

Stress-Aligned Hexahedral Lattice Structures

Bukenberger, D. R.; Wang, J.; Wu, J.; Westermann, R.

DOI

[10.1111/cgf.15265](https://doi.org/10.1111/cgf.15265)

Publication date

2025

Document Version

Final published version

Published in

Computer Graphics Forum

Citation (APA)

Bukenberger, D. R., Wang, J., Wu, J., & Westermann, R. (2025). Stress-Aligned Hexahedral Lattice Structures. *Computer Graphics Forum*, 44(1), Article e15265. <https://doi.org/10.1111/cgf.15265>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



Stress-Aligned Hexahedral Lattice Structures

D. R. Bukenberger,¹ J. Wang,¹ J. Wu² and R. Westermann¹

¹TU Munich, Munich, Germany
{dennis.bukenberger, junpeng.wang, westermann}@tum.de
²TU Delft, Delft, the Netherlands
J.Wu-1@tudelft.nl

Abstract

Maintaining the maximum stiffness of components with as little material as possible is an overarching objective in computational design and engineering. It is well-established that in stiffness-optimal designs, material is aligned with orthogonal principal stress directions. In the limit of material volume, this alignment forms micro-structures resembling quads or hexahedra. Achieving a globally consistent layout of such orthogonal micro-structures presents a significant challenge, particularly in three-dimensional settings. In this paper, we propose a novel geometric algorithm for compiling stress-aligned hexahedral lattice structures. Our method involves deforming an input mesh under load to align the resulting stress field along an orthogonal basis. The deformed object is filled with a hexahedral grid, and the deformation is reverted to recover the original shape. The resulting stress-aligned mesh is used as basis for a final hollowing procedure, generating a volume-reduced stiff infill composed of hexahedral micro-structures. We perform quantitative comparisons with structural optimization and hexahedral meshing approaches and demonstrate the superior mechanical performance of our designs with finite element simulation experiments.

Keywords: modelling, mesh generation, computational geometry, geometric modelling

CCS Concepts: • Computing methodologies → Shape analysis; Volumetric models; Mesh geometry models

1. Introduction

Lightweight design plays a key role in a sustainable future and is a critical task in a variety of industries, including aerospace, automotive and architecture. Lightweight designs, which also maximize mechanical performance, are often approached with topology optimization. Topology optimization is an established structural design method for optimizing the distribution of a given material budget within a specified domain under a given set of mechanical boundary conditions. It discretizes the design domain with a voxel grid and iteratively performs gradient-based material deposition to optimize for the stiffest layout. Especially in 2D scenarios, this process has been extensively studied [BS13]. The iterative optimization process involves intensive computation, especially for 3D structures as shown in Figure 2(a), thus demanding for dedicated GPU solvers [WDW16] or high-performance computing clusters [AAL15].

Lattice infills, also referred to as multi-scale approaches [WSG21], have shown a computationally less expensive alternative for creating high-resolution designs in 3D. It is well-established that in stiffness-optimal designs, material is aligned with orthogonal principal stress directions [Ped89]. In the limit of material

volume, this alignment forms micro-structures resembling quads or hexahedra.

Furthermore, as shown in recent work [SOG*22, JOB*24] and confirmed by our experiments (Table 1), *wall*-structures can significantly enhance the mechanical performance of 3D micro-structure designs. Achieving such *wall*-structures with a geometric approach necessitates a mesh comprising finite cells with regular vertices, making a pure hexahedral mesh a practical choice. While alternatives like tetrahedral or prismatic elements could be considered, the orthogonal nature of the principal stress directions in 3D favours the use of hexahedra, which also have orthogonal edges. This structure allows for easier alignment with the stress field and simplifies the embedding of subdivision micro-structures in each cell. Additionally, using non-hexahedral cells would complicate the subdivision process, as micro-structure geometries must be defined for each cell type, where general polyhedra introduce challenges due to irregular vertex valences. Although relaxing the constraint on *pure* hexahedral meshes is possible, it would increase the complexity of the overall procedure without offering clear benefits.




Figure 1: Stages of our pipeline: The input tetrahedral mesh (a) and applied loading conditions yield the stress tensor field (b). Our deformation aligns the tensor field to an orthogonal basis (c) on which the hexahedral lattice (d) is constructed. With the inverse deformation, we extract a stress aligned hexahedral mesh (e). This hexahedral lattice is materialized with beam (f) and wall (g) micro-structures, reducing the overall amount of material consumed by the object while maintaining an optimized mechanical performance under load.




Figure 2: A topology optimization result (a) compared to geometric structural optimization methods, guided by the induced stress field. All designs use only a fraction of the fully solid object's volume.

For given boundary shape and conditions, the challenge is to compile a globally consistent hex-lattice that follows the spatially varying stress directions. This is difficult due to the existence of degenerate points at which principal stress directions exchange their type and stress trajectories of the same type can cross. The stress field can be smoothed to filter out such regions [AJL*19], producing structures with a high degree of regularity yet significantly lower mechanical performance than topology optimization (see Figure 2b). Frame-aligned hex-dominant meshing [GPW*17] with the principal stress directions as guiding frame [WWG21] faces the same challenge. It cannot produce meaningful structures without significant distortions of the stress field (see Figure 2c) and, in general, the resulting structures are graphs which do not obey the composition rules of meshes. While in 2D, the trajectory-based approach [WWW22, WWW23] naturally forms a quad-dominant lattice, this does not translate to 3D such that intersecting stress trajectories would form a hex-lattice. In 3D space, there is no guarantee that such a stress line intersects even one other trajectory while transcending the object.

Therefore, we propose a technique that generates a 3D stress field-guided hex-lattice. As quads are the preferred 2D element shape for single load conditions (meaning multiple loads but acting at the same time), our structural design analogously comprises solely hexahedral elements with mutually orthogonal edges and faces that aim to align with the 3D stress tensor field (Figures 1f and 1g).

Our method builds upon the core concept of transforming an input tetrahedral mesh such that its inherent stress field aligns with an orthogonal basis. Therefore, we borrow concepts from cubification and object stylization methods [LJ19, LJ21, LZS*21, ZGL*23]. Our approach extends the deformation concept from triangulated manifold surfaces to tetrahedral meshes, aligning orthogonal stress tensors instead of surface normals. In the deformed state the resulting shape is discretized with a regular hexahedral grid. The inverse transformation retrieves a stress-aligned hexahedral mesh such that at each vertex, its edges are aligned with the orthogonal stress field. Eventually, hollow micro-structures replace the hexahedral cells, reducing the overall amount of material consumed by the object, while maintaining as much stiffness as possible. We evaluate our results by quantitative comparisons with related works, and can demonstrate superior mechanical performance of our designs with finite element simulation experiments.

2. Related Work

Despite some concepts in hexahedral meshing being related to ours, the primary goals we pursue are significantly different. Our focus on optimizing structural stiffness may sometimes compromise element quality and input fidelity in favour of achieving optimal stress field alignment within the object's interior. Consequently, our related work section does not aim to recapitulate the already thoroughly explored field of hexahedral meshing [PCS*22]. Instead, in the following, we collect and summarize work that has inspired or is closely related to our pursued goals.

Field alignment. A core principle in our approach is aligning the hexahedral mesh to a frame field. This is a complex and challenging problem because some field singularities cannot be resolved using hexahedra. In the concept proposed by Nieser *et al.* [NRP11], a coarse meta mesh is manually constructed, with singularities placed at favourable positions to ensure that the extracted field is suitable for generating a pure hexahedral mesh. However, the authors acknowledge that poorly placed singularities can lead to drastic deformations of the mesh. Since minimizing deformation is crucial for maintaining accurate stress field alignment in our approach, we opted for a method approaching singularities with implicit smoothing in the field deformation step. Similarly, the approach of Gao *et al.* [GJTP17] heavily smooths the generated frame field, resulting in a hex-dominant mesh that may include arbitrarily large polyhedra. At-most-hexa meshes [BTL22] also align to a frame field based on the input hull but by default may include elements topologically smaller than hexahedra, such as prisms or tetrahedra. Our proposed method constructs a pure hexahedral mesh aligned with a stress tensor field. Singularities in this stress field are not based on the object's shape but arise from the forces acting within the preceded simulation. Instead of smoothing, the deformation optimizes for the best compromises between extensive alignment of the hex-lattice to the stress field and accuracy in low-rank singularities.

Cubic maps. Constructing constrained shape deformations in form of PolyCube mappings [THCM04] is a core concept used in many hexahedral meshing frameworks. However, pre-computing PolyCube topology continues to pose a notable challenge. Progress in this area has been driven by advancements in labelling methods [GSZ11], exemplified by the utilization of graph cut segmentation techniques [LVS*13] and their advancements [DPM*22].

Recent trends in the field also adopt a similar cubical deformation and its inversion, as used in our work. Nevertheless, the intricate aspect of the PolyCube decomposition typically entails manual intervention, involving numerous smoothing steps and user-driven refinements [LZS*21] to reconstruct certain details. Moreover, this process lacks orientation invariance and results depend on the object's initial orientation. Similarly, approaches such as HexBox [ZGL*23] rely on surface projection techniques applied to a cubified representation of the mesh, also necessitating manual construction of the HexBox through user input.

In this context, the principle of cubification holds significant importance and is often tweaked and tailored for specific application requirements. In our approach, we utilize the cubic stylization principle [LJ19, LJ21], employing an as-rigid-as-possible (ARAP) deformation [IMH05, FSA23] with specific extensions to the initial core structure. Notably, we utilize only spokes within the rims-and-spokes configuration, as recently discussed in revised ARAP approaches [FSA23] and extend the alignment to full orthogonal matrices rather than only surface normals.

Multi-scale structures. Basic topology optimization approaches tend to produce strong single-truss-based structures carrying the majority of the load with thinner structures supporting the hull. However, compared to that, the micro-structure approach or porous materials have the advantage to be much more resilient to single defects [WAWS18]. Therefore, many recent approaches focus

on lattice infill, *i.e.* multi-scale, structures as the prevalent technique for converting the optimized density field into a binary material layout [WSG21]. There are also concepts that contrast the classical goal of improving stiffness, by utilizing micro-structured meta materials designed to enhance flexibility [SBR*15] and deformation [TTZ*20] in manufactured parts. The approach by Wang *et al.* [WWW22] for 2D designs proposes porous layouts based on mesh-like structures. However, in 3D domains, these structures do not necessarily conform to conventional meshes. Instead, they form graphs with irregular topology [WWG21], yet suitable to be produced with contemporary manufacturing technologies. Theoretical findings, on the other hand, advocate for the incorporation of *wall-like* structures in 3D designs in addition to solidified *beams*. To achieve such structures, a conforming hexahedral mesh is required.

3. Background

Compliance, *i.e.* the reciprocal of stiffness, is a well-established measure in structural design to express how much a structure deforms under a given load. Thus, minimizing compliance results in a structure with increased stiffness. Compliance is expressed as

$$c = \frac{1}{2} U^T \mathbf{K} U \quad (1)$$

where \mathbf{K} is the structural stiffness matrix. The external loads are expressed in the force vector F , and the displacement vector U is determined by solving the equilibrium

$$\mathbf{K} U = F. \quad (2)$$

In classical topology optimization, the compliance is minimized under a constraint on the amount of consumed material (volume). As compliance is not scale-invariant, the performance of the optimized design is always expressed as $\frac{c}{c_0}$ in relation to the compliance c_0 of the fully solid object. Furthermore, the volume of the optimized structure aims to be a fraction of the initial object and is thus given as $\alpha = \frac{v}{v_0}$, where v_0 is the volume of the solid object.

3.1. Stress field

We assume a given tetrahedral mesh (with vertices V and tetrahedra T) and specify loading conditions as sets of vertices that remain fixed in certain regions and others where forces are applied. In this work, we assume a *single-load* scenario, which may feature multiple forces but acting at the same time. Standard finite element methods yield deformation vectors at each vertex, which are interpolated across elements to compute the *Cauchy* stress tensors [GM20]. These stress tensors are then projected back to each vertex $v_i \in V$. Orthogonal principal stress tensors Ξ_i follow as their eigenbases, respectively. The eigenvalues of the *Cauchy* stress tensors further yield information on the magnitude of stress in certain regions of the object and serve as basis to derive the scalar *von Mises* stress norm σ_v . Our use of the *von Mises* stress is two-fold: We use it as an indicator of the edge thickness of the hexahedral design (Figure 8), and to prioritize the alignment of the design in regions with important mechanical properties (λ_i in Equation 4).

Aligned structures. The theoretical concept introduced by Michell in 1904 [Mic04], commonly summarized as *Michell*

Truss, serves as a cornerstone for designing stiffness-optimized lightweight structures. Michell demonstrated that an optimally stiff design under load experiences no shearing stress, leaving only tension and compression stresses. This motivates the widely established concept of micro-structures aligning with the principal stress directions, where shear stress is zero. The rationale behind this is that most materials perform significantly better under tensile and compressive stresses than under shear stresses.

4. Method

The core concept underlying many hexahedral meshing procedures is the computation of a robust mapping between a given frame field and a hexahedral domain subdivision. In classical hexahedral meshing, the goal is to achieve a final hexahedral layout which aligns with the domain boundaries and comprises elements satisfying shape and size constraints. The frame field needs to be selected with care, since it may contain local constellations which simply cannot be resolved with hexahedra. Our approach deviates from classical hexahedral meshing and strives for the best mapping under a global alignment objective, with the frame field being the principal stress directions in a 3D tensor field. Our proposed solution to achieve such a mapping covers the following principle steps:

- Firstly, a deformation aligns the stress field to an orthogonal basis by transforming an input tetrahedral mesh (Section 4.1).
- In the deformed state, hex-lattice structures are introduced and deformed back, using the inverse transformation (Section 4.2).
- Finally, elements in the stress-aligned hexahedral mesh are replaced by hollow micro-structures with adjustable fill rate to achieve a selected volume budget (Section 4.3).

4.1. Constrained field deformation

We introduce an optimization approach that computes a deformation of a tetrahedral mesh, such that the individual per vertex stress tensors Ξ_i align to a common orthogonal basis. We consider the set V of all vertices v_i of the input mesh ($\tilde{v}_i \in \tilde{V}$ in the deformed state, respectively) and $\mathcal{N}(i)$ are sets of locally adjacent neighbours. Notably, in these local neighbourhoods, we actually only consider directions to direct neighbours of vertex v_i (*spokes* only instead of the full *spokes-and-rims* structure), providing a more flexible basis for the deformation optimization. The system stiffness matrix $\mathbf{W} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ comprises standard 3D-cotan weights w_{ij} [Cra19], encoding the relation of vertices i and j connected with edge e_{ij} , derived from the tetrahedral mesh. Further denote $\mathbf{R}_i \in \mathbb{R}^{3 \times 3}$ individual rotation matrices per vertex. Local neighbourhood geometry is denoted with $d_{ij} = v_j - v_i$ at rest and $\tilde{d}_{ij} = \tilde{v}_j - \tilde{v}_i$ in the deformed state, respectively. v_i encodes the Voronoi cell volume of vertex v_i , and λ_i allows to weight the importance of individual vertices. The energy to be minimized is then expressed as

$$\underset{\tilde{V}, \{\mathbf{R}_i\}}{\text{minimize}} \sum_{v_i \in V} \sum_{v_j \in \mathcal{N}(i)} \frac{w_{ij}}{2} \|\mathbf{R}_i d_{ij} - \tilde{d}_{ij}\|_F^2 + \lambda_i v_i \|\mathbf{R}_i \Xi_i\|_1. \quad (3)$$

Similar to the *Cubic Stylization* approach by Liu *et al.* [LJ19], the first part of the objective function acts as an ARAP constraint. The use of the L_1 norm (with $\|\mathbf{X}\|_1 = \sum \|x_i\|_1$) in the second term encourages axis alignment because the L_1 norm sums the absolute val-




Figure 3: The stress field (+) on the left is formed for the given loading conditions (| ←). In the centre, the object is deformed, aligning the stress field to an orthogonal basis. In this state, the hex-lattice (□) is inserted and transformed back (right).

ues of the components of a vector across all dimensions. Minimizing this sum pushes the vector to align with one of the basis axes. For a unit vector, the L_1 norm reaches its minimum value of 1 when the vector is aligned with one of these axes, since only one component will be non-zero in this case.

Thus, the per-vertex optimization step can be formally expressed as

$$\mathbf{R}'_i = \arg \min_{\mathbf{R}_i \in \mathbb{R}^{3 \times 3}} \frac{1}{2} \|\mathbf{R}_i D_i - \tilde{D}_i\|_{\mathbf{W}_i}^2 + \lambda_i v_i \|\mathbf{R}_i \Xi_i\|_1 \quad (4)$$

where D_i and \tilde{D}_i ($\in \mathbb{R}^{|\mathcal{N}(i)| \times 3}$) stack all local spokes-edge vectors (not normalized) row-wise at rest and deformed state, respectively. The $|\mathcal{N}(i)| \times |\mathcal{N}(i)|$ sized matrix \mathbf{W}_i with 3D cotan-weights is employed in Equation (4) with $\|\mathbf{X}\|_{\mathbf{W}}^2 = \text{Tr}(\mathbf{X}\mathbf{Y}\mathbf{X}^T)$.

Optimization is performed with the *Alternating Direction Method of Multipliers* (ADMM) [BPC*11], where we use the common initialization parameters [LJ19]. A decisive difference in our application lies in the orthogonal Procrustes [GD04] formulation: Instead of considering only the normals of surface vertices, we expand the formulation with the principal stress tensors Ξ_i of the tetrahedral mesh vertices as formulated in Equation (5).

$$\mathbf{R}_i^{k+1} = \arg \max_{\mathbf{R}_i \in \mathbb{R}^{3 \times 3}} \text{Tr}(\mathbf{R}_i \mathbf{M}_i), \text{ with} \quad (5)$$

$$\mathbf{M}_i = \begin{bmatrix} D_i^T & \Xi_i^T \end{bmatrix} \begin{bmatrix} \mathbf{W}_i & \\ & \mathbf{P}^k \end{bmatrix} \begin{bmatrix} \tilde{D}_i \\ \mathbf{R}_i^k \Xi_i - \mathbf{U}_i^k \end{bmatrix}$$

Further, we have $\mathbf{P}_i, \mathbf{U}_i \in \mathbb{R}^{3 \times 3}$, where \mathbf{P}_i is a diagonal matrix comprising three penalty values $\rho_x^k, \rho_y^k, \rho_z^k$ and \mathbf{U}_i^k stacks scaled dual vectors to the aligned stress tensor. In each ADMM iteration, the optimal local rotation matrices are iteratively updated from the singular value decomposition $\mathbf{M}_i = \mathcal{U}_i \Sigma_i \mathcal{V}_i^T$ as $\mathbf{R}_i^{k+1} = \mathcal{V}_i \mathcal{U}_i^T$ (with signs flipped such that $\det(\mathbf{R}_i^{k+1}) > 0$). This local iteration determines the best alignment for each stress tensor Ξ_i independently, thus facilitating parallelized execution. A reduction to two dimensions, as exemplified in Figure 3, follows analogously.

In the global step, a single linear system is solved [SA07] for the deformation with the best local alignments. Unless explicitly specified, the vertices are not bound to any constraints beyond the neighbourhood geometry encoded in the global cotan-matrix \mathbf{W} .




Figure 4: 2D example of Equation (6) with the mapping of \tilde{u}_i , enclosed in the deformed element with vertices $\tilde{v}_0, \tilde{v}_1, \tilde{v}_2 \in \tilde{V}_t$, back to the original shape with vertices $v_0, v_1, v_2 \in V_t$ resulting in u_i .

4.2. Introducing hexahedral structures

Within the deformed structure, we generate a hex-lattice that is oriented on the same basis as the aligned stress tensors. The hex-lattice with vertices \tilde{u}_i is constructed within the deformed shape of the tetrahedral mesh. A hexahedral cell is included, only if all of its eight vertices lie within a tetrahedron $t \in T$ of the input mesh (however, not necessarily all in the same tetrahedron). This guarantees that all hexahedra are entirely included in the inputs outer hull. Transforming this hex-lattice back to the undeformed object's pose is realized by applying the mapping

$$u_i = b_t(\tilde{u}_i) [V_t - \tilde{V}_t]. \quad (6)$$

V_t and $\tilde{V}_t \in \mathbb{R}^{4 \times 3}$ denote the four vertices at rest and deformed state of the tetrahedron t , which includes the created vertex \tilde{u}_i of the hexahedral mesh. The corresponding barycentric coordinates of \tilde{u}_i in t are denoted by $b_t(\tilde{u}_i)$. These are analogously used to interpolate other field related quantities at the hexahedral vertices, such as the stress tensors σ_i from the tetrahedral vertices. Insertion of the hex-lattice and its back-deformation are included on the right in Figure 3 and shown in detail in Figure 4.

Hull layer strategy. The grid structure employed within the deformed object does not yet yield a conforming mapping to the object's outer faces. By design, the grid does align with object boundary features per default but only follows the induced stress field. Nevertheless, improved input fidelity can be approached with the two strategies presented in Figure 5. The progressive grid approach also features elements *partly* inside the input object, thus overlapping the boundary. A simple normal-projection of the grid hull does

not introduce new elements and preserves alignment well. However, as pointed out in Figure 5 (circled), this may cause inverted elements when vertices are projected into a concavity. Our conservative grid strategy incorporates only elements *fully* enclosed in the input hull. The resulting gap between the introduced structure and the object hull is filled with a padding layer (Figures 6c and 6d). Outer *open* quads of the hexahedral structure are projected along their normal directions onto the tetrahedral mesh's triangular hull. This spans up a padding layer between the inner hex-structure and the projected quads, again consisting solely of hex elements. When vertices of the aligned grid are projected onto the hull, the distance they travel is smaller than the spacing in the grid itself. Because the field is considered smooth and does not change significantly over such small proximities, in general, the resulting structures on the hull are also well aligned to the field. However, this excludes edges directly connecting the projected hull with the inner structure as these edges align with the projection directions, *i.e.* the normal directions. Further, in certain constellations, the projection may lead to self-intersecting projected quads. Thus, we employ a smoothing step on the new layer of hexahedra (Figures 6d and 6e), untangling degenerate hull-quads. Whereas feature alignment and accurate hull approximation are crucial for classic hexahedral meshing approaches, this step is more of a cosmetic feature for our approach where the internal field alignment matters most. Therefore, the padding could either be omitted or replaced by a more advanced solution as the scaffolding and deformation technique by Gao *et al.* [GSP19] or post-processed with dedicated hex-optimization techniques [LSVT15].

If the chosen resolution of the hex-lattice discretization is insufficient to capture certain details of the input geometry, two kinds of artifacts may arise: Fine details, thinner than a hexahedron, cannot be reconstructed and result in missing geometry. In other rare cases, the padding layer spanned up by the projection yields a non-untangleable mesh. This occurs if the quads of the *open* hull do not form a manifold mesh. This can be prevented by either increasing the resolution or performing a smoothing operation on an intermediate voxel representation. As shown in Figure 6(c), a structure is represented with a 3D voxel mask, indicating if a hex-cell is used or not. The mask is filtered using $3 \times 3 \times 3$ sized binary kernels, such that only cell configurations remain where adjacent hex-cells have to share a quad-face. In a configuration where cells only share an edge or single vertex, the hex-cell with the least number of direct neighbouring cells is removed. This usually affects no more than 0.1% of hex-cells from the most outer layer.

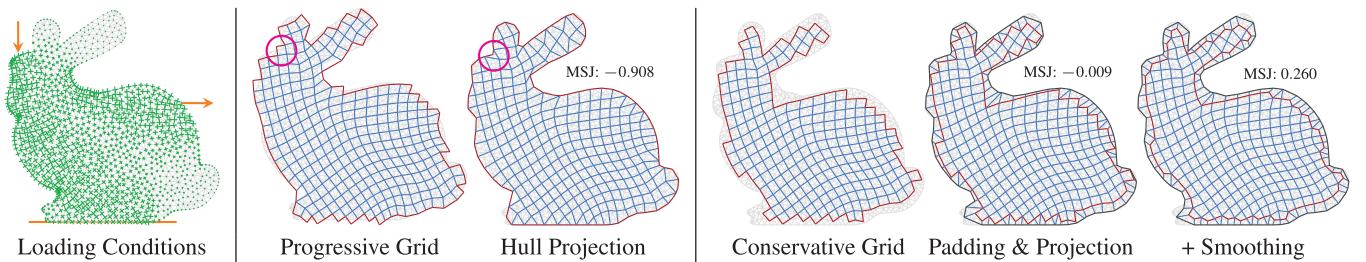


Figure 5: Progressive grids and their projection better preserve individual element alignments in the hull layer but can cause inverted vertices if projected into a concavity (circled). Affected elements cannot be untangled by smoothing. Our conservative strategy with padding offers better results by default as only a single edge (2D) or a quad (3D) per element is projected and can be further improved by smoothing.




Figure 6: The stress field in (a) emerges from loading conditions keeping the left face of the cube fixed and applying a tangential downwards force on the right. In (b), the tetrahedral mesh is deformed and aligns the stress field to the common basis, visualized by the inserted cubical frame. Further, the hex-lattice is inserted (c), the hull padded (d), smoothed (e) and transformed back to the original shape (f).




Figure 7: To reduce volume, hexahedra (left) are replaced by hollowed edge-inward-extruded micro-structures (centre, right). A volume fraction α is met by adjusting the extrusion parameter τ .

4.3. Micro-structures

To reduce the volume of a structure, *i.e.* the weight of a computed layout, the individual mesh cells are replaced by hollowed micro-structures. Therefore, as illustrated in Figure 7, the edges of each cell are extruded inwards. The micro-hexahedra spanned by the extrusion consume only a fraction of the volume of the original elements. The exact fraction is controlled via the parameter τ . Since all newly created vertices of a micro-structure are interpolations of their original edge, face or cell vertices, a conforming structure can be generated, *e.g.* two adjacent micro-structures share the same four interpolated vertices from the quad that separated their original hex-cells. For detailed interpolation matrices, see Appendix A. The reduction in volume is relative to each cell, with a global parameter τ analogously translating to the volume of the full structure.

The analogue to the common de-homogenization in 2D is to extrude only the edges of each element, forming solid *beams* as visualized in the centre of Figure 7. With 3D objects, however, the design is not limited to *beams* alone but may also feature *wall*-like structures by extruding the faces of each cell as well. Therefore, the *beam*-structure is simply augmented with six more hexahedra as shown on the right of Figure 7, creating a hollow cube.

Equation (7) defines the extrusion parameter τ as a function of the desired volume fraction $\alpha = \frac{v}{v_0}$, where v is the volume of the micro-structures, and v_0 is the volume of the solid object. We provide separate equations for *beam*-only (*b*) and *walled* (*w*) micro-structures, denoted by the respective subscripts. While parameter τ directly controls the thickness of *beams* and *walls* in the micro-structures, the volume fraction α is the more intuitive design parameter. Converting from τ to α is straightforward, as detailed in

Appendix B, but going the other way (expressing τ in terms of a given α) is more challenging. Specifically, for the *beam* micro-structure, the relationship between τ_b and α results in a *casus irreducibilis*, meaning there is no real-valued inverse function for arbitrary α . As a result, the expression for τ_b is an approximation, though it is exact when $\alpha = \frac{1}{2}$. More details on this can be found in Appendix B.

$$\tau_b(\alpha) = \frac{\cos^{-1}(1 - 2\alpha)}{2\pi} \quad \tau_w(\alpha) = \frac{1 - \sqrt[3]{1 - \alpha}}{2} \quad (7)$$

The insertion of micro-structures requires a conforming mesh with finite cells and, in particular, vertices with an in-cell valence of 3, *e.g.* hexahedra, triangular prisms and tetrahedra. General hex-dominant meshes, *e.g.* by Gao *et al.* [GJTP17], may feature arbitrary polyhedra with vertices of higher valence, thus are not suitable in general. At-most-hexa meshes [BTL22] solely consist of elements which are topologically smaller than hexahedra but are also encoded as such. Therefore, we can generate micro-structures for at-most-hexa meshes, but the featured collapsed edges cause the system matrix to become impractical for simulation. In general, however, any pure tetrahedral or hexahedral structure, featuring edges of predominantly similar lengths, is suitable for this micro-structure concept.

Targeted thickness. To approach bone-like porous structures [WAWS18], the global volume constraint is replaced by a local volume constraint to avoid accumulation of material. In topology optimization, this enforces that material is first distributed along the mechanically relevant stress directions, resulting in *beam*- and *wall*-like structures in 2D and 3D, respectively. In our context, we propose τ^* as extrusion parameter, determined individually for each vertex of the aligned hex-lattice structure. As link to the global stress field, we found the scalar *von Mises* stress σ_v a suitable heuristic. The stress tensors at the hexahedral vertices are interpolated according to Equation (6). Thus, a recomputation of the stress simulation on the new structure is not required. Further is the *von Mises* stress normalized by the global maximum value $\hat{\sigma}_v$ and squared to emphasize high-stress regions. However, individual extrusion parameters at each vertex yield non-uniform cell volume fractions, thus a global volume fraction α is no longer given.

$$\tau^* = \left(\left(\frac{\sigma_v}{\hat{\sigma}_v} \right)^2 \cdot (1 - |p|) + \max(p, 0) \right) \cdot (\tau_{\top} - \tau_{\perp}) + \tau_{\perp} \quad (8)$$

As a countermeasure, Equation (8) formulates vertex-individual parameters τ^* based on a global parameter $p \in [-1, 1]$. To meet




Figure 8: 2D micro-structures with uniform τ on the left, τ^* on the right and $\alpha = \frac{1}{2}$ for both. The σ_v stress field (middle, scaled for visualization) results from the same forces as shown in Figure 3.

a specified global α , all τ^* are p -tuned using standard numerical optimization. Further is τ^* limited by the natural maximum $\tau_{\top} = \frac{1}{2}$ for a fully solid cell. Our chosen non-zero lower boundary τ_{\perp} corresponds to a minimum of $\alpha = 5\%$ for an individual cell [WWW23]. Figure 8 illustrates a 2D example of the quad-mesh, extracted in Figure 3, with uniform τ_b cells on the left and individually extruded τ_b^* cells on the right. The effect of targeted thickness results in improved mechanical performance, expressed as the relative compliance $\frac{c}{c_0}$, which further decreases with τ^* . For reference, the same evaluation on the initial unoptimized triangle mesh input gives $\frac{c}{c_0} = 3.316$ for τ_b and $\frac{c}{c_0} = 1.646$ for τ_b^* , respectively. Equation (8) also offers a degree of freedom for using σ_v : In this example, the normalization was replaced by an inverse cumulative distribution function, yielding an equalized thickness distribution.

4.4. Extensions

In the following, we discuss modifications of the proposed method that allow for improving and fine-tuning the results.

Optimal alignment. Many hexahedral meshing approaches that rely on a polycube decomposition assume a suitable orientation of the input [LZS*21, ZGL*23]. With our method, the primary objective is the alignment of the hex-lattice to the stress field in the object's interior rather than surface features. Therefore, we can determine an optimal orientation automatically as the eigenbasis of a weighted average of all stress tensors σ_i . This option was used by default in all included results. Optionally, a specific orientation could be defined manually. Notably, if the optimal orientation is not given or automatically determined, the first few global optimization steps rotate the object to align with a suitable orientation.

Considering surface normals. Especially in regions of contact, *i.e.* input mesh vertices where the simulated forces are applied or which are kept fix, results can be improved if the hex-lattice also aligns to the outer hull. Let vertices $v_j \in V_{\perp}$ be the set of fixed or moved vertices. Then, the stress tensor Ξ_j is rotated (by minimal Euler rotation) such that its most normal-aligned axis actually becomes parallel to the surface normal n_j . Further, we can enforce stronger alignment in these points by selectively increasing the λ_j . This has a similar effect as cubification, *i.e.* regions of contact are flattened uniformly and become orthogonal to one of the coordinate axes. Thus, when introducing the hex-lattice (Section 4.2), the hexahedra closest to the hull tend to form a single-level layer rather than

staircases approximating a rounded hull surface. However, this additional alignment only makes sense if the forces act orthogonally on the surface, but not in a tangential configuration, *e.g.* as on the cube example in Figure 6.

Adjusting axis alignment. The λ_i parameter in the energy minimization Equation (3) allows one to weight the local alignment against the global ARAP-deformation, where higher values enforce stronger alignment. We have experimented with individual weights per vertex, *i.e.* scaled by the largest eigenvector from the local stress tensor computation. This allows for prioritizing regions with higher stress to have better alignment on the cost of regions with less stress, *i.e.* where sub-optimal alignment is less crucial. Furthermore, we have made a similar observations as Li *et al.* [LZS*21], which is that ramping up the *cubeness* over multiple iterations leads to overall better results.

5. Discussion and Evaluation

We evaluate the mechanical performance and stress-alignment of the structures generated with our approach in comparison to similar results of structural design and comparable hexahedral meshing methods. Shown comparisons are based on results available on HexaLab [BTP*19] and were selected as representative subsets of objects generated with alternative methods. The included numbers and plots are the results of FEM simulations using standard parameters (Young's modulus of 1000 and Poisson ratio of 0.3). Since the inputs are tetrahedral meshes and both our hex-lattice and the micro-structure results are pure hexahedral meshes, any common FEM library can be utilized for numerical stress simulation. The same forces are applied to the fully solid object and the different optimized lattice designs using micro-structures, to measure and compare the mechanical performance under load but with significantly reduced amounts of material.

5.1. Mechanical performance

Table 1 presents an evaluation of the compliance computed for different objects and hexahedral structures generated with different methods. This allows for comparing the performance of *beam*- and *wall*-structures introduced in Section 4.3, both with volume fractions of $\alpha = \frac{1}{2}$. All listed numbers are given as $\frac{c}{c_0}$, where c_0 is the compliance of the fully solid object, *i.e.* the lower the better. We tried to level the resolution of the compared structures as good as possible. An example is illustrated in Figure 9.

For each object in Table 1, the first row includes the TetWild [HSW*20] mesh that served as basis for our approach and also acts as a baseline in this comparison. Unsurprisingly the tetrahedral mesh performance is usually the weakest, as its internal structure is quite disturbed and not aligned to any meaningful direction. To highlight the effectiveness of *wall* micro-structures, the comparison also features other pure hexahedral meshes generated with an automatic block decomposition algorithm [LPP*20], state-of-the-art polycube mappings [DPM*22] and semi-manual approaches requiring user-designed volumetric functions [GMD*15], skeletons for tubular shapes [LAPS17] or specified blocks [ZGL*23] to guide

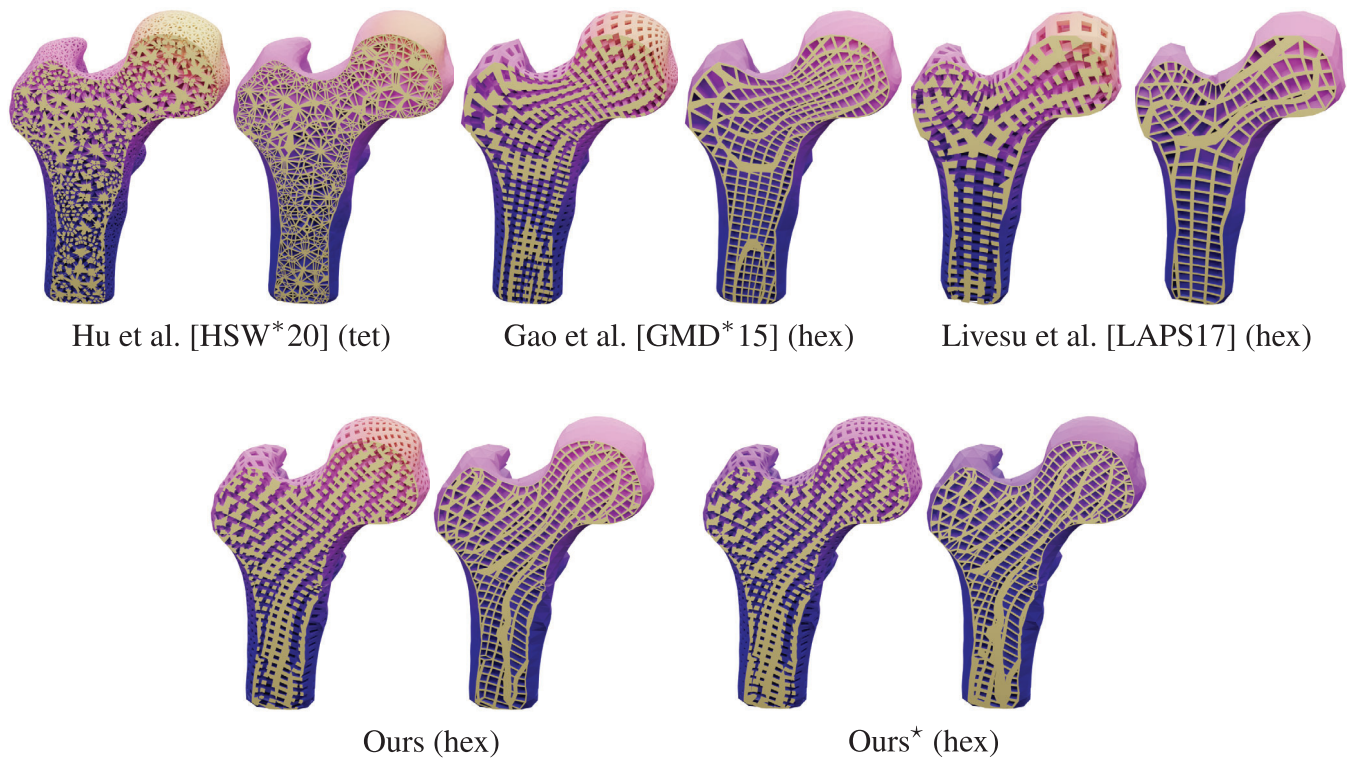


Figure 9: Results correspond to the numbers listed in Table 1 of the femur and meshes created by different methods. Each version is shown with beam (left) and wall (right) micro-structures in a cut-open view. Colouring visualizes the induced deformation under the given loading condition. The femur is fixed at the bottom and a force vector is applied at the top. Our results marked with a * employ the non-uniform τ^* extrusion to generate individually scaled micro-structures per cell. The locally varying thickness of extruded beams and walls in the τ^* version is noticeable in the lower parts of the femur.

the decomposition. These methods produce feature aligned meshes with high input fidelity and hexahedra of objectively high quality. We apply our micro-structure concept on these hexahedral meshes and directly use the results (also hexahedra) in the simulation. As our mesh is tailored to a specific loading scenario, for a reasonable comparison, we choose applied forces to be at least in favour of the competing structures or generate them using identical frame fields. As the results demonstrate, in all scenarios, our approach generates meshes that surpass the meshes generated with alternative methods in terms of mechanical performance. The observation that better alignment leads to stiffer structures is consistent with findings from classical topology optimization. Further, the measured numbers confirm that *wall*-like structures generally outperform *beam*-only structures, while consuming the same volume fraction. Table 1 also features our results with non-uniform extrusion rates τ^* , indicated with a *. This reduces the resulting compliance factor significantly, often by a factor of around two, and in some cases even lower. As shown in Figure 10 with the *fertility* and *kitten* model, also complex object topology is no limitation with our method. The only requirement for achieving a complete structural design that reproduces thin features is a sufficiently high resolution of the hex-lattice. Lower resolutions may result in defects, such as missing geometry. This issue can be partially mitigated by using the progressive grid method, as shown in Figure 5, with an additional padding layer.

However, this approach may generate some inverted elements due to the projection, necessitating more untangling operations. The *kitten* object is an example where this technique was used to include the tail despite the coarse resolution.

In Table 2, we compare the mechanical performance of infill structures generated with the structural optimization methods by Arora *et al.* [AJL*19] and Wu *et al.* [WWG21]. Notably both generate field-aligned edge-only designs, not supporting our micro-structure concept. To enable a fair comparison, we apply Arora *et al.*'s technique, solidifying the edge-graphs by extruding cylinders along the edges and computing a robust Boolean manifold surface for the extraction of a tetrahedral mesh (see Figure 2). The radius of the cylinders is chosen to meet the desired volume fraction $\alpha = \frac{1}{2}$. We apply the same technique to solidify our lattice structures for comparison in Table 2. This approach mitigates the bias introduced by using differently tessellated FEM structures [BPM*95, SHD*18, SHG*22], *i.e.* tetrahedralized *trusses* versus our hexahedral micro-structure *beams*. The compliance values also correlate with the observations in Figure 11, showing that Wu *et al.*'s lattices are generally well-aligned but feature arbitrary interconnecting diagonal edges. These consume material from the limited global budget, but do not significantly contribute to the overall stiffness and weaken the more structurally relevant edges.

Table 1: We evaluate the mechanical performance of different hexahedral meshes under identical loading conditions with a volume fraction of $\alpha = \frac{1}{2}$ using beam and wall micro-structures. Numbers on the right give the compliance increase $\frac{c}{c_0}$ with reduced volume compared to the fully solid object, i.e. the lower the better.

	Method	beam	wall
Cube	Hu et al. [HSW*20]	3.070992	2.880481
	Axis Aligned Hexahedra	3.486986	2.751497
	Ours	2.998545	2.531853
	Ours*	2.584732	2.185485
Femur	Hu et al. [HSW*20]	3.082209	2.921756
	Gao et al. [GMD*15]	2.933029	2.478329
	Livesu et al. [LAPS17]	2.927899	2.446718
	Ours	2.747245	2.435161
	Ours*	1.995508	1.624553
Fertility	Hu et al. [HSW*20]	3.302256	3.105790
	Gao et al. [GMD*15]	3.052227	2.491892
	Livesu et al. [LAPS17]	3.197407	2.607759
	Ours	2.920798	2.458964
	Ours*	2.522228	2.043973
Kitten	Hu et al. [HSW*20]	3.241617	3.050779
	Gao et al. [GMD*15]	3.313399	2.662459
	Livesu et al. [LPP*20]	3.074764	2.581898
	Ours	2.977281	2.497409
	Ours*	2.246125	1.804384
Spot	Hu et al. [HSW*20]	3.223385	3.032241
	Zoccheddu et al. [ZGL*23]	3.054186	2.484063
	Ours	2.846452	2.426191
	Ours*	2.487546	2.011845
Venus	Hu et al. [HSW*20]	3.327974	2.995528
	Gao et al. [GSP19]	2.942068	2.488273
	Dumery et al. [DPM*22]	2.912585	2.463419
	Ours	2.766663	2.363639
	Ours*	2.345045	1.895938

5.2. Field alignment

The concepts introduced by Arora et al. [AJL*19] and Wu et al. [WWG21] represent approaches for generating stress-aligned lattice structures, tailored towards sustainable manufacturing designs. Both approaches generate graph structures comprised of vertices that are connected via stress tensor aligned edges. While this results in closed triangular and quadrangular cells in 2D, the corresponding 3D structure does not generally comprise finite cells or faces. In Figure 11, we also compare our results and their aligned lattices in terms of adherence to the initial tensor field. Therefore, we interpret the input tetrahedral mesh and our hexahedral mesh also as edge graphs and compare their individual alignment to local field orientations. We measure the angle of an edge to the closest axis of its local principal stress tensor, and average at each vertex the measures for all its connected edges. In the tetrahedral mesh, edges do not follow a particular direction. This also reflects in the overall highest deviations in Figure 11. The average over all vertices in a structure is inset in the histograms, respectively. Averages of around 20° deviation in results of the *aligned* methods are unexpectedly high but can be explained as follows: The results by Arora et al.



Figure 10: Objects listed in Table 1 as a (solid) hexahedral mesh on the left, corresponding beam and wall micro-structure designs with $\alpha = \frac{1}{2}$ in the centre and on the right, respectively. Colour encodes deformation under load, normalized for visualization.

adhere more strictly to the initial grid, resulting in less alignment with the tensor field. In the conforming edge structure of Wu et al., many edges adhere to the local stress directions, but vertices are also connected by additional non-aligned diagonals, contributing to larger deviations on average. In our results, the highest per-vertex deviations are found at hull vertices. Since the hull padding

Table 2: Comparing relative compliance $\frac{c}{c_0}$ of field-aligned lattices using solidified trusses as shown in Figure 2 with $\alpha = \frac{1}{2}$. The same parameters and boundary conditions as in Table 1 apply.

	Method	truss
Cube	Arora et al. [AJL*19]	3.758369
	Wu et al. [WWG21]	3.405656
	Ours	3.298168
Femur	Arora et al. [AJL*19]	2.512450
	Wu et al. [WWG21]	2.188068
	Ours	2.509744
Fertility	Arora et al. [AJL*19]	2.716619
	Wu et al. [WWG21]	2.637801
	Ours	2.383448
Kitten	Arora et al. [AJL*19]	2.627741
	Wu et al. [WWG21]	2.692648
	Ours	2.521266
Spot	Arora et al. [AJL*19]	3.872281
	Wu et al. [WWG21]	3.267195
	Ours	2.668108
Venus	Arora et al. [AJL*19]	2.431033
	Wu et al. [WWG21]	3.546859
	Ours	2.389367

layer results from projections of inner vertices, local orientations are generally well preserved. However, if the projection creates inverted elements that require untangling, the smoothing process may move vertices and locally disturb the alignment. Therefore, when considering only inner vertices in our structures, the deviation averages are 18.0° for the *femur* and 16.2° for the *venus*. However, intersections of the histogram curves indicate that overall about 80% of our vertices have better aligned edges compared to the conforming lattices of Wu et al. This slight advantage of our hex-lattice is also reflected in the global averages over all vertices.




Figure 11: Comparing lattice alignments: The histograms on the left plot the angular deviation of vertex edges to the local stress tensor, the average deviation over all edges is listed in the inset. Corresponding lattice structures are shown on the right as full and cut-open versions, respectively. The colour-coded structures visualize the averaged edge deviation for individual vertices.

Singularities. Figure 12 illustrates loading conditions on a rectangular bar that provoke a singularity in the stress field. The method by Gao et al. [GJTP17] applies local smoothing to the frame field and extracts a quad- or hex-dominant mesh. This relaxed constraint on the mesh topology allows for more flexibility and better alignment of individual element edges to the field. However, this introduces many irregular vertices and cells of arbitrary degree. Arora et al. [AJL*19] employ global smoothing on the frame field before extracting the edge-graph. This approach tends to push problematic regions towards the object's boundary, which may result in artifacts and locally increased resolution in the extracted graph. Our proposed field deformation has a similar effect on the stress field as global smoothing, as singularities are approached by bending the field around them. Nevertheless, our extracted mesh has better alignment close to the boundary (e.g. upper and lower right corners of the bar), has no degenerate faces, guarantees inner vertices to be regular and generally consists of more evenly sized cells.

Varying loading conditions. In contrast to common hexahedral meshing options which optimize the hexahedral structure for best feature alignment and highest input fidelity, our method is tailored for the characteristics of a specific stress field. With different load scenarios, the stress field obtained from simulating on the tetrahedral mesh will be different. Figure 13 illustrates cut-open examples a *buddha* statue with varying fixed regions or forces applied. All three examples were generated with identical parameters, only varying the mesh resolution slightly to yield approximately the same number of hexahedral elements. Another significant difference to common hexahedral meshing methods becomes apparent in Figure 14: The *jet engine bracket* [WBM21] is faced with a static torsional force acting bidirectional on its handles. As the lattice follows the emerging stress field, resulting internal structures are wavy and curved. Untangling may introduce smoothing artifacts and limited resolution can lead to beveling of sharp features.




Figure 12: A rectangular bar is attached on the left with two forces applied on a single vertex each. As highlighted, the resulting stress field features a singularity where and denote the primary and secondary principle stress directions (in close proximity), respectively.




Figure 13: Cut-open visualizations of a buddha statue reveal hexahedral mesh variations under different loading scenarios. Orange boxes mark fixed regions, and arrows indicate the applied forces.




Figure 14: Static torsional forces are applied on a jet engine bracket [WBM21], shown on the left. The inside view on the right reveals the wavy structures that emerge in the resulting hex-lattice.

5.3. Element quality

The minimum scaled Jacobian (MSJ) is a measure of skewedness that is often used to express the quality of a hexahedral mesh and its suitability for robust computation. Although our hex-lattices are not primarily designed for this purpose, we can evaluate the MSJ to highlight the quality of our mesh and the individual cells. Additionally, to ensure applicability of our micro-structures, there should be no inverted elements ($MSJ < 0$). Plots in Figure 16 show the MSJ of our extracted hex-lattice and the resulting micro-structures. The inverse deformation applied on the inserted *perfect* hex-lattice only marginally decreases the quality of individual elements. Only the hull padding (Section 4.2) introduces elements of lower quality due to the projection. The targeted untangling of degenerate elements only assures to fix inverted elements but stops once it is accomplished. As shown in Figure 15, the SJ distribution in our meshes is dominated by the highest quality elements in the object's core

and features only a few poor elements in the most-outer layer. As the structure is a pure hexahedral mesh, common optimization techniques [LSVT15] trivially apply. More smoothing or optimization could eventually elevate the element quality but would ultimately also introduce more deviation from the initial stress field.

By construction, each hex-cell corner is also part of one of its micro-hexahedra. Thus, the MSJ of the overall structure does not change by replacing cells with micro-structures. However, introduced micro-vertices are determined as interpolations of the original cell vertices. This comes as a smoothing effect, noticeable in the histograms (Figure 16) as the MSJ distribution shifts upwards. This indicates more micro-hexahedra of higher quality in relation to the overall amount. The histograms also show the effect of non-uniform τ^* on the micro-hexahedra: Even in a perfectly rectangular hexahedron, varying τ^* values on each vertex will create slightly skewed and differently scaled micro-hexahedra. This can be seen in the diagrams with the rightmost bars (accounting for the highest quality elements) of τ^* falling slightly behind the τ bars.

5.4. Performance and implementation

As described in Section 4.1, the stress field deformation is computed with a standard energy optimization approach using a local and global step per iteration. In all experiments, the λ parameter was linearly scaled from 0 to 2 over the first 100 iterations, whereas the full optimization usually converges within 200 iterations. The computation- and memory-complexity of the procedure scales roughly linearly with the size of the input. Local steps are performed per vertex, while the global step involves highly sparse matrices. The time per iteration also decreases during the optimization, as with increasing alignment, fewer local ADMM iterations are required. For 10k tetrahedra, our single-core *Python* implementation performs around 10 iterations per second, whereas for larger inputs with 800k tetrahedra, one iteration (local + global step) may take up to 3 s. Extracting the aligned hex-lattice, padding the hull layer and introducing the micro-structures usually takes only a few seconds. However, once the most time-consuming deformation is computed, it can be reused to extract multiple hex-lattice structures, *e.g.* with different resolutions, *wall* or *beam* configurations and adapted τ^* values. An implementation of our method is available at <https://github.com/dbukenberger/HexahedralLattice>.

5.5. Future improvements

While our method demonstrates improved alignment with the stress field compared to certain alternatives, there are still cases




Figure 15: Hexahedral mesh quality. The scaled Jacobian measure (per vertex) from 0 (poor) to 1 (optimal) is colour-coded. Due to the hull projection and untangling, elements of the most-outer layer are usually more deformed, i.e. have lower quality.

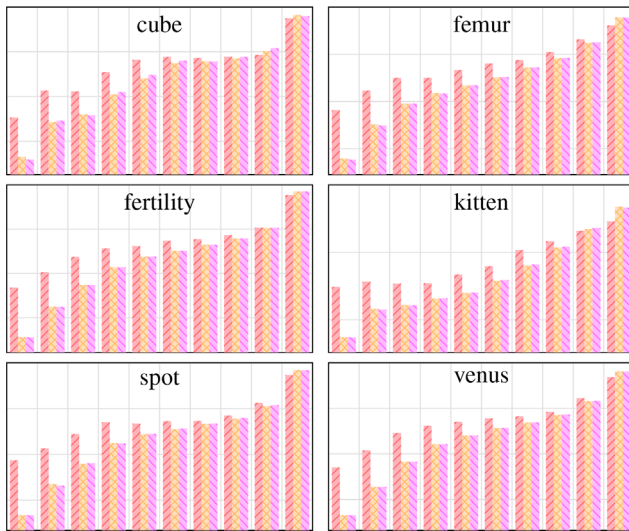


Figure 16: Histograms (0.1 bin width) plot the element quality, measured as minimum scaled Jacobian (log-scale, ranging from 0 (poor) to 1 (optimal)). The three bars include the extracted hex-lattice, our wall micro-structures using a uniform τ and τ^* .

where deviations from the field can be significant. Especially field singularities, which are not explicitly accounted for in the mesh structure, may introduce local distortions, as seen in Figure 12. Addressing these issues in future work could further enhance the performance and robustness of our approach.

The optimization to find the optimal deformation for aligning the stress field is constrained by a cotan-based ARAP energy. While this *rigidity* facilitates natural isometric deformations with consistent object proportions, the contrary could be a promising approach as well, i.e. with *conformal* deformations [BDS*12] (in this context meaning angle-preservation), altered proportions would lead to a varying resolution in the extracted lattice.

Our micro-structure concept is based on the single parameter τ to reduce the volume within each cell equally, thus uniformly for the whole object. This setting is easy to control, does not require any further optimization and is already sufficient to demonstrate the promising mechanical performance of the computed struc-

tural designs. However, our proposed non-uniformly scaled micro-structures using τ^* allow for more stiffness in stress-critical regions of the object at the expense of weaker non-critical regions. While this simple heuristic already improves the overall mechanical performance drastically, it is based on the *von Mises* stress of the fully solid structure. Updating τ^* in an iterative feedback loop with a compliance optimization objective could improve the performance further, but at the intensive computational costs of re-evaluating \mathbf{K} and \mathbf{U} in each iteration. Thus, investigating more advanced methods for balancing the volume distribution within the object for a given α could be a rewarding direction for further research.

While our micro-structures, by design, recreate the initial shape and orientation of the cells, relaxing this relation introduces another degree of freedom. Anisotropically scaled micro-structure geometry or individual alignment to their local stress tensors could further improve the object's stiffness.

In addition to stress-aligned hexahedral meshing, which we focus on in this paper, the optimization of spatially varying and direction-dependent thickness or porosity is a separate and important task in structural optimization. We leave the integration of our meshing method into such a design optimization routine as future work.

6. Conclusion

We introduce a novel method for constructing stress field-aligned lattice structures, yielding pure hexahedral meshes wherein the internal edge structures consider the major stress directions under given loading conditions. Our approach optimizes a deformation of the input object and its associated stress field to align with an orthogonal basis. The deformed object is filled with hexahedra and eventually transformed back to yield stress-aligned hex-lattices. The conforming nature of our result mesh with finite cells enables volume (=material) reduction via cell-inward edge-extrusion facilitating *wall* structures. Utilization of such micro-structures is a capability lacking in competing 3D approaches generating edge graphs. Our method achieves comparable alignment and mechanical performance as state-of-the-art 3D methods in *beam*-structure design. The loading conditions in our simulations were deliberately set in favour of the competing mesh structures, creating stress fields roughly following their inherent layout. Nevertheless, our approach surpasses common meshing methods by producing explicitly tailored stress-aligned structures. Scenarios where our method really shines are configurations like the cube (Figure 6f) or the jet engine bracket

(Figure 14), where common hexahedral meshing methods would produce completely different structures, *i.e.* perfectly axis- or feature edge aligned hex-cells. Through mechanical performance comparisons with state-of-the-art hexahedral meshing approaches, we have demonstrated the superiority of produced *wall*-like structures, affirming improved mechanical properties of the stress-aligned hexahedral infills. Employing the stress-aligned material distribution with non-uniform extrusion along the edges and faces of micro-structures further advances the mechanical performance of the infills. This sets new standards in mesh-based structural design and aims to inspire future work in the design and manufacturing of lightweight 3D objects.

Acknowledgements

The authors have nothing to report.

Open access funding enabled and organized by Projekt DEAL.

Conflicts of Interest

Authors state no conflict of interest.

Appendix A: Micro-Structure Design

The edge vertices for the micro-structure are linearly interpolated for every edge as $[v'_0, v'_1]^T = \mathbf{T}_e [v_0, v_1]^T$ using the weight matrix

$$\mathbf{T}_e = \begin{bmatrix} a & b \\ b & a \end{bmatrix} \quad \text{with} \quad \begin{aligned} a &= 1 - \tau, \\ b &= \tau. \end{aligned} \quad (\text{A.1})$$


The four vertices of each face are computed as $[v'_0, v'_1, v'_2, v'_3]^T = \mathbf{T}_f [v_0, v_1, v_2, v_3]^T$, using bilinear weights with the matrix

$$\mathbf{T}_f = \begin{bmatrix} a & b & c & b \\ b & a & b & c \\ c & b & a & b \\ b & c & b & a \end{bmatrix} \quad \text{with} \quad \begin{aligned} a &= \tau^2 - 2\tau + 1, \\ b &= \tau - \tau^2, \\ c &= \tau^2. \end{aligned} \quad (\text{A.2})$$

Inner micro-structure vertices within the cell compute as $[v'_0, v'_1, v'_2, v'_3, v'_4, v'_5, v'_6, v'_7]^T = \mathbf{T}_c [v_0, v_1, v_2, v_3, v_4, v_5, v_6, v_7]^T$, using trilinear weights in form of the matrix

$$\mathbf{T}_c = \begin{bmatrix} a & b & b & b & c & c & c & d \\ b & a & c & c & b & b & d & c \\ b & c & a & c & b & d & b & c \\ b & c & c & a & d & b & b & c \\ c & b & b & d & a & c & c & b \\ c & b & d & b & c & a & c & b \\ c & d & b & b & c & c & a & b \\ d & c & c & c & b & b & b & a \end{bmatrix} \quad \text{with} \quad \begin{aligned} a &= -(\tau - 1)^3, \\ b &= (\tau - 1)^2 \tau, \\ c &= \tau^2 - \tau^3, \\ d &= \tau^3. \end{aligned} \quad (\text{A.3})$$

The weighting matrix \mathbf{T}_c thereby corresponds to the vertices of a hexahedron ordered as shown on the right. τ s can be determined individually per vertex or averaged per edge, face or cell element, respectively.



Appendix B: Cell Volume Fraction

The volume fraction α for a hexahedral cell with a given parameter τ is formulated in Equation (B.1) as a polynomial for the *beam*-case. This follows the micro-structure elements introduced in Figure 7 with small cubes at the eight corners of the initial cell and elongated beams at the 12 edges, respectively.

$$\alpha_H^b(\tau) = 8\tau^3 + 12(1 - 2\tau)\tau^2 = 12\tau^2 - 16\tau^3 \quad (\text{B.1})$$

As formulated in Equation (B.2), the consumed volume of the *walled* micro-structure further includes six extruded faces.

$$\alpha_H^w(\tau) = 8\tau^3 + 12(1 - 2\tau)\tau^2 + 6(1 - 2\tau)^2\tau = (2\tau - 1)^3 + 1 \quad (\text{B.2})$$

Whereas the *wall* function $\alpha_H^w(\tau)$ is a 1-to-1 mapping and has an inverse, the *beam* function $\alpha_H^b(\tau)$ has no real-valued inverse.

$$\hat{\alpha}_H^b(\tau) = \frac{1 - \cos(2\pi\tau)}{2} \quad (\text{B.3})$$

However, for the domain $[0, \frac{1}{2}]$ and range $[0, 1]$, the polynomial α_H^b can be approximated with the simple term formulated in Equation (B.3). This trigonometric approximation has a maximum error of $\sim 1\%$ but is exact in $\alpha = 0, \frac{1}{2}$ and 1. Its trivial inverse exists and is formulated in Equation (7).

References

- [AAL15] AAGE N., ANDREASSEN E., LAZAROV B. S.: Topology optimization using PETSc: An easy-to-use, fully parallel, open source topology optimization framework. *Structural and Multidisciplinary Optimization* 51 (2015), 565–572.
- [AJL*19] ARORA R., JACOBSON A., LANGLOIS T. R., HUANG Y., MUELLER C., MATUSIK W., SHAMIR A., SINGH K., LEVIN D. I.: Volumetric Michell trusses for parametric design & fabrication. In *Proceedings of the 3rd Annual ACM Symposium on Computational Fabrication* (2019), pp. 1–13.
- [BDS*12] BOUAZIZ S., DEUSS M., SCHWARTZBURG Y., WEISE T., PAULY M.: Shape-up: Shaping discrete geometry with projections. *Computer Graphics Forum* 31, 5 (2012), 1657–1667.
- [BPC*11] BOYD S., PARIKH N., CHU E., PELEATO B., ECKSTEIN J.: Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning* 3, 1 (2011), 1–122.
- [BPM*95] BENZLEY S. E., PERRY E., MERKLEY K., CLARK B., SJAARDAMA G.: A comparison of all hexagonal and all tetrahedral finite element meshes for elastic and elasto-plastic analysis. In *Proceedings of the 4th International Meshing Roundtable* (1995), vol. 17, Sandia National Laboratories, Albuquerque, NM, pp. 179–191.

- [BS13] BENDSOE M. P., SIGMUND O.: *Topology Optimization: Theory, Methods, and Applications*. Springer Science & Business Media, Berlin, 2013.
- [BTL22] BUKENBERGER D. R., TARINI M., LENSCH H. P. A.: At-most-hexa meshes. *Computer Graphics Forum* 41, 1 (2022), 7–28.
- [BTP*19] BRACCI M., TARINI M., PIETRONI N., LIVESU M., CIGNONI P.: HexaLab.net: An online viewer for hexahedral meshes. *Computer-Aided Design* 110 (2019), 24–36.
- [Cra19] CRANE K.: The n-dimensional cotangent formula. Online note. <https://www.cs.cmu.edu/k~mcrane/Projects/Other/nDCotanFormula.pdf> (2019), pp. 11–32.
- [DPM*22] DUMERY C., PROTAIS F., MESTRALLET S., BOURCIER C., LEDOUX F.: Evocube: A genetic labelling framework for polycube-maps. *Computer Graphics Forum* 41, 6 (2022), 467–479.
- [FSA23] FINNENDAHU U., SCHWARTZ M., ALEXA M.: ARAP revisited discretizing the elastic energy using intrinsic Voronoi cells. *Computer Graphics Forum* 42 (2023), e14790.
- [GD04] GOWER J. C., DIJKSTERHUIS G. B.: *Procrustes Problems* (vol. 30). OUP, Oxford, 2004.
- [GJTP17] GAO X., JAKOB W., TARINI M., PANOZZO D.: Robust hex-dominant mesh generation using field-guided polyhedral agglomeration. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 114.
- [GM20] GUSTAFSSON T., MCBAIN G. D.: scikit-fem: A Python package for finite element assembly. *Journal of Open Source Software* 5, 52 (2020), 2369.
- [GMD*15] GAO X., MARTIN T., DENG S., COHEN E., DENG Z., CHEN G.: Structured volume decomposition via generalized sweeping. *IEEE Transactions on Visualization and Computer Graphics* 22, 7 (2015), 1899–1911.
- [GPW*17] GAO X., PANOZZO D., WANG W., DENG Z., CHEN G.: Robust structure simplification for hex re-meshing. *ACM Transactions on Graphics* 36, 6 (2017), 1–13.
- [GSP19] GAO X., SHEN H., PANOZZO D.: Feature preserving Octree-based hexahedral meshing. *Computer Graphics Forum* 38, 5 (2019), 135–149.
- [GSZ11] GREGSON J., SHEFFER A., ZHANG E.: All-hex mesh generation via volumetric polycube deformation. *Computer Graphics Forum* 30, 5 (2011), 1407–1416.
- [HSW*20] HU Y., SCHNEIDER T., WANG B., ZORIN D., PANOZZO D.: Fast tetrahedral meshing in the wild. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 117:1–117:18.
- [IMH05] IGARASHI T., MOSCOVICH T., HUGHES J. F.: As-rigid-as-possible shape manipulation. *ACM Transactions on Graphics (TOG)* 24, 3 (2005), 1134–1141.
- [JOB*24] JENSEN P. D. L., OLSEN T. F., BÆRENTZEN J. A., AAGE N., SIGMUND O.: Efficient inverse-designed structural infill for complex engineering structures. *Thin-Walled Structures* 195 (2024), 111427.
- [LAPS17] LIVESU M., ATTENE M., PATANÉ G., SPAGNUOLO M.: Explicit cylindrical maps for general tubular shapes. *Computer-Aided Design* 90 (2017), 27–36.
- [LJ19] LIU H.-T. D., JACOBSON A.: Cubic stylization. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–10.
- [LJ21] LIU H.-T. D., JACOBSON A.: Normal-driven spherical shape analogies. *Computer Graphics Forum* 40, 5 (2021), 45–55.
- [LPP*20] LIVESU M., PIETRONI N., PUPPO E., SHEFFER A., CIGNONI P.: LoopyCuts: Practical feature-preserving block decomposition for strongly hex-dominant meshing. *ACM TOG* 39, 4 (2020), 121:1–121:17.
- [LSVT15] LIVESU M., SHEFFER A., VINING N., TARINI M.: Practical hex-mesh optimization via edge-cone rectification. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 141.
- [LVS*13] LIVESU M., VINING N., SHEFFER A., GREGSON J., SCATENI R.: Polycut: Monotone graph-cuts for polycube base-complex construction. *ACM Transactions on Graphics (TOG)* 32, 6 (2013), 1–12.
- [LZS*21] LI L., ZHANG P., SMIRNOV D., ABULNAGA S. M., SOLOMON J.: Interactive all-hex meshing via cuboid decomposition. *ACM Transactions on Graphics (TOG)* 40, 6 (2021), 1–17.
- [Mic04] MICHELL A. G. M.: 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.
- [NRP11] NIESER M., REITEBUCH U., POLTHIER K.: CubeCover - parameterization of 3D volumes. *Computer Graphics Forum* 30, 5 (2011), 1397–1406.
- [PCS*22] PIETRONI N., CAMPEN M., SHEFFER A., CHERCHI G., BOMMES D., GAO X., SCATENI R., LEDOUX F., REMACLE J., LIVESU M.: Hex-mesh generation and processing: A survey. *ACM Transactions on Graphics* 42, 2 (2022), 1–44.
- [Ped89] PEDERSEN P.: On optimal orientation of orthotropic materials. *Structural Optimization* 1, 2 (June 1989), 101–106.
- [SA07] SORKINE O., ALEXA M.: As-rigid-as-possible surface modeling. In *Symposium on Geometry Processing* (2007), vol. 4, pp. 109–116.
- [SBR*15] SCHUMACHER C., BICKEL B., RYS J., MARSCHNER S., DARAIO C., GROSS M.: Microstructures to control elasticity in 3D printing. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–13.

- [SHD*18] SCHNEIDER T., HU Y., DUMAS J., GAO X., PANOZZO D., ZORIN D.: Decoupling simulation accuracy from mesh quality. *ACM Transactions on Graphics* 37 (2018), 1–14.
- [SHC*22] SCHNEIDER T., HU Y., GAO X., DUMAS J., ZORIN D., PANOZZO D.: A large-scale comparison of tetrahedral and hexahedral elements for solving elliptic PDEs with the finite element method. *ACM Transactions on Graphics (TOG)* 41, 3 (2022), 1–14.
- [SOG*22] STUTZ F. C., OLSEN T. F., GROEN J. P., TRUNG T. N., AAGE N., SIGMUND O., SOLOMON J., BÆRENTZEN J. A.: Synthesis of frame field-aligned multi-laminar structures. *ACM Transactions on Graphics (TOG)* 41, 5 (2022), 1–20.
- [THCM04] TARINI M., HORMANN K., CIGNONI P., MONTANI C.: Polycube-maps. *ACM Transactions on Graphics (TOG)* 23, 3 (2004), 853–860.
- [TTZ*20] TRICARD T., TAVERNIER V., ZANNI C., MARTÍNEZ J., HUGRON P.-A., NEYRET F., LEFEBVRE S.: Freely orientable microstructures for designing deformable 3D prints. *ACM Transactions on Graphics* 39, 6 (2020), 1–16.
- [WAWS18] WU J., AAGE N., WESTERMANN R., SIGMUND O.: Infill optimization for additive manufacturing-approaching bone-like porous structures. *IEEE Transactions on Visualization and Computer Graphics* 24, 2 (2018), 1127–1140.
- [WBM21] WHALEN E., BEYENE A., MUELLER C.: SimJEB: Simulated jet engine bracket dataset. *Computer Graphics Forum* 40, 5 (2021), 9–17.
- [WDW16] WU J., DICK C., WESTERMANN R.: A system for high-resolution topology optimization. *IEEE Transactions on Visualization and Computer Graphics* 22, 3 (2016), 1195–1208.
- [WSG21] WU J., SIGMUND O., GROEN J. P.: Topology optimization of multi-scale structures: A review. *Structural and Multidisciplinary Optimization* 63 (2021), 1455–1480.
- [WWG21] WU J., WANG W., GAO X.: Design and optimization of conforming lattice structures. *IEEE Transactions on Visualization and Computer Graphics* 27, 1 (2021), 43–56.
- [WWW22] WANG J., WU J., WESTERMANN R.: Stress trajectory guided structural design and topology optimization. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference* (2022), vol. 86212, American Society of Mechanical Engineers, pp. V002T02A012.
- [WWW23] WANG J., WESTERMANN R., WU J.: A streamline-guided dehomogenization approach for structural design. *Journal of Mechanical Design* 145, 2 (2023), 021702.
- [ZGL*23] ZOCCHEDDU F., GOBBETTI E., LIVESU M., PIETRONI N., CHERCHI G.: HexBox: Interactive box modeling of hexahedral meshes. *Computer Graphics Forum* 42, 5 (2023), e14899.

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Supporting Information