

Stenciling: Designing Structurally-Sound Surfaces with Decorative Patterns

Christian Schumacher^{1,2} Bernhard Thomaszewski² Markus Gross^{1,2}

¹ETH Zurich
²Disney Research



Figure 1: Our design tool allows users to create physical surfaces structured with decorative patterns. Given a base shape and pattern exemplar (left), our method creates a distribution of stencils that are used to remove material from an otherwise solid surface. The result (center, right) is a physical surface that combines user-controlled aesthetics and structural integrity.

Abstract

We present a novel method to design shells with artistic cutouts in a manner that produces a stable final result. The process of stenciling, removing material with a fixed shape, is a particularly appealing way to introduce a decorative pattern into the design of architectural structures, furniture, or household objects. However, removing material can easily weaken an object to the point where its integrity is compromised, while purely functional distributions of cutouts lack the desired aesthetic component. We tackle this problem by combining aesthetics, stability, and material efficiency in an optimization that determines the distribution and scaling of these stencils in a way that complies as much as possible with both pattern and stability objectives. We demonstrate the capabilities of our system on examples from architecture, furniture design, and decorative items, and show how user interaction can be integrated to guide the aesthetics of the final result.

1. Introduction

From curved pavilions to monocoque furniture and household items—thin shell structures integrate form and function at all scales of everyday life. Apart from overall shape, one particularly appealing way to impart style onto a shell structure is by removing material in decorative patterns. However, uninformed or too aggressive removal of material can easily degrade the structural performance of the shell and lead to failure. Topology optimization methods, by contrast, have been very successful at creating organic structures with optimal stability in a broad range of engineering applications [DG14], including shell structures [HTG13]. But although structurally-optimal, the aesthetics of the resulting patterns cannot be controlled.

Rather than a trade-off, we seek a design tool that combines aesthetics, stability, and material efficiency in a synergistic way. To this

end, we introduce the concept of *stencils*—parametrized masks that cut out material from an underlying solid surface. Inspired by recent work on discrete element distribution [MWT11, RÖM*15], we propose an example-based method for describing stencil patterns with custom shapes and arrangements. We formulate pattern creation as an energy minimization problem, allowing us to simultaneously optimize with respect to both aesthetic and structural goals. Using this approach, designers can quickly create structurally-sound shells with a broad range of decorative patterns, ranging from regular and homogeneous distributions to irregular and heterogeneous patterns.

The technical core of our method is an optimization algorithm that automatically determines stencil parameters in order to comply as much as possible with both pattern and stability objectives. Evaluating the stability objective and its derivatives requires solv-

ing for the equilibrium state of the corresponding patterned shell, and computing its derivatives with respect to the stencil parameters. When using meshes that conform to the boundaries of the stencil, solving for static equilibrium can easily take several minutes. In order to afford interactive design iterations, we lay aside computationally-expensive approaches based on conforming discretizations and remeshing operations. Inspired by material interpolation schemes in topology optimization [BS99], we instead perform all computations on the original surface mesh and scale the elastic energy of its (membrane and bending) elements according to the extent to which they are covered by stencils. Though approximate, this approach allows us to use a comparatively coarse mesh with constant topology, which is essential for reducing the time spent on parameter optimization to an acceptable level during design iterations. Once a satisfying design has been found, the user can verify its validity through simulation on a high-resolution conforming mesh.

The formulation that we propose in this work is general with respect to both aesthetic and structural goals, and it provides manifold possibilities for the user to guide the aesthetic appeal of the design. We illustrate this flexibility on a set of digital examples and physical prototypes that are representative of applications in industrial and architectural design.

2. Related Work

Fabrication-Oriented Design Recent research in computational design has looked into different ways to leverage state-of-the-art fabrication methods, augmenting or replacing traditional geometry-based design tools with fabrication-oriented methods. Several researchers have proposed approaches to change the deformation behavior of objects [BBO*10, PZM*15, SBR*15], introduce functional mechanical properties into prints [PWLSH13, BWBSH14], or design different forms of physical surfaces, including inflatables [STBG12], plush toys [MI07], bead-work [IIM12], and wire-mesh models [GSFD*14].

Texture & Distribution Synthesis Computer graphics has a long history of creating decorative patterns on digital surfaces. Example-based texture synthesis methods use a small input exemplar to create a seamless output in a larger domain; see for instance Wei et al. [WLKT09]. In contrast to traditional pixel-based or voxel-based textures, discrete example-based textures [IMIM08, LGH13] offer more flexibility to modify a texture without destroying its appearance, and is well suited for applications in digital fabrication.

We base our approach on the work of Ma et al. [MWT11], who introduced a texture synthesis approach with a neighborhood-based energy formulation, where shapes are represented by a collection of points. As a central difference to this, our method operates on rigid bodies, not points, and directly applies to curved surfaces.

Roveri et al. [RÖM*15] recently introduced a similar texture synthesis method that uses a more efficient matching approach, and can also handle continuous structures.

Thin Shell Simulation The simulation of thin shells has been extensively used in computer graphics to animate cloth, but has also

started to find applications in computational design tools in recent years, where it is used to predict and optimize for different types of surfaces. Grinspun et al. [GHDS03] have introduced a widely used shell model with hinge-based bending elements, which has been further investigated and extended [GSH*04, GGWZ07]. We employ a modified version of this shell model in our simulation, which scales the elastic energy with the fill ratio of individual elements to approximate the deformation behavior of small-scale details on a coarse simulation mesh. Another approach to account for fine details has been proposed by Kaufmann et al. [KMB*09], who use an XFEM method with enrichment textures to augment the basis function of a regular FEM simulation and allows for detailed cuts.

Structural Optimization & Surface Design Finding structures with good mechanical properties has long been a central topic in mechanical engineering. Topology optimization [BS03] has been used to find structures that provide the least compliance, given a fixed amount of material. Wu et al. [WDW16] have recently shown that even high resolution topology optimization problems can be solved efficiently.

In the computer graphics community, several papers have proposed efficient methods to increase the stability of 3D printed objects [SVB*12, WWY*13, LSZ*14], using various approaches to generate internal structures for closed objects, and Zhou et al. [ZPZ13] presented a method to determine likely points of failure in fabricated models.

Statics aware grid shells by Pietroni et al. [PTP*15] are an example of physics-based surface designs with an aesthetic component, though the appearance is restricted to Voronoi patterns.

Two recent publications investigate the combination of aesthetics and stability in objects with decorative patterns. Dumas et al. [DLL*15] use a voxel-based representation to remove material from an object by projecting a binary texture onto a mesh. They improve the stability of the object by inserting new connections into the structure using a heuristic approach. In contrast, our approach directly computes the effect of stencil parameters on the stability, and also considers the aesthetic objective while improving the stability.

Our work is most closely related to Martinez et al. [MDLW15], who directly combine appearance and compliance objectives in a unified optimization. Their approach is based on traditional topology optimization and tries to find an aesthetically-pleasing structure within a given stability bound, using a fixed amount of material. Our approach also uses a combination of objectives, but works on three-dimensional shells, and uses an energy-based stability criterion without material constraints that allows us to keep the desired texture appearance in non-critical regions of the model.

3. Overview

We propose a computational design tool for physical surfaces with decorative cutouts, fulfilling both aesthetic and structural objectives. Starting from a mesh representing the input surface, the user first creates a decorative pattern that defines the aesthetic objective

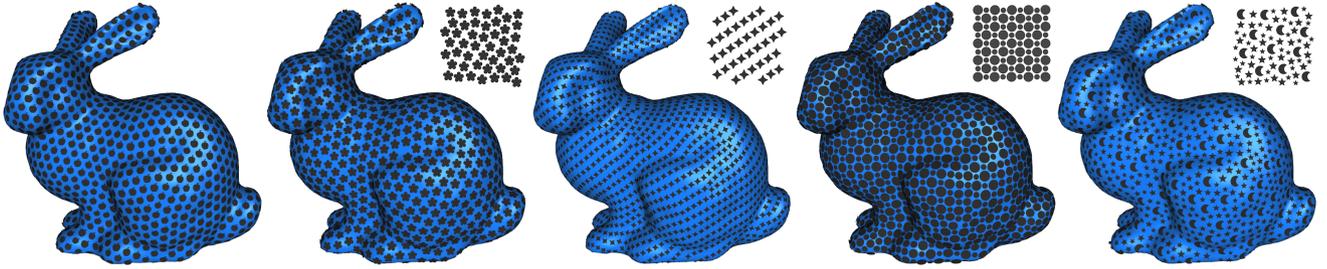


Figure 2: We show that our method can handle different types of textures. From left to right: (1) A distribution of apples using only our packing objective. (2) An exemplar with irregularly placed flowers creates a packing-like distribution, and offers the possibility to fill in missing stencils. (3)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

of the design (Sec. 4). This decorative pattern consists of a collection of stencils that are distributed across the surface according to a texture objective in the form of an energy function. We explore two options in this work: a packing scheme that places stencils on the surface according to a quasi-blue noise random distribution; and an example-based distribution scheme based on the principle of discrete element textures [MWT11]. We also provide a set of tools that allow the user to further control the distributions in terms of local scaling and orientation fields.

The user then defines a force distribution on the surface that is a conservative estimate of the expected loads. Furthermore, structural goals are defined to either enforce a stability criterion in terms of a fixed threshold on the maximum per-element energy density, or to minimize the compliance of the structure for a fixed amount of material. The aesthetic and structural goals provided by the user give rise to objectives that are jointly minimized in order to obtain optimal stencil parameters (Sec. 5).

4. Stencil Patterns

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

4.1. Stencil Representation & Projection

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

a vector from an underlying orientation field (computed with `libigl` [JP*16] using the approach described in [BZK09]) around the normal \mathbf{n}_i . The effective degrees of freedom \mathbf{p}_i of a stencil are thus defined by five parameters corresponding to its position, scale, and scalar rotation value. We concatenate the parameters of all stencils into a vector \mathbf{p} .

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

4.2. Pattern Objectives

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

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

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

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

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

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

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

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

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

$$O_{ij}^{\text{rep}} = \frac{w_{ij}}{d_{ij}^2}, \quad (3)$$

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

$$O_i^{\text{growth}} = -s_i. \quad (4)$$

The packing objective is then defined as

$$P^{\text{pack}} = \sum_{i,j} \left(O_{ij}^{\text{rep}} + O_{ij}^{\text{coll}} \right) + \sum_i O_i^{\text{growth}}. \quad (5)$$

Discrete Element Textures The packing objective can be used to create visually pleasing distributions as shown, e.g., in Fig. 2, (I). However, while the user can control the shape of the stencil, the distribution is homogeneous and its structure cannot be controlled. In order to provide more artistic control over the resulting distributions, we turn to an example-based approach inspired by *Discrete Element Textures* (DET) [MWT11]. DET was developed to synthesize and improve the distribution of points with arbitrary attributes. It works by matching points and their neighborhoods in an output distribution to neighborhoods in a user-provided exemplar, and then computes updates that aim to improve the match between output and input. In our case, we approximate each stencil as a single point, and encode the scale and type of the stencil as attributes. Letting $N(i)$ denote the neighborhood of stencil i , its corresponding energy is defined as

$$O_i^{\text{DET}} = \sum_{j \in N(i)} \left(\|\mathbf{q}_i - \hat{\mathbf{q}}_j\| - \mathbf{T}(\mathbf{q}'_i - \mathbf{q}'_j) \right)^2 + \|r_j - r'_j\|^2 + \xi(1 - \delta_{t_j t'_j}), \quad (6)$$

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

we augment the DET objective by a term that explicitly penalizes collisions.

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

4.3. Initialization

An adequate initial distribution of stencils is important to ensure that our optimization approach (Sec. 5) will converge to a desirable solution. We found that especially for the DET objective, a good initialization is crucial, and moreover, copying patches from the input exemplar to the output, as suggested in the original paper [MWT11], rarely led to good results in our examples.

We instead opted for an incremental initialization strategy similar to [IMM08]. We start by randomly selecting a point in the input exemplar, and then copy the neighborhood patch around this point onto a random location in the output domain. We then compute the DET matching from the output domain to the input exemplar, with the constraint that we do not match any patches in the input exemplar that are at the boundary of the domain. This is a one-way matching that does not include a penalty for stencils in the input neighborhood that are not matched to the output domain. However, we can use the neighborhood in the exemplar to estimate which stencils we could add to the output domain, by determining the stencils that are not matched. Adding unmatched stencils into the output domains, given that the overlap with existing stencils is not significantly larger than our threshold, allows us to expand the region covered by stencils. By iteratively applying this insertion scheme, we grow an initial stencil distribution that ultimately covers the whole model.

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

4.4. User Control

Apart from the shape of the stencils and their arrangement, we also provide additional tools that allow the user to control the aesthetics of the resulting pattern through sizing and orientation fields. In order to control the orientation of the stencils, we augment the computation of the orientation field by an objective that asks the resulting field to locally align with user-defined sketch directions (see Fig. 3, left).

Another way of stylizing the pattern is to prescribe an inhomogeneous scaling field across the surface; see Fig. 3, right. We provide a simple brush tool that allows the user to paint the desired stencil sizes directly onto the surface. For the DET objective, these sizing values are used to scale the exemplar when computing the matching energy (6), whereas for the packing objective, the scale of the stencil is simply set to the sizing value at the corresponding location.

5. Structural Optimization

Apart from aesthetic goals, a central requirement when designing functional surfaces is to satisfy structural constraints, i.e., con-

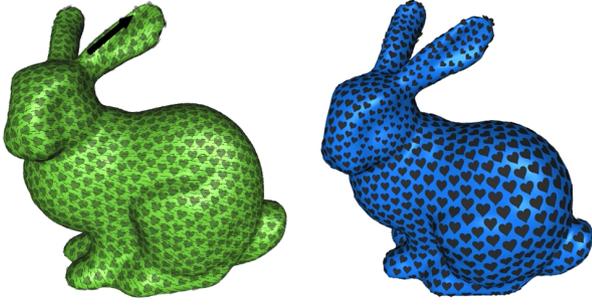


Figure 3: Users can control the orientation field through sketching (left) and locally adjust the size of the stencils using a brush tool (right).

ditions relating to the stability of the design. While the criteria that can be used to quantify stability are manifold, we focus on two common variants of structural optimization based on energy-density constraints and compliance minimization. For both of these approaches, we must compute *a*) the equilibrium state of the structure, *b*) the change in equilibrium state induced by a change in parameters, and *c*) the parameter values that lead to a desired equilibrium state. We explain these steps in the following.

5.1. Simulation

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

$$\Psi = \frac{\mu}{2} \text{tr}(\mathbf{E}^t \mathbf{E}) + \lambda \text{tr}(\mathbf{E})^2, \quad (7)$$

where μ and λ are (thickness-dependent) material parameters, $\mathbf{E} = \frac{1}{2}(\mathbf{F}^t \mathbf{F} - \mathbf{I})$ is the 2×2 Green strain tensor, and $\mathbf{F} \in \mathbb{R}^{3 \times 2}$ is the deformation gradient; refer to, e.g., Skouras et al. [STBG12] and the textbook by Bonet and Wood [BW97] for details. Using CST elements for discretization, the strain energy density is constant across the element and the membrane energy simply follows as $W_i^{\text{memb}} = \int_{\Omega_i} \Psi = \mathcal{A}_i \Psi$, where Ω_i is the parameter domain of element i and \mathcal{A}_i its undeformed area.

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

$$W_{ij}^{\text{bend}} = k_b (\theta(\mathbf{x}) - \bar{\theta})^2 \frac{\mathcal{A}_{ij}}{h_{ij}}, \quad (8)$$

where θ and $\bar{\theta}$ are the dihedral angles in the deformed and undeformed configuration, respectively, h_{ij} is a geometry factor, and $\mathcal{A}_{ij} = \frac{1}{3}(\mathcal{A}_i + \mathcal{A}_j)$ (see [GHDS03]). Note that k_b is a dependent coefficient that is computed from the material parameters λ and μ as well as the thickness of the element according to [GGWZ07].

5.2. Extension to Stenciled Shells

Perhaps the most obvious way to apply this computational model to the case of thin shells structured with decorative cutouts would be to use a triangle meshes that conforms to the boundaries of the stencils. However, this approach would require a high-resolution mesh that would drastically increase the time needed to compute equilibrium state—and it would be prohibitively expensive when used at the core of our optimization algorithm. We therefore turn away from conforming meshes and instead maintain a comparatively coarse mesh, each of whose triangles \mathcal{T}_i we endow with an additional fill-ratio variable α_i corresponding to the degree to which the element is overlapped by stencils. In order to compute the overlap between a given triangle and a stencil, we first project the triangle into the tangent space of the stencil. We then compute the intersection polygon using the Sutherland-Hodgman clipping algorithm [SH74]. The ratio between the area of the overlap and the area of the projected element is then subtracted from the triangle's fill ratio.

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

$$\hat{W}_{ij}^{\text{memb}}(\mathbf{x}, \alpha) = \int_{\hat{\Omega}_i} \Psi = \alpha_i \mathcal{A}_i W_i^{\text{memb}}(\mathbf{x}) \quad (9)$$

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

$$\hat{W}_{ij}^{\text{bend}}(\mathbf{x}) = \sqrt{\alpha_i \alpha_j} W_{ij}^{\text{bend}}(\mathbf{x}). \quad (10)$$

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

$$\hat{W}^{\text{el}} = \sum_i \hat{W}_i^{\text{memb}} + \sum_{(i,j) \in \mathcal{H}} \hat{W}_{ij}^{\text{bend}}, \quad (11)$$

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

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

$$\mathbf{f}_i(\mathbf{x}, \alpha) = \mathbf{f}_i^{\text{el}}(\mathbf{x}, \alpha) + \mathbf{f}_i^{\text{ext}} = \mathbf{0} \quad \forall i. \quad (12)$$

5.3. Optimization

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

to improve the structural objective, we consider its gradient with respect to the stencil parameters,

$$\frac{dS(\alpha, \mathbf{x}(\alpha))}{d\alpha} = \frac{\partial S(\alpha, \mathbf{x}(\alpha))}{\partial \alpha} + \frac{\partial S(\alpha, \mathbf{x}(\alpha))}{\partial \mathbf{x}} \frac{d\mathbf{x}}{d\alpha}. \quad (13)$$

It is clear from the above expression that computing the gradient requires the map between positions \mathbf{x} and material fill ratios α , which is given by Eq. (12): a combination of material fill ratios and deformed positions is admissible if and only if it corresponds to an equilibrium state, i.e., $\mathbf{f}(\mathbf{x}, \alpha) = \mathbf{0}$. Consequently, for an admissible change in material fill ratios it must hold

$$\frac{d\mathbf{f}}{d\alpha} = \frac{\partial \mathbf{f}}{\partial \alpha} + \frac{\partial \mathbf{f}}{\partial \mathbf{x}} \frac{d\mathbf{x}}{d\alpha} = \mathbf{0}, \quad (14)$$

and therefore

$$\frac{d\mathbf{x}}{d\alpha} = -\frac{\partial \mathbf{f}^{-1}}{\partial \mathbf{x}} \frac{\partial \mathbf{f}}{\partial \alpha}. \quad (15)$$

We can thus compute the derivative of \mathbf{x} wrt. α by solving a system of linear equations, whose matrix is given by the Hessian of the elastic energy of the surface. It is worth noting that Eqs (14-15) are interchangeably referred to as the *implicit function theorem* or *sensitivity analysis*. Finally, since we ultimately solve for stencil parameters, we compute $\frac{dS}{d\mathbf{p}}$ by applying the chain rule to (13) and use finite differences to numerically approximate the required derivatives $\frac{\partial \alpha}{\partial \mathbf{p}}$.

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

Energy-Density Objective A natural structural goal for a design is to ask that a given expected load should not lead to failure. Structural failure is typically indicated by a stability criterion that, depending on the type of material, depends on deformation, stress, or energy density. Without loss of generality, we opt for a criterion based on per-element energy density

$$\mathcal{W}_i = \frac{1}{\hat{\mathcal{A}}_i} \hat{\mathcal{W}}_i^{\text{memb}} + \sum_j \frac{\hat{\mathcal{A}}_i}{3\hat{\mathcal{A}}_{ij}^2} \hat{\mathcal{W}}_{ij}^{\text{bend}}, \quad (16)$$

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

$$S^{\text{stab}}(\alpha, \mathbf{x}(\alpha)) = \frac{1}{m} \sum_i \begin{cases} \frac{1}{2}(\mathcal{W}_i(\alpha, \mathbf{x}(\alpha)) - \beta)^2, & \text{if } \mathcal{W}_i > \beta \\ 0, & \text{else,} \end{cases} \quad (17)$$

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

In many cases, the materials used in a design will dictate specific threshold values that must not be exceeded. However, another way of optimizing the stability of a design is to minimize the maximum per-element energy density. We implement this strategy as an iterative scheme, each of whose steps update the current value of β to the average computed from the top $y\%$ of the elements with the highest energy density, where y can be chosen depending on the use case. See Sec. 6 for an example.

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

$$S^{\text{comp}}(\alpha, \mathbf{x}(\alpha)) = \sum_{i=0}^m \mathcal{A}_i \mathcal{W}_i(\alpha, \mathbf{x}(\alpha)), \quad (18)$$

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

5.4. Numerical Solution

With both the stability and pattern objectives in place, we compute optimal stencil parameters through minimization. Thanks to the equilibrium conditions (12), the positions \mathbf{x} are implicit functions of the material fill ratios α , which in turn are explicit functions of the stencil parameters \mathbf{p} . We thus define the joint objective as a function of the stencil parameters,

$$J(\mathbf{p}) = w_S S(\mathbf{x}(\mathbf{p}), \alpha(\mathbf{p})) + w_P P(\mathbf{p}), \quad (19)$$

where w_S and w_P are scaling parameters.

In order to minimize J , we use L-BFGS-B [BLNZ95], a Quasi-Newton method that combines analytically-computed gradients with a limited-memory BFGS approximation of the Hessian and bound constraints. We additionally employ line search in order to ensure monotonic decrease in the objective. This approach allows us to optimize for all continuous stencil parameters at the same time. For the stencil type, which is the only discrete stencil parameter, we use the update scheme proposed in [MWT11] that selects the stencil type based on a majority vote among the stencil's neighbors. Note that in order to use line search correctly, we need to ensure that our simulation is always in equilibrium, i.e., we need to update the simulation for any change in stencil parameters.

6. Results

6.1. Validation

Simulation with Scaled Energy Densities Our elastic simulation using scaled energy densities is inspired by simulation approaches from topology optimization [BS99], where the scaling of the energy density of each element is a function of its fill ratio. We check the validity of our simulation, which extends this approach to hinge elements that include two fill ratios, by ascertaining that for increasing mesh resolutions, the simulation converges to a regular simulation using a conforming mesh. Figure 4 shows that this convergence can indeed be observed, and moreover, that even for coarse resolutions,

our approach provides an adequate approximation of a simulation on a conforming mesh.

Physical Validation We validate that our optimization improves the stability of an object by performing a load test on a stenciled plate, shown in Figure 5. For the design and experiment, we fix one end of the plates, and apply a load on the center of the other end. The plate with the pattern-optimized stencil distribution shows the largest deflection, and the energy- and compliance-optimized stencil distributions show improved performance. The compliance-optimized distribution uses a fixed material constraint in this experiment, and the limited amount of material redistribution possible without any stencils overlapping explains the slightly better performance of the energy-optimized distribution.

Coarse Simulation vs. High-resolution Simulation For all our results that we fabricated, we compare our coarse simulation to a high-resolution conforming mesh. Figures 6, 7, 9, and 8 show that we capture the global energy density distribution well, while small local stress peaks along the boundaries of stencils might exceed the maximum energy density computed by our simulation.

6.2. Physical Results

We designed a number of results for different applications, which were 3D printed using selective laser sintering with a polyamide material (PA12). We show both the *pattern-optimized* results, which only use the pattern objective for the optimization, as well as the *stability-optimized* results that use both pattern and stability objectives. For all stability-optimized examples, we set the energy density threshold β to a fraction of the maximum energy density \mathcal{W}^{max} of the non-stenciled mesh to ensure that we create a meaningful scenario for our optimization. We show the statistics and performance of our examples in Table 1, and parameters in Table 2.

Chair and Table The chair example in Figure 6 and the table example in Figure 8 we both designed by fixing the bottom vertices of the model, and then applying a load on the seat and back rest for the chair, and on the table top for the table.

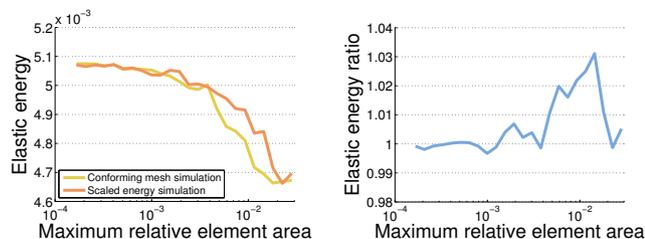


Figure 4: The convergence of elastic energy for simulations with our scaled energy density approach and simulations with a conforming mesh, for different mesh resolutions, specified by the maximum area of a mesh element. The simulation is performed on a thin, square, horizontally oriented, unit-sized sheet with a single circular stencil, with fixed outer boundaries and deformed by self-weight. Left: The elastic energy of both simulations. Right: The ratio between the elastic energies.

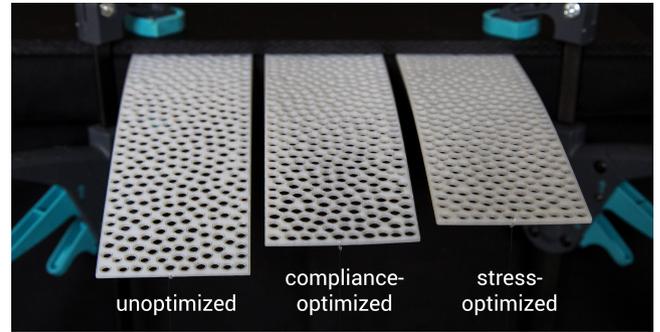


Figure 5: Comparison of the deflection of plates with circular stencils, before and after running the stability optimization. The plates were fixed on one side, and a weight was attached to the opposite side.

Fruit Bowl We also tested our method on alternative boundary conditions. The fruit bowl in Figure 9 has been designed to hang by a wire attached to its top, and the load case is a force pulling on the bottom of the bowl.

Pavilion Figure 1 and Figure 7 show the result of applying our method to an object under self-weight. In this case, the stencil distribution not only locally influences the stiffness, but also changes the load, and can have a global effect on the stability. Since the optimization considers the interplay between stencil parameters and the physical simulation, it automatically removes material from the top part of the pavilion in order to lighten the load on the base of the structure.

For this case of an object under self-weight, our approach of setting a fixed threshold for the energy density might be too restrictive. Rather, we would like to perform a minimization instead of a thresholding, with the assumption that any reduction in energy density will be beneficial. As an exercise, we performed an adaptive optimization: At the beginning of an optimization, we set the threshold to the median energy density. Whenever we reach an optimal solution, we check whether the median energy density changed. If so, we continue the optimization with the new threshold. Figure 7 shows the result of this adaptive optimization, which shows more extreme changes to the stencils than the regular optimization.

6.3. Interactive Design Session

The accompanying video contains an interactive design session, where a user adjusts the orientation and scaling fields during the design phase. This design session also shows that we can use our initialization approach (Sec. 4.3) to fill in holes in the pattern, which improved the appearance for this specific example. We run the initialization algorithm on the existing distribution whenever the optimization converges, and then restart the optimization if the number of stencils changed.

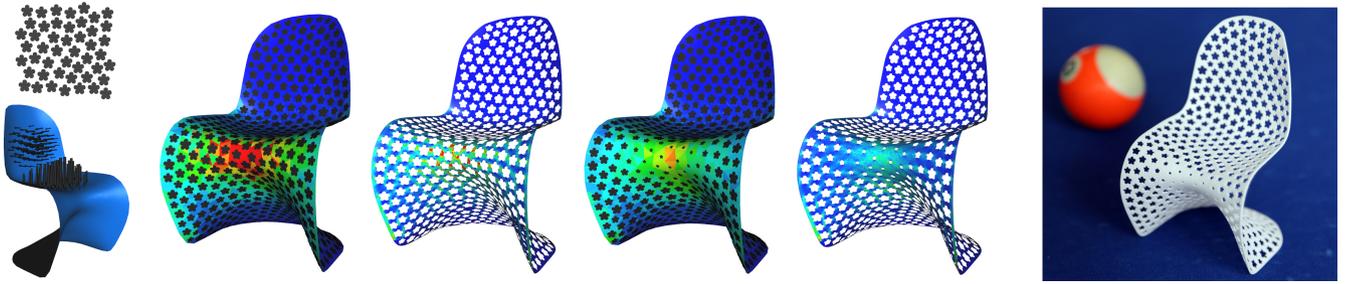


Figure 6: Comparison of energy densities between our coarse simulation mesh (second and fourth picture) and a high-resolution conforming mesh (third and fifth picture) for the texture-optimized chair (left pair) and the stability-optimized chair (right pair). Energy densities around the threshold β are shown in green, with higher densities in red and lower densities in blue. The pattern exemplar and load case are given on the left. The stability-optimized chair has been 3d printed (right).

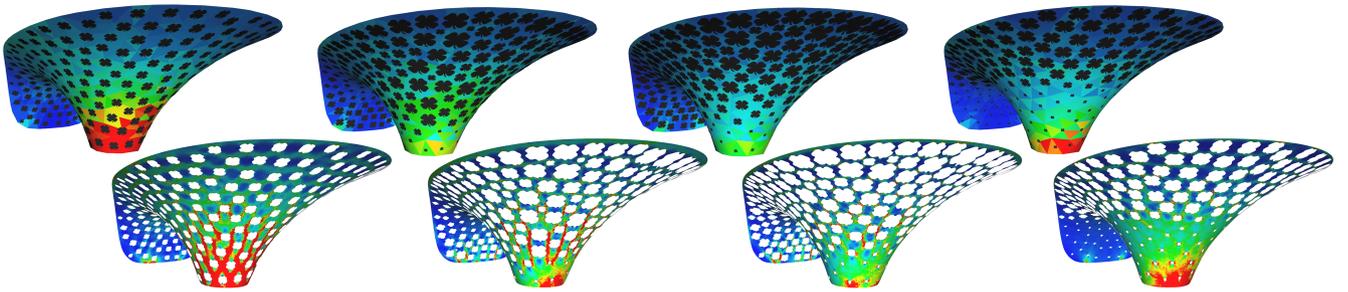


Figure 7: Comparison of energy densities for the pavilion model. The top row shows our coarse simulation mesh, the bottom row a high-resolution simulation on a conforming mesh. We show four different scenarios: The pattern-optimized pavilion (left), the energy density-optimized pavilion with a single energy density threshold (center left) and an adaptive energy density threshold (center right), and the compliance-optimized pavilion with a fixed material constraint (right). Energy densities around the threshold β are shown in green, with higher densities in red and lower densities in blue. The pattern exemplar and final fabricated model are shown in Figure 1.

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

Table 1: Performance statistics. The table shows the number of stencils, the number of vertices for the underlying mesh and the simulation mesh, and the runtime t for all our fabricated examples. The runtime is measured from initial distribution to convergence (for the pattern-only objective) or from a pattern-optimized solution to a stability-optimized solution (for all other objectives).

7. Conclusion

We introduced a novel method to design stenciled surfaces that combines aesthetics and stability. We have shown that we can effi-

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

Table 2: Parameters used for our examples: Scaling parameters w_S and w_P , energy density threshold β as a function of the maximum energy density \mathcal{W}^{max} , collision safety factor s , and the neighborhood radius r_N relative to the dimensions of the input exemplar.

ciently predict regions of failure and automatically optimize a stencil distribution to create stable objects with high visual fidelity.

7.1. Limitations and Future Work

While we can predict regions of failure, we cannot guarantee that our best solution will be able to avoid these, as either the load case or the texture objective could prevent a feasible solution. In such cases, we can provide feedback to the user that would ideally lead to adjustments in the load case or texture objective parameters, though including more sophisticated ways to decouple the appearance and stability objectives is an interesting direction for future work.

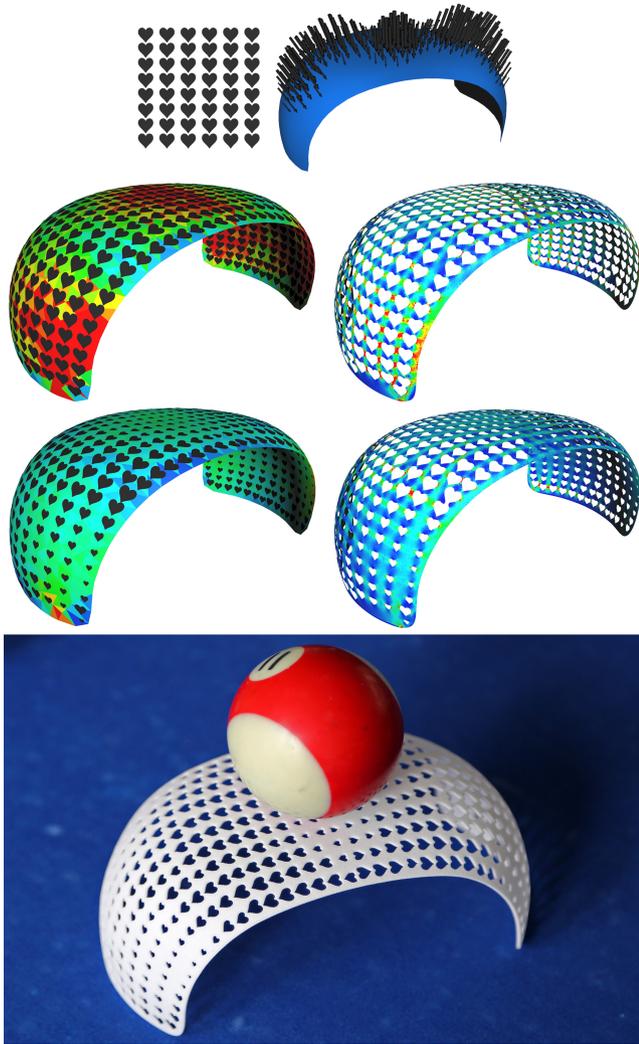


Figure 8: Comparison of energy densities between our coarse simulation mesh (left) and a high-resolution conforming mesh (right) for the pattern-optimized table (top) and the stability-optimized chair (middle). Energy densities around the threshold β are shown in green, with higher densities in red and lower densities in blue. The pattern exemplar and load case are given on the top. The stability-optimized table holds a billiard ball (bottom).

Our fill ratio-based simulation approach allows us to efficiently approximate the deformation behavior of a stenciled model. While we have shown that with increasing resolution, the simulation result converges to the solution obtained from a conforming mesh, for practical resolutions, our approach might not be able to capture any highly anisotropic deformation behavior introduced by stencils with very thin features. XFEM [KMB*09] or a homogenization method [KMOD09] could be used to achieve a better approximation in these cases, at the cost of computational efficiency.

All computations of positions and distances of stencils are performed in Euclidean space, and while this has not led to any prob-

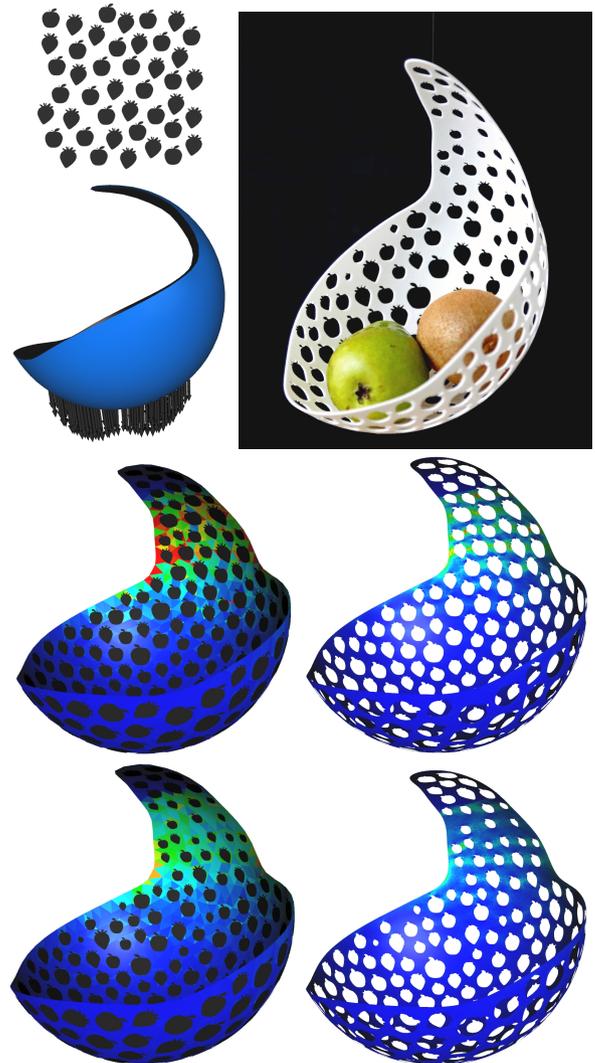


Figure 9: The pattern exemplar and load case of the bowl example (top left), and the final fabricated result (top right). In the bottom part, we compare the energy densities of our coarse simulation mesh (left) with a high-resolution simulation on a conforming mesh (right) for the pattern-optimized (middle) and stability-optimized bowl (bottom). Energy densities around the threshold β are shown in green, with higher densities in red and lower densities in blue.

lems in the examples we showed, our projection approach cannot guarantee injectivity, which might lead to unsatisfactory results in the case of folded meshes. Incorporating geodesic distances could prevent artifacts in such situations, though a sufficiently fast and accurate computation of geodesic distances is necessary to preserve our current level of efficiency.

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

to modify the stencils would increase the expressiveness of our method.

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

Acknowledgements

We would like to thank Loïc Ciccone for his work on this project's predecessor, and the code for cutting and projecting stencils. Many thanks also to Stelian Coros for his invaluable input. Finally, we would also like to thank the anonymous reviewers for their helpful comments.

References

- [BBO*10] BICKEL B., BÄCHER M., OTADUY M. A., LEE H. R., PFISTER H., GROSS M., MATUSIK W.: Design and fabrication of materials with desired deformation behavior. In *Proc. of ACM SIGGRAPH '10* (2010). 2
- [BLNZ95] BYRD R. H., LU P., NOCEDAL J., ZHU C.: A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing* 16, 5 (1995). 6
- [BS99] BENDSØE M. P., SIGMUND O.: Material interpolation schemes in topology optimization. *Arch. Applied Mech.* 69 (1999), 635–654. 2, 6
- [BS03] BENDSØE M. P., SIGMUND O.: *Topology optimization: Theory, methods and applications*. Springer, Berlin, 2003. 2
- [BW97] BONET J., WOOD R. D.: *Nonlinear Continuum Mechanics for Finite Element Analysis*. Cambridge Univ. Press, 1997. 5
- [BWBSH14] BÄCHER M., WHITING E., BICKEL B., SORKINE-HORNUNG O.: Spin-it: Optimizing moment of inertia for spinnable objects. In *Proc. of ACM SIGGRAPH '14* (2014). 2
- [BZK09] BOMMES D., ZIMMER H., KOBBELT L.: Mixed-integer quadrangulation. In *Proc. of ACM SIGGRAPH '09* (2009). 3
- [DG14] DEATON J. D., GRANDHI R. V.: A survey of structural and multidisciplinary continuum topology optimization: Post 2000. *Struct. Multidiscip. Optim.* 49, 1 (2014). 1
- [DLL*15] DUMAS J., LU A., LEFEBVRE S., WU J., DICK C.: By-example synthesis of structurally sound patterns. In *Proc. of ACM SIGGRAPH '15* (2015). 2
- [GGWZ07] GARG A., GRINSPUN E., WARDETZKY M., ZORIN D.: Cubic shells. In *Proc. of Symp. on Computer Animation '07* (2007), pp. 91–98. 2, 5
- [GHDS03] GRINSPUN E., HIRANI A., DESBRUN M., SCHRÖDER P.: Discrete shells. In *Proc. of Symp. on Computer Animation '03* (2003). 2, 5
- [GSFD*14] GARG A., SAGEMAN-FURNAS A. O., DENG B., YUE Y., GRINSPUN E., PAULY M., WARDETZKY M.: Wire mesh design. In *Proc. of ACM SIGGRAPH '14* (2014). 2
- [GSH*04] GINGOLD Y., SECORD A., HAN J. Y., GRINSPUN E., ZORIN D.: A discrete model for inelastic deformation of thin shells. In *Proc. of Symp. on Computer Animation '04* (2004). 2
- [HTG13] HASSANI B., TAVAKKOLI S. M., GHASEMNEJAD H.: Simultaneous shape and topology optimization of shell structures. *Struct. Multidiscip. Optim.* 48, 1 (2013). 1
- [IIM12] IGARASHI Y., IGARASHI T., MITANI J.: Beady: Interactive beadwork design and construction. In *Proc. of ACM SIGGRAPH '12* (2012). 2
- [IMIM08] IJIRI T., MĚCH R., IGARASHI T., MILLER G.: An example-based procedural system for element arrangement. In *Proc. of Eurographics '08* (2008). 2, 4
- [JP*16] JACOBSON A., PANOZZO D., ET AL.: libigl: A simple C++ geometry processing library, 2016. <http://libigl.github.io/libigl/>. 3
- [KMB*09] KAUFMANN P., MARTIN S., BOTSCH M., GRINSPUN E., GROSS M.: Enrichment textures for detailed cutting of shells. In *Proc. of ACM SIGGRAPH '09* (2009). 2, 9
- [KMOD09] KHAREVYCH L., MULLEN P., OWHADI H., DESBRUN M.: Numerical coarsening of inhomogeneous elastic materials. In *Proc. of ACM SIGGRAPH '09* (2009). 9
- [LGH13] LANDES P.-E., GALERNE B., HURTUT T.: A shape-aware model for discrete texture synthesis. In *Proc. of Symp. on Rendering '13* (2013). 2
- [LSZ*14] LU L., SHARF A., ZHAO H., WEI Y., FAN Q., CHEN X., SAVOYE Y., TU C., COHEN-OR D., CHEN B.: Build-to-last: Strength to weight 3d printed objects. In *Proc. of ACM SIGGRAPH '14* (2014). 2
- [MDLW15] MARTÍNEZ J., DUMAS J., LEFEBVRE S., WEI L.-Y.: Structure and appearance optimization for controllable shape design. In *Proc. of ACM SIGGRAPH Asia '15* (2015). 2
- [MI07] MORI Y., IGARASHI T.: Plushie: An interactive design system for plush toys. In *Proc. of ACM SIGGRAPH '07* (2007). 2
- [MWT11] MA C., WEI L.-Y., TONG X.: Discrete element textures. In *Proc. of ACM SIGGRAPH '11* (2011). 1, 2, 3, 4, 6
- [PTP*15] PIETRONI N., TONELLI D., PUPPO E., FROLI M., SCOPIGNO R., CIGNONI P.: Statics aware grid shells. In *Proc. of Eurographics '15* (2015). 2
- [PWSH13] PRÉVOST R., WHITING E., LEFEBVRE S., SORKINE-HORNUNG O.: Make it stand: Balancing shapes for 3d fabrication. In *Proc. of ACM SIGGRAPH '13* (2013). 2
- [PZM*15] PANETTA J., ZHOU Q., MALOMO L., PIETRONI N., CIGNONI P., ZORIN D.: Elastic textures for additive fabrication. In *Proc. of ACM SIGGRAPH '15* (2015). 2
- [RÖM*15] ROVERI R., ÖZTIRELI A. C., MARTIN S., SOLENTHALER B., GROSS M.: Example based repetitive structure synthesis. In *Proc. of Symp. on Geometry Processing '15* (2015). 1, 2
- [SBR*15] SCHUMACHER C., BICKEL B., RYS J., MARSCHNER S., DARAIO C., GROSS M.: Microstructures to control elasticity in 3d printing. In *Proc. of ACM SIGGRAPH '15* (2015). 2
- [SH74] SUTHERLAND I., HODGMAN G. W.: Reentrant polygon clipping. *Communications of the ACM* 17 (1974), 32–42. 5
- [STBG12] SKOURAS M., THOMASZEWSKI B., BICKEL B., GROSS M.: Computational design of rubber balloons. In *Proc. of Eurographics '12* (2012). 2, 5
- [SVB*12] STAVA O., VANEK J., BENES B., CARR N., MĚCH R.: Stress relief: improving structural strength of 3d printable objects. In *Proc. of ACM SIGGRAPH '12* (2012). 2
- [WDW16] WU J., DICK C., WESTERMANN R.: A system for high-resolution topology optimization. *IEEE Transactions on Visualization and Computer Graphics* 22, 3 (2016), 1195–1208. 2
- [WLKT09] WEI L.-Y., LEFEBVRE S., KWATRA V., TURK G.: State of the art in example-based texture synthesis. In *Eurographics '09 State of the Art Reports (STARs)* (2009). 2
- [WWY*13] WANG W., WANG T. Y., YANG Z., LIU L., TONG X., TONG W., DENG J., CHEN F., LIU X.: Cost-effective printing of 3d objects with skin-frame structures. In *Proc. of ACM SIGGRAPH Asia '13* (2013). 2
- [ZPZ13] ZHOU Q., PANETTA J., ZORIN D.: Worst-case structural analysis. In *Proc. of ACM SIGGRAPH '13* (2013). 2