

Worst-Case Stress Relief for Microstructures

JULIAN PANETTA, ABTIN RAHIMIAN, and DENIS ZORIN, New York University

Additive fabrication technologies are limited by the types of material they can print: while the technologies are continuously improving, still only a relatively small discrete set of materials can be used in each printed object. At the same time, the low cost of introducing geometric complexity suggests the alternative of controlling the elastic material properties by producing *microstructures*, which can achieve behaviors significantly differing from the solid printing material. While promising results have been obtained in this direction, fragility is a significant problem blocking practical applications, especially for achieving soft material properties: due to stress concentrations at thin joints, deformations and repeated loadings are likely to cause fracture.

We present a set of methods minimizing stress concentrations in microstructures by evolving their shapes. First, we demonstrate that the worst-case stress analysis problem (maximizing a stress measure over all possible unit loads) has an exact solution for periodic microstructures. We develop a new, accurate discretization of the shape derivative for stress objectives and introduce a low-dimensional parametric shape model for microstructures. This model supports robust minimization of maximal stress (approximated by an L_p norm with high p) and an efficient implementation of printability constraints. In addition to significantly reducing stresses (by a typical factor of 5 \times), the new method substantially expands the range of effective material properties covered by the collection of structures.

Additional Key Words and Phrases: additive fabrication, microstructures, deformable objects, homogenization, shape optimization, goal-based material design, stress minimization

ACM Reference format:

Julian Panetta, Abtin Rahimian, and Denis Zorin. 2017. Worst-Case Stress Relief for Microstructures. *ACM Trans. Graph.* 36, 4, Article 1 (July 2017), 16 pages.

<https://doi.org/http://dx.doi.org/10.1145/3072959.3073649>

1 INTRODUCTION

Most additive fabrication technologies share the distinctive feature that printing cost is primarily determined by the amount of material used and is generally independent of the object's complexity. These characteristics make it practical to fabricate small-scale structures, which can be used to fine-tune the object's deformation behavior and achieve more effective designs for classical problems such as maximizing a design's strength under particular loads while constraining its weight. More generally, spatially varying effective material properties, including exotic properties such as negative Poisson's ratio, can be achieved on even a single-material printer using fine-scale structures (microstructures).

This work is supported by the National Science Foundation, under grant DMREF-1436591.

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

© 2017 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

0730-0301/2017/7-ART1 \$15.00

<https://doi.org/http://dx.doi.org/10.1145/3072959.3073649>




Fig. 1. We design microstructures emulating a desired material while simultaneously experiencing minimized peak stresses under the worst-case load. Here we show two structures with identical macroscopic properties. Our optimization smoothes out the stress concentrations seen in red on the left, creating a robust, organic microstructure. The meshes are shaded with our novel worst-case micro-stress measure.

Geometric complexity is typically associated with high curvature variation and negatively-curved regions. These surface features are well-known to generate high stress concentrations, which may preclude practical applications: even moderate loads may result in cracks or plastic deformations. While the task of distributing stress evenly is essential to many mechanical design problems, stress distribution acquires particular importance for complex geometry, and it is more difficult to address on an ad hoc basis. Optimal designs, minimizing the maximal stress norm subject to a set of constraints, tend to be smooth, “organic” free-form structures, as small surface variations often result in significant changes in local stress.

Shape modeling tools have addressed the problem of creating low-stress designs in a variety of ways, primarily by providing engineers with techniques to add fillets and blends. These approaches are based on the intuition that eliminating sharp concave corners improves stress behavior, and engineers tend to make specific choices of geometry based on prior experience and trial and error.

Shape and topology optimization provide a principled approach to solving this problem and have been used with success to optimize various functionals, most commonly compliance. However, success has been limited so far in minimizing max stress accurately and efficiently. Moreover, the typical setting for stress minimization is to specify a load the structure is required to support. In many cases, especially the case of microstructure design we consider, the loads are not known in advance, so *worst-case analysis* is needed in the optimization loop, i.e. determining at each iteration the loads causing the highest maximal stress.

In this paper, we focus on optimizing *microstructures*: assemblies of small cells, each filled with a pattern that, when periodically tiled, produces a particular averaged (homogenized) elastic behavior. Many of the methods we propose are relevant to other stress optimization problems.

Contributions. In this paper, we describe a set of techniques for producing optimized microstructures that cover a broad range of isotropic material properties while at the same time experiencing minimized stress concentrations under the worst-case scenario. We also ensure these structures satisfy printability constraints for additive fabrication. Specifically, our contributions include:

- A novel exact microstructure worst-case stress analysis that expresses worst-case loads as eigenvectors of small tensors, incurring negligible computational cost beyond homogenization.
- An accurate and efficient approach to optimize a nearly nonsmooth large- p L_p norm stress objective approximating maximal stress. Combined with robust, adaptive meshing, this yields a method capable of reliably designing reduced-stress microstructures under the thousands of different target material property constraints from our microstructure database, using only tens of iterations per design.
- A parametric model significantly expanding material property coverage over previous work, while using a single topology and enabling substantial stress reduction (5x for most properties).
- A new blending operator for implicit surface modeling that eliminates bulging outside the convex hull. This is essential to efficiently enforce printability for our parametric model.

We demonstrate experimentally that our pipeline yields structures with significantly reduced stresses and, consequently, greater resilience to repeated loadings and stronger loads.

2 RELATED WORK

Stress minimization. Several works have focused on minimum stress design using shape and topology optimization. The closest work to ours is [Allaire et al. 2004], which applies topology optimization to design lightweight minimal-stress objects built from sequentially laminated composites. It considers the stress concentrations occurring in the microstructure, but only for a fixed loading scenario (not the worst-case), and only for sequential laminates, microstructures whose properties have nice closed-form expressions but are not manufacturable. Another similar work is [All 2008], which applies the level-set topology optimization method to minimize the p -norm of stress. [Xia et al. 2012] and [?] propose similar machinery but, to address numerical and convergence problems of p -norm minimization for high p , use different objective functions: the former adds a penalty term that activates for stresses above a specified threshold, and the latter switches to penalizing the deviation of stress from the average stress level. In the engineering literature, several works have considered fillet design with optimal shape, e.g., [Sonmez 2009; Van Miegroet and Duysinx 2007], minimizing the stress concentration factor.

Our work differs from traditional shape and topology optimization approaches in several important ways: (i) we use exact worst-case analysis to determine the loads; (ii) we optimize a parametric model for which formulating fabrication constraints is efficient and exact; and (iii) our formulation for the shape derivative in combination with adaptive meshing, absent in level-set and SIMP-based

formulations, yields much better accuracy when evaluating and differentiating stress.

A concurrent work, [?] uses a fully discrete approach to find 2D shapes with optimal L_p -norm of von Mises stress under known loads, treating all vertices as free variables (we use a low-dimensional parametric model instead). This method requires thousands of iterations in 2D compared to the 50 iterations typical in 3D for our method. It also cannot enforce our fabrication and tileability constraints.

In the computer graphics community, [Stava et al. 2012] presented several heuristic model correction techniques, such as thickening and strut insertion, to improve a structure’s resilience to specific loading scenarios. [Zhao et al. 2016] optimize the shell thickness to bound the von Mises stress on all vertices under a given load. This can be viewed as a similar type of optimization to ours, with the shell thickness serving as the shape parameter with a simple, direct relation to the effective material properties. Similarly, [Zhou et al. 2016] bound the von Mises stress under specified force, while penalizing deviation from the input structure, operating on coarse meshes with linear elements and without remeshing.

Worst-case stress analysis. In graphics, two papers have presented heuristics to analyze stress under unknown load. [Zhou et al. 2013] approximately determine the most efficient pressure distribution to break or severely deform a structure by solving a large eigenvalue problem and many linear programs. [Langlois et al. 2016] use a PCA-based Monte Carlo sampling to construct a stress probability distribution from a large number of force samples generated by a rigid body simulator. They use the failure probability as a constraint in a costly gradient-based topology optimization to reduce weight, but forbid the object’s surface from changing to avoid differentiating the full stress analysis pipeline. In contrast, we derive an *exact*, efficient way to determine worst-case loads for microstructures and compute its exact derivative.

Periodic homogenization. Homogenization is a central tool to our work, and our formulation is based on [Allaire 2002]. Homogenization has been used in graphics to reduce complexity of physical models in [Kharevych et al. 2009], which finds the constitutive parameters of a low resolution discretization that best approximates the behavior of the original, more complex object.

Microstructure design and optimization. There is a large body of literature dedicated to theoretical studies of composites’ effective moduli (our periodic structures are a limiting case of composites, combining a single material with free space), which was reviewed in [Panetta et al. 2015]. Recent monographs on this topic include [Allaire 2002; Cherkaev 2000; Cioranescu and Donato 1999; Milton 2002; Torquato 2002], which primarily focus on identifying microstructures with *extremal* behavior (having effective elasticity properties at the boundary of domain achievable by a given class of composites).

Microstructures have been designed using various forms of topology optimization [Bendsøe 1989; Bendsøe and Sigmund 2003; Nakasone and Silva 2010], seeking periodic structures minimizing, e.g., compliance for a fixed total volume fraction. The result is normally a single-scale structure, with scale controlled by the design grid resolution or other types of regularization.

[Schumacher et al. 2015] and [Panetta et al. 2015] propose methods to design tileable microstructures with printability constraints

to fabricate deformable objects with spatially varying elastic properties. Our method follows the general shape optimization approach of [Panetta et al. 2015], but introduces a different functional (making the elasticity tensor fitting a constraint), a different approach for computing shape derivatives, and a new, enriched parametric model.

In contrast to these works, [Martínez et al. 2016] construct aperiodic, printable structures with pointwise control over Young’s moduli, but cannot independently control the Poisson’s ratio.

Fabrication constraints. [Gaynor and Guest 2016; Langelaar 2016; Qian 2016] have made recent progress in incorporating undercut/overhang angle constraints in the topology optimization framework. However, these methods enforce the constraints only approximately, requiring parameter tuning, and add nonlinearities to the problem that hinder convergence ([Gaynor and Guest 2016]). We impose printability constraints natively on our model’s parameters similarly to [Panetta et al. 2015]; this is made possible in our smoothness-enriched parametric model by our novel blending algorithm that avoids overhangs if none are present in the pre-blended structure.

Fabrication. [Hollister 2005; Kang 2010; Lin et al. 2004a,b] have demonstrated the fabrication of optimized microstructures in the context of bone scaffold and fusion cage design, and [Andreassen et al. 2014; Bückmann et al. 2012; Greaves et al. 2011; Schwerdtfeger et al. 2011] have demonstrated the possibility of manufacturing auxetic materials. The idea of manufacturing objects with spatially varying properties using tileable structures also appears in [Hiller and Lipson 2009].

[Bickel et al. 2010] designs and fabricates objects satisfying an input deformation by optimizing for the best combination of stacked layers of their multi-material 3D printer’s base materials. [Skouras et al. 2013] applies discrete material optimization to achieve desired deformations of complex characters with actuation, fabricating the results with multi-material printing.

Blending. Both parametric and implicit blending have received extensive attention in geometry modeling. A survey of parametric blending methods can be found in [Vida et al. 1994], and [Lin et al. 2014] discusses more recent work. A broad range of techniques have been proposed for implicit blending, starting with the foundational works [Sabin 1968] and [Blinn 1982]. The problems of bulging and (lack of) local control—characteristic of the simplest implicit blending constructions—are key topics of research. Bulging is an important problem in our context, as it leads to violation of printability constraints.

Clean union operations and related work can be used to generate smooth surfaces as offsets to the union of input surfaces [Pasko and Adzhiev 2002], and can be made local by introducing blending volumes [Barthe et al. 2004; Bernhardt et al. 2010; Pasko et al. 2005]. However, these methods cannot easily control radius independently of blending for surfaces built from a 1-dimensional skeleton, and are restricted to binary operators. Convolution surfaces [Bloomenthal 1997] are generally bulge-free; however, the degree of blending at vertices is difficult to control independently at different joints, including in a recent advanced version discussed in [Zanni et al. 2013]. The approach of [Rockwood 1989] reduces the bulging problem but may not produce smooth surfaces, and is difficult to localize. More




Fig. 2. The stress a microstructure experiences depends strongly on the structure’s corner smoothness.

recent work modulating the blending based on the angles between primitive normals [Gourmel et al. 2013] produces high-quality results but does not fully control bulging in the sense our application requires (we consider gradient-based methods in more detail in Section 6). A gradient-based approach was extended to N-ary operations in [Zanni et al. 2015], but similar considerations apply. In our blending construction, we use the Kreisselmeier-Steinhauser function [Kreisselmeier and Steinhauser 1983], which is commonly used in optimization and has been applied to implicit surface blending in production [Quilez Quilez].

3 OVERVIEW

Our aim is to produce printable microstructures achieving a prescribed homogenized elasticity tensor (effective material properties) C^* with minimal pointwise worst-case stress; the general optimization problem can be written as

$$\operatorname{argmin}_{\substack{\omega \text{ printable} \\ \tilde{C}(\omega) = C^*}} J(s(\omega)), \quad (1)$$

where J is an integral stress measure, e.g., the L_p norm of a pointwise stress measure, ω is the microstructure shape, and $s(\omega)$ is the pointwise worst-case stress distribution. The pointwise worst-case stress distribution is a function of the shape only, obtained separately for each point by maximizing a stress norm over all possible unit macroscopic (averaged) loads applied to the structure.

To solve this optimization problem we need (i) a parametric shape description that can be meshed efficiently and differentiated with respect to its parameters; (ii) a way to impose printability and elastic tensor constraints; and (iii) an efficient and accurate method to compute shape derivatives of our stress objective.

Our shape representation is based on a skeleton graph, consisting of vertices connected by edges. We construct a smooth implicit surface from this skeleton by inflating each edge and joining them together smoothly with a custom blending operation (Section 6). The final inflated geometry is controlled by radius, position, and smoothing parameters assigned to each vertex. These shape parameters are the variables in the optimization.

To reduce our search space and improve performance, we focus on patterns with reflectional symmetry. These shapes are guaranteed to produce axis-aligned orthotropic material properties and can be designed by meshing and analyzing only the period cell’s positive octant (upper-right quadrant in 2D) with appropriate modifications to the homogenization boundary conditions. In the following, we refer to this portion of the period cell as the structure’s *symmetry*

cell. We enforce reflectional symmetries by assigning independent parameters only to vertices in the positive octant and constraining every vertex initially on a reflection plane to stay on that plane.

We enforce two types of printability constraints: minimal part thickness and a self-supporting constraint. We do this by first *enforcing printability on the pre-blended structure*, that is, the structure where all inflated edges are combined with an exact union. This can be done easily by applying inequality constraints on the position and radius variables during our optimization. In particular, we apply a lower bound constraint to the radius variables and ensure the sphere defined by each vertex is supported:

$$v.z - v.r \geq \min_u(u.z - u.r),$$

where z is the vertex's position along the printing axis, r is the radius, and the minimization is over neighboring vertices of v that can "support" v . Then we restrict our blending to avoid bulging that would violate the self-supporting constraint by creating overhanging features (Section 6). The result is a printable, smoothly blended structure.

We optimize a discretization of (1) using SLSQP [Kraft 1994] within the NLOpt package [Johnson 2016]. At each step of the optimization, we perform the following steps to evaluate the objective function and its gradient:

- (i) We use CGAL's [Alliez et al. 2016] 3D Mesh Generation package to create a high-quality tetrahedral mesh of the structure's symmetry cell from our smooth signed distance function.
- (ii) We solve the elasticity equations (4) within the symmetry cell to find the fluctuation displacements, then compute the effective macro-scale elasticity tensor and its shape derivative, as well as the worst-case stress measure at each point (by solving a small eigenvalue problem at each point). These steps are outlined in Section 4.
- (iii) We compute the perturbation velocity of the boundary vertices induced by perturbing each shape parameter using automatic differentiation of our signed distance function.
- (iv) We solve the adjoint problem for our stress objective and evaluate the shape derivative (Section 5): using the adjoint solution, we construct the discrete volume differential form (25) that computes the change in stress objective caused by perturbing each mesh vertex. Then from (25), we construct an accurate boundary differential form (discrete shape derivative) (16) which takes a boundary perturbation velocity, smoothly extends it to the interior and applies the volume differential form. Feeding in the velocities found in (iii), we obtain accurate partial derivatives for each parameter.

4 WORST-CASE STRESS ANALYSIS OF PERIODIC STRUCTURES

In this section, we obtain formulas for the worst-case stress at a point in the microstructure. In general, finding a load distribution (with the total load fixed) that maximizes the stress norm at a point is a difficult nonconvex problem—unlike minimization, maximization of a convex function on a constrained domain is not convex.

However, in our specific setting, the space of loads that we need to consider is low (6-dimensional) and periodic homogenization




Fig. 3. (Schematic) Periodic tiling of a domain Ω with base cell Y having geometry ω and length scale ϵ [Panetta et al. 2015].

imposes structure that greatly simplifies the optimization. We consider the idealized setting where the structure's length scale is much smaller than the object to be filled by the structure—a standard assumption in microstructure design—so that at the macro scale (the full object's scale) the structure can be viewed as a homogeneous material experiencing smooth stress and strain distributions varying gradually compared to the size of a single cell. Evaluating these smooth stress and strain distributions at a point in the full object obtains the *average* stress and strain tensor experienced by the microstructure cell at that point, and under the periodic homogenization assumptions either of these tensors completely determines the microstructure's deformation. Thus, the external loads transmitted to the cell are captured by a single averaged "macroscopic stress" tensor, reducing our space of loads to 6 dimensions.

Solving this optimization problem *for every point*, we obtain a distribution of per-point worst-case stress (and corresponding per-point worst-case load), which depends on the microstructure's shape only, not on the full object or choice of loading.

To derive the equations for the worst-case stress at a given point, we need to review the formulas for homogenization of linear elasticity (e.g., [Allaire 2002]). We employ Einstein summation notation, where summation over repeated indices is implied. When appropriate, we use coordinate free expressions with $A : B$ denoting double contraction, $A :: B$ denoting quadruple contraction, and \otimes denoting tensor products.

The homogenized elasticity tensor. We summarize the key steps in the homogenization process used to determine the effective material properties of a solid-void microstructure fabricated with printing material C^{base} . This process is presented in more detail in [Panetta et al. 2015]. The elastic response of a periodically tiled object, Ω (Figure 3), under macroscopic external load $\bar{\mathbf{f}}$ is governed by the linear elastostatic equation

$$-\nabla \cdot [\bar{\mathbf{C}} : \varepsilon(\bar{\mathbf{u}})] = \bar{\mathbf{f}} \quad \text{in } \Omega, \quad (2)$$

augmented with problem-dependent boundary conditions. Tensor $\bar{\mathbf{C}}$ is the homogenized elasticity tensor, vector $\bar{\mathbf{u}}$ denotes the macroscopic (averaged) displacement, and $\varepsilon(\mathbf{u}) := \frac{1}{2}(\nabla \mathbf{u} + (\nabla \mathbf{u})^T)$ is the Cauchy strain tensor.

Zooming in on a particular point, we see a periodic tiling of the structure's base cell Y where each cell undergoes the same average strain $\varepsilon(\bar{\mathbf{u}})$. Translational symmetry implies the microscopic strain and stress fields in the tiling are periodic, so the tiling's displacement field consists of a linear term (with constant strain $\varepsilon(\bar{\mathbf{u}})$) plus a

periodic “microscopic fluctuation” term \mathbf{w} (having zero average strain by periodicity). The average stress in the tiling is defined by $\bar{C} : \varepsilon(\bar{\mathbf{u}}) = \frac{1}{|\bar{Y}|} \int_{\omega} C^{\text{base}} : [\varepsilon(\mathbf{w}) + \varepsilon(\bar{\mathbf{u}})] \, dy$. Decomposing $\varepsilon(\bar{\mathbf{u}})$ in the basis of symmetric rank 2 tensors, i.e., $e^{kl} := \frac{1}{2}(\mathbf{e}_k \otimes \mathbf{e}_l + \mathbf{e}_l \otimes \mathbf{e}_k)$, we see that the homogenized elasticity tensor is given by

$$\bar{C}_{ijkl} = \frac{1}{|\bar{Y}|} \int_{\omega} C_{ijpq}^{\text{base}} [\varepsilon(\mathbf{w}^{kl}) + e^{kl}]_{pq} \, dx. \quad (3)$$

Finally, to find the microscopic fluctuation term, the force balance equation in the cell Y is solved for each of the six basis strain probing conditions e^{kl} :

$$-\nabla \cdot (C^{\text{base}} : [\varepsilon(\mathbf{w}^{kl}) + e^{kl}]) = 0 \quad \text{in } \omega, \quad (4a)$$

$$\hat{\mathbf{n}} \cdot (C^{\text{base}} : [\varepsilon(\mathbf{w}^{kl}) + e^{kl}]) = 0 \quad \text{on } \partial\omega \setminus \partial Y, \quad (4b)$$

$$\mathbf{w}^{kl}(\mathbf{x}) \text{ } Y\text{-periodic}, \quad (4c)$$

$$\int_{\omega} \mathbf{w}^{kl}(\mathbf{x}) \, dx = 0. \quad (4d)$$

The compliance tensor \bar{S} is defined as the symmetric rank four inverse of \bar{C} . Notice that \bar{C} is a function of the structure’s geometry alone; it is independent of the macroscopic details (shape Ω , forces $\bar{\mathbf{f}}$, and boundary conditions).

Macro- to micro-stress. After solving (4), we can construct a rank four tensor mapping the cell’s macroscopic strain to the microscopic strain at a point $\mathbf{x} \in \omega$: $G_{ijkl}(\mathbf{x}) \stackrel{\text{def}}{=} [\varepsilon(\mathbf{w}^{kl})(\mathbf{x}) + e^{kl}]_{ij}$. Then the microscopic stress at \mathbf{x} under a particular macroscopic loading $\bar{\sigma}$ is

$$\sigma(\mathbf{x}) = C^{\text{base}} : G(\mathbf{x}) : \bar{S} : \bar{\sigma} \stackrel{\text{def}}{=} F(\mathbf{x}) : \bar{\sigma}. \quad (5)$$

Tensor F is the linear map from the macroscopic stress to the microscopic stress at a point. Although the micro-stress can vary from point to point in complicated ways, it actually can be maximized explicitly due to its simple relationship to load $\bar{\sigma}$ at each point.




Fig. 4. Left: worst-case stress field with peak stress value circled. Right: macroscopic stress condition inducing this peak stress.

Various pointwise stress measures can be utilized. The worst-case maximum principal stress, predicting the failure of brittle material, is defined at a point in the microstructure as:

$$s_m(\mathbf{x}) = \max_{\bar{\sigma} : \bar{\sigma}=1} \lambda_{\max}(F(\mathbf{x}) : \bar{\sigma}), \quad (6)$$

where we maximize over macroscopic stresses with unit Frobenius norm. The worst-case microscopic stress Frobenius norm is:

$$s_f^2(\mathbf{x}) = \max_{\bar{\sigma} : \bar{\sigma}=1} \bar{\sigma} : \underbrace{F(\mathbf{x})^T : F(\mathbf{x})}_{T^F(\mathbf{x})} : \bar{\sigma}, \quad (7)$$




Fig. 5. The three worst-case microscopic stress measures from left to right: maximum principal stress, Frobenius norm, and von Mises stress. Each is efficient to compute exactly for every element in the mesh.

and the maximum von Mises microscopic stress (for predicting ductile failure) is given by

$$s_v^2(\mathbf{x}) = \max_{\bar{\sigma} : \bar{\sigma}=1} \bar{\sigma} : \underbrace{F^T(\mathbf{x}) : V^T : V : F(\mathbf{x})}_{T^V(\mathbf{x})} : \bar{\sigma}, \quad (8)$$

where V extracts the von Mises stress: $(V\sigma)_{ij} = \sqrt{\frac{3}{2}} \left(\sigma_{ij} - \frac{1}{3} \delta_{ij} \sigma_{kk} \right)$.

It turns out that each worst-case stress measure can be evaluated efficiently by solving a tensor eigenvalue problem. The expressions for s_f and s_v are already in the familiar eigenvalue form for symmetric tensors T^F and T^V : they can be computed by flattening these tensors into matrices and finding the maximum eigenvalue. The corresponding eigenvector, when unflattened, is the worst-case load. Worst-case maximum principal stress leads to a different type of eigenvalue problem (see Appendix A for more detail):

$$\mathbf{n}^* \cdot T^M(\mathbf{x}) : [\mathbf{n}^* \mathbf{n}^{*T}] = s_m^2 \mathbf{n}^*, \quad (9)$$

with $T^M(\mathbf{x}) = F(\mathbf{x}) : F(\mathbf{x})^T$. We denote the *per-point* macro-stress tensor attaining the peak stress measure by worst-case load $\bar{\sigma}^*$.

Surprisingly, all three of these measures are very similar, as the maximum stress tensor eigenvalue tends to dominate all others in the worst-case scenario (see Figure 5).

Optimization problem. We wish to find a microstructure that minimizes the L_p norm of the stress measure (either s_m , s_f , or s_v) while achieving a particular homogenized elasticity tensor C^* :

$$\underset{\substack{\omega \text{ printable} \\ \bar{C}(\omega) = C^*}}{\text{argmin}} J(\omega), \quad J(\omega) = \int_{\omega} j(s(\mathbf{x})) \, dx, \quad j(s) \stackrel{\text{def}}{=} s^{p/2}. \quad (10)$$

Here, s is the *squared* stress measure of interest, e.g. s_v^2 . In our experiments, we use p up to 16. For simplicity, the remainder of this paper will consider s_f , but the other cases are nearly identical.

Discretization. We discretize the cell problems (4) with quadratic tetrahedral FEM, which we found essential for accurate stress and homogenized tensor evaluation. We use straight-edged elements (subparametric FEM) for representing the geometry to simplify meshing and the shape derivative formulas (edge nodes are placed at the edge midpoints).

The integral in the objective function (10) is computed with numerical quadrature, and the quadrature points determine where worst-case stress must be evaluated. We use piecewise constant quadrature, meaning *a single worst-case stress quantity is computed*

for each mesh element (from the element's averaged fluctuation strains).

5 SHAPE OPTIMIZATION

We found the equality-constrained optimization formulation (10) necessary to reliably reduce stress while preserving macroscopic behavior. For instance, naively smoothing to reduce stress dramatically changes macroscopic properties and is undone by re-fitting via [Panetta et al. 2015]. Further, though stress-optimal designs generally exhibit smooth features, predicting how smooth is difficult. In fact, [?] demonstrates that theoretically optimal structures can have sharp corners. While we have not encountered this in our experiments, we see significant variation in optimal curvature.

To solve (10) with SLSQP, we need partial derivatives of the objective and constraints with respect to each shape parameter. We compute these derivatives in two stages: first, we compute the quantities' *shape derivatives*, i.e., how they change when the domain ω is perturbed by an arbitrary velocity field \mathbf{v} . Then, to differentiate with respect to a parameter, we feed into these shape derivatives the velocity field induced by changing the parameter.

5.1 Volumetric vs. boundary shape derivatives

The shape derivative of the objective functional J in the direction of a shape perturbation \mathbf{v} is defined as the Gâteaux derivative

$$dJ[\mathbf{v}] \stackrel{\text{def}}{=} \lim_{t \rightarrow 0} \frac{J(\omega(t, \mathbf{v})) - J(\omega)}{t}, \quad (11)$$

where $\omega(t, \mathbf{v}) = \{\mathbf{x} + t\mathbf{v}(\mathbf{x}) : \mathbf{x} \in \omega\}$.

Shape derivatives have been used to minimize a wide variety of objective functions, and they are typically expressed as a boundary integral depending only on the boundary's normal velocity, e.g. [All 2008]. However, we found that in our setting (nearly nonsmooth objective functionals resulting from using the L_p norm for a high p), these standard formulas give poor accuracy (Figure 6). To address this issue, we have developed a different formulation using volume integrals instead. In this section, we use a simplified model problem to clarify the source of inaccuracy in the standard method and to introduce the main idea of our approach. Our expression for the shape derivative for worst-case stress is described in Section 5.2 and derived in the supplementary document.

Our model problem is to differentiate $J = \int_{\omega} j(\nabla u) \, dx$ with respect to changes of the domain ω , where u is a scalar field solving

$$\nabla \cdot (\nabla u + g) = 0 \quad \text{in } \omega, \quad \hat{\mathbf{n}} \cdot (\nabla u + g) = 0 \quad \text{on } \partial\omega. \quad (12)$$

Here, g is analogous to the macroscopic strain.

The standard expression for J 's shape derivative is:

$$dJ[\mathbf{v}] = \int_{\partial\omega} \left(j(\nabla u) - \nabla p \cdot (\nabla u + g) \right) \mathbf{v} \cdot \hat{\mathbf{n}} \, dx,$$

where p is the adjoint scalar field solving

$$\nabla \cdot (\nabla p - j'(\nabla u)) = 0 \quad \text{in } \omega, \quad \hat{\mathbf{n}} \cdot (\nabla p - j'(\nabla u)) = 0 \quad \text{on } \partial\omega.$$

For optimizing relatively smooth functionals like compliance, this method yields good results with standard discretization approaches. To illustrate the problem with it in our setting, we briefly review the main steps of its derivation.




Fig. 6. Comparisons of different techniques for differentiating our L_p worst-case stress objective with respect to a pattern parameter (in this case, an offset variable whose value is plotted on the horizontal axis). For low p , derivatives are computed reasonably accurately by all methods. However, for L_{12} (which is needed to approximate L_∞ as shown in Figure 15) the traditional formula for the shape derivative is wildly inaccurate. Using our discrete differential form with interpolated boundary velocities maintains good accuracy.

We consider the initial domain ω and slightly perturbed domain ω_t , with the deformation from ω to ω_t given by map $\mathbf{x}_t = \mathbf{x} + t\mathbf{v}$ with Jacobian $I + t\nabla\mathbf{v}$. Using the Reynolds Transport Theorem, the derivative of J can be computed in terms of the *Eulerian derivative* \dot{u} of u_t (the PDE solution at time t) evaluated at $t = 0$:

$$dJ[\mathbf{v}] = \int_{\omega} j'(\nabla u) \cdot \nabla \dot{u} \, dx + \int_{\partial\omega} j(\nabla u) \mathbf{v} \cdot \hat{\mathbf{n}} \, dA(\mathbf{x}),$$

We obtain \dot{u} starting from the weak form of (12) in domain ω_t :

$$\int_{\omega_t} \nabla \phi \cdot (\nabla u_t + g) \, dx = 0, \quad \forall \phi. \quad (13)$$

We apply the Reynolds Transport Theorem for the second time to obtain an equation for the Eulerian derivative, \dot{u} . As detailed in the supplement, an intermediate step in the derivation of the Reynolds Transport Theorem—when applied to the weak form (13)—yields an

equation for the material derivative, δu :

$$\int_{\omega} \nabla \phi \cdot (\nabla \delta u - \nabla \mathbf{v}(\nabla u + g) - (\nabla \mathbf{v})^T \nabla u + (\nabla u + g) \nabla \cdot \mathbf{v}) \, dx = 0 \quad (14)$$

holding for all ϕ . The derivation then proceeds by substituting the definition $\dot{u} \stackrel{\text{def}}{=} \delta u - \mathbf{v} \cdot \nabla u$. Applying integration by parts to the $\nabla \cdot \mathbf{v}$ term, we obtain:

$$\int_{\omega} \nabla \phi \cdot \nabla \dot{u} \, dx = \underbrace{\int_{\omega} \nabla(\nabla \phi \cdot \mathbf{v}) \cdot (\nabla u + g) \, dx}_{I_1} - \int_{\partial \omega} \nabla \phi \cdot (\nabla u + g)(\mathbf{v} \cdot \hat{\mathbf{n}}) \, dx \quad \forall \phi. \quad (15)$$

Note that I_1 is exactly of the form (13), with $\nabla \phi \cdot \mathbf{v}$ instead of ϕ . As ϕ is assumed to be arbitrary (from a sufficiently smooth space of functions), (13) implies that this term vanishes, leaving us with a simple equation for \dot{u} (in this case, a Laplace equation with a Neumann boundary condition).

Our key observation is that this term may not vanish identically in the discrete setting, with the error being particularly large for low-order finite element discretizations. We found that for rapidly varying functionals like our worst-case stress, this leads to large errors in derivatives. To solve this problem, one could choose the finite element basis and test functions for the problem in a way that guarantees that I_1 is exactly zero. Alternatively—and this is the approach we follow—one can skip integration by parts and use (14) to obtain the discretization. This approach ends up obtaining our objective’s *exact discrete derivative*: it computes the exact change in the objective if the tetrahedral mesh nodes are infinitesimally advected by the velocity field without remeshing.

Extending boundary perturbation into the domain. Unlike the traditional shape derivative formulas, this volume form will depend on the perturbations \mathbf{v} on the entire domain ω , rather than only the boundary. This should not be surprising—finite element analysis is sensitive to the interior meshing—but it is problematic: perturbing the microstructure parameters induces a velocity \mathbf{v}^b on ω ’s boundary *only*, leaving the interior perturbation undefined.

Although we know that, in the limit of refinement, the shape derivative is independent of \mathbf{v} inside the domain (interior perturbations do not impact shape), setting the interior velocity to 0 hurts accuracy in practice (Figure 6). This is because perturbing only boundary vertices distorts the incident elements and increases the discretization error as a result. Smoothly extending velocities into the interior (so that elements advect with the boundary, distorting less) improves accuracy.

We perform this extension by solving the Laplace equation with Dirichlet condition \mathbf{v}^b on the microstructure boundary and periodic conditions on the cell boundary. Denoting the velocity at the interior vertices by \mathbf{v}^i , we compute $\mathbf{v}^i = -L_{ii}^{-1} L_{ib} \mathbf{v}^b$, where L denotes the linear finite element laplacian matrix. Using this extension, we can define a differential form acting on the boundary velocity \mathbf{v}^b :

$$dJ^b[\mathbf{v}^b] \stackrel{\text{def}}{=} dJ[\mathbf{v}] = dJ \left[\begin{pmatrix} -L_{ii}^{-1} L_{ib} \\ I \end{pmatrix} \mathbf{v}^b \right]. \quad (16)$$

This differential form can be constructed explicitly by applying the transpose of the interpolation matrix to the explicit representation




Fig. 7. Visualization of shape derivative quantities. Left: the normal velocity induced by changing a position variable, as computed by automatic differentiation of the signed distance function. On the right, the steepest descent velocity for our worst-case stress objective and the compliance tensor-fitting objective from [Panetta et al. 2015].

of the volume differential ((25) below). This accelerates gradient computation for patterns with many parameters and recovers an accurate steepest descent boundary velocity, as visualized in Figure 7.

5.2 Shape derivative in volume form for worst-case stress

We summarize the formulas for stress shape derivatives, in the form we use in our implementation. The derivations of these are somewhat lengthy and appear in the supplementary document.

The worst-case stress objective’s shape derivative is, in volume form:

$$dJ[\mathbf{v}] = \int_{\omega} j \nabla \cdot \mathbf{v} + \tau^{kl} : D[\varepsilon(\mathbf{w}^{kl})] + \gamma :: d\bar{C}[\mathbf{v}] \, dx, \quad (17)$$

where $D[\bullet]$ denotes the material derivative, and \mathbf{v} is the perturbation velocity field. Tensors τ^{kl} and γ are partial derivatives of $j(s(\mathbf{x})) = j(s(\varepsilon^{kl}, \bar{C}, \mathbf{x}))$ with respect to the fluctuation strains and homogenized elasticity tensor:

$$\delta j = (j') \frac{\partial s}{\partial \varepsilon^{kl}} : \delta \varepsilon^{kl} + (j') \frac{\partial s}{\partial \bar{C}} :: \delta \bar{C} \stackrel{\text{def}}{=} \tau^{kl} : \delta \varepsilon^{kl} + \gamma :: \delta \bar{C},$$

$$\tau^{kl} = (2j' C^{\text{base}} : F : \bar{\sigma}^*) [\bar{S} : \bar{\sigma}^*]_{kl}, \quad (18)$$

$$\gamma = (-2j' F^T : F : \bar{\sigma}^*) \otimes (\bar{S} : \bar{\sigma}^*). \quad (19)$$

(Recall that the $\bar{\sigma}^*$ value for a point \mathbf{x} is the unit macroscopic stress maximizing microscopic stress at \mathbf{x}). These formulas do not include the derivative of the worst-case load $\bar{\sigma}^*$ since the derivative of an eigenvalue (pointwise worst-case stress) does not depend on its eigenvector’s derivative; see Appendix B. In other words, the worst-case load can be considered constant when differentiating.

Because we use straight-edged elements (i.e. the fluctuation displacement fields are piecewise quadratic, but the geometry representation is piecewise linear), the perturbation velocity \mathbf{v} is piecewise linear. Thus \mathbf{v} is represented as a perturbation vector $\delta \mathbf{q}_i$ on each mesh vertex:

$$\mathbf{v} = \sum_i \lambda_i \delta \mathbf{q}_i, \quad (20)$$

where λ_i is vertex i ’s linear shape function.

The first term in the integrand of (17) can be computed directly, and the third term is the homogenized elasticity tensor’s shape derivative, which can be computed from the fluctuation displacements \mathbf{w}^{kl} . The second term includes the unknown material derivative of the fluctuation strains; we re-express this term using the solution

to the adjoint cell problems, which are in the same form as (4) but with different right hand sides.

Adjoint equation. We need to solve an adjoint equation to express the term involving the fluctuation displacements' material derivatives in a computationally tractable form. The weak form of the kl^h adjoint cell problem PDE is:

$$\int_{\omega} \tau^{kl} : \varepsilon(\phi) \, d\mathbf{x} = \int_{\omega} \varepsilon(\mathbf{p}^{kl}) : C^{\text{base}} : \varepsilon(\phi) \, d\mathbf{x}, \quad (21)$$

for all test functions ϕ , where \mathbf{p}^{kl} is the adjoint state vector field. \mathbf{p}^{kl} is discretized using the same piecewise quadratic basis functions as used for fluctuation displacements.

The expression for the discrete shape derivative is:

$$\begin{aligned} dJ[\mathbf{v}] = \int_{\omega} \left[(j - \varepsilon(\mathbf{p}^{kl}) : \sigma^{kl}) \nabla \cdot \mathbf{v} + (\nabla \mathbf{p}^{kl} \nabla \mathbf{v}) : \sigma^{kl} \right. \\ \left. + (\varepsilon(\mathbf{p}^{kl}) : C^{\text{base}} - \tau^{kl}) : (\nabla \mathbf{w}^{kl} \nabla \mathbf{v}) \right] d\mathbf{x} \\ + \left(\int_{\omega} \gamma \, d\mathbf{x} \right) :: d\bar{C}[\mathbf{v}], \end{aligned} \quad (22)$$

where $\sigma^{kl} \stackrel{\text{def}}{=} C^{\text{base}} : [\varepsilon(\mathbf{w}^{kl}) + e^{kl}]$ is the microscopic stress corresponding to \mathbf{w}^{kl} .

For computation, it is convenient to express $dJ[\mathbf{v}]$ explicitly as a 1-form acting on the per-vertex perturbation vector field $\delta\mathbf{q}$. To do this, we re-express \mathbf{v} in terms of $\delta\mathbf{q}$. Using (20),

$$\nabla \cdot \mathbf{v} = \sum_m \delta\mathbf{q}_m \cdot \nabla \lambda_m, \quad \nabla \mathbf{v} = \sum_m \delta\mathbf{q}_m \otimes \nabla \lambda_m. \quad (23)$$

We can write $\nabla \mathbf{p}^{kl}$ in terms of each scalar-valued finite element shape function φ_n and its vector-valued coefficient \mathbf{p}_n^{kl} as:

$$\nabla \mathbf{p}^{kl} = \sum_n \mathbf{p}_n^{kl} \otimes \nabla \varphi_n. \quad (24)$$

Substituting these into $dJ[\mathbf{v}]$ and simplifying:

$$\begin{aligned} dJ[\lambda_m \delta\mathbf{q}_m] = \delta\mathbf{q}_m \cdot \int_{\omega} a \nabla \lambda_m + (\nabla \lambda_m \cdot \mathbf{b}_n) \nabla \varphi_n \, d\mathbf{x} \\ + \left(\int_{\omega} \gamma \, d\mathbf{x} \right) :: d\bar{C}[\lambda_m \delta\mathbf{q}_m], \quad (25) \\ a = j - \varepsilon(\mathbf{p}^{kl}) : \sigma^{kl}, \\ \mathbf{b}_n = \sigma^{kl} \mathbf{p}_n^{kl} + (\varepsilon(\mathbf{p}^{kl}) : C^{\text{base}} - \tau^{kl}) \mathbf{w}_n^{kl}, \end{aligned}$$

where summation over repeated subscripts is implied.

The exact discrete shape derivative of homogenized tensor \bar{C} is $d\bar{C}_{ijkl}[\lambda_m \delta\mathbf{q}_m] =$

$$\begin{aligned} \frac{\delta\mathbf{q}_m}{|Y|} \cdot \int_{\omega} \left(\sigma^{ij} : C^{\text{base}-1} : \sigma^{kl} \right) \nabla \lambda_m \\ - \left[\nabla \lambda_m \cdot \left(\sigma^{kl} \mathbf{w}_n^{ij} + \sigma^{ij} \mathbf{w}_n^{kl} \right) \right] \nabla \varphi_n \, d\mathbf{x}, \end{aligned}$$

where again summation over m and n is implied.

To summarize, we compute the shape derivative in three stages:

- (i) solve the adjoint problem to obtain solutions \mathbf{p}^{kl} ; (ii) compute the shape derivative of the homogenized elasticity tensor \bar{C} ; and
- (iii) substitute these quantities into (25) and evaluate the integrals.

6 CONVEX HULL-RESTRICTED BLENDING

The parametric microstructure model is an essential part of our method. It is designed to have few parameters (ensuring optimization is stable and fast), while at the same time providing enough degrees of freedom to achieve our stress reduction goal. Our model consists of a skeleton graph with position, radius, and smoothing parameters on each vertex. These parameters describe a sphere for each vertex, and we define our edge geometry *primitives* as the convex hull of the endpoint spheres for each edge. We then combine every edge primitive incident on a vertex to form a smooth joint, with smoothness determined by the vertex's smoothing parameter, and finally blend all the joints together into a single signed distance function.

A variety of techniques, discussed briefly in Section 2, have been developed to control smoothness both for parametric and implicit surfaces. Unfortunately, existing techniques do not meet some criteria our method needs to satisfy. Our list of requirements includes: (i) *Smooth geometry*: the surface should be smooth for almost all parameter values. (ii) *Smooth dependence on parameters*: the dependence of surface points on shape parameters should be differentiable, as (efficiently computable) derivatives are needed for shape optimization. (iii) *Locality*: smoothing needs to be controlled locally, with each smoothness parameter affecting a part of the geometry (in our case, a joint). (iv) *No bulging*: the surface does not create unnecessary protrusions, which are a common side effect of smoothed boolean operations; this aspect is particularly important for printability. (v) *Constraint-compatibility*: printability constraints can be efficiently expressed in terms of shape parameters.

As all existing methods we considered appeared difficult to adapt directly to our application, we have developed a new, simple blending technique, which works quite well for our application. We believe that the basic idea may be useful in other contexts, but we have not evaluated its utility for general-purpose implicit modeling.

Our method works on a collection of smooth shapes; for each we need a signed distance function, and for a group of shapes we need to compute signed distances to the convex hull efficiently. We define a *joint* as a maximal set of primitives that have a nonempty intersection. In our setting, this is a set of edge geometry primitives incident on a vertex.

Bulging. Intuitively, bulging is a behavior of shape blending methods where material is added to a solid in places unnecessary to create a smoothed shape (e.g., Figure 8). It is of special importance in the context of 3D printing, as bulging is likely to violate the printability of certain types of structures. To clarify: we enforce printability constraints on the edge primitives during optimization, so the boolean (nonsmooth) union of edges is guaranteed to be printable. Our blending operator is always additive and thus cannot violate the edge primitives' minimum thickness constraint, but bulging can easily add overhanging material that violates the self-supporting constraint for SLA printers (Figure 9).

While no formal definition of bulging is likely to capture the intuition completely, we use the following definition, based on the intuition for blended sets of convex primitives, which matches our goals well: *a bulge is the difference between the exact union of the*




Fig. 8. *Bulging. Left: bulging exhibited by smoothed distance function; Center: blending region defined by convex hull; Right: smooth, bulge-free joint.*




Fig. 9. *Unprintability due to bulging (overhanging features bulging below the red line) is resolved by the convex hull-restricted blending.*

constituent primitives and the blended union that is not contained in the convex hull of the primitives forming a joint.

The rationale for this definition is that for smooth primitives, the parts on the surface of the convex hull are smooth, so there is no need for further smoothing.

Summary. Our shape’s boundary is defined as the zero level-set of a scalar volume function, which is constructed from the distance functions to the convex, tangent-plane continuous primitives. As a basic component of our method, we use the *Kreisselmeier-Steinhauser (KS) function*,

$$KS(y_1, \dots, y_n; \rho) = -\frac{1}{\rho} \ln \left(\sum_{i=1}^n e^{-\rho y_i} \right), \quad (26)$$

where variables y_i are the distances to be blended. This function is a smoothed version of the minimum function, with smoothness controlled by parameter ρ . For the purposes of blending, it is more intuitive—and yields better-scaled optimization variables—to apply the change of parameters $s = 1/\rho$. Then, increasing smoothing parameter s increases the amount of blending.

The KS function has two important properties: (i) if the difference between maximal y_i and the rest exceeds approximately $7s$ it vanishes to 10^{-12} accuracy, so it is effectively local; and (ii) it is always smaller than the non-smoothed min function, with the maximal difference proportional to s .

We use a two-stage blending process: first, we perform local blending at joints, using position-dependent smoothing (s drops to zero for points outside the joint’s convex hull). This yields larger smooth, non-convex building blocks for the structure: the blended joints, which are connected by shared primitives (the single common edge primitive for a pair of adjacent joints in our setting). Second, we blend all joints together with a spatially-varying smoothing amount

that avoids sharp creases while preventing unnecessary bulging on the shared primitives.

Blending algorithm. The algorithm computes the level-set function $F(\mathbf{p})$ at a point $\mathbf{p} \in \mathbb{R}^3$ given a collection of convex primitives P_i ($i = 1 \dots N_p$) with signed distance functions $F_i(\mathbf{p})$. The primitives form joints J_m ($m = 1 \dots N_J$), each defined by a subset of the primitives. We define R_m to be the minimal distance from the medial axis of any primitive in J_m to the convex hull boundary.

Blending is controlled by smoothing parameters s_m , defined *per joint* (skeleton vertex), and the smoothing amounts s_{ab} , defined *per overlapping joint pair*. For reasons we explain below, we use only s_m as optimization variables (in addition to the primitives’ parameters). *Constructing blended joints.* We use the smoothed minimum function KS to construct joint blends, but with a spatially varying smoothing parameter $s(\mathbf{p})$. For points inside the convex hull of J_m , we use the maximal smoothing amount, equal to joint’s blending parameter s_m . Outside the convex hull, s quickly decays to zero: we need to leave a small buffer zone outside the convex hull to keep the blended function differentiable and to avoid nearly non-smooth surfaces when an intersection curve of two primitives approaches the boundary. To apply this blending modulation, we define $\eta_J(t)$, where t is the (nondimensionalized) distance from \mathbf{p} to the convex hull boundary, with the following requirements: (i) $\eta_J(t) = 1$ for $t \leq -1$, (well inside the convex hull); and (ii) $\eta_J(t) = 0$ at $t \geq \epsilon$, a user specified distance outside the convex hull.

We define two signed distance functions per joint, F_m^s , representing the blended joint, and F_m^h , representing the (nonsmooth) exact union of the primitives:

$$F_m^s(\mathbf{p}) = KS \left(\{F_i(\mathbf{p}), i \in J_m\}; s_m \eta_J \left(\frac{\text{dist}_{\text{hull}}(J_m)(\mathbf{p})}{R_m} \right) \right) \quad (27)$$

$$F_m^h(\mathbf{p}) = \min_{i \in J_m} F_i(\mathbf{p}).$$

Combining joints. The exact union of the smoothed joints in general is non-smooth: creases can appear on the shared primitives at the intersection of two joint surfaces.

The obvious solution is to combine all joints using the smoothed minimum function KS ; however, applying it in all cases will lead to bulging (Figure 10). We apply it for a point \mathbf{p} only to the two joints closest to \mathbf{p} in the sense of smoothed distances F_m^s , and only when we detect that the joints “conflict” in a way that would cause a crease. Specifically, we observe that $|F_m^s(\mathbf{p}) - F_m^h(\mathbf{p})|$, $m = a, b$, measures the degree to which each joint’s surface differs from the original primitive near \mathbf{p} . A crease forms along the shared primitive only if both of these quantities are nonzero; if one of the joints coincides with the shared primitive near \mathbf{p} , an exact union of the joint surfaces simply adds the smooth blend of the other joint (if any) atop the primitive. Thus we apply no blending when either quantity is zero and rapidly but smoothly transition to a slight blend as both quantities become nonzero. We achieve this with a smoothing modulation function, $\eta_F(t)$, of these quantities’ squared geometric mean. We design this function to transition from $\eta_F(0) = 0$ quickly to 1 as t increases.

$$F(\mathbf{p}) = KS(F_a^s(\mathbf{p}), F_b^s(\mathbf{p}); \tilde{s}_{ab}(\mathbf{p}))$$

$$\tilde{s}_{ab}(\mathbf{p}) = s_{ab} \eta_F \left((F_a^h(\mathbf{p}) - F_a^s(\mathbf{p}))(F_b^h(\mathbf{p}) - F_b^s(\mathbf{p})) \right) \quad (28)$$




Fig. 10. *Top left: when two joints’ blending regions overlap, we must smoothly union them to avoid a sharp crease. Top right: blending using KS with spatially constant parameter s_{ab} dilates the shared edge even when no overlap exists, introducing bulging that violates printability. Blending with smoothing modulation η_F (bottom) solves both problems.*

Transition functions. Our specific choice of transition functions η_J and η_F is heuristic but yields good results based on our experience. Other functions could be used as long as the specifications stated above are satisfied. Both functions are constructed using tanh, which can be viewed as constant outside a local interval similarly to the KS function. We prefer this choice to, e.g., splines, which can be made truly constant outside an interval, primarily due to the more compact formulas. The specific choices we found to work well are $\eta_J(t) = 1 - \tanh(C \max(1 + t, 0))^r$, with $r = 10$ and $C = 1.025$, which transitions from 1 to 0 roughly in the range -0.5 to 0.1 , and $\eta_F(t) = \tanh(1000t)$, which ensures smoothing is applied if even a slight crease appears. We did not attempt to optimize s_{ab} , which we fix to 0.02 for all joint pairs, because transition regions between joints do not typically concentrate stresses in our structures.

Gradient-based implicit blending. Gradient-based implicit blend-




Fig. 11. *Two surfaces’ normals intersect forming a right angle at a point outside the convex hull (translucent gray). Blending at this point would lead to potential unprintability.*

ing ([Gourmel et al. 2013]) is a popular method to avoid bulging in implicit modeling. However, it is not guaranteed to restrict blending to the convex hull of the input primitives and therefore cannot

preserve the primitives’ printability in our setting. One example is demonstrated in Figure 11: evaluating at the arrows’ intersection point outside the convex hull, the two primitives’ gradients are orthogonal. Since [Gourmel et al. 2013] uses only the gradients’ inner product to modulate blending, the method is unable to distinguish this situation—where blending would violate printability—from an evaluation point near a 90° intersection of two cylinders for which the maximum amount of blending would be desired.

From parameters to structure. We use CGAL’s 3D Mesh Generation package, which robustly constructs high-quality adaptive tetrahedral meshes from smooth signed distance functions. For accurate stress evaluation, we found adaptive refinement essential. We achieved this using CGAL’s `facet_distance` parameter, which ensures the mesh closely approximates the isosurface, leading to automatic refinement in regions of high curvature where stress concentrations are likely. One could instead drive refinement with a sizing field constructed from the worst-case stress field evaluated on a coarse initial mesh.

Creases in the isosurface lead to low-quality, occasionally non-manifold meshes with many tiny elements created to approximate the sharp features. However, CGAL can still produce high-quality meshes if it is provided all such creases explicitly as polygonal curves. Our tiled structures are smooth everywhere, but to mesh the symmetry cell, we must perform a boolean intersection with the cell’s bounding cube. This intersection operation creates sharp curves along the cube faces, which we extract with marching squares and pass to CGAL as feature lines.

The vertices CGAL creates on the symmetry cell boundary do not lie perfectly on the cell’s faces and must be snapped. However, points originating from the feature curves we extract *are* placed perfectly on the boundary. We use this fact to snap vertices in a robust way that is not too sensitive to the snapping threshold parameter: we use the feature curves to segment the mesh’s boundary vertices into connected components and then decide whether each whole component lies on the boundary.

Shape velocities. When a parameter p (controlling vertex position, radius, or smoothing) changes, each point \mathbf{x} on the microstructure’s boundary evolves with some velocity. We need this induced velocity field to evaluate our objective and constraints’ partial derivatives. We compute these fields directly from the signed distance function using automatic differentiation (with Eigen’s `Auto Diff` module).

Boundary point \mathbf{x} ’s motion in the normal direction is determined by differentiating the level set equation:

$$\hat{\mathbf{n}} \cdot \frac{d\mathbf{x}}{dp} = -\frac{1}{\|\nabla\phi\|} \frac{\partial\phi}{\partial p}, \quad (29)$$

where ϕ is the signed distance function and $\hat{\mathbf{n}} = \frac{\nabla\phi}{\|\nabla\phi\|}$ is the surface normal. The tangential velocity is left undefined, but it is natural to define it as zero.

7 RESULTS

In this section, we summarize the results obtained from our method.

Isotropic elasticity dictionary. We apply our framework to design structures with each of the six topologies chosen by [Panetta et al. 2015], rather than performing a new topology search.




Fig. 12. Comparison of material properties covered by [Panetta et al. 2015] (top) and our optimization (bottom). Each colored region shows a single topology’s coverage, and the dashed curve at the bottom outlines the region from [Panetta et al. 2015].

First, for each topology, we explore our new parametric model’s coverage of isotropic material space using a breadth-first-style search: we draw an infinite grid in the Young’s modulus, Poisson’s ratio plane (with grid points spaced logarithmically in Young’s modulus) and plot the point corresponding to the default parameter settings. We then design structures for every unattained grid point neighboring known structures by running our optimizer without the stress objective. We repeat until no more progress is made.

Next, starting from each structure found above, we run a worst-case stress minimization while holding fixed the homogenized material properties to design a more robust structure.

As can be seen from Figure 12, although we consider the more difficult problem of minimizing stress while simultaneously fitting the elasticity tensor, we achieve far greater coverage of isotropic elastic moduli. This is primarily due to two factors: a different shape parametrization enriched with additional parameters (smoothing factors at the joints) and a greater robustness of the meshing algorithm. Mesh robustness—particularly to topology changes as thicknesses grow and parts merge—helped extend the Young’s modulus range considerably (by a factor of 10), and the additional degrees of freedom controlling joint smoothness expanded the Poisson’s ratio range much closer to the theoretical limits (-1 to $1/2$).




Fig. 13. Stress reduction (top) and \log_{10} of optimized max stress (bottom) for each microstructure in our database.

The total range covered is still significantly less than theoretically possible, as given by Hashin-Shtrikman bounds [Hashin and Shtrikman 1963], however it is currently unknown if these bounds are realizable by printable structures. We conjecture that a more complete coverage requires finer-scale topology than provided by the type of structures we consider. We also note that the narrowing range of achieved Poisson’s ratios for higher E is predictable, since a fully solid structure can only produce the printing material’s ν . One surprising outcome of these experiments is that a single topology (number “0646” in [Panetta et al. 2015]’s enumeration) covers essentially the entire space reached by all six topologies (with the caveat that patterns can change topology during the optimization by merging, especially as the thickness parameters increase).

Stress reduction. Figure 13 shows the magnitude of stress reduction achieved by our algorithm is generally $5\times$ to $6\times$, but rapidly decays towards the boundary of the covered area. This is to be expected, as the subspace of shape parameters satisfying the elasticity tensor and printability constraints shrinks as we approach the boundary. We also observe a growth of maximal stress as Young’s modulus decreases, since thinner trusses are needed to make the material sufficiently flexible. Because we perform optimization in a reduced parameter space, we do not claim that our structures are optimal. However, the relatively uniform stress distributions over their surfaces suggest that much greater reductions are unlikely.




Fig. 14. Significant stress reduction is achieved for a diverse set of structures. The top row visualizes structures in our pattern subspace found by fitting to C^* without penalizing stress. The moduli are reported assuming a base printing material, $C^{base} = (200\text{MPa}, 0.35)$. The bottom row displays structures achieving the same C^* but with dramatically reduced peak stress. Histograms of the element counts at each stress level (reported in MPa) in the unoptimized (top) and optimized (bottom) designs are plotted on a log scale.

Note that the initial coverage exploration and worst-case stress minimization are performed over the same design space: the initial designs against which we compare stress levels already have smooth joints, but our worst-case stress objective is needed to exploit the smoothing parameters to significantly reduce stress.

Figure 14 plots histograms of per-element stresses. We observe that a large number of elements with high stress disappear after optimization. We also note that the total number of elements has

decreased because stress optimization generally eliminates the surface's high curvature regions, which require significant refinement. **Choice of p in L_p .** The dependence of maximal stress achieved by L_p optimization on the chosen objective norm is shown in Figure 15. We observe that $p = 12$ is generally adequate for approximating the L_∞ stress norm for optimization purposes. At this value of p , the objective is highly non-smooth, necessitating the accurate computation of shape derivatives described in Section 5.




Fig. 15. The peak worst-case stress level and corresponding structures designed by minimizing the L_p worst-case stress norm.

Resolution dependence. We have also verified that the optimization’s effectiveness at minimizing stress is insensitive to mesh resolution after a reasonable level of refinement is reached (Figure 16). We based our meshing settings for the sweep in Figure 13 on refinement experiments like this one.




Fig. 16. Optimizations run at increasing mesh resolution from right to left (controlled by adjusting CGAL’s facet_distance parameter to adaptively refine the surface mesh based on curvature), with the L_p norm of the stress of the final iterate evaluated on the fine mesh.

Experimental validation. Validating stress reduction experimentally is far more difficult than, e.g., measuring Young’s modulus. First, stress concentrations are difficult to measure directly and are affected by small geometry variations caused by printing inaccuracies. Second, though we design our structures for robustness under worst-case loads, these loads are often difficult to apply in a laboratory setting: our testing setup limited us to compression tests, which do not reveal the full benefits of our optimization. Finally, testing small tilings ($2 \times 2 \times 1$) leads to errors, as our compression test does not match the periodic homogenization boundary conditions. These

errors often manifest themselves as early fractures in the dangling bars touching the compression plates.




Fig. 17. We compressed unoptimized and optimized structures designed for the same material tensor until breakage. Despite the unoptimized pattern ending up slightly stiffer due to printing inaccuracies, it fails under much lower deformation.

We have, however, consistently observed that our optimized patterns can withstand deformations much larger than the deformations breaking the corresponding unoptimized patterns (Figure 17, Figure 18). Homogenized moduli are reported based on an assumed isotropic printing material with a Young’s modulus of 200MPa and a Poisson’s ratio of 0.35. Note that curing time significantly affects the true stiffness; we ensured the unoptimized and optimized structures of each tested pair were post-processed identically.

Figure 19 shows the results of loading to a fixed deformation, unloading, and repeating to test for fractures and plastic deformation; we have observed a substantial decrease in strength for unoptimized patterns and no change for the optimized ones.

8 CONCLUSIONS AND FUTURE WORK

We have demonstrated that the problem of determining worst-case loads, considered by numerous previous works, admits an exact solution for the specific case of periodic structures. We have developed a robust set of algorithms to minimize worst-case stresses, achieving a $5\times$ reduction for a significant share of patterns and substantially expanding the covered region of (E, ν) space versus previous work.

Our work has several limitations, which are all worth exploring further in future work. First, it is unclear how much further stress can be reduced while maintaining particular elasticity properties—a more in-depth study would be illuminating, but requires L_p stress optimization in the full-dimensional shape space. In our experience, this optimization is unstable, and it is not obvious how to constrain the optimization to printable designs. Second, it may be possible to cover a greater region of the material property space by considering finer topologies for the cells. Finally, one can apply our code to other physics-based design problems involving L_∞ stress optimization.




Fig. 18. Additional compression test results for a variety of topologies and material properties. A photo of the optimized pattern of each pair is inset. The optimized structure in the bottom left example did not actually break: it simply buckled into a lower energy configuration.




Fig. 19. We compressed one optimized and unoptimized structure at 0.5mm/s to 2mm. We then unloaded each structure and repeated the experiment twice. The optimized pattern traced out three overlapping lines at the bottom: its moduli were unchanged. The unoptimized structure broke on the first compression (top curve), causing weaker behavior on the two subsequent compressions.

A WORST-CASE MAXIMUM STRESS

The maximum principal stress at a point is given by (6):

$$s_m = \max_{\bar{\sigma}: \bar{\sigma}=1} \lambda_{\max}(F : \bar{\sigma}) = \max_{\bar{\sigma}: \bar{\sigma}=1} \max_{\|\mathbf{n}\|=1} \mathbf{n}^T [F : \bar{\sigma}] \mathbf{n}.$$

The Lagrangian for this maximization problem is

$$\mathcal{L}(\bar{\sigma}, \mathbf{n}, \lambda_{\bar{\sigma}}, \lambda_{\mathbf{n}}) = \mathbf{n}^T [F : \bar{\sigma}] \mathbf{n} + \lambda_{\bar{\sigma}} (\bar{\sigma} : \bar{\sigma} - 1) + \lambda_{\mathbf{n}} (\|\mathbf{n}\|^2 - 1).$$

The stationary condition with respect to $\bar{\sigma}$ is

$$\frac{\partial \mathcal{L}}{\partial \bar{\sigma}} = [\mathbf{nn}^T] : F + 2\lambda_{\bar{\sigma}} \bar{\sigma} = 0, \quad (30)$$

showing that, for arbitrary \mathbf{n} , principal stress is maximized by load:

$$\bar{\sigma}^* = -\frac{[\mathbf{nn}^T] : F}{2\lambda_{\bar{\sigma}}} = \pm \frac{[\mathbf{nn}^T] : F}{\sqrt{[\mathbf{nn}^T] : F : F^T : [\mathbf{nn}^T]}}$$

(since $\bar{\sigma}^*$ must have unit Frobenius norm).

This allows us to rewrite the (squared) worst-case principal stress computation as a maximization over \mathbf{n} only:

$$s_m^2 = \max_{\|\mathbf{n}\|^2=1} [\mathbf{nn}^T] : \underbrace{F : F^T}_{T^M} : [\mathbf{nn}^T], \quad (31)$$

which is a symmetric rank 4 tensor eigenvalue problem. It is straightforward to show that it is maximized by a \mathbf{n}^* such that

$$\mathbf{n}^* \cdot T^M : [\mathbf{n}^* \mathbf{n}^{*T}] = s_m^2 \mathbf{n}^*. \quad (32)$$

B EIGENVALUE DERIVATIVES

Assume matrix A has a non-repeated maximum eigenvalue λ with corresponding unit eigenvector v . Then

$$\dot{\lambda} = \frac{d}{dt}(v^T A v) = v^T \dot{A} v + 2\lambda v^T \dot{v} = v^T \dot{A} v,$$

where we used the fact that $v(t)$ is a unit vector ($v^T \dot{v} = 0$). A similar argument holds for rank four tensor T with non-repeated maximum eigenvalue λ and corresponding unit eigenvector n :

$$\dot{\lambda} = \frac{d}{dt}(T_{ijkl} n_i n_j n_k n_l) = \dot{T}_{ijkl} n_i n_j n_k n_l + 4\lambda \dot{n}_i n_i = \dot{T}_{ijkl} n_i n_j n_k n_l.$$

In practice, elements with high stress have one dominant maximum eigenvalue (Figure 20), so we do not need to worry about a repeated eigenvalue breaking our L_p objective's differentiability




Fig. 20. Regions of high worst-case stress (left) have the greatest difference in eigenvalues (1.0, on right).

ACKNOWLEDGEMENTS

We thank Bob Kohn for illuminating discussions, Francisca Gil-Ureta for helping create the paper figures, and Chelsea Tymms for help with printing. We also thank Yu Zhang and Minglei Zhao (NYU College of Dentistry) for help with compression testing.

REFERENCES

2008. Minimum stress optimal design with the level set method. *Engineering Analysis with Boundary Elements* 32, 11 (2008), 909–918. <https://doi.org/10.1016/j.enganabound.2007.05.007> Shape and Topological Sensitivity Analysis: Theory and Applications.
- Grégoire Allaire. 2002. *Shape optimization by the homogenization method*. Vol. 146. Springer.
- G Allaire, F Jouve, and H Maillot. 2004. Topology optimization for minimum stress design with the homogenization method. *Structural and Multidisciplinary Optimization* 28, 2-3 (2004), 87–98.
- Pierre Alliez, Clément Jamin, Laurent Rineau, Stéphane Tayeb, Jane Tournois, and Mariette Yvinec. 2016. 3D Mesh Generation. In *CGAL User and Reference Manual* (4.9 ed.). CGAL Editorial Board. http://doc.cgal.org/4.9/Manual/packages.html#PkgMesh_3Summary
- Erik Andreassen, Boyan S. Lazarov, and Ole Sigmund. 2014. Design of manufacturable 3D extremal elastic microstructure. *Mechanics of Materials* 69, 1 (2014), 1–10. <https://doi.org/10.1016/j.mechmat.2013.09.018>
- Loïc Barthe, Brian Wyvill, and Erwin De Groot. 2004. Controllable Binary CSG operators for “soft objects”. *International Journal of Shape Modeling* 10, 02 (2004), 135–154.
- Martin Philip Bendsoe. 1989. Optimal shape design as a material distribution problem. *Structural optimization* 1, 4 (1989), 193–202.
- Martin Philip Bendsoe and Ole Sigmund. 2003. *Topology optimization: theory, methods and applications*. Springer.
- Adrien Bernhardt, Loïc Barthe, Marie-Paule Cani, and Brian Wyvill. 2010. Implicit blending revisited. In *Computer Graphics Forum*, Vol. 29. Wiley Online Library, 367–375.
- Bernd Bickel, Moritz Bäcker, Miguel A. Otaduy, Hyunho Richard Lee, Hanspeter Pfister, Markus Gross, and Wojciech Matusik. 2010. Design and Fabrication of Materials with Desired Deformation Behavior. *ACM Trans. Graph.* 29, 4, Article 63 (July 2010), 10 pages. <https://doi.org/10.1145/1778765.1778800>
- James F Blinn. 1982. A generalization of algebraic surface drawing. *ACM transactions on graphics (TOG)* 1, 3 (1982), 235–256.
- Jules Bloomenthal. 1997. Bulge elimination in convolution surfaces. In *Computer Graphics Forum*, Vol. 16. Wiley Online Library, 31–41.
- Tiemo Bückmann, Nicolas Stenger, Muamer Kadic, Johannes Kaschke, Andreas Frölich, Tobias Kennerknecht, Christoph Eberl, Michael Thiel, and Martin Wegener. 2012. Tailored 3D Mechanical Metamaterials Made by Dip-in Direct-Laser-Writing Optical Lithography. *Advanced Materials* 24, 20 (2012), 2710–2714.
- Andrej Cherkav. 2000. *Variational methods for structural optimization*. Vol. 140. Springer.
- Doina Cioranescu and Patrizia Donato. 1999. *An introduction to homogenization*. Oxford University Press.
- Andrew T Gaynor and James K Guest. 2016. Topology optimization considering overhang constraints: Eliminating sacrificial support material in additive manufacturing through design. *Structural and Multidisciplinary Optimization* 54, 5 (2016), 1157–1172.
- Olivier Gourmel, Loïc Barthe, Marie-Paule Cani, Brian Wyvill, Adrien Bernhardt, Mathias Paulin, and Herbert Grasberger. 2013. A gradient-based implicit blend. *ACM Transactions on Graphics (TOG)* 32, 2 (2013), 12.
- G. N. Greaves, A. L. Greer, R. S. Lakes, and T. Rouxel. 2011. Poisson’s ratio and modern materials. *Nature Materials* 10, 11 (2011), 823–837. <https://doi.org/10.1038/nmat3134>
- Zvi Hashin and Shmuel Shtrikman. 1963. A variational approach to the theory of the elastic behaviour of multiphase materials. *Journal of the Mechanics and Physics of Solids* 11, 2 (1963), 127–140.
- Jonathan Hiller and Hod Lipson. 2009. Design and analysis of digital materials for physical 3D voxel printing. *Rapid Prototyping Journal* 15, 2 (2009), 137–149.
- Scott J. Hollister. 2005. Porous scaffold design for tissue engineering. *Nature Materials* 4, 7 (2005), 518–524. <https://doi.org/10.1038/nmat1421>
- Steven G. Johnson. 2016. The Nlopt nonlinear-optimization package. (2016). <http://ab-initio.mit.edu/nlopt>
- Hee Suk Kang. 2010. *Hierarchical design and simulation of tissue engineering scaffold mechanical, mass transport, and degradation properties*. Ph.D. Dissertation. The University of Michigan.
- Lily Kharevych, Patrick Mullen, Houman Owahdi, and Mathieu Desbrun. 2009. Numerical Coarsening of Inhomogeneous Elastic Materials. *ACM Trans. Graph.* 28, 3, Article 51 (July 2009), 8 pages. <https://doi.org/10.1145/1531326.1531357>
- Dieter Kraft. 1994. Algorithm 733: TOMP—Fortran modules for optimal control calculations. *ACM Transactions on Mathematical Software (TOMS)* 20, 3 (1994), 262–281.
- Gerhard Kreiselmeier and Reinhold Steinhauser. 1983. Application of vector performance optimization to a robust control loop design for a fighter aircraft. *Internat. J. Control* 37, 2 (1983), 251–284.
- Matthijs Langelaar. 2016. An additive manufacturing filter for topology optimization of print-ready designs. *Structural and Multidisciplinary Optimization* (2016), 1–13.
- Timothy Langlois, Ariel Shamir, Daniel Dror, Wojciech Matusik, and David IW Levin. 2016. Stochastic structural analysis for context-aware design and fabrication. *ACM Transactions on Graphics (TOG)* 35, 6 (2016), 226.
- Chia-Ying Lin, Chun-Ching Hsiao, Po-Quan Chen, and Scott J Hollister. 2004a. Interbody fusion cage design using integrated global layout and local microstructure topology optimization. *Spine* 29, 16 (2004), 1747–1754. PMID: 15303018.
- Cheng Yu Lin, Noboru Kikuchi, and Scott J. Hollister. 2004b. A novel method for biomaterial scaffold internal architecture design to match bone elastic properties with desired porosity. *Journal of Biomechanics* 37, 5 (2004), 623–636. <https://doi.org/10.1016/j.jbiomech.2003.09.029>
- Hongwei Lin, Yunyang Xiong, and Hongwei Liao. 2014. Semi-structured B-spline for blending two B-spline surfaces. *Computers & Mathematics with Applications* 68, 7 (2014), 706–718.
- Jonàs Martínez, Jérémie Dumas, and Sylvain Lefebvre. 2016. Procedural Voronoi foams for additive manufacturing. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 44.
- Graeme W Milton. 2002. *The theory of composites*. Cambridge University Press.
- PH Nakasone and ECN Silva. 2010. Dynamic design of piezoelectric laminated sensors and actuators using topology optimization. *Journal of Intelligent Material Systems and Structures* 21, 16 (2010), 1627–1652.
- 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.
- Alexander Pasko and Valery Adzhiev. 2002. Function-based shape modeling: mathematical framework and specialized language. In *International Workshop on Automated Deduction in Geometry*. Springer, 132–160.
- Galina I Pasko, Alexander A Pasko, and Tosiyasu L Kunii. 2005. Bounded blending for function-based shape modeling. *IEEE Computer Graphics and Applications* 25, 2 (2005), 36–45.
- Xiaoping Qian. 2016. Undercut and overhang angle control in topology optimization: a density gradient based integral approach. *Internat. J. Numer. Methods Engrg.* (2016).
- Inigo Quilez. smooth minimum. (????). <http://www.iquilezles.org/www/articles/smin/smin.htm>
- Alyn P Rockwood. 1989. The displacement method for implicit blending surfaces in solid models. *ACM Transactions on Graphics (TOG)* 8, 4 (1989), 279–297.
- Malcolm Sabin. 1968. The use of potential surfaces for numerical geometry. *British Aircraft Corporation, Weybridge, UK, Technical Report No. VTO/MS/153* (1968).
- Christian Schumacher, Bernd Bickel, Jan Rys, Steve Marschner, Chiara Darai, and Markus Gross. 2015. Microstructures to control elasticity in 3D printing. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 136.
- J. Schwerdtfeger, F. Wein, G. Leugering, R. F. Singer, C. KÄurner, M. Stingl, and F. Schury. 2011. Design of Auxetic Structures via Mathematical Optimization. *Advanced Materials* 23, 22 (2011), 2650–2654. <https://doi.org/10.1002/adma.201004090>
- Mélina Skouras, Bernhard Thomaszewski, Stelian Coros, Bernd Bickel, and Markus Gross. 2013. Computational design of actuated deformable characters. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 82.
- FO Sonmez. 2009. Optimal shape design of shoulder fillets for flat and round bars under various loadings. *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science* 223, 8 (2009), 1741–1754.
- Ondrej Stava, Juraj Vanek, Bedrich Benes, Nathan Carr, and Radomir Mvech. 2012. Stress Relief: Improving Structural Strength of 3D Printable Objects. *ACM Trans. Graph.* 31, 4, Article 48 (July 2012), 11 pages.
- Salvatore Torquato. 2002. *Random heterogeneous materials: microstructure and macroscopic properties*. Vol. 16. Springer.
- Laurent Van Mieghroet and Pierre Duysinx. 2007. Stress concentration minimization of 2D fillets using X-FEM and level set description. *Structural and Multidisciplinary Optimization* 33, 4-5 (2007), 425–438.
- Janos Vida, Ralph R Martin, and Tamas Varady. 1994. A survey of blending methods that use parametric surfaces. *Computer-Aided Design* 26, 5 (1994), 341–365.
- Qi Xia, Tielin Shi, Shiyuan Liu, and Michael Yu Wang. 2012. A level set solution to the stress-based structural shape and topology optimization. *Computers & Structures* 90 (2012), 55–64.
- Cédric Zanni, Adrien Bernhardt, Maxime Quiblier, and M-P Cani. 2013. SCALE-invariant Integral Surfaces. In *Computer Graphics Forum*, Vol. 32. Wiley Online Library, 219–232.
- Cédric Zanni, Michael Gleicher, and M-P Cani. 2015. N-ary implicit blends with topology control. *Computers & Graphics* 46 (2015), 1–13.
- Haiming Zhao, Weiwei Xu, Kun Zhou, Yin Yang, Xiaogang Jin, and Hongzhi Wu. 2016. Stress-Constrained Thickness Optimization for Shell Object Fabrication. In *Computer Graphics Forum*. Wiley Online Library.
- Qingnan Zhou, Julian Panetta, and Denis Zorin. 2013. Worst-case structural analysis. *ACM Trans. Graph.* 32, 4, Article 137 (July 2013), 12 pages.
- Yahan Zhou, Evangelos Kalogerakis, Rui Wang, and Ian R Grosse. 2016. Direct shape optimization for strengthening 3D printable objects. In *Computer Graphics Forum*, Vol. 35. Wiley Online Library, 333–342.