

Co-Optimization of Structure and Manufacturable Semi-Continuous Layers for Laminated Composites

TAO LIU, The University of Manchester, United Kingdom
 AORAN LYU, The University of Manchester, United Kingdom
 YONGXUE CHEN, The University of Manchester, United Kingdom
 YU JIANG, The University of Manchester, United Kingdom
 MICHAEL PETTY, The University of Manchester, United Kingdom
 CHARLIE C.L. WANG*, The University of Manchester, United Kingdom

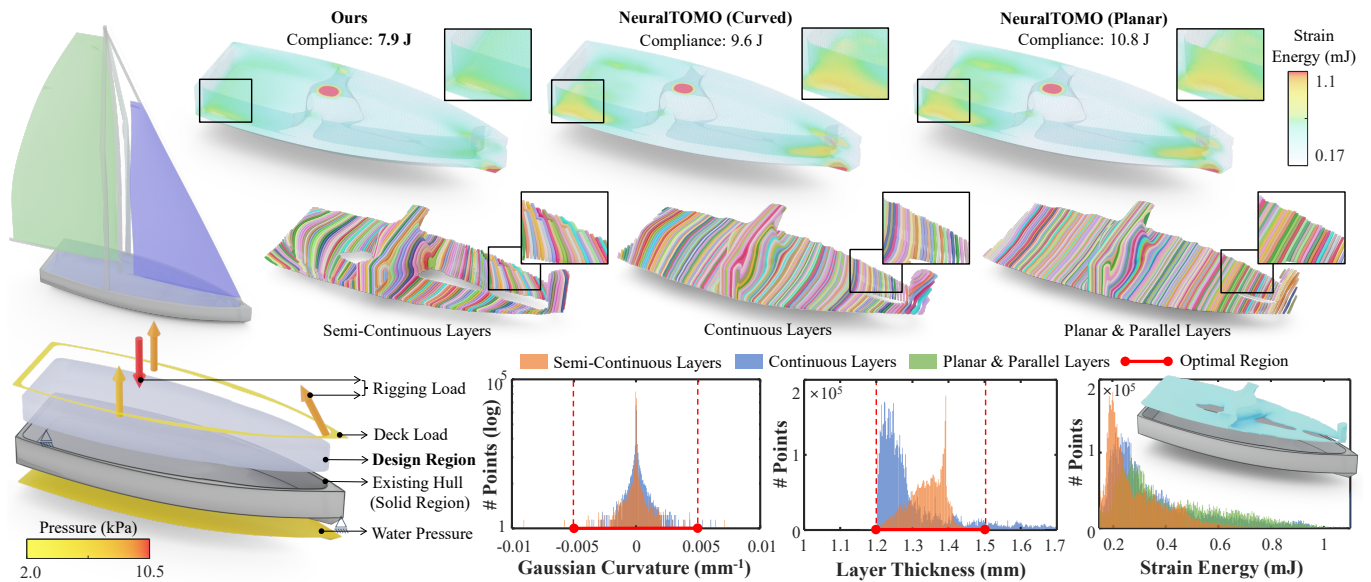


Fig. 1. Co-optimized topology and curved semi-continuous layers for a woven fabric-reinforced laminated composite boat structure. The proposed optimization pipeline supports both complete and partial layers through the computation of semi-continuous periodic scalar fields, resulting in improved mechanical performance (reduced strain energy) and enhanced manufacturability, with well-controlled Gaussian curvature and layer thickness. For comparison, co-optimization results generated by NeuralTOMO [Liu et al. 2025] using both (i) curved continuous layers and (ii) planar & parallel layers are also shown.

To enable the design and manufacturing of optimized composite structures using fabric plies, we propose a field-driven optimization framework that

*Corresponding author: charlie.wang@manchester.ac.uk (Charlie C.L. Wang).

Authors' Contact Information: Tao Liu, The University of Manchester, Manchester, United Kingdom; Aoran Lyu, The University of Manchester, Manchester, United Kingdom; Yongxue Chen, The University of Manchester, Manchester, United Kingdom; Yu Jiang, The University of Manchester, Manchester, United Kingdom; Michael Petty, The University of Manchester, Manchester, United Kingdom; Charlie C.L. Wang, The University of Manchester, Manchester, United Kingdom.

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

ACM SIGGRAPH, Los Angeles, CA

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-x-xxxx-xxxx-x/YYYY/MM

<https://doi.org/10.1145/3811393>

jointly optimizes structural topology and manufacturable layers. A central challenge in this setting is the modeling and optimization of partial fabric layers with near-uniform thickness, which we formulate as a semi-continuous periodic scalar field parameterized by a continuous implicit neural vector field. Within this concurrent structure-layer optimization framework, we further derive a formulation of inter-layer anisotropic mechanical behavior that enables effective modeling of mechanical property transitions induced by partial-layer boundaries, together with additional objectives for manufacturability and field regularization. We validate the effectiveness of our approach through both numerical simulations and physical experiments, demonstrating that the optimized fabric-reinforced laminated composites achieve up to 43.8% higher stiffness compared to counterparts fabricated using planar fabric plies.

CCS Concepts: • **Computing methodologies** → **Shape modeling**; • **Applied computing** → **Engineering**.

Additional Key Words and Phrases: Partial Layers, Inter-Layer Anisotropy, Topology Optimization, Fabric-Reinforced, Laminated Composites

ACM Reference Format:

Tao Liu, Aoran Lyu, Yongxue Chen, Yu Jiang, Michael Petty, and Charlie C.L. Wang. 2026. Co-Optimization of Structure and Manufacturable Semi-Continuous Layers for Laminated Composites. In *Proceedings of (ACM SIGGRAPH)*. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3811393>

1 Introduction

Composite structures have attracted significant attention in the design and manufacturing of high-value products (such as aircraft, boats and electric vehicles) due to their outstanding mechanical performance and exceptional strength-to-weight ratio, making them particularly suitable for applications that demand extreme lightweight designs. Among various composites, woven fabric-reinforced laminated composites are widely adopted because they offer ease of manufacturing, high damage tolerance, and reliable structural performance [Hasan et al. 2021]. Despite their broad industrial use, relatively little effort has been devoted to simultaneously optimizing both the structural topology and the manufacturable layers of such composites. This gap largely stems from the difficulty of modeling fabric layers with near-uniform thickness within the computational loop of structural topology optimization driven by *Finite Element Analysis* (FEA).

Recent research has begun to address the co-optimization of structural topology and fiber orientations for composites fabricated via robot-assisted 3D printing (e.g., [Cheng et al. 2023; Liu et al. 2025; Xu et al. 2024]). While promising, these approaches typically rely on complex manufacturing processes and specialized hardware, posing a substantially higher barrier to adoption. In this work, we investigate a new computational framework that enables the co-optimization of structural topology and manufacturable semi-continuous layers for woven fabric-reinforced laminated composites – see Fig. 1 for an example.

1.1 Challenges and Motivation

Requiring a structural solid to be formed exclusively by complete layers – i.e., enforcing that every layer boundary coincides with the boundary of the solid – significantly limits the design space for optimizing the mechanical performance of composite structures. In particular, when layers are constrained to have (nearly) uniform thickness, such a requirement greatly restricts the flexibility of generating diverse curved layers to realize a given solid geometry. In contrast, as shown in Fig. 2, allowing partial layers substantially enlarges the design space, enabling more effective exploration of optimum¹.

In prior work, partial layers are typically introduced as a post-processing step to enforce thickness constraints after optimization (e.g., [Zhang et al. 2022]). Consequently, the locations of partial-layer boundaries that introduce locally weakened mechanical behavior are not explicitly modeled or optimized during the structural design of composites. Directly modeling and optimizing composite structures with partial layers needs to overcome the following challenges.

- **Partial Layers:** A central challenge lies in how to effectively represent partial layers within topology optimization

¹Note that it is also important to minimize the extent of interior boundaries introduced by partial layers, as these regions exhibit relatively weak mechanical performance.

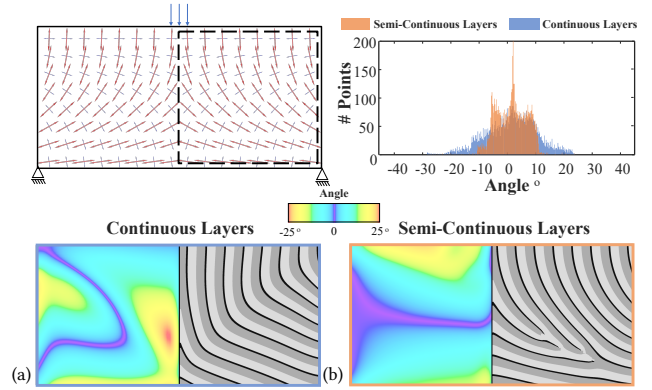


Fig. 2. For a structural solid subjected to prescribed loads, a laminated composite achieves optimal mechanical performance when its layers are aligned with the directions of maximum principal stress. However, this requirement cannot generally be satisfied when only complete layers are allowed, which can be evidenced by the large deviations between the tangents of continuous layers and the maximum stress directions – see the color map in the left of (a) for angle deviation distribution and also the histogram in blue color for the statistics. By introducing partial layers as shown in (b), the desired layer orientations can be better accommodated – see also the deviation distribution as the color map and the histogram in orange color.

pipelines, which typically rely on field-based representations such as density fields [Sigmund 2001]. When layers are modeled as iso-surfaces of a scalar field (e.g., [Liu et al. 2025]), existing formulations naturally generate complete layers spanning the entire structure. Enabling partial layers therefore requires a new computational paradigm that supports semi-continuous fields.

- **Mechanical Property:** In fabric-reinforced composites, regions fully covered by continuous layers exhibit significantly higher mechanical performance than regions containing layer boundaries (i.e., singularities) introduced by partial layers. A new formulation is required for effectively capturing the transition of mechanical properties across partial-layer boundaries within FEA-based topology optimization frameworks.
- **Manufacturability:** Manufacturability imposes additional constraints on optimized composite structures. When fabrication is performed using fabric plies, the solid must be decomposable into layers with nearly uniform thickness, and each layer should approximate a developable surface. These constraints must be explicitly incorporated into the optimization process rather than by post-processing corrections.

All these difficulties prevent the generation of optimized layer structures for laminated composites by existing approaches.

1.2 Our Method and Contributions

To address the above challenges, we propose a field-based computational framework that jointly optimizes structural topology and semi-continuous developable layers with uniform thickness for laminated composites. The optimization is conducted over a domain Ω and parameterized by two implicit neural fields:

- (1) a density field $\rho(\mathbf{x}) \in [0, 1]$ represented by a neural network \mathcal{N}_ρ ;
- (2) a vector field $\mathbf{z}(\mathbf{x}) \in \mathbb{R}^3$ represented by another neural network \mathcal{N}_z .

The density field $\rho(\mathbf{x})$ indicates the solid ($\rho \approx 1$) and the void ($\rho \approx 0$) regions, while the vector field $\mathbf{z}(\mathbf{x})$ specifies demanded local layer-normal directions. From $\mathbf{z}(\mathbf{x})$, we derive a semi-continuous periodic scalar field $\alpha(\mathbf{x})$ to represent both complete and partial layers by enforcing $\nabla\alpha(\mathbf{x}) \approx \mathbf{z}(\mathbf{x})$, while allowing $\mathbf{z}(\mathbf{x})$ to be non-integrable (Sec. 4.1). The domain Ω is discretized into a regular grid (Fig. 3), and the phase values $\alpha(\mathbf{x})$ are computed indirectly at grid nodes. Our scalar-field construction extends the stripe-field formulation of [Knöppel et al. 2015] to a 3-manifold setting. During the process of scalar-field construction, we evaluate a singularity metric for each voxel element e , which determines whether partial-layer boundaries pass through the element and accordingly selects different element stiffness matrices (Sec. 4.2). The optimization objectives are defined as a total loss combining compliance energy, Gaussian curvature requirement for developability, helicoid-free constraints, and harmonic regularization (Sec. 5). The total loss is minimized via backpropagation to update \mathcal{N}_ρ and \mathcal{N}_z , producing optimized solid topology together with manufacturable semi-continuous fabric layers.

Our work introduces the following technical contributions:

- A computational pipeline for generating semi-continuous periodic scalar fields from non-integrable vector fields, enabling the modeling of partial layers that satisfy manufacturing constraints – such as uniform thickness and developable surfaces – to guide the optimization;
- A field-driven co-optimization framework that couples structural topology and manufacturable layers, where the structural solid is represented by a neural density field and the layers are represented by a semi-continuous periodic scalar field parameterized by a neural vector field;
- A formulation of inter-layer anisotropic mechanical behavior within the coupled structure-layer topology optimization framework, enabling modeling of mechanical property transitions induced by partial layer boundaries, together with the losses for manufacturability and regularization.

To the best of our knowledge, no existing approach in the literature provides a unified capability for modeling and optimizing fabric-reinforced composite structures with partial layers. The effectiveness of our approach has been verified by both numerical and physical experiments.

2 Related Works

2.1 Topology Optimization for Composites

Topology optimization (TO) has been widely adopted in engineering design to determine optimal material distributions within a prescribed domain for enhanced mechanical performance, originally developed for isotropic materials [Sigmund 2001]. Recent advances have extended TO to explicitly account for anisotropy across diverse applications, including fluidic systems [Li et al. 2022], multi-material structures [Yan et al. 2020], and microstructure design [Arora et al.

2019; Wu et al. 2021; Zhu et al. 2017]. For fiber-reinforced composites, incorporating anisotropic mechanical analysis into TO is particularly critical, as load-bearing behavior is strongly governed by fiber orientations and their spatial continuity [Luo et al. 2023]. Enabled by GPU acceleration [Qi et al. 2025], recent work has further explored the concurrent optimization of 3D structural topology and fiber orientations [Liu et al. 2025; Qiu et al. 2022].

For laminated composites, most existing TO studies focus on designing in-plane fiber angle distributions. Representative examples include constant-stiffness laminate design, which optimizes fiber angles and ply shapes to meet prescribed laminate stiffness targets [Fang et al. 2025; Wang et al. 2023c]. Beyond constant stiffness, variable-stiffness laminate design allows spatially varying fiber orientations, with discrete variable-stiffness laminates and related formulations explored in [Gandhi et al. 2025; Smith and Norato 2023]. However, these approaches are inherently two-dimensional and cannot generate optimized designs for loading conditions governed by fully three-dimensional stress fields.

For 3D laminated composite designs, manufacturability from planar plies must be explicitly considered. Recent TO researches have incorporated various manufacturing constraints, including support structures [Langelaar 2016; Wang et al. 2023a], printing-induced distortion [Miki and Yamada 2021], printing orientation [Ye et al. 2023], feature accessibility [Yamada and Noguchi 2022], and concurrent optimization of topology and manufacturing sequences [Wang et al. 2024, 2023b; Wu et al. 2024a]. However, no existing approach addresses the manufacturability of laminated composites in TO.

2.2 Process Planning and Optimization for Fabrication

Process planning has been studied in computer graphics for years. A widely adopted strategy is domain decomposition, where a given model is partitioned into manufacturable regions, such as height fields [Muntoni et al. 2018], pyramid primitives [Hu et al. 2014], convex domains [Wu et al. 2020], or more general layered representations [Dai et al. 2018]. Recently, decomposition strategies have been extended to volume-surface co-decomposition [Zhong et al. 2023] and voxel-based inverse nullification [Chen et al. 2025], for hybrid manufacturing scenarios. Related work has also explored fabrication beyond conventional planar slicing, including the methods of using curved and non-planar layers [Etienne et al. 2019; Fang et al. 2020; Mitropoulou et al. 2020], periodic-function-based infill structures [Chermain et al. 2023], and volumetric atom-based toolpaths [Chermain et al. 2025]. However, these deterministic algorithms are difficult to integrate into optimization loops.

Field-based optimization approaches have attracted increasing attention, as they allow the integration of multiple manufacturing and design objectives within a unified framework. A variety of field representations have been explored in prior work, including height-based deformation fields [Etienne et al. 2019], vector fields [Fang et al. 2020], and quaternion fields coupled with scaling fields [Liu et al. 2024; Zhang et al. 2022]. These field-based strategies have further been extended to applications involving 3D-printed parts reinforced with continuous carbon fibers [Fang et al. 2024; Zhang et al. 2025]. A key advantage of field-based optimization is its natural compatibility with density-based topology optimization, enabling

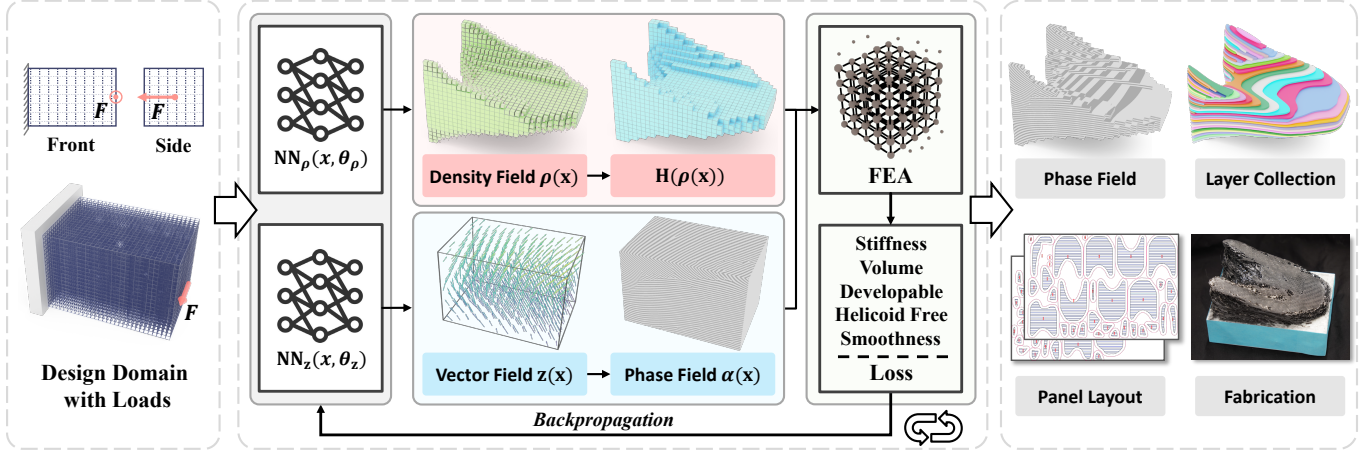


Fig. 3. Computational pipeline of the proposed framework. Two continuous fields – a density field and a vector field – are parameterized by neural networks (NN) and jointly optimized in a self-learning loop via backpropagation. Mechanical performance is evaluated through FEA, while manufacturability is enforced using geometry-based loss terms. The optimized fields are then used to extract a collection of curved semi-continuous layers, which are flattened into planar panel layouts for the fabrication of laminated composites.

the simultaneous optimization of structural topology and curved manufacturing layers [Liu et al. 2025].

2.3 Singularity-Aware Field Computing

Singularities and periodic parameterizations have been widely studied for field and pattern design. Ray et al. [Ray et al. 2006] introduced globally periodic functions for surface parameterization, laying the foundation for stripe-based representations, which were later formulated as a minimum-eigenvalue optimization problem by Knöppel et al. [2015]. Subsequent work has focused on improving the control and optimization of stripe fields, including efficient editing of singularity configurations [Noma et al. 2022], circle-bundle-based lifting of directional fields [Palmer et al. 2024], and neural-network-based singularity-aware field computation [Dong et al. 2025].

Singularity-based fields have also been applied to tasks of computational fabrication, such as knitting pattern generation [Mitra et al. 2023; Zhang et al. 2021], layered toolpath planning for self-shaping shells [Jourdan et al. 2022], differentiable stripe-based inverse design coupled with FEA [Maestre et al. 2023], and the fabrication of orientable microstructures [Tricard et al. 2020]. Extensions to 3-manifold domains have also been explored in [Corman and Crane 2019; Zhang et al. 2020]. Nevertheless, existing methods do not explicitly address the computation of fields that jointly model complete and partial layers with near-uniform thickness.

2.4 Developable Surface Optimization

Developable surfaces have long been employed to approximate complex shapes (e.g., [Julius et al. 2005; Solomon et al. 2012; Tang et al. 2016; Wang and Tang 2004]), primarily because they admit well-established manufacturing processes, such as forming with paper [Ion et al. 2020; Verhoeven et al. 2022; Zhao et al. 2023], woven fabrics [Julius et al. 2005], or glass panels [Jiang et al. 2020; Liu et al. 2006]. A comprehensive review of developable surface modeling is beyond the scope of this paper; therefore, only representative

works are discussed here. Among existing approaches, a common strategy is to minimize Gaussian curvature through surface deformation to obtain approximately developable geometries (e.g., [Kilian et al. 2008; Wang 2008]). More recently, neural-network-based parameterizations have been introduced to model piecewise developable surfaces via deformation [Wu et al. 2024b]. Woven fabrics can form non-zero Gaussian curvature through shearing deformation by Chebyshev Net [Aono et al. 1994; Oehri et al. 2024], a phenomenon that has been studied in textile simulation research. However, laminate fabrication usually restricts this sort of deformation as it produces unwanted wrinkles, which lead to stress concentrations and consequently reduce the mechanical strength of the final structure [Krogh et al. 2023]. However, no existing method formulates the generation of a stack of developable surface layers as isosurfaces of a scalar field, which is the problem addressed in this paper in the context of optimizing laminated composites.

3 Overview

3.1 Domain Discretization and Representation

To facilitate FEA and layer computation, we discretize the computational domain $\Omega \subset \mathbb{R}^3$ into a regular grid of $L \times M \times N$ voxels collected in the set \mathcal{C} . Each voxel serves as an element for FEA and is assigned homogenized mechanical properties under a piecewise formulation. The voxel corners are referred to as grid nodes and stored in the set \mathcal{V} , where the values of the semi-continuous periodic scalar field are evaluated locally on a per-voxel basis. The voxel edges are referred to as grid edges and stored in the set \mathcal{E} . The boundary faces of each voxel are constructed on-the-fly during computation according to the ordering of its grid nodes. Overall, the computational domain Ω is represented as a regular grid mesh $\mathcal{M} = (\mathcal{V}, \mathcal{E}, \mathcal{C})$ (see the leftmost in Fig. 3).

Despite this discrete representation for geometric and mechanical computation, the density field $\rho(\mathbf{x}) \in [0, 1]$ and the normalized vector field $\mathbf{z}(\mathbf{x}) \in \mathbb{S}^2$ are defined continuously and implicitly over \mathbb{R}^3

by two *neural networks* (NN), $\mathcal{N}_\rho(\mathbf{x}, \theta_\rho)$ and $\mathcal{N}_z(\mathbf{x}, \theta_z)$, respectively parameterized by network coefficients θ_ρ and θ_z that serve as the design variables in our optimization framework.

3.2 Computational Pipeline

Guided by the design and manufacturing objectives, we define a total loss function $\mathcal{L}_{\text{total}}$ to drive the optimization (Sec. 5).

Optimization is performed via backpropagation, which updates θ_ρ and θ_z to minimize $\mathcal{L}_{\text{total}}$. Specifically, our computational pipeline consists of the following steps:

- (1) Discretize the design domain Ω into a regular grid mesh \mathcal{M} ;
- (2) Initialize the networks $\mathcal{N}_\rho(\mathbf{x}, \theta_\rho)$ and $\mathcal{N}_z(\mathbf{x}, \theta_z)$ (Sec. 6.3);
- (3) To construct the periodic scalar field $\alpha(\mathbf{x})$, we compute the values of $\cos(\alpha_i)$ and $\sin(\alpha_i)$ at each grid node v_i by enforcing $\nabla\alpha(\mathbf{x}) \approx \mathbf{z}(\mathbf{x})$ (Sec. 4.1.1);
- (4) For each element $e \in C$, perform the following operations:
 - (a) Apply the phase unwrapping scheme to update the nodal phase values α_i (Sec. 4.1.2);
 - (b) Evaluate the singularity indicator I_e to determine whether element e contains a singularity (Sec. 4.1.3);
 - (c) Construct the element stiffness matrix \mathbf{K}_e accordingly (Sec. 4.2.1 & 4.2.2);
- (5) Assemble the global stiffness matrix and solve the equilibrium equations to obtain the displacement field \mathbf{U} (Sec. 4.2.3);
- (6) Evaluate the loss terms defined by the design and manufacturing objectives (Sec. 5);
- (7) Compute gradients $\partial\mathcal{L}_{\text{total}}/\partial\theta_\rho$ and $\partial\mathcal{L}_{\text{total}}/\partial\theta_z$ via backpropagation;
- (8) Update the network parameters θ_ρ and θ_z using the gradients and the optimizer;
- (9) Repeat Steps (3)–(8) until the learning convergences;
- (10) Construct a higher-resolution volumetric mesh \mathcal{M}^* for layer extraction;
- (11) Extract isosurfaces of $\alpha(\mathbf{x})$ on \mathcal{M}^* and trim them by the solid defined by $H(\rho(\mathbf{x}))$ to obtain curved layers (Sec. 6.4);
- (12) Generate planar panels and perform packing for fabrication.

To balance computational efficiency and geometric fidelity, we use a higher resolution for \mathcal{M}^* than for the grid \mathcal{M} employed in FEA.

4 Laminated Composites with Partial Layers

Modeling of fabric-reinforced laminated composites needs to jointly address the geometric representation (i.e., semi-continuous periodic scalar field) and the mechanical evaluation (i.e., anisotropic stiffness). In this section, we present our formulations for both aspects.

4.1 Geometric Modeling

The structural solid of a laminated composite is implicitly defined by the density field as

$$\mathcal{H} = \{\mathbf{x} \in \Omega \mid H(\rho(\mathbf{x})) = 1\}, \quad (1)$$

where $H(\cdot)$ denotes the Heaviside function, a formulation widely adopted in density-based topology optimization (e.g., [Liu et al. 2025; Sigmund 2001]). We focus next on modeling the laminate layers by generating a semi-continuous periodic scalar field from the vector field $\mathbf{z}(\mathbf{x})$.

4.1.1 Semi-continuous Periodic Scalar Field. For a regular grid mesh $\mathcal{M} = (\mathcal{V}, \mathcal{E}, C)$, we denote grid nodes by $v_i \in \mathcal{V}$ with positions $\mathbf{v}_i \in \mathbb{R}^3$, and oriented grid edges by $(v_i, v_j) \in \mathcal{E}$. Each edge is associated with a geometric vector

$$\mathbf{e}_{ij} = \mathbf{v}_j - \mathbf{v}_i, \quad \|\mathbf{e}_{ij}\| = h, \quad \hat{\mathbf{e}}_{ij} = \mathbf{e}_{ij}/h \quad (2)$$

with h being the grid spacing – constant on the regular grid.

Layers, including both complete and partial ones, are encoded by a periodic phase field $\alpha(\mathbf{x}) : \mathbb{R}^3 \mapsto \mathbb{R}/2\pi\mathbb{Z}$ defined on grid nodes, where the level sets of $\alpha(\mathbf{x})$ determine the layer assignment. To avoid branch cuts induced by the 2π -periodicity of α , we embed its computation on the unit circle by representing it via its sine and cosine components, yielding a globally continuous formulation. Specifically, we define

$$\mathbf{p}_i(\alpha_i) := \begin{bmatrix} \cos \alpha_i \\ \sin \alpha_i \end{bmatrix} \in \mathbb{S}^1 \quad (3)$$

at each grid node v_i , where $\mathbf{p}_i(\alpha_i) = \mathbf{p}_i(\alpha_i + 2k\pi)$ for all $k \in \mathbb{Z}$.

The phase field $\alpha(\mathbf{x})$ is determined with the help of the above unit circle embedding by enforcing its gradients to align with the vector field $\mathbf{z}(\mathbf{x})$ – i.e., letting $\nabla\alpha(\mathbf{x}) \approx \nu\mathbf{z}(\mathbf{x})$, where $\nu \in \mathbb{R}$ controls the phase growth speed and is chosen according to the laminate thickness. The computation is taken in a discrete manner edge by edge. For each edge e_{ij} , the enforcement is based on the directional derivative using a forward finite difference to let

$$(\alpha_j - \alpha_i)/h \approx \nu\mathbf{z}_{ij} \cdot \hat{\mathbf{e}}_{ij}. \quad (4)$$

In a discrete implementation, we simply evaluate \mathbf{z}_{ij} at the middle point of an edge as $\mathbf{z}_{ij} = \mathbf{z}((\mathbf{v}_i + \mathbf{v}_j)/2)$. This gives the target (unwrapped) phase increment per edge as

$$Z_{ij} = h\nu\mathbf{z}_{ij} \cdot \hat{\mathbf{e}}_{ij} = \nu\mathbf{z}_{ij} \cdot \mathbf{e}_{ij}, \quad (5)$$

where reversing the edge orientation changes the sign, i.e., $Z_{ji} = -Z_{ij}$. This edge-based discretization therefore requires the phase values to satisfy $\alpha_j = \alpha_i + Z_{ij}$.

Adding Z_{ij} to α_i corresponds to a planar rotation of the embedded vector such as (ref. [Ray et al. 2006])

$$\begin{bmatrix} \cos(\alpha_i + Z_{ij}) \\ \sin(\alpha_i + Z_{ij}) \end{bmatrix} = \mathbf{R}(Z_{ij}) \begin{bmatrix} \cos \alpha_i \\ \sin \alpha_i \end{bmatrix} \quad (6)$$

with

$$\mathbf{R}(Z_{ij}) = \begin{bmatrix} \cos Z_{ij} & -\sin Z_{ij} \\ \sin Z_{ij} & \cos Z_{ij} \end{bmatrix}.$$

Perfect edge consistency ideally satisfies $\mathbf{p}_j = \mathbf{R}(Z_{ij})\mathbf{p}_i$. Therefore, the phase field can be determined in a least-squares sense by minimizing the following energy:

$$E_\alpha = \sum_{e_{ij} \in \mathcal{E}} \|\mathbf{R}(Z_{ij})\mathbf{p}_i - \mathbf{p}_j\|_2^2. \quad (7)$$

To avoid degenerated solutions, we follow the strategy introduced in [Knöppel et al. 2015] to solve it via eigenvector analysis. Specifically, each embedding $\mathbf{p}_i = [\cos \alpha_i, \sin \alpha_i]^\top$ is treated as 2 independent variables $[c_i, s_i]^\top$ and collected into an unknown vector as

$$\mathbf{u} = [c_1, \dots, c_n, s_1, \dots, s_n]^\top \in \mathbb{R}^{2|\mathcal{V}|} \quad (8)$$

with $|\cdot|$ giving the number of elements in a set. Since E_α is in the quadratic form of \mathbf{u} , its minimum can be achieved by solving

$$\min_{\mathbf{u}} \frac{1}{2} \mathbf{u}^\top \mathbf{A} \mathbf{u}, \quad \text{s.t.} \quad \|\mathbf{u}\|_2^2 = 1, \quad (9)$$

where \mathbf{A} depends only on the mesh connectivity of \mathcal{M} and the guiding field $\mathbf{z}(\mathbf{x})$, and the constraint is introduced to avoid the trivial solution. After introducing a Lagrange multiplier λ , the first-order optimality condition gives the eigenproblem as (ref. [Sun 1990])

$$\mathbf{A} \mathbf{u} = \lambda \mathbf{u}, \quad (10)$$

and the eigenvector associated with the smallest eigenvalue gives the solution minimizing the energy E_α .

After determining the value of \mathbf{p}_i on every grid node as $\mathbf{p}_i = [c_i, s_i]^\top$, we have the initial phase on grid node as

$$\alpha_i = \text{atan2}(s_i, c_i). \quad (11)$$

4.1.2 Phase Unwrapping. The phase values obtained by the above procedure are restricted to the interval $(-\pi, \pi]$, which may violate the edge-wise consistency condition $\alpha_j = \alpha_i + Z_{ij}$ even when the equivalent relations $\cos(\alpha_j) = \cos(\alpha_i + Z_{ij})$ and $\sin(\alpha_j) = \sin(\alpha_i + Z_{ij})$ are satisfied. To resolve this issue, we apply a face by face phase unwrapping scheme akin to field-combing [Diamanti et al. 2014].

First of all, we define a wrapping operator $\varpi(\cdot)$ as

$$\varpi(\theta) := \text{atan2}(\sin \theta, \cos \theta) \in (-\pi, \pi]. \quad (12)$$

For each oriented edge (v_i, v_j) , the wrapped residual between $\alpha_i + Z_{i,j}$ and α_j is defined as

$$\delta_{ij} := \varpi(\alpha_j - (\alpha_i + Z_{ij})). \quad (13)$$

The unwrapped phase increment on edge e_{ij} can then be obtained as

$$\sigma_{ij} := Z_{ij} + \delta_{ij} = Z_{ij} + \varpi(\alpha_j - (\alpha_i + Z_{ij})). \quad (14)$$

Using the unwrapped phase increment defined in the above equation, the unwrapped phase values α_i^* , α_j^* , α_k^* and α_l^* on a quad face f_{ijkl} (with boundary orientation $i \rightarrow j \rightarrow k \rightarrow l \rightarrow i$) can be obtained by propagation:

$$\begin{aligned} \alpha_i^* &:= \alpha_i, \\ \alpha_j^* &:= \alpha_i^* + \sigma_{ij}, \\ \alpha_k^* &:= \alpha_j^* + \sigma_{jk}, \\ \alpha_l^* &:= \alpha_k^* + \sigma_{kl}. \end{aligned} \quad (15)$$

4.1.3 Singularity Detection. For a quad face f_{ijkl} , the winding number is evaluated as (ref. [Noma et al. 2022])

$$n(f_{ijkl}) := \frac{1}{2\pi} \oint_{\partial f_{ijkl}} d\alpha = \frac{1}{2\pi} (\sigma_{ij} + \sigma_{jk} + \sigma_{kl} + \sigma_{li}), \quad (16)$$

where σ_{ij} denotes the unwrapped phase increment along edge e_{ij} . If $n(f_{ijkl}) = 0$, the vector field $\mathbf{z}(\mathbf{x})$ is regular over the face. Otherwise, $n(f_{ijkl}) \approx k \in \mathbb{Z} \setminus \{0\}$ indicates the presence of a field singularity within the face, where such a face is called a singular face. For a grid cell e , the cell is singular if any face is singular. Without loss of generality, a singular cell has even number of singular faces. The

following metric evaluates whether e is singular by the Heaviside function $H(\cdot)$:

$$I_e = H\left(\sum_{f \in e} |n(f)|\right), \quad (17)$$

where a soft approximation as sigmoid is employed in our differentiable pipeline.

4.2 Physical Modeling

In this subsection, we formulate the mechanical model of our approach, focusing on the treatment of transient mechanical properties and inter-layer anisotropy with varying orientations over the computational domain Ω , both of which evolve dynamically during optimization.

4.2.1 Transient Mechanical Properties. Mechanical properties are evaluated via FEA on all elements within the computational domain Ω . An element with density $\rho = 1$ is treated as solid and assigned an appropriate Young's modulus together with the corresponding constitutive matrix. For void elements (i.e., $\rho = 0$), a small Young's modulus $E_{\min} = 1.0\text{e-}3$ is employed to ensure numerical stability.

Since partial layers are permitted, we further distinguish between elements that intersect partial-layer boundaries and those fully filled by complete laminate layers. Elements containing partial-layer boundaries are modeled as isotropic material (e.g., resin), with Young's modulus $E_{\text{iso}} = (E_r, E_r, E_r)$. In contrast, elements fully occupied by laminated layers exhibit pronounced inter-layer anisotropy, characterized by an orthotropic Young's modulus $E_{\text{lam}} = (E_x, E_y, E_z)$ with $E_x \approx E_y \gg E_z$. This modeling choice is motivated by prior studies on laminated composites with stacking sequences such as $[0^\circ/90^\circ, 45^\circ/135^\circ]$ and their symmetric counterparts. Such layups are nearly isotropic in the laminate plane: the effective in-plane stiffness is largely insensitive to rotational orientation, while stiffness along the laminate normal (inter-layer direction) is significantly weaker [Wisnom 1999]. These coefficients of Young's modulus can be obtained from mechanical tests. Specifically, $E_r = E_z = 4.2$ GPa and $E_x = E_y = 32.7$ GPa are observed experimentally and correspond to a resin-dominated response and a significantly stiffer in-plane behavior respectively; these values are used in all examples presented in this paper.

By incorporating the above considerations and following the density interpolation scheme of the SIMP method [Sigmund 2001], the interpolated Young's modulus for an element $e \subset \Omega$ is defined as

$$E_e = E_{\min} + H^p(\rho(\mathbf{x}_e^c))(((1 - I_e)E_{\text{lam}} + I_e E_{\text{iso}}) - E_{\min}), \quad (18)$$

where the density field $\rho(\mathbf{x})$ is evaluated at the element center \mathbf{x}_e^c , and $p = 3$ is used as the penalization exponent for the Heaviside function $H(\cdot)$. The transition between regular elements and singular elements (i.e., those containing partial-layer boundaries) is smoothly handled through the singularity metric I_e defined in Eq. (17).

4.2.2 Element Stiffness. Inter-layer anisotropy of laminated composites is explicitly considered when formulating the element stiffness matrix for FEA. This results in a corresponding 6×6 constitutive matrix in Voigt notation, denoted by $\mathbf{C}(E_e)$, which is assembled using the interpolated element Young's modulus E_e together with orthotropic mechanical properties. For simplicity, we assume constant




Fig. 4. The material frame, defined by the local layer orientation \mathbf{n}_e at the element center \mathbf{x}_e^c , is aligned with the Cartesian coordinate frame through two successive rotations. First, a rotation by angle θ brings the material frame's Z -axis onto the xy -plane of the Cartesian frame. This is followed by a rotation by angle ϕ , which aligns the material frame's Z -axis with the Cartesian z -axis.

Poisson's ratios and set both the in-plane and inter-layer Poisson's ratios to 0.25 throughout the domain [Jones 1999]. Shear moduli are computed from the Young's moduli and the constant Poisson's ratio.

For each regular element e , the layer orientation is computed as the average of the vector field $\mathbf{z}(\cdot)$ over the element domain and denoted by \mathbf{n}_e . A local material frame (X, Y, Z) is then constructed by aligning its Z -axis with \mathbf{n}_e and placing its origin at the element center \mathbf{x}_e^c . Two rotation angles, θ about the X -axis and ϕ about the Y -axis, are subsequently determined to align the material frame with the global Cartesian coordinate frame of the computational domain Ω (see Fig. 4).

With the above preparation, the element stiffness matrix is defined as

$$\mathbf{K}_e = \mathbf{T}(\mathbf{x}_e^c) \mathbf{C}(E_e) \mathbf{T}(\mathbf{x}_e^c)^T, \quad (19)$$

where $\mathbf{T}(\mathbf{x}_e^c) = \mathbf{T}_2(I_e \phi) \mathbf{T}_1(I_e \theta) \in \mathbb{R}^{6 \times 6}$ is a transformation matrix associated with two successive rotations (see Fig. 4). Specifically, we employ

$$\mathbf{T}_1(\theta) = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & c_1^2 & s_1^2 & -2s_1c_1 & 0 & 0 \\ 0 & s_1^2 & c_1^2 & 2s_1c_1 & 0 & 0 \\ 0 & s_1c_1 & -s_1c_1 & c_1^2 - s_1^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & c_1 & s_1 \\ 0 & 0 & 0 & 0 & -s_1 & c_1 \end{bmatrix} \quad (20)$$

$$\mathbf{T}_2(\phi) = \begin{bmatrix} c_2^2 & 0 & s_2^2 & 0 & -2s_2c_2 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ s_2^2 & 0 & c_2^2 & 0 & 2s_2c_2 & 0 \\ 0 & 0 & 0 & c_2 & 0 & -s_2 \\ s_2c_2 & 0 & -s_2c_2 & 0 & c_2^2 - s_2^2 & 0 \\ 0 & 0 & 0 & s_2 & 0 & c_2 \end{bmatrix} \quad (21)$$

with $c_1 = \cos \theta$, $s_1 = \sin \theta$, $c_2 = \cos \phi$, and $s_2 = \sin \phi$. By incorporating the singularity metric I_e , Eq. (19) naturally accommodates both regular elements and elements intersecting partial-layer boundaries, enabling a smooth transition between anisotropic and isotropic constitutive responses.

4.2.3 Global Displacement. Assembling all element stiffness matrices yields the global stiffness matrix \mathbf{K} . The global displacement

field \mathbf{U} is then obtained by solving the linear system

$$\mathbf{K}\mathbf{U} = \mathbf{F}, \quad (22)$$

where \mathbf{F} denotes the applied load vector. The resulting displacement field is defined over all elements and is subsequently used to evaluate the compliance energy for stiffness optimization.

5 Optimization Formulation

We adopt a widely used design objective in topology optimization, namely minimizing the compliance energy subject to a prescribed volume fraction of the resulting structure relative to the computational domain Ω . Following the SIMP framework for topology optimization [Sigmund 2001], the co-optimization of fabric-reinforced laminated composites can be formulated as

$$\begin{aligned} \min_{\theta_\rho, \theta_z} \quad & \mathbf{F}^T \mathbf{U} \\ \text{s.t.} \quad & \mathbf{K}\mathbf{U} = \mathbf{F}, \\ & \int_{\Omega} H(\rho(\mathbf{x})) \, d\mathbf{x} \leq V^*, \\ & \Pi_j(\theta_\rho, \theta_z), \quad j = 1, 2, \dots, M. \end{aligned} \quad (23)$$

The network parameters θ_ρ and θ_z of the two implicit fields $\rho(\cdot)$ and $\mathbf{z}(\cdot)$ serve as the design variables. The first constraint in Eq. (23) is the equilibrium equation (i.e., Eq.(22)) for FEA computed on a discrete mesh to determine the displacements \mathbf{U} . The total volume of the structure is then evaluated in the second constraint and controlled by the maximally allowed volume V^* . The manufacturing requirements are defined as the constraints $\Pi_j(\cdot)$ in the last row, the details of which will be presented in Sec. 5.2.

The optimization objectives in our framework are formulated as loss functions. In this section, we present them in three categories: design, manufacturing, and regularization.

5.1 Design Losses

Two design objectives are formulated as loss functions, corresponding to the compliance energy (i.e., mechanical stiffness) and the total material volume.

5.1.1 Stiffness. For a given loading case with force vector \mathbf{F} and the corresponding displacement field \mathbf{U} obtained by solving Eq. (22), the structural compliance is defined as

$$J = \mathbf{F}^T \mathbf{U} = \sum_e \mathbf{U}_e^T \mathbf{K}_e \mathbf{U}_e, \quad (24)$$

where \mathbf{U}_e denotes the vector of nodal displacements associated with element e , and \mathbf{K}_e is the element stiffness matrix defined in Eq. (19). The compliance J represents the work done by external forces and is inversely related to the overall stiffness of the structure under the applied loading.

The stiffness loss is defined as

$$\mathcal{L}_{\text{stf}} := J/J_0, \quad (25)$$

where J_0 is a reference compliance used to non-dimensionalize the objective. Without loss of generality, we compute J_0 using a fully solid structure with planar layers, i.e., $\rho(\mathbf{x}) = 1$ and $\mathbf{z}(\mathbf{x}) = (0, 0, 1)$ for all $\mathbf{x} \in \Omega$.

5.1.2 *Volume.* To enforce the volume constraint in Eq. (23) with respect to the prescribed maximum volume V^* , we introduce the following volume loss as

$$\mathcal{L}_{\text{vol}} := \text{ReLU} \left(\frac{\sum_e H(\rho(\mathbf{x}_e^e)) V_e}{V^*} - 1 \right), \quad (26)$$

where V_e denotes the volume of element e . The $\text{ReLU}(\cdot)$ function ensures that the loss is activated only when the total material volume exceeds the prescribed limit V^* .

5.2 Manufacturing Losses

All manufacturability losses are evaluated at a set of points \mathcal{P} uniformly sampled over the design domain Ω . These loss terms are computed only within the solid region, i.e., at sample points satisfying $H(\rho(\mathbf{x}_p)) \approx 1$. Accordingly, we define the following normalization factor:

$$\Psi = \sum_{\mathbf{x}_p \in \mathcal{P}} H(\rho(\mathbf{x}_p)), \quad (27)$$

which measures the total number (or accumulated weight) of samples located inside the solid region. This normalization ensures that the manufacturability losses are scale-invariant with respect to the sampling density.

5.2.1 *Gaussian Curvature.* Considering the developable surface requirement for each layer, the Gaussian curvature within the solid region should ideally vanish (i.e., $\kappa_G \approx 0$). We evaluate the Gaussian curvature on the implicit level sets of the scalar field $\alpha(\mathbf{x})$ using the formulation of [Goldman 2005], under the assumption that the constraint $\nabla\alpha(\mathbf{x}) = \mathbf{z}(\mathbf{x})$ has been enforced. For a point $\mathbf{x} \in \mathbb{R}^3$, the Gaussian curvature of the implicit surface is given by

$$\kappa_G(\mathbf{x}) = \frac{\mathbf{z}(\mathbf{x})^T \mathbf{H}_\alpha^*(\mathbf{x}) \mathbf{z}(\mathbf{x})}{\|\mathbf{z}(\mathbf{x})\|^4}, \quad (28)$$

where $\mathbf{H}_\alpha^*(\mathbf{x})$ denotes the adjoint matrix of $\mathbf{H}_\alpha(\mathbf{x})$ – i.e., the Hessian matrix of $\alpha(\mathbf{x})$. Note that the Hessian $\mathbf{H}_\alpha(\mathbf{x})$ can be evaluated directly from the spatial derivatives of $\mathbf{z}(\mathbf{x})$.

We then define a Gaussian curvature loss to penalize the non-developable regions as

$$\mathcal{L}_{\text{gc}} := \frac{1}{\Psi} \sum_{\mathbf{x}_p \in \mathcal{P}} H(\rho(\mathbf{x}_p)) \text{ReLU}(|\kappa_G(\mathbf{x}_p)| - \kappa_{gc}), \quad (29)$$

where \mathcal{P} denotes the set of sampled points, $H(\rho(\cdot))$ is an indicator function that restricts the loss to points within the structural solid, and $\kappa_{gc} = 0.005$ is a user-defined curvature threshold determined empirically.

5.2.2 *Helicoid Free.* Although the topology of the isosurfaces of the phase field $\alpha(\mathbf{x})$ within a singular element can be highly complex, we observe that enforcing the local constraint $\kappa_G \approx 0$ restricts the admissible surface configurations to only two types: (1) a collection of developable surfaces, or (2) a helicoid surface, which are not manufacturable using fabric lamination. If the set of singular points within a singular element is traced as a singular curve, two characteristic cases can be observed (also illustrated in Fig. 5):

- When the singular curve forms the boundary of a partial layer, it is nearly orthogonal to the surface normal;

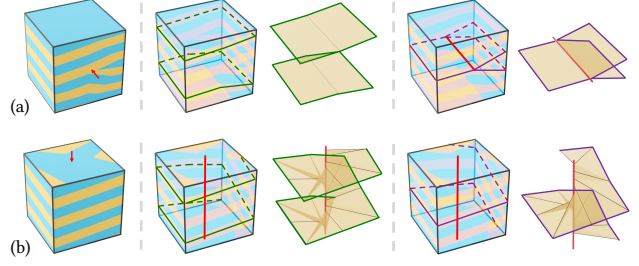


Fig. 5. Two different cases of singular curves (displayed in red color) within a singular element: (a) the boundary curve of a partial layer, and (b) the rational axis of a helicoid surface.

- When the singular curve coincides with the axis of a helicoid surface, the curve (i.e., the axis) is nearly parallel to the surface normal.

These geometric observations are exploited to construct a loss function that penalizes the formation of helicoid surfaces.

For ease of implementation, given an element e that contains a singularity (i.e., $I_e = 1$), we first identify all of its singular faces and store them in the set \mathcal{F}_e^s . For any pair of faces $f_i, f_j \in \mathcal{F}_e^s$, the line segment $\mathbf{c}_i \mathbf{c}_j$ is used as a simplified approximation of the singular line, where \mathbf{c}_i and \mathbf{c}_j denote the centers of f_i and f_j respectively. Based on this construction, we define the helicoid-free loss that softly enforces that all soft singularity directions remain at least $\frac{\pi}{4}$ away from the local layer normal. That gives

$$\mathcal{L}_{\text{hef}} := \frac{1}{\Psi} \sum_{e \in \mathcal{C}} I_e \sum_{\mathbf{x}_p \in \mathcal{P}_e} H(\rho(\mathbf{x}_p)) \left(\frac{|\mathcal{F}_e^s|}{2} \right)^{-1} \sum_{f_i, f_j \in \mathcal{F}_e^s} \text{ReLU} \left(\frac{|\mathbf{z}(\mathbf{x}_p) \cdot (\mathbf{c}_j - \mathbf{c}_i)|}{\|\mathbf{z}(\mathbf{x}_p)\| \|\mathbf{c}_j - \mathbf{c}_i\|} - \cos \frac{\pi}{4} \right), \quad (30)$$

where $\mathcal{P}_e \subset \mathcal{P}$ denotes the set of sampling points that lie inside element e , and $|\mathcal{F}_e^s|$ gives the number of faces in \mathcal{F}_e^s . By this loss, the optimization is steered away from helicoid-like configurations, yielding helicoid-free layers as the resultant isosurfaces of the phase field $\alpha(\cdot)$.

5.3 Regularization Loss

Even when the global developability and helicoid-free constraints are satisfied, the phase field under optimization may still exhibit high-frequency oscillations. We therefore introduce an edge-based regularization loss to encourage smoothness of the vector field $\mathbf{z}(\mathbf{x})$. This regularization is applied only within the solid region as

$$\mathcal{L}_{\text{reg}} := \frac{\sum_{(v_i, v_j) \in \mathcal{E}} H(\rho(\frac{1}{2}(\mathbf{x}_i + \mathbf{x}_j))) \|\mathbf{z}(\mathbf{x}_i) - \mathbf{z}(\mathbf{x}_j)\|_2^2}{\sum_{(v_i, v_j) \in \mathcal{E}} H(\rho(\frac{1}{2}(\mathbf{x}_i + \mathbf{x}_j)))} \quad (31)$$

This loss encourages neighboring vectors – equivalently, the gradients of the phase field $\alpha(\cdot)$ – to vary smoothly across adjacent regions.

5.4 Total Loss

The total loss function for the co-optimization of fabric-reinforced laminated composites, corresponding to the problem formulated in

Eq. (23), is defined as a weighted sum of the individual loss terms.

$$\mathcal{L}_{\text{total}} := \omega_{\text{stf}} \mathcal{L}_{\text{stf}} + \omega_{\text{vol}} \mathcal{L}_{\text{vol}} + \omega_{\text{gc}} \mathcal{L}_{\text{gc}} + \omega_{\text{hef}} \mathcal{L}_{\text{hef}} + \omega_{\text{reg}} \mathcal{L}_{\text{reg}}. \quad (32)$$

The weights of the loss terms are adjusted dynamically to balance the design objective with the other losses. Details of the weighting strategy are provided in Sec. 6.2.

6 Implementation Details

6.1 Differentiable Pipeline

Based on the geometric and physical modeling framework introduced in Sec. 4, all loss functions defined in Sec. 5 are differentiable, enabling the use of gradient-based optimizers (e.g., Adam [Kingma and Ba 2015]) to efficiently solve the proposed optimization problem. We analyze the differentiability of the overall pipeline below.

First, the Heaviside function appearing in several formulations is approximated using a sigmoid function, ensuring smoothness and differentiability. With this approximation, it is straightforward to verify that the loss terms \mathcal{L}_{vol} , \mathcal{L}_{gc} , and \mathcal{L}_{reg} are differentiable with respect to the neural network parameters θ_ρ and θ_z , as these losses are directly formulated in terms of the scalar field $\rho(\cdot)$ and the vector field $\mathbf{z}(\cdot)$. Both fields are parameterized by the coefficients θ_ρ and θ_z of neural networks.

The differentiability of \mathcal{L}_{stf} and \mathcal{L}_{hef} is less immediate. In particular, the displacement field \mathbf{U} and stiffness matrix \mathbf{K} in \mathcal{L}_{stf} , as well as the singularity metric I_e in \mathcal{L}_{hef} , depend on the phase field $\alpha(\cdot)$. The phase field is obtained by solving an eigenvector problem defined on the vector field $\mathbf{z}(\cdot)$. Theoretical results on the differentiability of eigenvector problems can be found in [Sun 1990]. Following similar techniques to those described in Sec. 6.1 of [Montes Maestre et al. 2023], the differentiability of $\alpha(\cdot)$ with respect to $\mathbf{z}(\cdot)$ (consequently the network parameters θ_z) can be established.

In summary, all loss terms appearing in Eq. (32) are differentiable with respect to the learnable parameters θ_ρ and θ_z , allowing the entire optimization pipeline to be trained end-to-end using gradient-based methods.

6.2 Network Architecture, Learning and Weighting

The architecture of *multilayer perceptrons* (MLPs) is adopted to parameterize both the density field $\rho(\cdot)$ and the vector field $\mathbf{z}(\cdot)$. In our implementation, each field is represented by an independent MLP consisting of five hidden layers, each with 256 neurons. This network configuration is chosen empirically to effectively capture the geometric complexity of the corresponding fields. To ensure third-order continuity, the *Sigmoid Linear Unit* (SiLU) activation function is employed throughout all layers. The networks take the spatial position vector $\mathbf{x} \in \mathbb{R}^3$ as input and generate the corresponding values of the scalar density field $\rho(\mathbf{x})$ and the vector field $\mathbf{z}(\mathbf{x})$ as output respectively.

The optimization problem is solved using the NN-based automatic differentiation framework provided by PyTorch [Paszke et al. 2019]. The proposed loss function is minimized via backpropagation to update the network parameters. After initializing the field values (as detailed in the next subsection), the optimization is carried out using an initial learning rate of 1.0×10^{-3} . The learning rate is adaptively adjusted during training using the 'ReduceLROnPlateau'

scheduler, which monitors the optimization progress and reduces the learning rate when a plateau is detected, thereby promoting stable and efficient convergence. The minimum learning rate is set to 1.0×10^{-6} to ensure robust convergence control.

We adopt a weighting strategy similar to [Liu et al. 2025], where \mathcal{L}_{stf} serves as the primary objective and all other loss terms are treated as constraints. A hybrid scheme combining the penalty method of TOuNN [Chandrasekhar and Suresh 2021] and the dynamic constraint handling of DC3 [Donti et al. 2021] is employed. The weight ω_{stf} is initialized as $10/\mathcal{L}_{\text{stf}}$ based on the initial solution and is dynamically adjusted to balance $\omega_{\text{stf}} \mathcal{L}_{\text{stf}}$ with the remaining loss terms, while remaining within $[0, 10]$. The weights of the constraint losses are initialized to zero and increased by 0.05 per iteration. When $\mathcal{L}_{\text{stf}} < 1.0 \times 10^{-5}$, a constraint-correction step is applied by temporarily setting $\omega_{\text{stf}} = 0$. This hybrid scheme is found to be effective in practice. More details and discussions for balancing the competition among loss terms during training can be found in [Wang et al. 2021].

6.3 Field Initialization

The density field $\rho(\cdot)$ is initialized as a uniform field with a value slightly above 0.5, which is realized by assigning random coefficients for the first four MLP layers and then computing the coefficients of the last layer to make this uniform output.

The vector field $\mathbf{z}(\cdot)$, which specifies the layer orientations, cannot be initialized in a trivial manner analogous to the density field. This is because the phase field $\alpha(\cdot)$ in our formulation is a semi-continuous periodic field. As a result, the condition $\nabla \alpha(\mathbf{x}) = \mathbf{z}(\mathbf{x})$ generally admits multiple solutions, some of which lead to fragmented layers in the resulting isosurfaces of $\alpha(\cdot)$ (as illustrated in Fig. 6). To mitigate this ambiguity induced by the semi-continuous periodic formulation, we introduce a more robust initialization strategy, named *continuous-field gradient* (CFG) scheme, which is applied to generate the initial guess of our self-learning based optimization. Specifically, the CFG scheme consists of the following two steps:

- Computing a continuous scalar field $\tilde{\alpha}(\cdot)$ parameterized on a neural network \mathcal{N}_β having the same architecture as \mathcal{N}_ρ by minimizing $\sum_{v_i \in \mathcal{V}} \|\nabla \tilde{\alpha}(\mathbf{x}_i) - \mathbf{z}(\mathbf{x}_i)\|^2$ over all grid nodes, using the current estimate of $\mathbf{z}(\cdot)$;
- Updating the network parameters θ_z such that $\mathbf{z}(\mathbf{x}) = \nabla \tilde{\alpha}(\mathbf{x})$ is satisfied for all grid nodes in \mathcal{V} .

By explicitly enforcing global consistency through the continuous field $\tilde{\alpha}(\cdot)$, this initialization scheme can effectively capture the global structure of the layer orientations and significantly reduce the occurrence of fragmented layers in subsequent optimization.

6.4 Post-processing

After optimizing the density field $\rho(\cdot)$ and the vector field $\mathbf{z}(\cdot)$ via the proposed NN-based self-learning framework, the surface patches corresponding to laminated layers are generated in a post-processing stage. This stage consists of two main steps: (1) layer extraction and (2) layer flattening.

For layer extraction, we first compute the phase field $\alpha(\cdot)$ corresponding to the optimized vector field $\mathbf{z}(\cdot)$. Isosurfaces of $\alpha(\cdot)$ with isovalues $(2\mathbb{Z} + 1)\pi$ are extracted and treated as layer patches for

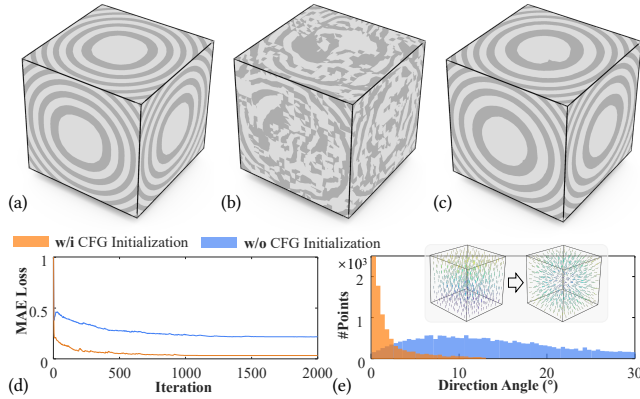


Fig. 6. For a semi-continuous periodic phase field $\alpha(\cdot)$ (as shown in (a)) whose gradient $\nabla\alpha(\cdot)$ is constrained to align with a radial vector field $\mathbf{v}(\cdot)$ pointing outward from the domain center, a naïve optimization may lead to fragmented layers, as shown in (b). By applying the proposed CFG initialization, the optimization better captures the global structure of the vector field, resulting in coherent and non-fragmented layers as illustrated in (c). The comparisons are given for (d) the convergence curves of $E_{\text{avg}} = \frac{1}{|\mathcal{P}|} \sum_{\mathbf{x}_p \in \mathcal{P}} \|\mathbf{v}(\mathbf{x}_p) - \nabla\alpha(\mathbf{x}_p)\|$ and (e) the histogram of angle deviations between $\mathbf{v}(\cdot)$ and $\nabla\alpha(\cdot)$, which both demonstrate the effectiveness of CFG initialization.

laminated composites. For each element e containing at least one point with density value $\rho > 0.5$, a marching-cross-style method presented in Sec. 3.4.2 of [Sharp et al. 2019] is applied to extract connected isocurves on the faces of e . For regular elements, these isocurves form closed loops. In contrast, for singular elements, a closed loop can only be obtained by adding the corresponding singular line, as described in Sec. 5.2.2. Triangular surface meshes are then generated using the hole-filling triangulation method of [Kazhdan et al. 2007]. The resolution of the grids used for mesh surface extraction is set as $4\times$ of the FEA grids. Finally, the resulting meshes are trimmed by the implicit solid \mathcal{H} defined in Eq. (1) to produce the final geometry of each layer.

Although the Gaussian curvature of the extracted isosurfaces has been optimized to be nearly zero, the resultant surface patches are generally not perfectly developable due to numerical errors introduced during isocurve extraction and subsequent triangulation. We apply the as-rigid-as-possible surface flattening method [Liu et al. 2008] to each triangular mesh to finally obtain planar panel shapes for fabrication.

7 Results and Discussion

7.1 Computational Experiments

Our computational pipeline is implemented in Python. The source code of our pipeline will be released upon the acceptance of this paper. All computational experiments are performed on a desktop PC with an Intel Core i5-12600K CPU (10 cores @ 3.6GHz) with 96GB RAM and an NVIDIA RTX4080 GPU with 16GB VRAM, running Ubuntu 20.04 LTS.

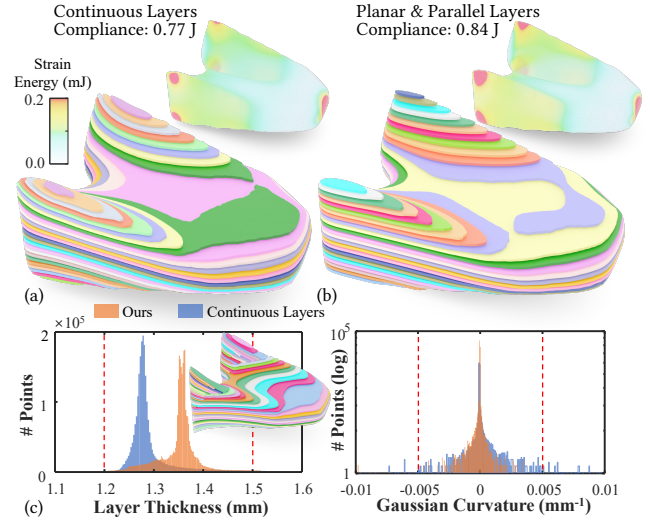


Fig. 7. Result comparison for the cantilever beam example. Results obtained using (a) curved continuous layers and (b) planar & parallel layers generated by NeuralTOMO are shown, with strain energy distributions visualized using color maps. Histograms of layer thickness and Gaussian curvature are provided in (c) to evaluate manufacturability.

7.1.1 Examples and Comparisons. Our approach has been tested on a variety of examples and are compared with existing approaches, where the computational statistics are listed in Table 1.

The first example considers the optimization of a Boat structure under the loading conditions illustrated in Fig. 1, which is similar to [Mancuso et al. 2022]. For comparison, we generate co-optimization results using NeuralTOMO [Liu et al. 2025]. Specifically, curved continuous layers are produced by allowing 5-axis motion, while planar & parallel layers are obtained by restricting the motion to 2.5-axis printing. The same maximum allowable structural volume is imposed in both methods. In addition, the NeuralTOMO implementation is modified to incorporate the Gaussian curvature loss formulated in Eq. (28). Statistical analyses of the resultant designs are visualized using the histograms shown in Fig. 1. Overall, the designs generated by NeuralTOMO exhibit weaker mechanical performance compared to our results. Although the Gaussian curvature of NeuralTOMO’s layers can be partially controlled through the above modification, it fails to produce curved layers with well-controlled layer thickness.

The second example considers the Cantilever Beam design shown in Fig. 3. Following the same comparison protocol as in the Boat example, we compute NeuralTOMO results using both curved and planar layers, as shown in Fig. 7. The compliance energy J , defined in Eq. (24), is reported for quantitative comparison. While our optimized composite achieves a compliance of 0.67 J (Fig. 3), the designs generated by NeuralTOMO exhibit 14.5% and 25.3% higher compliance – i.e., lower stiffness respectively. In addition, histograms of layer thickness and Gaussian curvature are provided for manufacturability evaluation. In summary, although NeuralTOMO can effectively control layer thickness, the curved continuous layers it

Table 1. Statistics of computational and physical experiments

Model	Fig.	FEA Grid Resolution	# Singular Chains	Peak Memory (GB) [†]			Computing Time (sec.)					Fabrication Info.	
				FEA	Phase	All	FEA [‡]	Phase [‡]	Optm. [‡]	Post-proc.	Total	#Layers	Weight (g)
Boat	1	45 × 15 × 150	18	13.57	3.27	30.35	12.52	3.43	2.55	405.35	19,012	492	–
Cantilever Beam	3,7,9	45 × 30 × 30	12	8.35	2.01	14.37	5.32	1.17	0.94	147.62	7,598	206	657
A-Bracket	8	54 × 48 × 61	26	20.41	5.22	43.75	19.51	4.62	3.55	638.16	33,912	134	–
Shell (Ours)	10,12	90 × 13 × 90	4	13.83	3.73	33.48	15.93	3.82	2.48	362.33	22,642	92	424
Shell (Planar)	10,12	90 × 13 × 90	0	10.83	–	28.48	15.93	–	2.03	324.15	18,364	84	453
Propeller (Grid Size 1.0)	11,15	20 × 72 × 64	15	12.04	3.20	26.52	15.04	3.75	2.24	307.57	21,572	230	–
Propeller (Grid Size 1.5)	11,15	13 × 48 × 43	6	4.72	1.75	9.62	3.93	0.64	0.89	130.89	6,631	162	–
Propeller (Grid Size 2.0)	11,15	10 × 36 × 32	2	3.16	1.38	5.22	2.35	0.26	0.84	75.67	4,329	120	–

[†] FEA and Stripe Generation work on CPU, and Neural Network works on GPU, and all the VRAM usage is less than 16Gb.

[‡] The computing time of FEA, Phase, and the rest of the optimization is reported as the average value of each step.

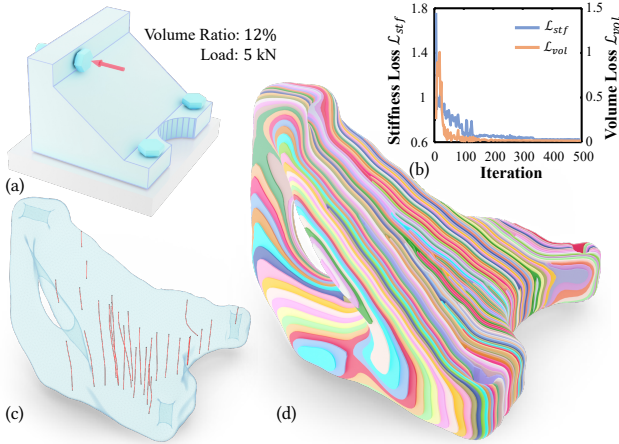


Fig. 8. Result of the A-Bracket example. (a) Loading conditions and the design domain, (b) the convergence curves of the stiffness loss \mathcal{L}_{stf} and the volume loss \mathcal{L}_{vol} , (c) the singular chains visualized as red curves, and (d) the optimized structure and curved semi-continuous layers.

generates – even after incorporating the Gaussian curvature constraint – still contain localized regions with large Gaussian curvature.

To demonstrate the importance of incorporating semi-continuous periodic field-based layer optimization that supports partial layers within the optimization loop, we conduct an additional study on the Cantilever Beam example. Specifically, we apply Neural Slicer [Liu et al. 2024] to the optimized solid produced by our approach to generate non-uniform layers, which are subsequently processed using the adaptive slicing scheme proposed in the S^3 -slicer [Zhang et al. 2022] to obtain ‘nearly uniform’ layer thickness (see Fig. 9). This post-processing strategy substantially degrades mechanical performance, resulting in compliance values that are 19.4% and 22.3% higher respectively. Moreover, this sequential strategy fails to jointly control layer thickness and Gaussian curvature. While the adaptive slicing step can enforce a minimum layer thickness, it does not regulate the maximum thickness, and the resulting layers still contain regions with large Gaussian curvature.

The third example concerns the topology optimization of an A-bracket with a prescribed volume fraction of 12% subjected to a load of 5 kN (see Fig. 8(a)). Our method simultaneously optimizes the

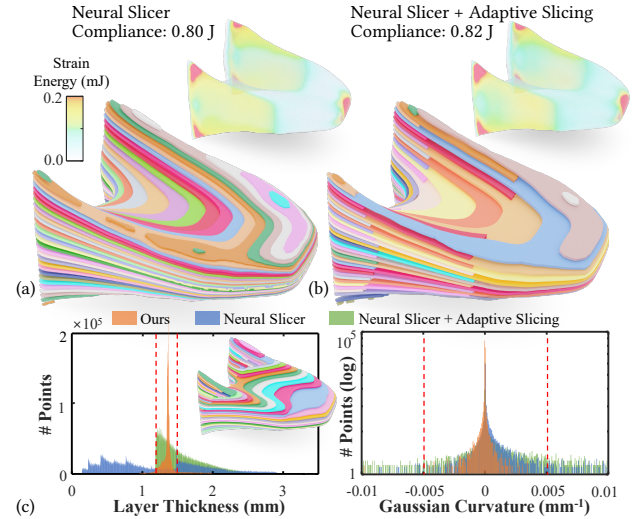


Fig. 9. Study of an alternative solution. (a) Application of Neural Slicer to the structural topology generated by our approach, followed by (b) a post-processing step using adaptive slicing. Histograms of two manufacturability metrics – layer thickness and Gaussian curvature – are shown in (c).

topological structure and the curved semi-continuous layers. As shown in Fig. 8(b), the optimization process converges in terms of both stiffness and volume. The final design gives an optimized load-bearing structure that can be fabricated using curved developable layers tailored to the geometry (Fig. 8(c)). The identified singular chains, which define the boundaries of the partial layers, are sparse and spatially localized within the optimized solid (Fig. 8(d)).

The fourth example considers a shell model subjected to a compressive force applied at its center, as shown in Fig. 10. The result generated by our approach is compared with those obtained by NeuralTOMO using planar & parallel layers. Comparisons of the strain energy distribution, the stiffness loss \mathcal{L}_{stf} , and the volume loss \mathcal{L}_{vol} are reported. The composite optimized by our approach achieves a compliance of 22.61 J, which is substantially lower than that of the planar-layer laminated composite generated by NeuralTOMO, with a compliance of 30.49 J.

The fifth example considers the topological design of a manufacturable interior core structure for the KP505 controllable-pitch




Fig. 10. Results of the Shell example. (a.1) The resultant structure and curved layers generated by our approach with its strain energy distributions shown in (a.2). (b.1) The resultant structure with planar layers generated by NeuralTOMO [Liu et al. 2024]. The convergence curves are given for (c.1) the stiffness loss \mathcal{L}_{stf} and (c.2) the volume loss \mathcal{L}_{vol} .

propeller [Korsmik et al. 2020], operating in open water at an advance coefficient of 0.5. The loading conditions and design domain are shown in Fig. 11(a); additional details can be found in [Korsmik et al. 2020]. The topology optimization result for fabric-reinforced laminated composite construction is shown in Fig. 11(b), achieving a compliance of 28.76 J. Notably, all generated layers exhibit very low Gaussian curvature, ensuring manufacturability.

7.1.2 Ablation Study for Gaussian Curvature. The ablation study for the Gaussian curvature loss \mathcal{L}_{gc} has been taken on the propeller example. As can be seen in Fig. 11(c, d), the resultant layers will have many regions with large Gaussian curvatures – i.e., not developable after turning off the Gaussian curvature loss by setting $\omega_{gc} = 0.0$ although the compliance energy can be further reduced from 28.76 J to 23.72 J. This test verifies the effectiveness of the \mathcal{L}_{gc} loss.

7.1.3 Ablation Study for Helicoid-Free. We further conduct an ablation study to validate the effectiveness of the helicoid-free loss \mathcal{L}_{hef} using the shell example. To disable this loss, we set $\omega_{hef} = 0$ during optimization. After convergence, the resulting layers are extracted and the induced singularity curves (shown in red) on the shell and across layers are visualized, as summarized in Fig. 12(b). Without the helicoid-free loss, the optimized layers exhibit a helical structure (see the zoom view in Fig. 12(b)): the underlying phase field drifts with a twist-like behavior, producing singularity curves that act as twisting axes – i.e., the layer surface normals become nearly parallel to the singular curves. Notably, layers with such helical structures cannot be made from fabric pieces and are therefore unsuitable for laminated composite fabrication. In contrast, when the full set of

loss terms is applied, the normals of the resulting layers are nearly perpendicular to the singular curves (see Fig. 12(a)) and all layers are manufacturable.

7.2 Physical results

Physical experiments were conducted to validate the optimized laminated composite designs. The specimens fabricated in our work follow a standard wet hand-layups workflow for woven laminated composites [Campbell 2004]. All specimens were fabricated by first 3D printing curved moulds that serve as geometric supports similar to support structures used beneath overhangs in additive manufacturing – see Fig.13(a) for an example of fabrication process. After that, each layer is first cut as a 2D panel, then warped onto the target geometry using a 3D mold, and finally stacked and consolidated. To improve assembly accuracy, alignment markers in the form of small holes are introduced on the panels to ensure consistent layer registration during layup. The laminated layers were fabricated using heavy twill-weave carbon fabric (Twill 12k Black Stuff Carbon Fibre, 650 g/m²). The nominal fabric thickness is 0.45 mm, and the consolidated single-ply thickness is approximately 0.65 mm, with an effective thickness range of 0.60–0.75 mm.

To prevent fraying of the carbon fabric during cutting, a prepregging step was adopted. The dry fabric was first impregnated on a flat plate using EL2 epoxy resin with the fast AT30 hardener and left at room temperature for approximately 1.5 hours.

The prepared plies were then laid up to follow the target curved geometry and consolidated under vacuum using an EL2 epoxy laminating system, with resin, FAST AT30 hardener, and SLOW AT30 hardener mixed at a weight ratio of 100 : 15 : 15. All lay-up and curing steps were performed at a temperature of 25°C with details following [Campbell 2004].

The results of all physically fabricated models are shown in Fig.13(b).

Mechanical tests were conducted to evaluate the stiffness of laminated composites fabricated using layers that are generated by our approach. All tests were performed on an INSTRON 5969 universal testing machine with a load capacity of 30 kN. Compression–bending tests were carried out on shell composites fabricated using our optimized semi-continuous layers (Fig. 10(a.1)) and the conventional planar & parallel layers (Fig. 10(b.1)). Force–displacement curves were recorded during the tests (Fig. 14), showing a 42.38% increase in stiffness – i.e., from 2.35 kN/mm to 3.38 kN/mm – for the specimens fabricated using our optimized structures and layers.

7.3 Discussion

7.3.1 Gaussian Curvature. Although woven fabrics can accommodate non-zero Gaussian curvature through in-plane shear, such deformation is much more restricted in laminated composite fabrication, where excessive shear may induce wrinkles, stress concentrations, and degraded structural performance. This limitation becomes more pronounced after resin impregnation or prepregging, which improves handling stability but reduces drapability. To reflect the physical behavior of impregnated carbon fabrics, our method constrains Gaussian curvature to a narrow band around zero, allowing limited geometric flexibility while keeping the optimized layers

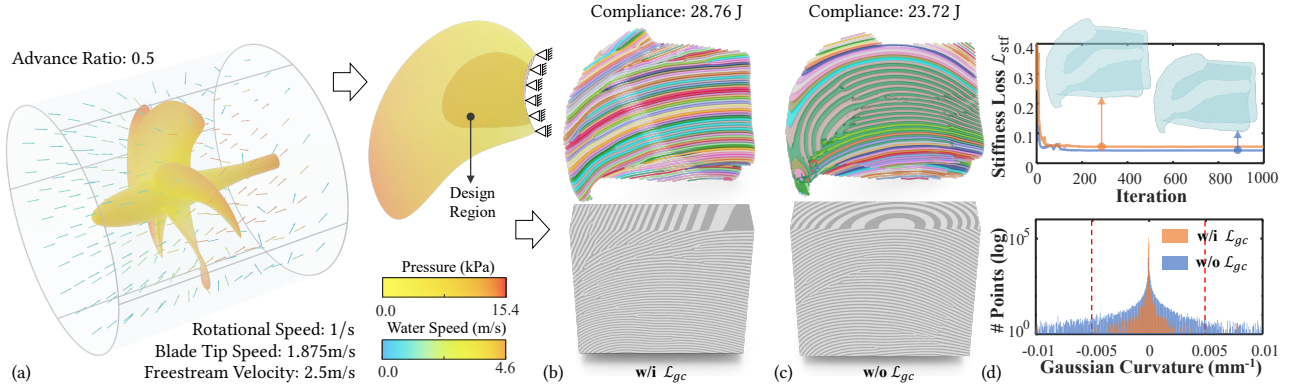


Fig. 11. Co-optimization results for the interior core structure of a laminated-composite propeller: (a) Loading conditions and design domain of the propeller, (b) optimized interior structure and curved semi-continuous layers generated by our approach, (c) the ablation study result obtained by disabling the Gaussian curvature loss \mathcal{L}_{gc} (i.e., using $\omega_{gc} = 0.0$), and (d) convergence curves of the stiffness loss \mathcal{L}_{stf} and histograms of Gaussian curvature for cases with and without the \mathcal{L}_{gc} loss.

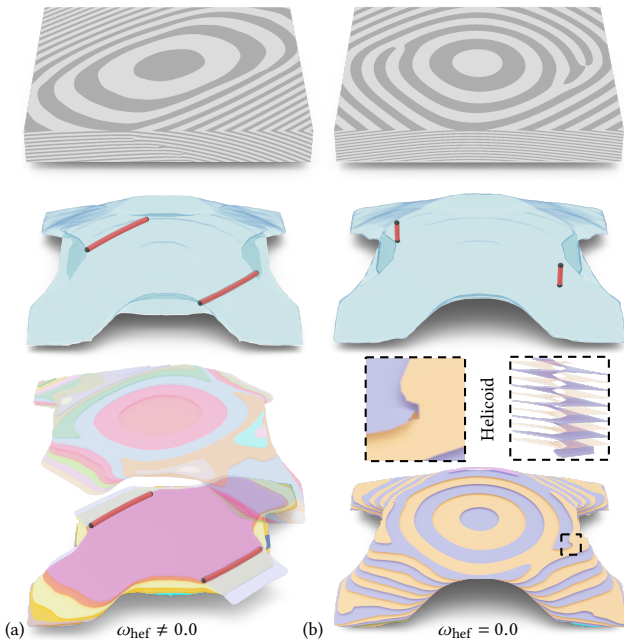


Fig. 12. Ablation study of the helicoid-free loss \mathcal{L}_{hef} : (a) the resultant structure and layers of the Shell example by using all losses in \mathcal{L}_{total} and (b) the result by using $\omega_{hef} = 0.0$ to disable the \mathcal{L}_{hef} loss. The singular curves are displayed in red color.

close to developable configurations. This manufacturability-driven relaxation acknowledges that small deviations from zero curvature improve geometric adaptability, whereas large non-zero Gaussian curvature is difficult to realize reliably without introducing laminate defects.



Fig. 13. The fabrication of laminated composites are taken on 3D printed moulds – see an example of fabrication sequence as illustrated in (a.1-a.6), and all the physically fabricated models are shown in (b.1-b.3).

7.3.2 Vector-Field Representation. We introduce a unit vector field² to compute the phase function rather than modeling it directly as a scalar field in order to implicitly represent a semi-continuous function. In other words, this formulation allows the phase function to tolerate value ‘jumps’. Specifically, Eq. (10) is solved in an intermediate step to convert the vector field into phase function values defined on the grid nodes. For the same reason, the smoothness regularization (Eq. (31)) is applied only to the vector field but not to the phase function itself; imposing smoothness directly on the phase

²The unit vector field $\mathbf{z}(\cdot)$ is obtained by normalizing the output of a neural network.

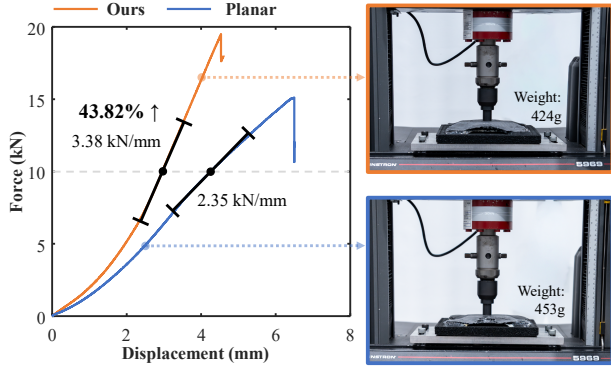


Fig. 14. Tensile test results for the Shell model. Specimens fabricated using our curved semi-continuous layers and the conventional planar & parallel layers shown in Fig. 10 are tested under compression. The measured force-displacement curves are shown, and the corresponding stiffness values are reported. Our method results in a structure with 43.8% stronger stiffness that has similar and slightly lighter weight.

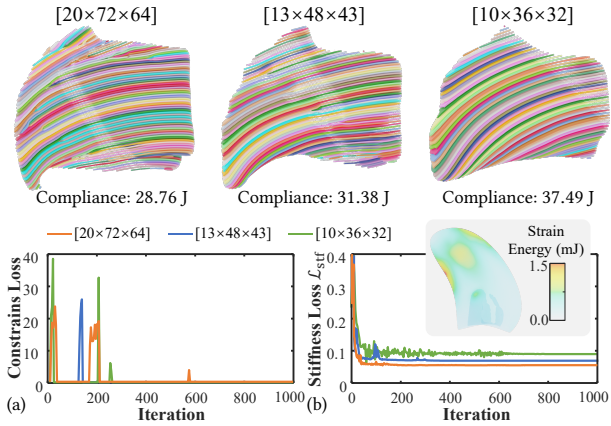


Fig. 15. Optimization results of the Propeller example by using grids in different resolutions: the convergence curves of both (a) the sum of loss terms considered as constraints and (b) the stiffness loss \mathcal{L}_{stf} .

function would prevent the optimization from accommodating such jumps.

7.3.3 Limitations. Several limitations of the proposed approach are discussed below, including the resolution dependency, the types of singularity and the layer extraction.

First, the optimization results are influenced by the grid resolution used in the computational pipeline. As illustrated by the comparison study on the propeller model shown in Fig. 15, we compute optimization results using grids at different resolutions. While the constraint losses consistently converge to near-zero values across all grid resolutions, the stiffness loss \mathcal{L}_{stf} converges to different values. The compliance energies of the resulting designs are evaluated using the *same* finite-element mesh in the commercial software Abaqus to ensure a fair comparison. We observe that employing a denser grid resolution enables the generation of laminated composite structures

with lower compliance, owing to the enlarged design space available for exploring improved structural configurations.

Secondly, the singularities permitted in our layer representation are currently restricted to the boundaries of surface patches, which can be effectively captured by the proposed semi-continuous periodic phase field. However, this representation does not support the generation of developable surfaces with darts or the assembly of connected developable surface patches as enabled in [Skouras et al. 2014; Zhang et al. 2019], not to mention non-smooth developable geometries with creases and folds [Schreck et al. 2016]. Although such geometries are less commonly employed in laminated composites – since layer discontinuities can induce stress concentrations and consequently weaken mechanical performance, their inclusion remains an interesting problem for future research. Moreover, strength-based topology optimization involving semi-continuous layers requires more advanced mechanical formulations to accurately model material behavior in the vicinity of singular regions, which has not been covered by our current work.

Lastly, the current layer mesh extraction procedure used in the post-processing stage is relatively complex and non-robust. It involves multiple steps, including strip segment generation, sorting and connection, hole filling and triangulation, reconstruction of complete mesh connectivity, and Boolean trimming with the implicit solid. A subsequent surface-flattening step is also required. Developing a mesh-free approach that can directly map or deform the isosurfaces of the phase field onto a planar domain would be an interesting direction for future work, as it could enable the direct extraction of planar layer panels in 2D and significantly simplify the post-processing pipeline.

8 Conclusions

We present a field-driven computational pipeline for co-optimizing structural topology and manufacturable layers with near-uniform thickness that can be warped from planar fabric plies. Design variables are parameterized by two neural networks representing a density field for structural topology and a vector field for layer orientations. A key component of the pipeline is the construction of semi-continuous periodic scalar fields from non-integrable vector fields, enabling the modeling of partial layers. Optimization objectives are achieved through a formulation of inter-layer anisotropic mechanical behavior that captures mechanical property transitions induced by partial-layer boundaries, together with additional loss terms enforcing manufacturability and field regularization. The framework is validated through numerical examples under diverse loading conditions, as well as physical fabrication and mechanical testing, demonstrating improved stiffness and manufacturability of fabric-reinforced laminated composites.

Acknowledgments

This project was supported by the Chair Professorship Fund at the University of Manchester and the UK Engineering and Physical Sciences Research Council (EPSRC) Fellowship Grant (Ref. #: EP/X032213/1).

References

- Masaki Aono, David E. Breen, and Michael J. Wozny. 1994. Fitting a Woven-Cloth Model to a Curved Surface: Mapping Algorithms. *Computer-Aided Design* 26, 4 (1994), 278–292.
- Rahul Arora, Alec Jacobson, Timothy R. Langlois, Yijiang Huang, Caitlin Mueller, Wojciech Matusik, Ariel Shamir, Karan Singh, and David I.W. Levin. 2019. Volumetric michell trusses for parametric design & fabrication. In *Proceedings of the 3rd Annual ACM Symposium on Computational Fabrication*. 1–13.
- F.C. Campbell. 2004. Manufacturing Processes for Advanced Composites. In *Manufacturing Processes for Advanced Composites*, F.C. Campbell (Ed.), 223–239.
- Aaditya Chandrasekhar and Krishnan Suresh. 2021. TOuNN: Topology Optimization using Neural Networks. *Structural and Multidisciplinary Optimization* 63, 3 (2021), 1135–1149.
- Yongxue Chen, Tao Liu, Yuming Huang, Weiming Wang, Tianyu Zhang, Kun Qian, Zikang Shi, and Charlie C. L. Wang. 2025. Can Any Model Be Fabricated? Inverse Operation Based Planning for Hybrid Additive-Subtractive Manufacturing. *ACM Transactions on Graphics* 44, 6 (2025), 1–14.
- Ping Cheng, Yong Peng, Shixian Li, Yanni Rao, Antoine Le Duigou, Kui Wang, and Said Alhi. 2023. 3D printed continuous fiber reinforced composite lightweight structures: A review and outlook. *Composites Part B: Engineering* 250 (2023), 110450.
- Xavier Chermain, Giovanni Cocco, Cédric Zanni, Eric Garner, Pierre-Alexandre Hugron, and Sylvain Lefebvre. 2025. Atomizer: Beyond Non-Planar Slicing for Fused Filament Fabrication. *Computer Graphics Forum* 44, 5 (2025).
- Xavier Chermain, Cédric Zanni, Jonàs Martínez, Pierre-Alexandre Hugron, and Sylvain Lefebvre. 2023. Orientable Dense Cyclic Infill for Anisotropic Appearance Fabrication. *ACM Transactions on Graphics* (2023).
- Etienne Corman and Keenan Crane. 2019. Symmetric moving frames. *ACM Transactions on Graphics* 38, 4 (2019), 1–12.
- Chengkai Dai, Charlie C.L. Wang, Chenming Wu, Sylvain Lefebvre, Guoxin Fang, and Yong-Jin Liu. 2018. Support-free volume printing by multi-axis motion. *ACM Transactions on Graphics* 37, 4 (2018), 1–14.
- Olga Diamanti, Amir Vaxman, Daniele Panozzo, and Olga Sorkine-Hornung. 2014. Designing N-PolyVector Fields with Complex Polynomials. *Computer Graphics Forum* 33, 5 (2014), 1–11.
- Qiujie Dong, Huibiao Wen, Rui Xu, Shuangmin Chen, Jiaran Zhou, Shiqing Xin, Changhe Tu, Taku Komura, and Wenping Wang. 2025. NeurCross: A Neural Approach to Computing Cross Fields for Quad Mesh Generation. *ACM Transactions on Graphics* 44, 4 (2025), 1–12.
- Priya Donti, David Rolnick, and J Zico Kolter. 2021. DC3: A learning method for optimization with hard constraints. In *International Conference on Learning Representations*. Vienna, Austria.
- Jimmy Etienne, Nicolas Ray, Daniele Panozzo, Samuel Hornus, Charlie C.L. Wang, Jonàs Martínez, Sara McMains, Marc Alexa, Brian Wyvill, and Sylvain Lefebvre. 2019. Curvilinear: Slightly curved slicing for 3-axis printers. *ACM Transactions on Graphics* 38, 4 (2019), 1–11.
- Guoxin Fang, Tianyu Zhang, Yuming Huang, Zhizhou Zhang, Kunal Masania, and Charlie C.L. Wang. 2024. Exceptional mechanical performance by spatial printing with continuous fiber: Curved slicing, toolpath generation and physical verification. *Additive Manufacturing* 82 (2024), 104048.
- Guoxin Fang, Tianyu Zhang, Sikai Zhong, Xiangjia Chen, Zichun Zhong, and Charlie C. L. Wang. 2020. Reinforced FDM: Multi-axis filament alignment with controlled anisotropic strength. *ACM Transactions on Graphics* 39, 6 (2020), 1–15.
- Pingchu Fang, Tong Gao, Yongbin Huang, Longlong Song, Hongquan Liu, Pierre Duysinx, and Weihong Zhang. 2025. Uniform multiple laminates interpolation model and design method for double-double laminates based on multi-material topology optimization. *Computer Methods in Applied Mechanics and Engineering* 433 (2025), 117492.
- Yogesh Gandhi, Alejandro M. Aragón, Julián Norato, and Giangiacomo Minak. 2025. A geometry projection method for the topology optimization of additively manufactured variable-stiffness composite laminates. *Computer Methods in Applied Mechanics and Engineering* 435 (2025), 117663.
- Ron Goldman. 2005. Curvature formulas for implicit curves and surfaces. *Computer Aided Geometric Design* 22, 7 (2005), 632–658.
- K.M.F. Hasan, P.G. Horváth, and T. Alpar. 2021. Potential fabric-reinforced composites: a comprehensive review. *Journal of Materials Science* 56 (2021), 14381–14415.
- Ruizhen Hu, Honghua Li, Hao Zhang, and Daniel Cohen-Or. 2014. Approximate pyramidal shape decomposition. *ACM Transactions on Graphics* 33, 6 (2014), 1–12.
- Alexandra Ion, Michael Rabinovich, Philipp Herholz, and Olga Sorkine-Hornung. 2020. Shape approximation by developable wrapping. *ACM Transactions on Graphics* 39, 6 (2020), 1–12.
- Caigui Jiang, Cheng Wang, Florian Rist, Johannes Wallner, and Helmut Pottmann. 2020. Quad-mesh based isometric mappings and developable surfaces. *ACM Transactions on Graphics* 39, 4 (2020), 1–13.
- Robert M. Jones. 1999. *Mechanics of Composite Materials* (2nd ed.).
- David Jourdan, Mélina Skouras, Etienne Vouga, and Adrien Bousseau. 2022. Computational Design of Self-Actuated Surfaces by Printing Plastic Ribbons on Stretched Fabric. *Computer Graphics Forum* 41, 2 (2022), 493–506.
- Dan Julius, Vladislav Kraevoy, and Alla Sheffer. 2005. D-Charts: Quasi-Developable Mesh Segmentation. *Computer Graphics Forum* 24, 3 (2005), 581–590.
- Michael Kazhdan, Allison Klein, Ketan Dalal, and Hugues Hoppe. 2007. Unconstrained isosurface extraction on arbitrary octrees. In *Proceedings of the Fifth Eurographics Symposium on Geometry Processing*. 125–133.
- M. Kilian, S. Flöry, Z. Chen, N. J. Mitra, A. Sheffer, and H. Pottmann. 2008. Curved Folding. *ACM Transactions on Graphics* 27, 3 (2008), 1–9.
- Diederik P. Kingma and Jimmy Lei Ba. 2015. Adam: A method for stochastic optimization. *3rd International Conference on Learning Representations, ICLR 2015-Conference Track Proceedings* (2015).
- Felix Knöppel, Keenan Crane, Ulrich Pinkall, and Peter Schröder. 2015. Stripe patterns on surfaces. *ACM Transactions on Graphics* 34, 4 (2015), 1–11.
- R.S. Korsmik, A.A. Rodionov, V.A. Korshunov, D.A. Ponomarev, I.S. Prosychev, and V.V. Promakhov. 2020. Topological optimization and manufacturing of vessel propeller via LMD-method. *Materials Today: Proceedings* 30 (2020), 538–544.
- Christian Krogh, Sebastian M. Hermansen, Erik Lund, Jørgen Kepler, and Johnny Jakobsen. 2023. A matter of course: Generating optimal manufacturing instructions from a structural layout plan of a wind turbine blade. *Composites Part A: Applied Science and Manufacturing* 172 (2023), 107599.
- Matthijs Langelaar. 2016. Topology optimization of 3D self-supporting structures for additive manufacturing. *Additive Manufacturing* 12 (2016), 60–70.
- Yifei Li, Tao Du, Sangeetha Grama Srinivasan, Kui Wu, Bo Zhu, Eftychios Sifakis, and Wojciech Matusik. 2022. Fluidic Topology Optimization with an Anisotropic Mixture Model. *ACM Transactions on Graphics* 41, 6 (2022), 1–14.
- Ligang Liu, Lei Zhang, Yin Xu, Craig Gotsman, and Steven J. Gortler. 2008. A local/global approach to mesh parameterization. In *Proceedings of the Symposium on Geometry Processing (SGP '08)*. 1495–1504.
- Tao Liu, Tianyu Zhang, Yongxue Chen, Yuming Huang, and Charlie C. L. Wang. 2024. Neural Slicer for Multi-Axis 3D Printing. *ACM Transactions on Graphics* 43, 4 (2024), 1–15.
- Tao Liu, Tianyu Zhang, Yongxue Chen, Weiming Wang, Yu Jiang, Yuming Huang, and Charlie C. L. Wang. 2025. Neural Co-Optimization of Structural Topology, Manufacturable Layers, and Path Orientations for Fiber-Reinforced Composites. *ACM Transactions on Graphics* 44, 4 (2025), 1–17.
- Yang Liu, Helmut Pottmann, Johannes Wallner, Yong-Liang Yang, and Wenping Wang. 2006. Geometric modeling with conical meshes and developable surfaces. *ACM Transactions on Graphics* 25, 3 (2006), 1–9.
- Yi-Rong Luo, Robert Hewson, and Matthew Santer. 2023. Spatially optimised fibre-reinforced composites with isosurface-controlled additive manufacturing constraints. *Structural and Multidisciplinary Optimization* 66, 6 (2023), 130.
- Juan Montes Maestre, Yinwei Du, Ronan Hinchet, Stelian Coros, and Bernhard Thomaszewski. 2023. Differentiable Stripe Patterns for Inverse Design of Structured Surfaces. *ACM Transactions on Graphics* 42, 4 (2023), 1–14.
- Antonio Mancuso, Antonio Saporito, and Davide Tumino. 2022. Designing the internal reinforcements of a sailing boat using a topology optimization approach. *Applied Ocean Research* 129 (2022), 103384.
- Takao Miki and Takayuki Yamada. 2021. Topology optimization considering the distortion in additive manufacturing. *Finite Elements in Analysis and Design* 193 (2021), 103558.
- Rahul Mitra, Liane Makatura, Emily Whiting, and Edward Chien. 2023. Helix-Free Stripes for Knit Graph Design. *Proceedings - SIGGRAPH 2023 Conference Papers* (2023), 1–9.
- Ioanna Mitropoulou et al. 2020. Print Paths Key-framing: Design for Non-Planar Layered Robotic FDM Printing. In *Proceedings of the ACM Symposium on Computational Fabrication*.
- Juan Sebastian Montes Maestre, Yinwei Du, Ronan Hinchet, Stelian Coros, and Bernhard Thomaszewski. 2023. Differentiable Stripe Patterns for Inverse Design of Structured Surfaces. *ACM Transactions on Graphics* 42, 4 (2023), 1–14.
- Alessandro Muntoni, Marco Livesu, Riccardo Scateni, Alla Sheffer, and Daniele Panozzo. 2018. Axis-Aligned Height-Field Block Decomposition of 3D Shapes. *ACM Transactions on Graphics* 37, 5 (2018), 1–15.
- Yuta Noma, Nobuyuki Umetani, and Yoshihiro Kawahara. 2022. Fast Editing of Singularities in Field-Aligned Stripe Patterns. *Proceedings - SIGGRAPH Asia 2022 Conference Papers* (2022), 1–8.
- Annika Oehri, Aviv Segall, Jing Ren, and Olga Sorkine-Hornung. 2024. Chebyshev Parameterization for Woven Fabric Modeling. *ACM Transactions on Graphics* 43, 6 (2024), 1–15. doi:10.1145/3687928
- David Palmer, Albert Chern, and Justin Solomon. 2024. Lifting Directional Fields to Minimal Sections. *ACM Transactions on Graphics* 43, 4 (2024), 1–20.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: an imperative style, high-performance deep learning library. *Advances in neural information processing systems* 32 (2019), 1–12.

- Tianyuan Qi, Junpeng Zhao, and Chunjie Wang. 2025. An efficient GPU solver for 3D topology optimization of continuous fiber-reinforced composite structures. *Computer Methods in Applied Mechanics and Engineering* 435 (2025), 117675.
- Zheng Qiu, Quhao Li, Yunfeng Luo, and Shutian Liu. 2022. Concurrent topology and fiber orientation optimization method for fiber-reinforced composites based on composite additive manufacturing. *Computer Methods in Applied Mechanics and Engineering* 395 (2022), 114962.
- Nicolas Ray, Wan Chiu Li, Bruno Lévy, Alla Sheffer, and Pierre Alliez. 2006. Periodic global parameterization. *ACM Transactions on Graphics* 25, 4 (2006), 1460 – 1485.
- Camille Schreck, Damien Rohmer, Stefanie Hahmann, Marie-Paule Cani, Shuo Jin, Charlie C. L. Wang, and Jean-François Bloch. 2016. Nonsmooth Developable Geometry for Interactively Animating Paper Crumpling. *ACM Transactions on Graphics* 35, 1 (2016), 1–18.
- Nicholas Sharp, Yousuf Soliman, and Keenan Crane. 2019. Navigating intrinsic triangulations. *ACM Transactions on Graphics* 38, 4 (2019), 1–16.
- Ole Sigmund. 2001. A 99 line topology optimization code written in matlab. *Structural and Multidisciplinary Optimization* 21, 2 (2001), 120–127.
- Mélina Skouras, Bernhard Thomaszewski, Peter Kaufmann, Akash Garg, Bernd Bickel, Eitan Grinspun, and Markus Gross. 2014. Designing inflatable structures. *ACM Transactions on Graphics* 33, 4 (2014), 1–10.
- Hollis Smith and Julián Norato. 2023. Simultaneous material and topology optimization of composite laminates. *Computer Methods in Applied Mechanics and Engineering* 404 (2023), 115781.
- Justin Solomon, Etienne Vouga, Max Wardetzky, and Eitan Grinspun. 2012. Flexible Developable Surfaces. *Computer Graphics Forum* 31, 5 (2012), 1567–1576.
- Ji-Guang Sun. 1990. Multiple eigenvalue sensitivity analysis. *Linear Algebra Appl.* 137–138 (1990), 183–211.
- Chengcheng Tang, Pengbo Bo, Johannes Wallner, and Helmut Pottmann. 2016. Interactive Design of Developable Surfaces. *ACM Transactions on Graphics* 35, 2 (2016), 1–12.
- Thibault Tricard, Vincent Tavernier, Cédric Zanni, Jonàs Martínez, Pierre-Alexandre Hugron, Fabrice Neyret, and Sylvain Lefebvre. 2020. Freely orientable microstructures for designing deformable 3D prints. *ACM Transactions on Graphics* 39, 6 (2020), 1–16.
- Floor Verhoeven, Amir Vaxman, Tim Hoffmann, and Olga Sorkine-Hornung. 2022. Dev2PQ: Planar Quadrilateral Strip Remeshing of Developable Surfaces. *ACM Transactions on Graphics* 41, 3 (2022), 1–18.
- Charlie C. L. Wang. 2008. Towards flattenable mesh surfaces. *Computer-Aided Design* 40, 1 (2008), 109–122.
- Charlie C. L. Wang and Kai Tang. 2004. Achieving developability of a polygonal surface by minimum deformation: a study of global and local optimization approaches. *Visual Computer* 20, 8–9 (2004), 521–539.
- Sifan Wang, Yujun Teng, and Paris Perdikaris. 2021. Understanding and Mitigating Gradient Flow Pathologies in Physics-Informed Neural Networks. *SIAM Journal on Scientific Computing* 43, 5 (2021), A3055–A3081.
- Weiming Wang, Dongwei Feng, Li Yang, Shan Li, and Charlie C.L. Wang. 2023a. Topology optimization of self-supporting lattice structure. *Additive Manufacturing* 67 (2023), 103507.
- Weiming Wang, Kai Wu, Fred van Keulen, and Jun Wu. 2024. Regularization in space–time topology optimization for additive manufacturing. *Computer Methods in Applied Mechanics and Engineering* 431 (2024), 117202.
- Yan Wang, Dan Wang, Yucheng Zhong, David W. Rosen, Shuxin Li, and Stephen W. Tsai. 2023c. Topology optimization of Double-Double (DD) composite laminates considering stress control. *Computer Methods in Applied Mechanics and Engineering* 414 (2023), 116191.
- Ziqi Wang, Florian Kennel-Maushart, Yijiang Huang, Bernhard Thomaszewski, and Stelian Coros. 2023b. A Temporal Coherent Topology Optimization Approach for Assembly Planning of Bespoke Frame Structures. *ACM Transactions on Graphics* 42, 4 (2023), 1–13.
- M.R. Wisnom. 1999. Size effects in the testing of fibre-composite materials. *Composites Science and Technology* 59, 13 (1999), 1937–1957.
- Chenming Wu, Chengkai Dai, Guoxin Fang, Yong-Jin Liu, and Charlie C. L. Wang. 2020. General Support-Effective Decomposition for Multi-Directional 3-D Printing. *IEEE Transactions on Automation Science and Engineering* 17, 2 (2020), 599–610.
- Jun Wu, Weiming Wang, and Xifeng Gao. 2021. Design and Optimization of Conforming Lattice Structures. *IEEE Transactions on Visualization and Computer Graphics* 27, 1 (2021), 43–56.
- Kai Wu, Weiming Wang, Fred van Keulen, and Jun Wu. 2024a. Space–time topology optimization for anisotropic materials in wire and arc additive manufacturing. *International Journal of Mechanical Sciences* 284 (2024), 109712.
- Kang Wu, Zheng-Yu Zhao, Zheng Zhang, Ligang Liu, and Xiao-Ming Fu. 2024b. Piecewise Developable Modeling via Implicit Neural Deformation and Feature-Guided Cutting. *IEEE Transactions on Visualization and Computer Graphics* 30, 9 (2024), 5993–6004.
- Yanan Xu, Zhaoxuan Feng, Yunkai Gao, Chi Wu, Jianguang Fang, Guangyong Sun, Na Qiu, Grant P Steven, and Qing Li. 2024. Topology optimization for additive manufacturing of CFRP structures. *International Journal of Mechanical Sciences* 269 (2024), 108967.
- T. Yamada and Y. Noguchi. 2022. Topology optimization with a closed cavity exclusion constraint for additive manufacturing based on the fictitious physical model approach. *Additive Manufacturing* 52 (2022), 102630.
- Xin Yan, Cong Rao, Lin Lu, Andrei Sharf, Haisen Zhao, and Baoquan Chen. 2020. Strong 3D Printing by TPMS Injection. *IEEE Transactions on Visualization and Computer Graphics* 26, 10 (2020), 3037–3050.
- Jun Ye, Qichen Guo, Hongjia Lu, Pinelopi Kyvelou, Yang Zhao, Leroy Gardner, and Yi Min Xie. 2023. Topology optimisation of self-supporting structures based on the multi-directional additive manufacturing technique. *Virtual and Physical Prototyping* 18, 1 (2023), 2271458.
- Paul Zhang, Josh Vekhter, Edward Chien, David Bommers, Etienne Vouga, and Justin Solomon. 2020. Octahedral Frames for Feature-Aligned Cross Fields. *ACM Transactions on Graphics* 39, 3 (2020).
- Tianyu Zhang, Xiangjia Chen, Guoxin Fang, Yingjun Tian, and Charlie C. L. Wang. 2021. Singularity-Aware Motion Planning for Multi-Axis Additive Manufacturing. *IEEE Robotics and Automation Letters* 6, 4 (2021), 6172–6179.
- Tianyu Zhang, Guoxin Fang, Yuming Huang, Neelotpal Dutta, Sylvain Lefebvre, Zekai Murat Kilic, and Charlie C. L. Wang. 2022. S3-Slicer: A General Slicing Framework for Multi-Axis 3D Printing. *ACM Transactions on Graphics* 41, 6 (2022), 1–15.
- Tianyu Zhang, Tao Liu, Neelotpal Dutta, Yongxue Chen, Renbo Su, Zhizhou Zhang, Weiming Wang, and Charlie C.L. Wang. 2025. Toolpath generation for high density spatial fiber printing guided by principal stresses. *Composites Part B: Engineering* 295 (2025), 112154.
- Xiaoting Zhang, Guoxin Fang, Melina Skouras, Gwenda Gieseler, Charlie C. L. Wang, and Emily Whiting. 2019. Computational design of fabric formwork. *ACM Transactions on Graphics* 38, 4 (2019), 1–13.
- Zheng-Yu Zhao, Mo Li, Zheng Zhang, Qing Fang, Ligang Liu, and Xiao-Ming Fu. 2023. Evolutionary Piecewise Developable Approximations. *ACM Transactions on Graphics* 42, 4 (2023), 1–14.
- Fanchao Zhong, Haisen Zhao, Haochen Li, Xin Yan, Jikai Liu, Baoquan Chen, and Lin Lu. 2023. VASCO: Volume and Surface Co-Decomposition for Hybrid Manufacturing. *ACM Transactions on Graphics* 42, 6 (2023), 1–17.
- Bo Zhu, Mélina Skouras, Desai Chen, and Wojciech Matusik. 2017. Two-Scale Topology Optimization with Microstructures. *ACM Transactions on Graphics* 36, 4 (2017), 1–16.