Speedy Stress-Scaled Sheet Structure Synthesis

Dennis R. Bukenberger[*]
Technical University Munich
Munich, Germany
dennis.bukenberger@tum.de

Jun Wu
Delft University of Technology
Delft, The Netherlands
j.wu-1@tudelft.nl

Rüdiger Westermann
Technical University Munich
Munich, Germany
westermann@tum.de

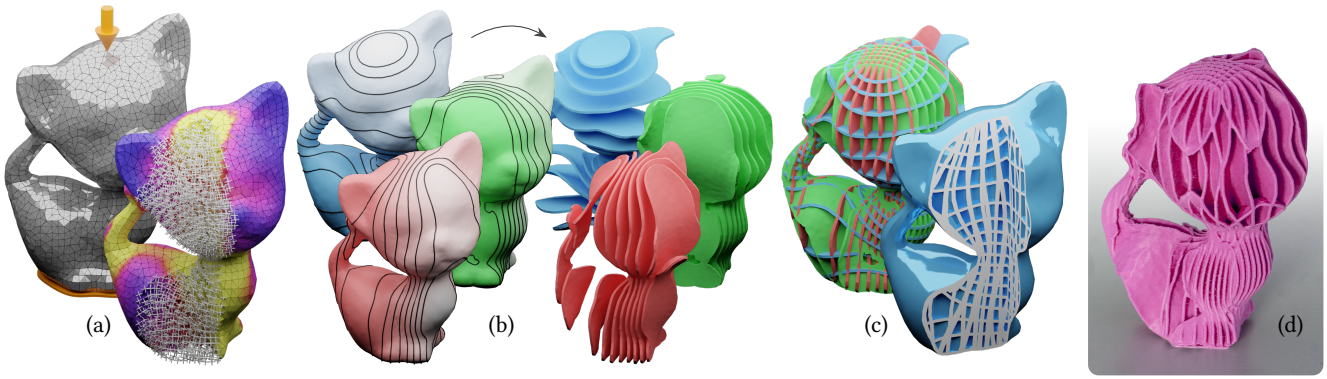

Figure 1: *Generating lightweight and stiff infill sheet structures in seconds: (a) Loading conditions on the solid object, represented with a tet mesh, and 3D stress field. (b) Scalar stress potential fields are recovered from the principal stress components and isocontours consuming a specified fraction of the solid's volume are extracted. (c) The resulting infill design exhibits improved stiffness compared to alternative methods (Figure 2). (d) Implicit modulations of the guiding field allow for the extraction of print-ready self-supporting infill structures (shown without the hull).*

Abstract

This paper presents a highspeed method for constructing load-resistant infill patterns for additively manufactured parts. Owing to its computational efficiency, the method supports interactive design editing and optimization and can be integrated as an alternative infill-generation strategy in 3D printer slicer software. Due to its implicit formulation, it further facilitates the incorporation of manufacturing constraints, such as overhang angle limits or user-defined design preferences. The method constructs a sheet structure within 3D objects whose morphology dynamically adapts to the stress field induced by a specified load setup. Starting from a stress-aligned tensor field, stress potential fields are recovered as an implicit representation of the structural design. Modulating these potential fields enables the generation of stress-specific material concentrations, improving overall mechanical performance. The final sheet structures are extracted as isosurfaces, with spacing and thickness adjusted to satisfy a prescribed material budget. Results demonstrate superior stiffness and substantially lower construction times compared to common sheet- or graph-based infill designs. The 3D printability of the generated designs is validated through the physical fabrication of multiple parts, confirming that all geometries are self-supporting and free of manufacturing artifacts.

[*]Also with Delft University of Technology.

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.
SCF '26, Tokyo, Japan
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2821-1/2026/12
https://doi.org/10.1145/3828611.3845636

CCS Concepts

• **Applied computing** → **Computer-aided design**; • **Computing methodologies** → **Mesh geometry models**; **Parametric curve and surface models**.

Keywords

Volumetric Models, Shape Analysis, 3D Printing

ACM Reference Format:
Dennis R. Bukenberger, Jun Wu, and Rüdiger Westermann. 2026. Speedy Stress-Scaled Sheet Structure Synthesis. In *Proceedings of the 11th ACM Symposium on Computational Fabrication (SCF '26), December 06–08, 2026, Tokyo, Japan.* ACM, New York, NY, USA, 16 pages. https://doi.org/10.1145/3828611.3845636

1 Introduction

In additive manufacturing (AM), infill configuration is a pivotal prefabrication step, typically performed interactively within the slicing software's user interface (UI). The infill choice is crucial because the internal structure directly determines the part's final mechanical performance and material consumption. Conventional slicing methods predominantly rely on uniform geometric patterns for infill designs, ranging from regular grids over honeycomb structures to continuous fractal space-filling curves. Enhancing the mechanical properties of a component typically necessitates higher infill volume, inducing a direct trade-off between structural integrity and material use.

Designing lightweight designs reducing material usage while maintaining their stiffness under mechanical stress therefore is a multi-objective problem: internal layouts must provide sufficient stiffness

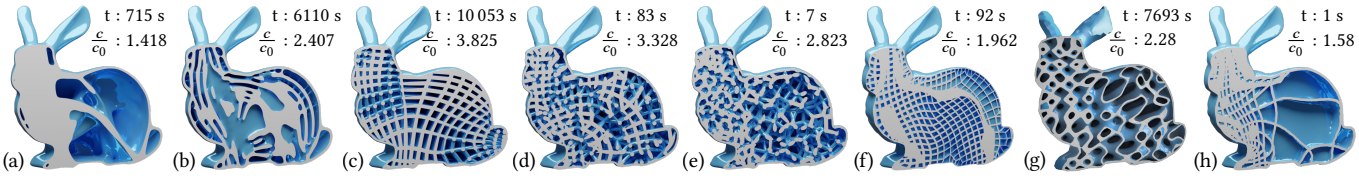

Topology Opt. [2025]	Porous Infill [2018] (wall)	Michell Truss [2019] (beam)	Conform. Lattice [2021b] (beam)	Graded Voronoi [2025] (beam)	Hex. Lattice [2024] (wall)	TPMS [2022] (wall)	S^6 [Ours] (wall)

Figure 2: *The bunny is fixed at the bottom and a down force acts on its head. Shown examples consume 45 % of the solid object's volume. Numbers give the computation time in seconds and the relative compliance compared to the solid object (c_0), thus the lower the better.*

and robustness under expected loading conditions and a given material budget, remain manufacturable on consumer-grade hardware (respecting overhang constraints) and be configurable within an interactive slicing workflow. In practice, preparing a print job in the slicer UI involves interactively tweaking parameters to satisfy user and hardware constraints, often requiring multiple slicing iterations. Existing pattern-based infills offer simplicity and predictable printing behavior but typically lack stress awareness. They are unable to adapt their material distribution to local load demands and can be either overly conservative (wasting material) or insufficiently robust in critical regions.

Our goal is to provide a fast infill design method that yields space-filling internal structures with a favorable stiffness-to-volume ratio, is directly printable without additional geometric post-processing and supports interactive user modifications. Therefore, we employ lightweight geometric operations that translate stress tensor fields into domain-filling infill layouts. The resulting stress-driven material distributions capture essential structural logic of mechanically sensible designs while remaining compatible with standard slicing workflows. Furthermore, our concept exposes parameters that allow users to balance load-tailored structural performance against global resilience, adjust infill volumes, wall thicknesses and maximum overhang angles.

We present S^6, a novel interactive method and tool for generating stress-adaptive, self-supporting, space-filling infill patterns on-the-fly. Within the slicer context, S^6 enables users to explore feasible design states under different mechanical and manufacturing constraints, providing a practical bridge between stress-aware structural design and everyday AM workflows.

Our specific contributions can be summarized as:

- A smoothing- and optimization-based scheme that reduces singularities while preserving a smooth, best-compromise tensor field for guiding the structure synthesis.
- A novel concept that recovers scalar fields from the stress tensor fields and extracts their isocontours as load-bearing sheets, with field modulation to tune between stiffness and resilience.
- Stress-aligned structures outperforming related methods in terms of mechanical performance, while exhibiting higher robustness to variations in material properties and loading conditions.
- Enforcing 3D-printing constraints via projection in the implicit field modulation stage, yielding a pipeline that outputs print-ready designs compatible with consumer-grade printers.
- A computationally lightweight pipeline that enables interactive structural design workflows and seamless integration into 3D printer slicer software.

2 Related Work

Lightweight Design. One way to generate lightweight designs is with topology optimization (TO), which optimizes a given material budget within a design space under specified mechanical constraints. In its simplest form, density-based TO [Bendsøe and Kikuchi 1988; Bendsøe and Sigmund 2013; Sigmund 1994, 2001] relies on a first-order finite element method (FEM) discretization of the design space. The material is iteratively distributed in this space using a gradient-based strategy to achieve the stiffest structure. This process involves repeated solutions of sparse linear systems, and even with performance-tuned implementations [Wang et al. 2025] and for medium-resolution 3D problems interactive rates cannot be achieved. Furthermore, the core idea of density-based TO relies on a theoretically unconstrained design space. Practically, however, the design domain is usually limited to a given shape. When employed as an infill design strategy, TO can also incorporate the hull of this given input shape (Figure 2a), yielding a few strong load-bearing paths plus thinner structures that support the hull. Under ideal conditions, this is a highly effective strategy, but it is vulnerable to local defects or alterations of the loading conditions [Wu et al. 2018]: any crack or flaw in a critical branch can significantly compromise overall performance or lead to structural failure. Recent work has started to incorporate manufacturability constraints directly into TO to better align it with industrial AM workflows, but the added constraints typically introduce further computational overhead [Liu et al. 2018].

Space Filling Infill Designs. Greater resilience to varying loading conditions or structural defects is offered by microstructures or porous infill (PI) materials (Figure 2b). This motivated many recent concepts using multi-scale structures or lattice-based infills to materialize the optimized density field into binary material maps [Wu et al. 2021a]. Other concepts also perform optimization of conventional infills patterns [Wang et al. 2023b; Wu et al. 2016]. Such grid-like structures show improved resilience to changing loading scenarios [Wang et al. 2025], a crucial trait for applications in engineering. Many common concepts are well understood and established in 2D, but direct translations to 3D are non-trivial. Simple extensions of 2D edge structures into 3D often yield *graph-* or *beam*-based networks (Figures 2c, 2d, 2e) that do not fully exploit the advantages of volumetric designs. Therefore, recent trends favor *sheet-* or *wall*-like structures [Bukenberger et al. 2024; Stutz et al.

2022] due to their improved bending and shear stiffness at a given material budget (Figure 2f). Examples include honeycomb-like structures [Lu et al. 2014; Zhang et al. 2015], porous infills [Wu et al. 2018], and triply periodic minimal surfaces (TPMS), which provide smooth, continuous space-filling geometries (Figure 2g). However, the inherently periodic nature and barely changeable surface orientation of TPMS make them poorly suited for stress-adaptive designs, where local alignment with principal stress directions is desirable [Hu et al. 2022]. Recent advancements also address manufacturability constraints [Zheng et al. 2023] or involve iterative optimization routines [Ning et al. 2026] but list computation times of over 100 hours for some designs.

Although infill design and TO serve different purposes - infill design focuses on domain-filling internal structures with a preserved outer shell whereas TO usually targets globally optimal load paths — the compliance of TO results is often used as a performance reference for infill approaches.

Field Aligned Meshing. Classic Michell theory [Michell 1904] states that minimum-volume truss structures are obtained by arranging members along the principal stress directions. In such layouts, each member carries only axial tension or compression, eliminating bending and thereby achieving the prescribed load-bearing capacity with the least material. Since the principal stress directions form smooth families of curves, Michell theory naturally motivates stress-field-aligned layouts such as trusses, lattices, and sheet structures in our setting. Meshing methods focused on direction-fields such as Mixed-Integer Quadrangulation [Bommes et al. 2009] construct smooth cross fields from sparse directional constraints via a global smoothness energy, then derive field-aligned parameterizations and meshes. A recent approach for fabric formwork [Sati et al. 2026] picks up this concept for strain field segmentation. They formulate a smoothing energy also promoting alignment, and use the resulting guidance field to organize fabric patches and seams, without eliminating all singularities. In contrast, our method separates these steps: we first deliberately oversmooth the stress-derived direction field to obtain a robust guidance field and in a second stage realign and modulate it to balance stress alignment with robustness or manufacturing constraints. In 2D, the orthogonal principal stress directions layouts naturally form quad meshes [Wang et al. 2022]. An analog for 3D scenarios are hexahedral meshes, which are notorious for the complications that arise during their field alignment [Pietroni et al. 2022]: Some forms of degenerate points in the tensor field, called singularities, can simply not be resolved with hexahedra. Thus, common hexahedral meshing approaches require manual guidance for the decomposition [Nieser et al. 2011], introduce non-hexahedral elements [Bukenberger et al. 2022; Gao et al. 2017], smoothen the field or revert to edge-graph based structures [Arora et al. 2019; Bukenberger et al. 2024; Wu et al. 2021b].

Manufacturability. Current research in structural design increasingly focuses on manufacturability and robustness [Liu et al. 2018; Mirzendehdel and Suresh 2016; Qian 2017; Wu et al. 2016; Zheng et al. 2024]. Worst-case analyses are used to improve stiffness and reduce stress concentrations, which often lead to the failure of additively manufactured parts [Liu et al. 2022; Panetta et al. 2017]. Other concepts integrate manufacturing constraints such as overhang angles, length limits and build direction directly into the TO process to create self-supporting structures [Wang et al. 2023a; Wu et al. 2016]. Advancements in manufacturing hardware also gave rise to new design options like multi-axis slicing [Liu et al. 2024; Zhang et al. 2022] or mechanically adaptive metamaterials [Han et al. 2024]. However, our method aims for compatibility with currently available consumer 3D printing technologies and delivers readily printable results without additional post-processing or manual support structure configuration.

Isocontouring. While isocontouring is a standard technique for visualizing scalar fields, there is, to the best of our knowledge, no prior work employing isocontours themselves to generate the final load-bearing structural system. Isocontours are curves or surfaces of constant value revealing salient features of an underlying function; in our case stress potentials, recovered from the stress tensor field. In density-based TO, the final design is typically obtained as the solid region defined by an isocontour extracted from the optimized material density field [Bendsøe 1989]. I.e., marching cubes [Lorensen and Cline 1987] can resolve the solid-void ambiguity inherent in intermediate density fields [Bendsøe and Sigmund 2013]. In the 2D isoline topology design (ITD) method [Victoria et al. 2009], the structural layout similarly emerges as the region enclosed by an isoline of a stress magnitude scalar field. In homogenization-based TO, a key open challenge is to convert a homogenized design into a manufacturable, single-scale solid-void structure. For 2D fields, a proposed projection strategy uses scalar stress fields to modulate periodic functions, thereby generating stress-aligned lattices [Groen and Sigmund 2018], similar to our example discussed in Appendix A. In a concept for advanced multi-axis printing, isocontours provided the basis for toolpath optimization, where isosurfaces of a governing field define curved working surfaces for aligned filament deposition [Fang et al. 2020]. Our work builds on this lineage by treating isocontours of stress-derived scalar fields not just as visualization devices, but as the scaffolding for a space-filling infill layout that can be generated and updated interactively (Figure 2h).

3 Background

As input we consider a tessellated object with vertices V and triangles or tetrahedra T for the 2D or 3D case, respectively. Boundary conditions are specified as sets of vertices with all (or some) degrees of freedom fixed and sets of vertices where force vectors are applied. Other forms of background meshes such as Cartesian voxel grids are also possible, our concept translates analogously from triangular to quadrangular or tetrahedral to hexahedral elements.

3.1 Stiffness

The stiffness of a design is measured, based on its inverse: compliance, the amount of deformation under load. Minimizing compliance therefore gives stiffer structures. It computes as

$$c = \frac{1}{2}\mathbf{U}^T\mathbf{K}\mathbf{U} \tag{1}$$

with the global stiffness matrix $\mathbf{K}$ assembled from the element stiffness matrices, the external force vector $\mathbf{F}$ and the displacement

vector $\mathbf{U}$ resulting from

$$\mathbf{KU} = \mathbf{F}. \tag{2}$$

The objective is to minimize compliance, while adhering to a constrained total amount of material (volume). As the compliance value is not invariant of the object's scale, structural performance is measured as $\frac{c}{c_0}$, relative to the fully solid object's compliance c_0. Further is the volume fraction $\alpha : \frac{v}{v_0}$ expressed as fraction of the optimized volume v and the original solid volume v_0.

3.2 Stress Field

With the given mesh and boundary conditions, standard FEM [Gustafsson and McBain 2020] gives deformation vectors per vertex and *Cauchy* stress tensors $\mathbf{C}$ per element, then projected to each vertex $\mathrm{v}_i \in V$. The eigenbases of tensors $\mathbf{C}$ give the orthogonal principal stress tensors Ξ. While isolated degenerate or nearly isotropic stress states may occur, they are typically confined to small regions or single points. In our pipeline, the subsequent smoothing of Ξ regularizes these cases and yields a practically coherent field. Eigenvalues are interpreted as local principal stresses and are used for computing the scalar *von Mises* stress norm σ_v, an established heuristic for stress magnitudes. The 2D case trivially follows by using triangles instead of tetrahedra.

3.3 Additive Manufacturing

Slicer software UIs for AM typically provide many adjustable parameters, allowing users to fine-tune underlying algorithms to meet specific manufacturing constraints: *Fused filament fabrication (FFF)* is one of the most popular and cost-efficient technologies for 3D printing. Therein, thermoplastic filaments are deposited layer-by-layer through a heated nozzle to build the object. One key limitation is the filament extrusion width, which constrains the minimum printable layer height and line width. To ensure structural integrity, walls must have a minimum thickness τ, being an integer multiple of the extrusion width. This guarantees, that each wall consists of several adjacent filament lines rather than single, non-load-bearing strands. Another limitation are steep overhang angles, i.e., new layers would be deposited mid-air as there is no more layer below supporting it. These constraints are addressed in Sections 4.2 & 4.3, with further experiments discussed in Section 5.2 and 5.3. With *stereolithography (SLA)*, objects are manufactured by incrementally raising a build plate out of a vat of liquid photopolymer resin, selectively cured layer-by-layer with a precise laser. As partially formed structures remain submerged in resin, a critical design constraint is the avoidance of hollow structures that can form suction cups - closed cavities concealing resin during platform elevation. This requirement is difficult to satisfy with any wall- or sheet-based infill strategy, which inherently tend to form closed cellular structures. Nevertheless, in Section 5.2 we demonstrate how our method can also produce SLA-compatible high-performance graph structures.

4 Method

Our lightweight structural design method builds on Michell theory's field-aligned principles, while balancing stiffness-oriented material distribution with sparse architectures for resilience. To achieve this, our method implements a pipeline with three principle steps:

4.1 Smoothing and Optimization (Figure 4)
4.2 Field reconstruction and Iso-Contouring (Figures 5 & 6)
4.3 Structure Synthesis and Extraction (Figure 9)

4.1 Field Preparation

For clearer visualization, we use a 2D example to illustrate concepts introduced later, though the full pipeline applies equally to 2D and 3D. The following examples use the 2D bunny shown in Figure 3, fixed at the bottom with two forces applied to its head and back. From a single FEM simulation on the solid, we obtain the principal stress tensor field. This field is used to define the final structure layout, a key difference from methods that alternate numerical simulations and optimization. This shortcut is motivated by the observation that principal stress directions in an isotropic solid closely match those of an equivalent orthotropic solid with aligned material axes. Thus, optimized structures largely follow patterns already evident in the initial simulation. However, since the assumption that the once-computed stress field fully defines the final optimized structure is not universally valid, i.e., the optimization process can alter the stress field's topology, our method cannot, in general, produce structures as stiff as those produced with TO.

Geometric Smoothing. Singularities, i.e., so called degenerate points [Delmarcelle and Hesselink 1994] (or degenerate regions in 3D domains [Hergl et al. 2021; Hung et al. 2023; Wang et al. 2020]) where the stress tensor has repeating eigenvalues, make the principal stress directions indeterminate. This prohibits the computation of consistent domain parameterizations [Stutz et al. 2020] and complicates the recovery of a smooth potential field, as required by our approach. Thus, in a first step the initial stress field is smoothed, largely removing such structures. We call this process oversmoothing. In a second step, a field as close as possible to the initial one is restored, by considering the initial and oversmoothed field.

First, the symmetric field tensors $\mathbf{C}$ are converted using the reversible half-vectorization operator $\mathrm{vech} : \mathbb{R}^{n\times n} \rightarrow \mathbb{R}^{n\frac{n+1}{2}}$, i.e., for $\mathbb{R}^3$ it is defined as

$$\mathrm{vech}\begin{pmatrix} \mathrm{a} & \mathrm{f} & \mathrm{e} \\ \mathrm{f} & \mathrm{b} & \mathrm{d} \\ \mathrm{e} & \mathrm{d} & \mathrm{c} \end{pmatrix} = [\mathrm{a,b,c,d,e,f}]^T. \tag{3}$$

This stacks solely the unique entries of a symmetric tensor into a vector shape. For simplified notation, we use $\overline{\mathbf{X}} = \mathrm{vech}\,(\mathbf{X})$ and may refer to either as tensor. Since vech is a linear isomorphism from the space of symmetric matrices $\mathbb{R}^{n\times n}$ to $\mathbb{R}^{n\frac{n+1}{2}}$, employing the $\overline{\mathbf{C}}$ representation for smoothing is equivalent to smoothing the symmetric tensor field under the constraint that $\mathbf{C} = \mathbf{C}^T$.

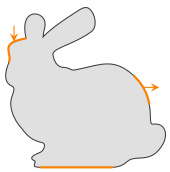

Figure 3: *This 2D bunny serves as an example object for upcoming illustrations, using the loading conditions (shown here in orange): A fixation at the bottom and two forces, one acting downwards on its head and one pulling to the right on its back.*

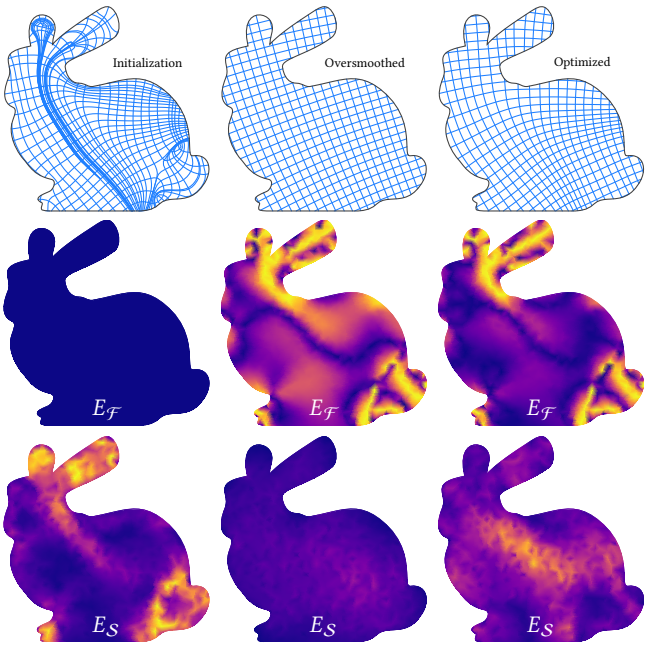

Figure 4: *Top: Principal stress lines in the initial, oversmoothed and optimized stress fields. Middle & Bottom: $\mathcal{F}$ Fidelity (deviation from input) and $\mathcal{S}$ smoothness energy of the stress fields (per vertex, scaled for visualization).*

Smoothing is then realized by employing the implicit Laplacian method with the n-dimensional [Crane 2019] cotan Laplacian matrix $\mathbf{L}$ and mass matrix $\mathbf{M}$. Following standard heat-diffusion heuristics [Crane et al. 2013], we choose λ proportional to the squared mean edge length of the background tessellation, linking the smoothing strength to the mesh resolution. The smoothed field tensors $\overline{\mathbf{C}'}$ then result by solving the well-posed linear system

$$(\mathbf{M} - \lambda \mathbf{L})\,\overline{\mathbf{C}'} = \mathbf{M}\overline{\mathbf{C}}. \tag{4}$$

The overall amount of smoothing is controlled by the combination of λ and the number of smoothing iterations. In our experiments, input tessellations are generated with a common target edge length using fTetWild [Hu et al. 2020]. Further, we keep the number of iterations fixed (default $\#s : 10$) across all results so that the total smoothing strength remains comparable.

Objective and Optimization. When smoothing the stress field too much, field singularities are largely removed but at the cost of introducing deviation from the input field. To recover the initial field's orientation, we formulate an optimization realigning the smoothed field to the input. This optimization is an optional refinement of the standard pipeline, requiring a few additional seconds of computation but providing improved mechanical performance. The optimization is initialized with the original and oversmoothed tensors $\mathbf{C}$ and $\mathbf{C}'$, respectively. The objective function to be minimized consists of three components, elaborated in the following. First we recall that Ξ are the eigenbases of $\mathbf{C}$, thus consisting of orthonormal eigenvectors v_i. Then, with $v_i^{\otimes d}$ being a d-wise tensor power, we can use the orthogonal decomposable order-d tensor $\mathbf{T} = \sum_i v_i^{\otimes d}$. Tensor $\mathbf{T}'$ corresponds analogously to tensor $\mathbf{C}'$. This ODECO [Robeva 2016; Zhu et al. 2025] representation $\mathbf{T}$ is invariant to axis permutations, axis sign-flips or $\pm 90°$ rotations of Ξ. Using these unique Cartesian coefficients of the symmetric tensor and the corresponding Euclidean coefficient-space metric, we first formulate the fidelity-term, expressing the deviation of the smoothed field from the input as

$$E_{\mathcal{F}} = \sum_{\mathrm{v}_i \in V} \left(\overline{\mathbf{T}}_i - \overline{\mathbf{T}'}_i\right)^2. \tag{5}$$

Secondly, we express the smoothness of the tensor field as

$$E_{\mathcal{S}} = \sum_{\mathrm{v}_i \in V} \sum_{\mathrm{v}_j \in \mathcal{N}(i)} \frac{w_{ij}}{2} \left(\overline{\mathbf{T}'}_i - \overline{\mathbf{T}'}_j\right)^2, \tag{6}$$

where w_{ij} are standard n-dimensional cotan weights for edges connecting a vertex v_i with a direct neighbor v_j. Eventually, we also employ the term

$$E_{\mathcal{N}} = \sum_{\mathrm{v}_i \in V} \left| |\mathbf{C}_i| - |\mathbf{C}'_i| \right|, \tag{7}$$

to keep optimized field vectors at a reasonable scale. Otherwise, the optimization might converge on a bad local optimum by collapsing tensor components to zero length.

Putting all together, we formulate the objective function as

$$\underset{\mathbf{C}'_i}{\text{minimize}}\; w_{\mathcal{F}} E_{\mathcal{F}} + w_{\mathcal{S}} E_{\mathcal{S}} + w_{\mathcal{N}} E_{\mathcal{N}}, \tag{8}$$

with adjustable weights $w_{\mathcal{F}}, w_{\mathcal{S}}$ and $w_{\mathcal{N}}$. While $w_{\mathcal{F}}$ and $w_{\mathcal{N}}$ can be 1, we found $w_{\mathcal{S}} : 20$ to allow for the best compromise between field recovery and smoothness. The optimization problem is nonlinear due to the eigenvalue dependence of $\mathbf{C}$, and it is solved using the Adam optimizer [Kingma and Ba 2014] (see Section 5.4 for details). Figure 4 illustrates the process of oversmoothing and field recovery: The initial, intermediate and final stress fields are visualized with Principal Stress Lines, i.e. trajectories along the minor and major stress directions derived from the stress tensor. For the initial field (left column), the fidelity energy $E_{\mathcal{F}}$ is zero everywhere and the smoothness energy $E_{\mathcal{S}}$ still contains problematic regions. After smoothing (center column), the fidelity term highlights deviations whereas the smoothness term is almost zero. Eventually, the optimization result (right column) gradually balances out both terms.

Combing. The smoothed and optimized field of principal stress frames is only defined up to permutations and sign flips of the three orthogonal eigenvector directions. Therefore, a combing step is performed on the background mesh to obtain globally consistent orientations. Starting from an arbitrary seed element, neighboring elements are traversed in a flood-fill manner. This determines, for each visited neighbor, the permutation and sign configuration of the local eigenvectors that minimizes the rotation to already visited adjacent frames. Analogous to standard combing procedures for cross and frame fields in geometry processing libraries [Jacobson et al. 2018] and related frame field design methods [Corman and Crane 2019; Ray et al. 2016], this yields a combed frame field with consistently ordered directions.

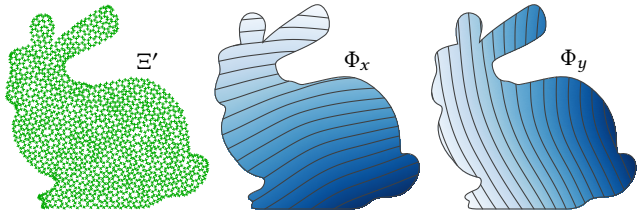

Figure 5: *x and y components of the cross field tensors Ξ_i are used to solve for the scalar field potentials Φ_x and Φ_y. The resulting scalar fields are shown as color gradients along with isocontours.*

4.2 Scalar Field Recovery

The combed directional components of Ξ', being the smoothed and optimized orthonormal cross field derived from the principal stress tensors of $\mathbf{C}'$, are now utilized to recover the underlying scalar potentials Φ [Crane et al. 2013]. This can be done by minimizing $\sum_{v_i \in V} \left|\nabla\Phi_i - \Xi'_i\right|^2$. However, with $\mathbf{D}$ being the tessellation's divergence matrix, we can equivalently [Schwarz 1995] also solve

$$-\mathbf{L}\Phi = \mathbf{D}\begin{bmatrix}\Xi'_0 & \Xi'_1 & \dots & \Xi'_{n-1}\end{bmatrix}^T. \tag{9}$$

Up to an additive constant, the solution for Φ is unique and for further handling can be shifted and scaled to unit range. Figure 5 illustrates the solved scalar potentials Φ_x and Φ_y corresponding to the x and y components of the cross field tensors Ξ', respectively.

Field Modulation. As described in the next section, the final geometry of the structural design is extracted as the isocontours of the scalar potential fields. Thus, while in this implicit state, we can modulate the field, to provoke certain features to emerge in the final structures. The isocontours on Φ in Figure 5 are quite evenly spaced since the components of the cross field Ξ' are all unit length. However, eventually these isocontours are thickened and interpreted as load bearing structures. Thus, for structural performance it is often beneficial to distribute more material to regions experiencing higher stress. We can achieve this, i.e., by introducing γ_i for locally scaling corresponding Ξ'_i vector components. Then, when interpreted as gradients in the field recovery, longer vectors cause steeper field potentials Φ and, thus, spatially compressed isocontours. Figure 6 illustrates an example using scalar stress indicator σ_v to modulate the cross field. Due to the steeper field deviations, this causes more densely spaced structures in regions with higher stress and fewer or none in regions with lower stress (bunny ears or tail). Effects of the compression on structural performance and robustness are illustrated in Figure 11 and analyzed in detail in Section 5.2.



Figure 6: *Local stress indicators σ_v are normalized and employed as γ to modulate field Ξ'. In high stress regions, this results in steeper Φ field deviations and compressed isocontours.*

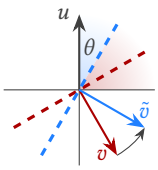

Figure 7: *The isocontour (dashed) corresponding to vector v forms too large an angle θ with the up-vector u. To reduce θ, v is therefore rotated onto a cone around u resulting in $\tilde{v}$.*

Besides influencing spatial density, this stage also allows for superimposing manufacturing constraints regarding overhang angles of the resulting sheets: For a surface to be 3D printed, we therefore define overhang as the angle θ enclosed by the surface and the printing direction u, usually the positive z-axis. As elaborated in the following section, sheets are eventually extracted as isocontours of the recovered scalar potential field. By definition, these isocontours transcend the field orthogonally to the unit vector components of Ξ', where the potential was recovered from, respectively. In Figure 7 such a vector component is shown as v with the corresponding isocontour as a dashed line. Thus, to correct for too large overhangs, all vectors $v \in \Xi'$ are restricted (if necessary) to a corresponding $\tilde{v}$. Affected vectors are therefore rotated onto a cone around u, now enclosing the reciprocal angle $\frac{\pi}{2} \pm \theta$ with u while maintaining their individual azimuth. Others, with angles between $\frac{\pi}{2} - \theta$ and $\frac{\pi}{2} + \theta$, remain unchanged. This rotation may introduce local deviations from the smoothed and optimized fields while still enabling robust sheet extraction and satisfying the manufacturing constraints.

Singularities and Integrability. The effect of residual singularities on sheet extraction is illustrated in Figure 8 for the 2D bunny and femur. In the bunny case, the remaining singularity after combing does not prevent coherent sheet extraction and isocontours can be constructed directly from the unsmoothed field. In contrast, the more pronounced singularity in the femur results in a non-integrable field and consequently in broken isocontours. These examples show that the presence of a singularity does not necessarily lead to a failure, but that sufficiently severe singularities can propagate into the final infill pattern. The proposed smoothing and optimization are therefore introduced as a robustness measure. While smoothing substantially regularizes the field and in practice removes the

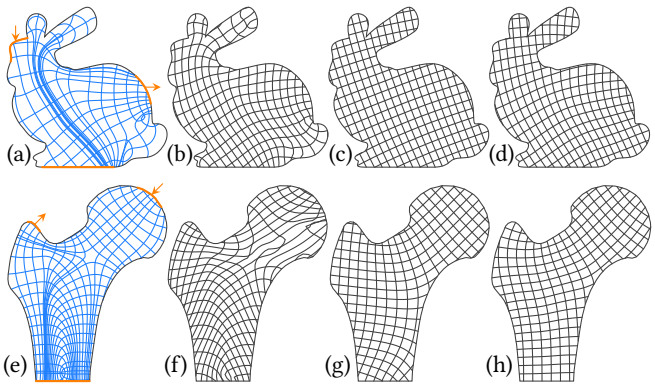

Figure 8: *Boundary conditions and principal stress lines on the 2D bunny and femur (a,e) featuring singularities. Sheet structures extracted from three cases: (b,f) unsmoothed field with a singularity, (c,g) smoothed field and (d,h) smoothed field after optimization.*

problematic singularities observed in our examples, it does not provide a theoretical guarantee that all singularities are eliminated. The optimization step then bends the smoothed field back toward the original input field while preserving its smoothness, thereby recovering alignment with the input field as far as possible without sacrificing the smoothness of the regularized field.

4.3 Geometry Generation

The recovered field Φ provides per vertex scalar values for the background mesh of the input object. First, sets of isocontours are extracted for the x, y and z components of Φ, respectively. This results in evenly spaced sheets, transcending the three dimensional object along critical stress directions. These sheets are eventually thickened to become the load bearing structure of the final design. Dependent on the type of tessellation one may use the marching tets [Treece et al. 1999] or cubes [Lorensen and Cline 1987] algorithm to extract triangulated surface sheets corresponding to certain isovalues.

Sheet Extraction. As field Φ may already incorporate compressive modulations, the isovalues are designed to be equally spaced $\in (0, 1)$. Therefore, the Φ field components are normalized to unit range and n gives the number of sheets. Isovalues are defined as

$$s_k = \frac{k}{n+1}, \; k = 1, \ldots, n. \tag{10}$$

This excludes the first ($= 0$) and last ($= 1$) isovalues: On a normalized field, their isocontour would coincide with the object boundary anyway. The number of sheets n can be set as an input parameter or derived from manufacturing constraints, i.e., considering the fixed and limited material budget, a minimum wall thickness and the absolute size of the final print.

Materialization. Extracted isocontours are simple manifold triangular mesh sheets. To generate the final load bearing geometry, every sheet is extruded by a small amount $\frac{\tau}{2}$ in positive as well as in negative local normal direction, thus resulting in a total sheet thickness τ. The object's hull is extruded by τ but only inwards. Therefore, we can modulate the volume of the final design with this parameter τ to meet a specific fraction α of the solid object's volume v_0. As visualized in Figure 9, we let $A_\triangle$ be the accumulated surface area of all extracted isocontour sheets as well as the outer boundary hull. Figure 10 exemplifies the merge of two simple extruded sheets in a boolean union operation. This highlights that estimating τ based on the desired target volume and extrudeable surface area alone ($\alpha v_0 = A_\triangle \cdot \tau$) does not account for sheets intersecting each other:

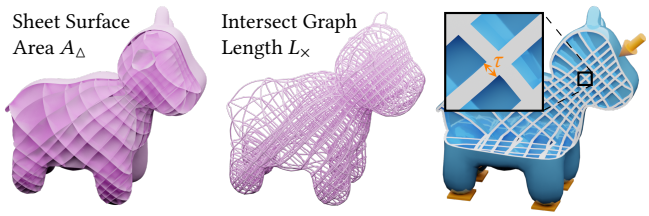

Figure 9: *Accumulated sheet surface area $A_\triangle$ and the length of their intersection graph $L_\times$ are used to compute the sheet thickness τ (inset) to meet a given volume fraction. The cut-open result on the right also shows initial fixations and acting forces.*

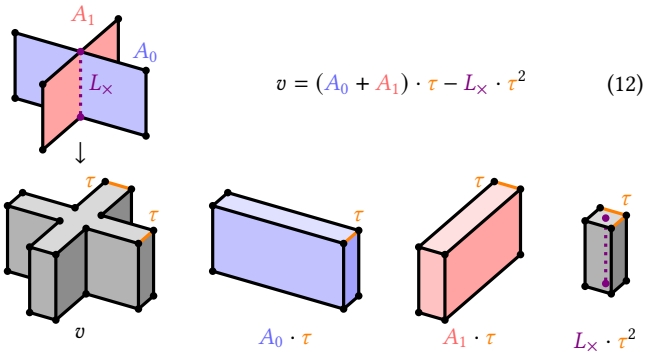

Figure 10: *Two simple sheets are thickened by τ and merged in a boolean union. Equation* (12) *gives the result volume, dependent on τ.*

small parts of extruded volume would naively be included twice. As expressed in Equation (12), the countermeasure is to subtract $L_\times$, as the total length of the intersection graph resulting from sheets of different sets crossing each other, times τ^2. This also includes the length of their open boundary, accounting for their intersection with the outer hull. Eventually, the simple quadratic Equation (12) with $v := \alpha v_0$ is solved as

$$\tau = \frac{A_\triangle - \sqrt{A_\triangle^2 - 4\,(L_\times \cdot \alpha v_0)}}{2 \cdot L_\times}. \tag{11}$$

While this very robustly meets the desired target volume with high enough precision (< 0.1 % relative error), it still neglects small portions of volume where three sheets intersect in one point. This can be mitigated by adding $N_\times \cdot \tau^3$ in Equation (12), with $N_\times$ as the number of three sheets intersecting and then solving the cubic term. We also found a small additional optimization of τ may further reduce the relative error by a factor of 10+ within only 5-6 steps if necessary.

As shown, extruded sheets and hull are eventually combined in a boolean union operation, resulting in a watertight manifold mesh. This is suitable for direct use in 3D printer slicer software or as base for tetrahedralization [Hu et al. 2020] or voxelization [Wang et al. 2025] for further FEM simulation tasks. An alternative geometry extraction approach and its drawback are discussed in Appendix A.

5 Results and Discussion

Our results are evaluated in comparison to other state-of-the-art concepts using FEM simulations with common standard parameters (Young's Modulus of 1000 and Poisson's ratio of 0.3). We utilize the SGLDBench framework [Wang et al. 2025] to compare our approach against established methods with available implementations and standard settings (voxel grid resolution of 2^8). Furthermore, our results are closed manifold surfaces and ready for analysis as tetrahedralization, i.e., using fTetWild [Hu et al. 2020] and standard FEM libraries [Gustafsson and McBain 2020]. We also fabricate representative designs and perform physical tests to validate the simulated behavior and confirm practical manufacturability. As implementations and datasets of related methods are rarely made available, we reproduced the reported boundary conditions and compared our results with the published numbers. If not explicitly stated otherwise, results include the object's boundary surface.

5.1 Mechanical Performance

Table 1 lists computation times in seconds and relative compliance values (the lower the better) of different models in comparison to related methods. Therein, density based TO is employed as a baseline for optimal mechanical performance, although being known for representing "only" a local optimum. Results produced while incorporating overhang constraints for additive manufacturing are marked in the ≧̃ column. What clearly stands out from these numbers is that our proposed method consistently produces the best compliance values in every comparison against other lightweight infill methods, even rivaling the TO results. As discussed in the following, this can mostly be credited to the efficient material agglomeration in high stress regions through the compressive field modulation. By design, neighboring sheets of the same type are not limited by a minimum distance. Thus in extreme cases, distances between sheets, smaller than the wall thickness τ, will cause them to merge, forming solid chunks of material.

5.2 Simulation Experiments

Field Modulation. In Section 4.2 we introduced parameter γ to modulate the vector field, allowing to compress or relax the density of extracted isocontours as shown in Figure 11. While the principal stress eigenvalues encode information about the local stress state, they can vary substantially in magnitude and are therefore more difficult to normalize and handle consistently. The scalar *von Mises* stress σ_v, however, is an established indicator for differences in local stress. Therefore, we propose to use $\gamma = \text{icdf}(\sigma_v)^p$ where $\text{icdf} : \mathbb{R} \to [0,1]$ is the inverse cumulative distribution function for a mapping to the unit interval and p an exponent to exaggerate γ's influence. While p is continuous in general, in this experiment we evaluate seven discrete settings $p \in [0,\dots,6]$. The resulting effect is shown with examples in Figure 11 and plotted in the graph on the right. With $p : 0$, the modulation has no effect and the resulting isocontours are evenly spaced with regular space filling. Increasing p leads to more densely compressed structures in high-stressed regions and sparser distributions in low-stress areas. This effect is directly measurable: for the target loading case, higher values of p improve mechanical performance (lower compliance) from $\frac{c}{c_0} : 1.983$ to 1.162 because material is distributed more effectively. However, the opposite trend appears when we probe arbitrary loading conditions on the optimized structures. Figure 11

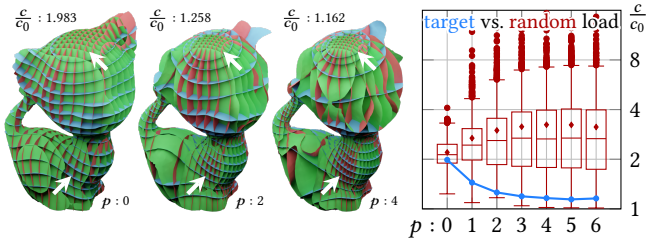

Figure 11: *Exponent p exaggerates the field compression γ. Without compression (p:0), structures are regular, space filling and respond well to arbitrary loads. Higher p improve mechanical performance for the target load by concentrating material (α:0.5) in high-stress regions (arrows) but lowers the performance for arbitrary load cases.*

Table 1: *Comparison of the computation time (in seconds) and mechanical performance as relative compliance (the lower the better). Results meeting AM constraints, are marked in the ≧̃ column, rows denoted with a ⋆ also include field optimization. The numbers next to the object names refer to illustrations showing corresponding examples.*

	method	≧̃	$t(s)$	v/v_0	c/c_0
buddha (25)	Top. Opt. [Wang et al. 2025]		550.0	0.50	1.117
	Porous Infill [Wu et al. 2018]	✓	4930.0	0.50	1.693
	Conf. Lattice [Wu et al. 2021b]		482.5	0.50	3.362
	Voronoi [Wang et al. 2025]		5.5	0.50	2.124
	S^6 [Ours] ($p:4$)		2.2	0.50	1.316
	S^6 [Ours] ($p:4$)⋆	✓	19.4	0.50	1.299
bunny (2)	Top. Opt. [Wang et al. 2025]		1490.0	0.45	1.418
	Porous Infill [Wu et al. 2018]	✓	6180.0	0.45	2.407
	Michell Truss [Arora et al. 2019]		10 052.5	0.46	3.825
	TPMS [Hu et al. 2022]		7693.2	0.45	2.280
	Conf. Lattice [Wu et al. 2021b]		83.0	0.45	3.328
	Hexa [Bukenberger et al. 2024]		91.7	0.45	1.962
	Voronoi [Wang et al. 2025]		7.5	0.45	2.823
	S^6 [Ours] ($p:4$)		1.1	0.45	1.580
	S^6 [Ours] ($p:4$)⋆	✓	13.0	0.45	1.602
femur (17)	Top. Opt. [Wang et al. 2025]	✓	143.0	0.41	1.089
	Conf. Lattice [Wu et al. 2021b]	✓	160.0	0.41	2.976
	Hexa [Bukenberger et al. 2024]		39.2	0.41	1.748
	TO-TPMS [Ning et al. 2026]	✓	365 148.0	0.41	2.720
	S^6 [Ours] (Graph, $p:0$)	✓	2.0	0.41	3.124
	S^6 [Ours] (Graph, $p:2$)	✓	2.0	0.41	1.222
	S^6 [Ours] ($p:4$)		0.3	0.41	1.315
	S^6 [Ours] ($p:4$)⋆		11.4	0.41	1.299
fertility (18)	Top. Opt. [Wang et al. 2025]		515.0	0.40	1.407
	Build-to-Last [Lu et al. 2014]		-	0.37	3.189
	TPMS [Hu et al. 2022]		2455.8	0.45	1.510
	Hexa [Bukenberger et al. 2024]		59.2	0.40	1.962
	TO-TPMS [Ning et al. 2026]	✓	189 504.0	0.40	4.110
	S^6 [Ours] ($p:2$)		0.5	0.40	1.619
	S^6 [Ours] ($p:4$)		0.5	0.40	1.573
kitten (1)	Top. Opt. [Wang et al. 2025]		1020.0	0.41	1.238
	Build-to-Last [Lu et al. 2014]		-	0.41	2.888
	Porous Infill [Wu et al. 2018]	✓	2680.0	0.41	2.125
	TPMS [Hu et al. 2022]		5893.8	0.45	2.310
	TO-TPMS [Ning et al. 2026]	✓	239 616.0	0.34	3.000
	Hexa [Bukenberger et al. 2024]		59.9	0.41	2.000
	S^6 [Ours] ($p:4$)		0.9	0.41	1.471
	S^6 [Ours] ($p:4$)⋆		8.5	0.41	1.399
	S^6 [Ours] ($p:4$)⋆	✓	8.6	0.41	1.403
molar (16)	Top. Opt. [Wang et al. 2025]		936.0	0.35	1.795
	TPMS [Hu et al. 2022]		2346.0	0.45	1.820
	TO-TPMS [Ning et al. 2026]	✓	172 872.0	0.35	3.000
	Hexa [Bukenberger et al. 2024]		41.8	0.35	2.147
	S^6 [Ours] ($p:2$)		0.4	0.35	1.844
	S^6 [Ours] ($p:2$)⋆		3.5	0.35	1.817
	S^6 [Ours] ($p:2$)⋆	✓	3.6	0.35	1.767

summarizes the compliance of the seven designs over random loading scenarios using boxplots. For each design, we generate 1000 trials, each defined by a randomly selected pair of fixation and load application regions with a random force direction applied to the load region. Boxes indicate the interquartile range (IQR), the central line denotes the median, the ♦ marks the mean, whiskers extend to 1.5 × IQR with outliers shown as • outside the whisker range. Load cases that are, by chance, similar to the target loading also benefit from higher p, as indicated by the lower whisker ends tracking the target-case compliance curve. In contrast, the rest of the distribution shows that, as the structure becomes more compressed, thus less regular and space filling, both the average and worst-case compliance increase substantially. Quantitatively, in the uncompressed case ($p : 0$) the worst-case relative compliance over all 1000 trials is $\frac{c}{c_0} : 4.099$, whereas for highly compressed designs this worst-case value can double or even quadruple, confirming the substantial loss in robustness for large p. Thus, parameter p can effectively act as a slider that trades off regular, space-filling structures with high robustness to arbitrary loading conditions against more concentrated material distributions optimized for the target case. I.e., the example on the right is generated analog to the results shown in Figure 12 but with $p : 4$ instead of $p : 0$. With a fixation at the bottom and a down force acting on the bunny's head, this causes massive material concentration in the front with only a few sheets in its back and even completely empty ears. This drastically targeted material distribution is reflected with the best compliance result over the uncompressed results in Figure 12 but also compared to other methods, in Table 1.

Field Alignment. As elaborated in Section 4.1 the initial stress field may contain singularities, making it unsuitable for a straightforward use in subsequent steps with a scalar potential field recovery and isocontour extraction. Thus, we introduce a two-step preprocessing by first oversmoothing the field to reduce and regularize singularities and then recover the best local alignment to the input without breaking smoothness. In Figure 12 we experimentally analyze the influence of structural alignment with the initial stress field on the resulting mechanical performance. Therefore, we compare results using an increasing number of smoothing iterations (1, 5 & 10) and with applied optimization, respectively, denoted with a ★. The illustration also lists the field deviation angle $\overline{\theta}_n$, computed as area weighted mean of angles enclosed by sheet surface normals and principal stress directions in field Ξ. As can be

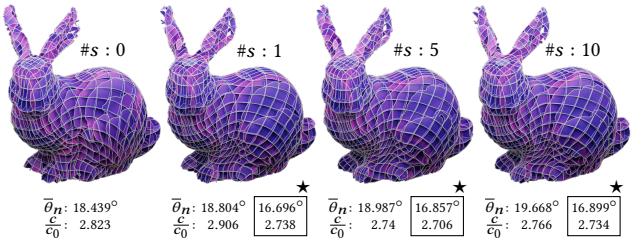

Figure 12: *Comparing the impact of field smoothing plus alignment optimization (★) using an increasing number of smoothing steps #s. No field compression is applied ($p : 0$) and compliance is evaluated on materialized results using a volume fraction of $\alpha : 0.45$.*

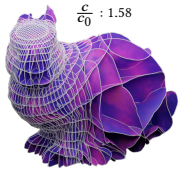

Figure 13: *Applying a compression of $p : 4$ to the bunny shown in Figure 12 results in a substantial improvement in mechanical performance: material is concentrated in highly stressed regions, while the ears and back are left almost empty.*

observed in the figure, the deviation from the initial field increases with the number of smoothing iterations, as expected. The second observation aligns with Michell theory and proves the optimization strategy: For each smoothing level, however, the subsequent optimization consistently reduces the deviation from the input field, thus improves the mechanical performance. This indicates that the optimization is able to partially compensate for the changes introduced by smoothing. For the present case, the mechanically optimized structures obtained after 1, 5, and 10 smoothing iterations exhibit very similar compliance, with the result obtained after five smoothing steps being marginally better. This suggests that only limited smoothing is required for this particular field. Nevertheless, the amount of smoothing required is inherently dependent on the complexity of the principal stress tensor field Ξ. The field can contain swirls and singularities, which complicate the recovery of the scalar potential Φ. As expressed in Equation (9), Φ is recovered as the most suitable smooth scalar function in a least-squares sense. Consequently, insufficient smoothing may leave local singularities and swirls in the field, which can be reflected in the extracted sheet structures. For other scenarios with more pronounced swirls or singularities, a higher number of smoothing iterations may be beneficial before optimization, which is why our pipeline includes a configurable smoothing stage.

Smoothing alone often already provides a sufficiently well aligned field for extracted sheets resulting in good mechanical performance. The optimization step aims to harness most of the structures potential by realigning the field and resulting structure to the input stress [Michell 1904]. Numbers listed in Figure 12 and Table 1 show improvements of 3 to 6 % better compliance when the optimization is applied. Parameters such as the smoothing steps, number of sheets, compression factor or volume fraction can be adjusted interactively, allowing for a responsive and explorative design workflow. The more computationally involved optimization can be deferred until the desired configuration has been established. It can then be applied within a few seconds to finalize the design for subsequent processing or export. Tuning these other parameters also has a measurable influence on the results and can lead to moderate performance gains. However, the overall biggest improvements on stiffness is achieved by material concentration using the aforementioned parameter p. This indicates that the approach overall is robust to parameter variations in preprocessing, reducing the need for extensive manual optimization.

Number of Sheets. The number of sheets n can be chosen interactively by the user or derived from manufacturing constraints such as minimum printable wall thickness. Figure 14 shows, without field compression, sheets are spaced evenly where more and thinner sheets disperse material more homogeneously, thus degrade mechanical performance for a target load scenario. With

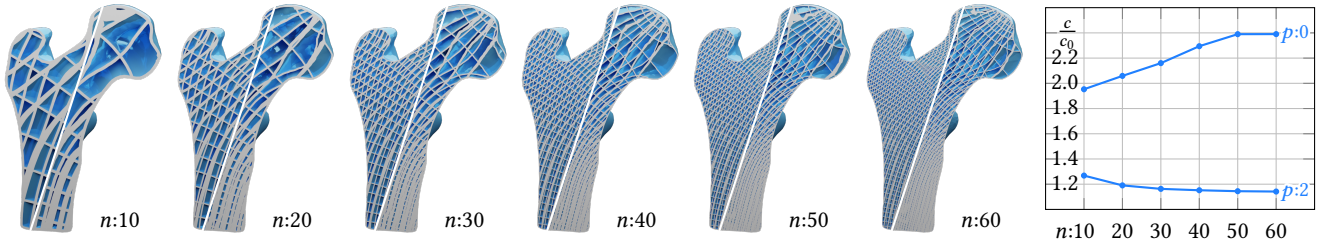

Figure 14: *The split views show the $p : 0$ (left) and $p : 2$ (right) versions of the femur ($\alpha : 0.5$) for increasing numbers of sheets n. The plot on the right indicates that mechanical performance can be improved by increasing the number of sheets when compression is applied. Without compression, however, additional sheets distribute the material more homogeneously, resulting in lower performance for a given target load.*

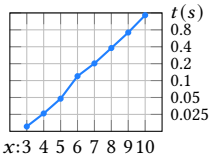

Figure 15: *Measurements for the sheet extraction time on the femur with am exponential increase in the number of sheets $n = 2^x$. The plot shows a linear relationship between the number of sheets and the computation time.*

compression, however, increasing numbers of sheets improve the structure's mechanical performance. In practice, n should be selected with the intended print scale in mind so that the resulting wall thickness τ remains manufacturable (e.g., above the minimum feature size and nozzle-dependent limits). For our examples, we set a default of $n : 20$, which yields an approximate wall thickness of four times the nozzle extrusion width for objects of about 100 mm height. In the log-log plot shown in Figure 15, the number of sheets is given by $n = 2^x$, with x ranging from 3 to 10, corresponding to 8 to 1024 sheets per dimension. Although such a large number of sheets would become impractical to manufacture, these values serve as an example to illustrate the linear relationship between the number of sheets and computation time.

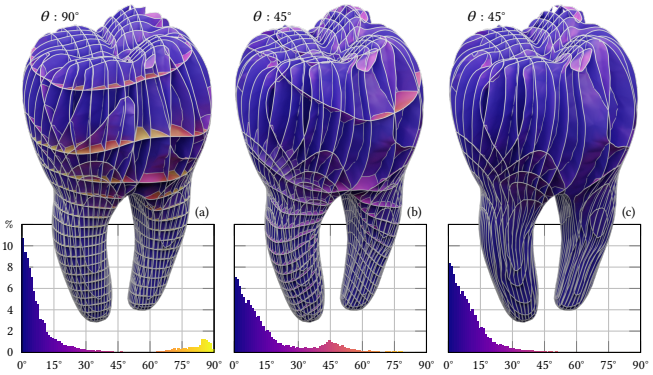

Figure 16: *Examples of extracted isocontours and corresponding histograms, showing percentages of sheet area with a certain overhang angle. Tweaking field Ξ' with θ angle limitation allows for extracted surfaces to adhere manufacturability constraints. In the rightmost image, sheets oriented predominantly orthogonal to the print direction have been experimentally removed.*

Manufacturability. FFF is one of the most popular AM techniques and produces objects by stacking thin layers of thermoplastic, deposited linewise through a heated nozzle. Capabilities to print such lines *free floating* are still limited, thus common hardware requires new layers to be deposited onto previous layers. This necessitates either support structures (to be removed after the print) or surfaces not exceeding a certain overhang angle. Internal support structures are wasteful, time consuming and can't be removed from closed shapes. Thus, our goal is to produce readily printable geometry with limited overhang angles. Figure 16 illustrates the effect introduced in Section 4.2 for tweaking components of Ξ' before the geometry extraction, adhering to given angular constraints. This modulation necessarily departs from exact stress alignment in favor of manufacturability, allowing the extracted sheets to respect FFF overhang limits and thus avoiding the need for non-removable support structures. We therefore treat this step as a pragmatic manufacturability option rather than a change in design intent: it slightly perturbs the stress-guiding field only where needed to satisfy overhang limits, while preserving the overall structural performance benefits. Histograms in Figure 16 plot the percentages of sheet surface areas with certain angles, respectively. The default result without constraints is shown in 16a, allowing for overhang angles up to 90°. As the force acts on the object vertically, so does the first principal stress component of Ξ' and, therefore, the corresponding sheets transcend the object almost horizontally. This is directly reflected with the bump above 75° in the histogram. The enforced θ limit is noticeable in 16b as these sheets are bent upwards. Although their histogram bump shifted to around 45°, the histogram reveals that some small regions still exhibit angles over 45°. This can be explained as the field is constrained only locally. The remaining overhang regions can be attributed to ridges forming between the steepened slopes on either side. Regions where the overhang angle remains critical are, however, very tiny and can be bridged by the filament extrusion nozzle. However, the sheets in question do not significantly contribute to mechanical performance for the given load case and are, therefore, experimentally omitted in 16c. The freed material budget is relocated to thicken the remaining sheets, then further significantly improving overall stiffness, as listed in Table 1. A manufactured example of these self-supporting infill structures without the hull is included in Figure 24.

Another AM constraint comes with the concept of SLA, where objects emerge from a bath of photopolymer resin. Therein, the complication is to avoid closed hollow regions in the final structure. Those can ultimately ruin a print by forming suction cups and ripping off the unfinished print from the build plate during platform

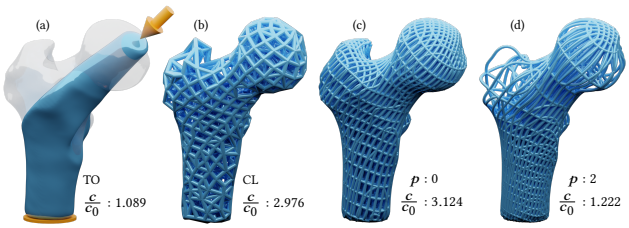

Figure 17: *Comparing void-free results without the object hull. Fixation and force on the femur bone are shown on the left with a TO result [2025]. Others (left to right) are a CL result [2021b] and two versions of ours, one space-filling (p : 0) and one concentrated (p : 2).*

elevation. Further, they may even conceal uncured resin that can not be rinsed out and will over time degrade the print from the inside. Avoiding cavities will always be challenging with sheet-based structures. Graph-based methods like Michell trusses [Arora et al. 2019] or conforming lattice structures (CL) [Wu et al. 2021b] are perfectly permeable for resin and rinsing liquids. Our approach can be easily adapted to form graph-structures as well by extracting the intersection graph network of sheets intersecting each other. Examples on the femur (α : 0.4) are compared in Figure 17: Without the constraint to feature the input object's closed hull, the TO result is also free of voids. Further, we compare a CL result and two versions of our approach. Our first result with p : 0 matches the resolution, space-fillingness and mechanical performance of the CL result. Naturally, our field compression concept can also be applied when extracting only graph based structures. This is shown with p : 2 in the second result, better material concentration in high-stress regions drastically improve the design compliance.

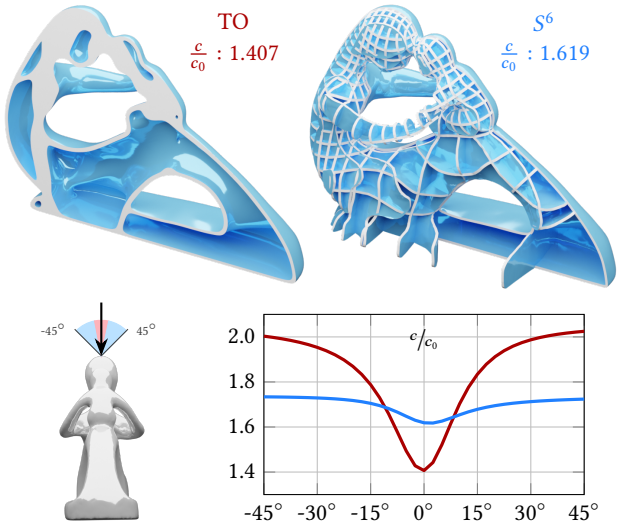

Figure 18: *Compliance responses of TO (left) and our S^6 method (right), optimized for the bottom loading case. Curves show the relative compliance of designs when the applied force deviates from the initial force for which the design was optimized. The red cone shows the narrow range in which TO performs best. For wider angel variations (blue), our result has the best compliance.*

Resilience. Common structural design strategies optimize material distribution for stiffness under a single specific loading scenario. In real-world applications, load variations may occur that must be accommodated. Such variations may include angular deviations of acting forces, unanticipated loads or damage from application or manufacturing. Figure 18 shows an experimental compliance analysis under rotated loading conditions of ±45°. This illustrates how TO designs' are particular vulnerability to such changes: While exhibiting the best stiffness results in the initial loading setup (±0°), variations lead to rapid mechanical performance degradation of up to 44 % higher compliance. Our design starts with a slightly higher compliance for the initial setup but proves quite resilient to induced load deviations with only a 5 % compliance increase. Marked by the two curves crossing, there is only a ±10° range for variations in which the TO result actually performs best. This narrow cone is also highlighted in the illustration at the bottom left.

The jet engine bracket (JEB) challenge [GE 2013] specifies four different loading scenarios for the engine bracket, shown in Figure 19. This experiment focuses on single-load designs: both our method and the TO baseline are optimized for one prescribed load case (4). The other load configurations are then applied on this design to compare their off-design behavior under identical information and assumptions. Although TO methods exist for multiple or uncertain load cases, giving TO access to all load cases simultaneously would define a different, more informed setting. For the target scenario, the mechanical performance of our design already comes very close to the performance of the TO design. When applying scenario (2), the TO result still performs better than ours, which can likely be attributed to the similarity of the scenarios (2) and (4), both with horizontal forces. However, for larger variations as with scenarios (1) and (3) our result significantly outperforms the TO result. The same behavior occurs analogously when optimizing for any of the other loading cases and evaluating performance against the remaining ones.

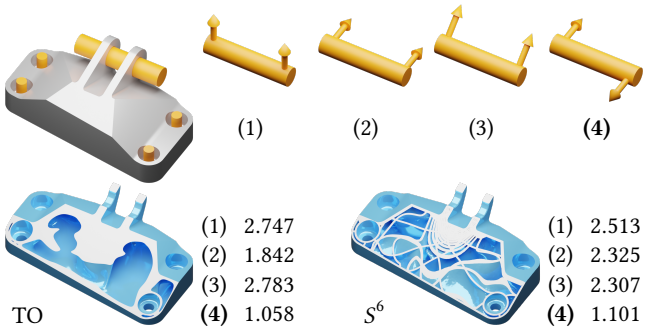

Figure 19: *The jet engine bracket challenge [GE 2013] specifies four different load scenarios (top). The designs are optimized for scenario (4) with α : 0.5. Numbers give the relative compliance $\frac{c}{c_0}$ of the resulting designs to all four loading scenarios, respectively.*

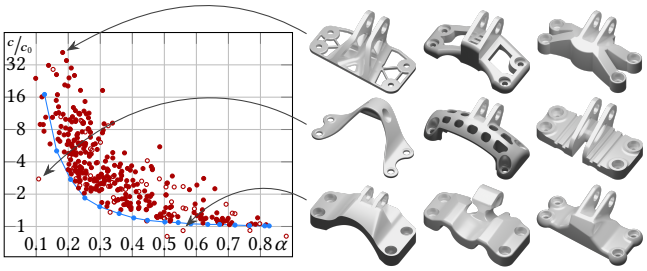

Figure 20: *Handcrafted submissions from the GE JEB challenge [GE 2013] have different relative volumes (α) and are not constrained by the shape of the initial bracket. Red data points mark the compliance response for each design to load case **(4)**. The blue line connects responses of our S^6 design, realized at different volumes.*

Varying Volume. Figure 20 shows examples of the many handcrafted designs, submitted in the original GE JEB challenge [GE 2013]. The plot on the left illustrates the designs' individual compliance responses to load case **(4)**, respectively. Our design from Figure 19 is adapted to different volume fractions by increasing the sheet thickness τ. Compliance responses, corresponding to the different volume versions, are connected with the blue line. This demonstrates the versatility of our method: A single result design can be adapted to a continuous range of volumes. With increasing volume, the compliance response effectively converges to that of the solid object and practically beats all manual designs. The JEB challenge specifications only describe the loading conditions but do not restrict or constrain the design space. This mainly explains the data points in the plot below our blue curve. I.e., the bottom left design even has a relative compliance below 1, thus is stiffer than the reference bracket (top-left in Figure 19), although consuming only about 55 % of its volume. That is because it practically exceeds the given shape and is much bulkier in the top center than the original object. Empty dots in the diagram mark responses of designs that exceed the convex hull of the reference bracket. Our method, however, is designed to produce stiff infills for a given shape without changing or exceeding its boundary. Nevertheless, the plotted results position it as one of the mechanically most performant options for optimizing the compliance to volume ratio.

Dehomogenization. Our comparisons with TO designs all refer to results from density-based TO, a concept optimizing only the density values of the finite elements in a high-resolution mesh. A different approach is homogenization-based TO, yielding not only a continuous density field but also an orthotropic tensor field, corresponding to periodic microstructures on a coarse mesh. The result is rather a relaxed description than a manufacturable solid-void geometry.

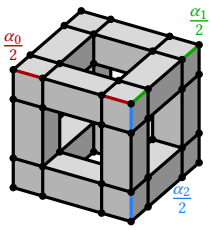

Figure 21: *Homogenization-based TO optimizes a density field, which is materialized as a multi-scale structure using microstructures. The local density determines the configuration of the corresponding microstructure ($\alpha_{0,1,2}$), translating the optimized material distribution into fine-scale geometry.*

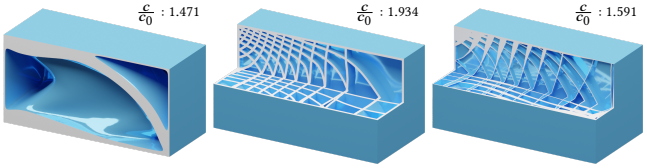

Figure 22: *Comparing TO and two S^6 versions. A single-simulation stress field is used as basis for the center result. The right one is based on an optimized direction field and anisotropic material properties from a homogenization process.*

Therein optimized parameters include not only density but also orientations and corresponding anisotropic scaling parameters of the microstructure elements [Geoffroy-Donders et al. 2020], as shown in Figure 21. The optimization results $\alpha_{0,1,2}$ for the microstructures correspond to the individual components of the principal stress direction tensors. Dehomogenization methods then translate the continuous homogenized design into a manufacturable binary material field. The right of Figure 22 illustrates a result obtained by incorporating our method as dehomogenization. Therefore, optimized microstructure orientations are employed for the field potential recovery instead of principal stress directions. Further, anisotropic material properties $\alpha_{0,1,2}$ locally scale the sheet extrusion τ of their corresponding sheets. The directional field and material properties are obtained through a separate, preceding optimization process. Using this *optimal* data as the foundation for our approach enables the generation of even stiffer designs. This improvement is reflected in the reduced compliance of the dehomogenized results compared to one based on a single-simulation stress field with uniform, isotropic sheet thickness. This represents a more experimental use case of our method, as the separate homogenization process involves substantial additional computational effort. Nevertheless, the results highlight the versatility and potential of our approach, suggesting its applicability in future dehomogenization frameworks.

5.3 Physical Experiments

Commonly available lightweight infill methods, including TO or microstructure-based designs, do not inherently satisfy manufacturing constraints, are thus, not readily printable or may require manual post-processing. Although the PI method [Wu et al. 2018] does not explicitly enforce overhang limits either, its cellular inner structure can still be printed reliably on standard FFF printers when scaled appropriately. For this reason, we compare our designs



Figure 23: *Examples of the bunny, buddha & kitten objects being printed on a FFF printer using different infill designs but with identical reduced material budgets, respectively.*



Figure 24: *Top: Printed objects with a height of 100 mm from bottom to loading area. Bottom: Inner structures are fully self-supporting and can be printed without the hull or other additional support.*

against the default material-reduction infill strategy available in common 3D printer software, namely a rectilinear grid of sparsely spaced filament strands, as well as against the strength-optimized PI method, across a range of test objects. All shown results were fabricated on an *Ultimaker 2* printer using standard polylactic acid (PLA) filament and default printer settings, without additional support material (Figure 23). Infills are adjusted so that all designs consumed the same amount of material per object. The objects are printed with a height of 100 mm from the base to the loading point. As shown in Figure 24, loading points are either flattened (buddha) or extended with a small solid pedestal (bunny & kitten) to provide an even surface for uniform load application. The test setup is shown on the right in Figure 25, with the objects placed flat on a surface and subjected to a compressive load from above. The force was increased gradually, while the resulting deformation was measured until a termination threshold of 2 % of the original object height was reached [DIN 2003]. The corresponding force-deformation curves are shown in Figure 25. As can be seen, our sheet structures exhibited lower deformation than standard infill and even porous infill designs under identical loading conditions. The standard and porous infill curves of the bunny exhibit slight discontinuities near the end. This behavior could be attributed to eccentric loading on the asymmetric specimen, causing local material yielding or buckling, similar to the fold formation observed in damaged carbon-fiber reinforced

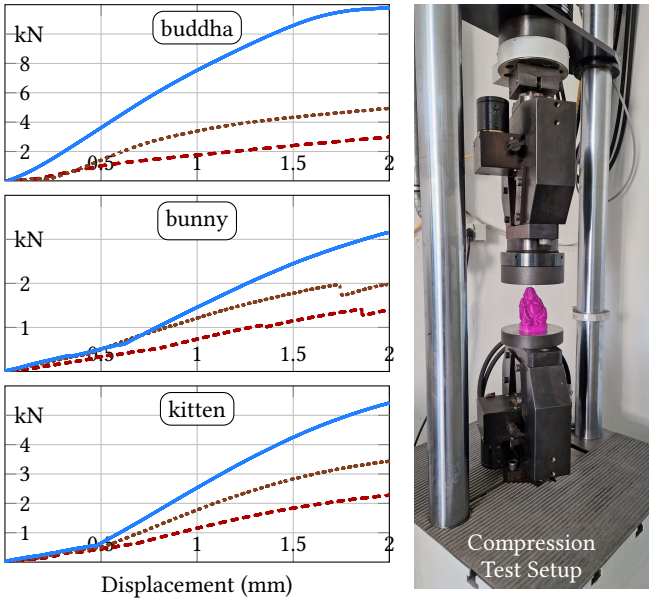

Figure 25: *Graphs on the left show deformation response curves from the compression test on the right, comparing - - - Standard Infill, ····· Porous Infill and our —— S^6 Infill on different objects (Figure 23). Our designs withstand 2–4× the force of standard infills before reaching the 2 % deformation termination criterion.*

polymer components at high compressive stresses. Nevertheless, on average our infills require 2 to 4 times larger forces to reach the deformation termination threshold, depending on the given object geometry. Furthermore, the bottom examples in Figure 23 shows only the inner sheets without the outer hull. This structure was printed successfully without the need for any additional support, further demonstrating that the overhang-handling strategy introduced in Section 4.3 also practically translates into physical manufacturing. This experiment confirms that our proposed structures can be manufactured reliably using standard consumer-grade 3D printers. The results of the compression test also show that they provide a highly efficient alternative to conventional infill strategies for improving the mechanical performance of printed parts under a fixed material budget.

5.4 Computation and Implementation

Our method is implemented in Python with the code available at https://github.com/dbukenberger/SpeedySheets. Timings listed in Table 1 are generated with a i9-12900k CPU and measure the time from input fields to sheet-structures, which includes smoothing, combing, the scalar field recoveries and isocontour extraction. The plot in Figure 26 shows the field preparation time in seconds as a function of the background tessellation size. The data points correspond to our test objects, with the number of tetrahedral elements $|\mathcal{T}|$ ranging from 8k to over 60k. The time required to extract the isocontour sheets follows a similar trend, but is approximately one order of magnitude faster. Furthermore, the extraction time scales linearly with the number of sheets. The resulting meshes can subsequently be evaluated using the SGLDBench framework [Wang et al. 2025], with the FEM simulations taking less than 20 s

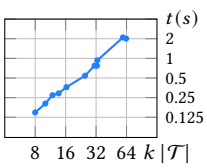

Figure 26: *Measurements of the field preparation time in relation to the background tessellation size for the objects used in this paper. The preparation time scales approximately linearly with the number of tetrahedra $|\mathcal{T}|$, given here in thousands.*

in our experiments, using a grid resolution of 2^8. Thus, while the current implementation is not integrated into a 3D-printing slicer, these timings indicate that the method could support an interactive design workflow, providing immediate visual feedback when adjusting parameters such as the number of sheets or material concentration. Based on the experiments presented in this paper, the parameter p can be selected to balance general structural performance against targeted loading scenarios, while n should primarily be chosen according to manufacturing constraints. During interactive exploration, smoothing is sufficient to rapidly evaluate different designs and parameter settings. Once a satisfactory design has been found, the optional optimization can be applied to further improve the last few percent of performance, requiring less than 20 s in our experiments. Experimentally, we were able to further reduce the optimization times reported in Table 1 by approximately half using a CUDA-based JAX-backend running on a 3090Ti GPU. The optimization is realized using Adam [Kingma and Ba 2014] with gradient clipping and an adaptive learning-rate schedule. The optimizer is initialized with a learning rate of 10^{-2} and we apply global-norm clipping with a threshold of 1.0. If the objective does not improve over 5 iterations, we halve the learning rate; after several such reductions or when the learning rate and loss fall below predefined thresholds, the optimization is terminated. Across all our experiments, we observed stable convergence without manual tuning of any hyperparameters beyond these settings. A practical integration into slicer software would additionally require a tetrahedralization of the input geometry, rather than only its surface hull; this preprocessing can, however, be performed within a few seconds using existing tools such as fTetWild [Hu et al. 2020]. The slicer interface would further need to expose load specifications, extending existing interaction tools for operations such as measuring, translating and scaling to allow users to define relevant forces and loading scenarios. An FEM solver would then be required to evaluate the resulting structures. These components are external to our method and are likewise not part of our current implementation. Consequently, while a complete slicer integration is left to future work, the computational requirements and required supporting components suggest that such an integration is practically feasible.

Numerical Stability. All results are based on volumetric meshes generated with fTetWild [Hu et al. 2020], which produces high-quality tetrahedralizations without sliver elements, in contrast to some legacy meshing tools. On these inputs we did not encounter crashes or non-terminating runs: the smoothing stage behaves well numerically and the subsequent optimization consistently converges to a stable solution. Isosurface extraction and FEM evaluation also proved robust across different solver and contouring backends (e.g., standard sparse linear or multiprocessing solvers and common Python mesh libraries).

6 Conclusion

In this work we present a novel concept for rapid stress-tailored computational design. Componentwise potential fields are recovered from an orthogonal principal stress field. Extracted isocontours therein align orthogonally with the principal stress directions and, when materialized as extruded sheets, serve as stiff structural components. The implicit representation as vector fields allow for easy manipulation of the structure before its extraction. I.e., overhang limitations of FFF methods can be met by introducing angular constraints on the individual vector components. Further, the normalized vector field can be modulated by a stress magnitude indicator. This effectively compresses the field in high stress regions, resulting in optimized material concentration and, thus, significantly improves stiffness. Involved computations are lightweight and allow for interactive design iterations. Further mechanical performance improvements can be obtained by an optional realignment optimization step, finalizing the design. As a combination of interactive design synthesis and simple adaptation of FFF and SLA manufacturing constraints, our concept can be readily included in an AM slicer software workflow. It therein offers a realistic alternative to common infill pattern strategies, usually invariant of any applied forces on the resulting object. In extensive simulation based experiments we demonstrate superior mechanical performance over several established lightweight design methods. As our results are space-filling by design, they are resilient to multiple or varying loading conditions. Our method establishes a new paradigm for lightweight structural design, particularly suitable for real-world loading scenarios, rapid AM prototyping cycles and large-scale 3D applications where the computational complexity of other methods prove prohibitive.

While our method achieves interactive layout design, the subsequent refinement step to optimize volume fraction and field alignment remains a sequential bottleneck that disrupts the interactive workflow. Future research will focus on eliminating this two-stage process by embedding the refinement directly into the interactive loop. Promising directions include (i) training a lightweight surrogate model (e.g., a neural network) to predict near-optimal volume fractions on-the-fly from the current stress field and user edits, (ii) developing a continuous optimization that incrementally updates both topology and thickness during interaction, and (iii) leveraging GPU-accelerated implicit differentiation to perform the refinement in the background without blocking user input. Ultimately, we aim for a fully interactive pipeline where also the printable geometry is instantly available, enabling a user guided exploration of the stiffness-resilience trade-off.

Acknowledgments

This work was partially funded by the German Science Foundation (DFG project 577775355).

References

Rahul Arora, Alec Jacobson, Timothy R Langlois, Yijiang Huang, Caitlin Mueller, Wojciech Matusik, Ariel Shamir, Karan Singh, and David IW Levin. 2019. Volumetric michell trusses for parametric design & fabrication. In *Proceedings of the 3rd Annual ACM Symposium on Computational Fabrication*. Association for Computing Machinery, New York, NY, USA, 1–13.

Martin P Bendsøe. 1989. Optimal shape design as a material distribution problem. *Structural optimization* 1, 4 (1989), 193–202.

Martin Philip Bendsøe and Noboru Kikuchi. 1988. Generating optimal topologies in structural design using a homogenization method. *Computer Methods in Applied Mechanics and Engineering* 71, 2 (1988), 197–224.

Martin Philip Bendsøe and Ole Sigmund. 2013. *Topology optimization: theory, methods, and applications*. Springer Science & Business Media, 233 Spring Street„ New York, NY, 10013-1578, USA.

David Bommes, Henrik Zimmer, and Leif Kobbelt. 2009. Mixed-integer Quadrangulation. *ACM Transactions on Graphics (TOG)* 28, 3, Article 77 (July 2009), 10 pages. doi:10.1145/1531326.1531383

Dennis R. Bukenberger, Marco Tarini, and Hendrik P. A. Lensch. 2022. At-Most-Hexa Meshes. *Computer Graphics Forum* 41, 1 (2022), 7–28. doi:10.1111/cgf.14393

Dennis R. Bukenberger, Junpeng Wang, Jun Wu, and Rüdiger Westermann. 2024. Stress-Aligned Hexahedral Lattice Structures. *Computer Graphics Forum* 44.1 (2024), e15265. doi:10.1111/cgf.15265

Etienne Corman and Keenan Crane. 2019. Symmetric Moving Frames. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–16.

Keenan Crane. 2019. The n-dimensional cotangent formula. 11–32 pages.

Keenan Crane, Clarisse Weischedel, and Max Wardetzky. 2013. Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Transactions on Graphics (ToG)* 32, 5 (2013), 1–11.

Thierry Delmarcelle and Lambertus Hesselink. 1994. The topology of symmetric, second-order tensor fields. In *Proceedings Visualization'94*. IEEE Computer Society Press, Washington, DC, United States, 140–147.

Normenausschuss Kunststoffe DIN, Deutsches Institut für Normung e. V. 2003. Plastics - Determination of compressive properties.

Guoxin Fang, Tianyu Zhang, Sikai Zhong, Xiangjia Chen, Zichun Zhong, and Charlie CL Wang. 2020. Reinforced FDM: Multi-axis filament alignment with controlled anisotropic strength. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–15.

Xifeng Gao, Wenzel Jakob, Marco Tarini, and Daniele Panozzo. 2017. Robust hex-dominant mesh generation using field-guided polyhedral agglomeration. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 114.

GE. 2013. *GE Jet Engine Bracket Challenge*. GrabCAD. https://grabcad.com/challenges/ge-jet-engine-bracket-challenge

Perle Geoffroy-Donders, Grégoire Allaire, and Olivier Pantz. 2020. 3-d topology optimization of modulated and oriented periodic microstructures by the homogenization method. *J. Comput. Phys.* 401 (2020), 108994.

Jeroen P Groen and Ole Sigmund. 2018. Homogenization-based topology optimization for high-resolution manufacturable microstructures. *Internat. J. Numer. Methods Engrg.* 113, 8 (2018), 1148–1163.

Tom Gustafsson and G. D. McBain. 2020. scikit-fem: A Python package for finite element assembly. *Journal of Open Source Software* 5, 52 (2020), 2369. doi:10.21105/joss.02369

Jiacheng Han, Xiaoya Zhai, Lili Wang, Di Zhang, Junhao Ding, Winston Wai Shing Ma, Xu Song, Wei-Hsin Liao, Ligang Liu, Jun Wu, et al. 2024. Inverse-designed 3D sequential metamaterials achieving extreme stiffness. *Materials & Design* 247 (2024), 113350.

Chiara Hergl, Christian Blecha, Vanessa Kretzschmar, Felix Raith, Fabian Günther, Markus Stommel, Jochen Jankowai, Ingrid Hotz, Thomas Nagel, and Gerik Scheuermann. 2021. Visualization of tensor fields in mechanics. In *Computer Graphics Forum*, Vol. 40.6. Wiley Online Library, The Eurographics Association and John Wiley & Sons Ltd., Groene Loper 3, 5612AE Eindhoven, The Netherlands, 135–161.

Jiangbei Hu, Shengfa Wang, Baojun Li, Fengqi Li, Zhongxuan Luo, and Ligang Liu. 2022. Efficient Representation and Optimization for TPMS-Based Porous Structures. *IEEE Transactions on Visualization and Computer Graphics* 28, 7 (2022), 2615–2627. doi:10.1109/TVCG.2020.3037697

Yixin Hu, Teseo Schneider, Bolun Wang, Denis Zorin, and Daniele Panozzo. 2020. Fast tetrahedral meshing in the wild. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 117–1.

Shih-Hsuan Hung, Yue Zhang, and Eugene Zhang. 2023. Global topology of 3d symmetric tensor fields. *IEEE Transactions on Visualization and Computer Graphics* 30, 1 (2023), 1282–1291.

Alec Jacobson, Daniele Panozzo, et al. 2018. libigl: A simple C++ geometry processing library. https://libigl.github.io/.

Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. arXiv:1412.6980 [cs.LG] https://arxiv.org/abs/1412.6980

Jikai Liu, Andrew T Gaynor, Shikui Chen, Zhan Kang, Krishnan Suresh, Akihiro Takezawa, Lei Li, Junji Kato, Jinyuan Tang, Charlie CL Wang, et al. 2018. Current and future trends in topology optimization for additive manufacturing. *Structural and multidisciplinary optimization* 57, 6 (2018), 2457–2483.

Tao Liu, Tianyu Zhang, Yongxue Chen, Yuming Huang, and Charlie CL Wang. 2024. Neural slicer for multi-axis 3D printing. *ACM Transactions on Graphics (TOG)* 43, 4 (2024), 1–15.

Zhenyuan Liu, Jingyu Hu, Hao Xu, Peng Song, Ran Zhang, Bernd Bickel, and Chi-Wing Fu. 2022. Worst-Case Rigidity Analysis and Optimization for Assemblies with Mechanical Joints. In *Computer Graphics Forum*, Vol. 41.2. Wiley Online Library, The Eurographics Association and John Wiley & Sons Ltd., Groene Loper 3, 5612AE Eindhoven, The Netherlands, 507–519.

William E Lorensen and Harvey E Cline. 1987. Marching cubes: A high resolution 3D surface construction algorithm. *ACM Transactions on Graphics (TOG)* 21, 4 (1987), 163–169.

Lin Lu, Andrei Sharf, Haisen Zhao, Yuan Wei, Qingnan Fan, Xuelin Chen, Yann Savoye, Changhe Tu, Daniel Cohen-Or, and Baoquan Chen. 2014. Build-to-Last: Strength to Weight 3D Printed Objects. *ACM Transactions on Graphics (ToG)* 33, 4 (2014), 1–10.

Anthony George Maldon Michell. 1904. The Limits of Economy of Material in Frame-Structures. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 8, 47 (1904), 589–597.

Amir M Mirzendehdel and Krishnan Suresh. 2016. Support structure constrained topology optimization for additive manufacturing. *Computer-Aided Design* 81 (2016), 1–13.

Matthias Nieser, Ulrich Reitebuch, and Konrad Polthier. 2011. CubeCover - Parameterization of 3D Volumes. In *Computer Graphics Forum*, Vol. 30.5. Wiley Online Library, The Eurographics Association and John Wiley & Sons Ltd., Groene Loper 3, 5612AE Eindhoven, The Netherlands, 1397–1406.

HY Ning, WS Huang, and GH Tang. 2026. Data-driven lightweight and robust design of hybrid TPMS structures. *Additive Manufacturing* 118 (2026), 105106.

Julian Panetta, Abtin Rahimian, and Denis Zorin. 2017. Worst-case stress relief for microstructures. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–16.

Nico Pietroni, Marcel Campen, Alla Sheffer, Gianmarco Cherchi, David Bommes, Xifeng Gao, Riccardo Scateni, Franck Ledoux, Jean Remacle, and Marco Livesu. 2022. Hex-mesh generation and processing: a survey. *ACM Transactions on Graphics (TOG)* 42, 2 (2022), 1–44.

Xiaoping Qian. 2017. Undercut and overhang angle control in topology optimization: a density gradient based integral approach. *Internat. J. Numer. Methods Engrg.* 111, 3 (2017), 247–272.

Nicolas Ray, Dmitry Sokolov, and Bruno Lévy. 2016. Practical 3D frame field generation. *ACM Transactions on Graphics (TOG)* 35, 6 (2016), 233.

Elina Robeva. 2016. Orthogonal decomposition of symmetric tensors. *SIAM J. Matrix Anal. Appl.* 37, 1 (2016), 86–102.

Abhinit Sati, Tiffany Bao, Jeff Tedi, Edward Chien, and Emily Whiting. 2026. Strain-Field Based Segmentation for Fabric Formwork. *Computer Graphics Forum* 25.2 (2026), e70367. doi:10.1111/cgf.70367

Günter Schwarz. 1995. *Hodge Decomposition-A method for solving boundary value problems*. Springer, Springer Berlin, Heidelberg.

Ole Sigmund. 1994. Materials with prescribed constitutive parameters: an inverse homogenization problem. *International Journal of Solids and Structures* 31, 17 (1994), 2313–2329.

Ole Sigmund. 2001. A 99 line topology optimization code written in Matlab. *Structural and Multidisciplinary Optimization* 21, 2 (2001), 120–127.

FC Stutz, JP Groen, O Sigmund, and JA Bærentzen. 2020. Singularity aware dehomogenization for high-resolution topology optimized structures. *Structural and Multidisciplinary Optimization* 62, 5 (2020), 2279–2295.

Florian Cyril Stutz, Tim Felle Olsen, Jeroen Peter Groen, Tuan Nguyen Trung, Niels Aage, Ole Sigmund, Justin Solomon, and Jakob Andreas Bærentzen. 2022. Synthesis of frame field-aligned multi-laminar structures. *ACM Transactions on Graphics (TOG)* 41, 5 (2022), 1–20.

Graham M Treece, Richard W Prager, and Andrew H Gee. 1999. Regularised marching tetrahedra: improved iso-surface extraction. *Computers & Graphics* 23, 4 (1999), 583–598.

Mariano Victoria, Pascual Martí, and Osvaldo M Querin. 2009. Topology design of two-dimensional continuum structures using isolines. *Computers & structures* 87, 1-2 (2009), 101–109.

Junpeng Wang, Dennis R. Bukenberger, Simon Niedermayr, Christoph Neuhauser, Jun Wu, and Rüdiger Westermann. 2025. SGLDBench: A Benchmark Suite for Stress-Guided Lightweight 3D Designs. *IEEE Transactions on Visualization and Computer Graphics* -, - (2025), 1–13. doi:10.1109/TVCG.2025.3573774

Junpeng Wang, Jun Wu, and Rüdiger Westermann. 2020. A globally conforming lattice structure for 2d stress tensor visualization. In *Computer graphics forum*, Vol. 39.3. Wiley Online Library, The Eurographics Association and John Wiley & Sons Ltd., Groene Loper 3, 5612AE Eindhoven, The Netherlands, 417–427.

Junpeng Wang, Jun Wu, and Rüdiger Westermann. 2022. Stress trajectory guided structural design and topology optimization. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, Vol. 86212. American Society of Mechanical Engineers, American Society of Mechanical Engineers, 2 Park Ave, New York, NY 10016, USA, V002T02A012.

Shengfa Wang, Zheng Liu, Jiangbei Hu, Na Lei, and Zhongxuan Luo. 2023b. An Efficient Self-supporting Infill Structure for Computational Fabrication. In *Computer Graphics Forum*, Vol. 42.7. Wiley Online Library, The Eurographics Association

and John Wiley & Sons Ltd., Groene Loper 3, 5612AE Eindhoven, The Netherlands, e14953.

Weiming Wang, Dongwei Feng, Li Yang, Shan Li, and Charlie CL Wang. 2023a. Topology optimization of self-supporting lattice structure. *Additive Manufacturing* 67 (2023), 103507.

Jun Wu, Niels Aage, Rüdiger Westermann, and Ole Sigmund. 2018. Infill optimization for additive manufacturing-approaching bone-like porous structures. *IEEE Transactions on Visualization and Computer Graphics* 24, 2 (2018), 1127–1140.

Jun Wu, Ole Sigmund, and Jeroen P Groen. 2021a. Topology optimization of multi-scale structures: a review. *Structural and Multidisciplinary Optimization* 63 (2021), 1455–1480.

Jun Wu, Charlie CL Wang, Xiaoting Zhang, and Rüdiger Westermann. 2016. Self-supporting rhombic infill structures for additive manufacturing. *Computer-Aided Design* 80 (2016), 32–42.

Jun Wu, Weiming Wang, and Xifeng Gao. 2021b. Design and Optimization of Conforming Lattice Structures. *IEEE Transactions on Visualization and Computer Graphics* 27, 1 (2021), 43–56.

Qiancheng Zhang, Xiaohu Yang, Peng Li, Guoyou Huang, Shangsheng Feng, Cheng Shen, Bin Han, Xiaohui Zhang, Feng Jin, Feng Xu, et al. 2015. Bioinspired engineering of honeycomb structure–Using nature to inspire human innovation. *Progress in Materials Science* 74 (2015), 332–400.

Tianyu Zhang, Guoxin Fang, Yuming Huang, Neelotpal Dutta, Sylvain Lefebvre, Zekai Murat Kilic, and Charlie CL Wang. 2022. S3-slicer: A general slicing framework for multi-axis 3D printing. *ACM Transactions on Graphics (TOG)* 41, 6 (2022), 1–15.

Nan Zheng, Xiaoya Zhai, and Falai Chen. 2023. Topology optimization of self-supporting porous structures based on triply periodic minimal surfaces. *Computer-Aided Design* 161 (2023), 103542.

Nan Zheng, Xiaoya Zhai, Jingchao Jiang, and Falai Chen. 2024. Topology optimization of self-supporting structures for additive manufacturing via implicit b-spline representations. *Computer-Aided Design* 175 (2024), 103745.

Haikuan Zhu, Hongbo Li, Hsueh-Ti Derek Liu, Wenping Wang, Jing Hua, and Zichun Zhong. 2025. Designing 3D Anisotropic Frame Fields with Odeco Tensors. *ACM Transactions on Graphics (TOG)* 44, 4 (2025), 1–14.

A Implicit Geometry

As our concept builds on the idea of implicit geometry, we have also experimented with an approach to take the implicit representation one step further. In Section 4.3 the final geometry results as the boolean union of the object hull and isocontour sheets extruded with a fixed thickness τ. Therefore, n isocontours are extracted per potential field $\Phi_{x,y,z}$. Alternatively, the fields may also drive superimposed periodic functions, then compiled into one field

$$\Phi_s = \max_{d \in \{x,y,z\}} \frac{\cos^{-1}\left(\cos\left(2\pi\Phi_d n\right)\right)}{\pi}. \tag{13}$$

The final structure is then extracted as a single isocontour to isolate peaks of the superimposed triangle wave functions, thus requires adjusting for the correct isovalue to meet a specified target volume fraction. Although this approach seems more in line with the concept of implicit geometry representation, it comes with some drawbacks: The triangle wave functions included in Equation (13) introduce artificially high frequencies on the domain. Extracting reasonably smooth isocontours, then requires a drastically increased resolution of the background tessellation [Groen and Sigmund 2018], i.e., for the example in Figure 27 the background triangulation is recursively subdivided 3 ×, thus increased the number of triangles by factor 4^3. Furthermore, widths of the triangle waves naturally stretch with the wavelength. Thus, in case of modulated fields Φ, isolated peaks narrow in compressed regions and widen in relaxed regions as highlighted on the right. This not only complicates guaranteeing a minimum sheet thickness for manufacturing but also effectively contrasts the initial concept for better material concentration through field compression. Lastly, including the object hull is non-trivial as it requires incorporating another simulated and normalized distance field in Φ_s from Equation (13).

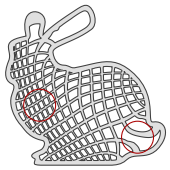

Figure 27: *Trigonometric modulation of the scalar field (Equation* (13)*) enables individual field components to be combined and sheet structures extracted simultaneously. However, compression then introduces artifacts that undermine its intended purpose.*