Field Optimization for Scalable and Distortion-Aware Process Planning in Hybrid Additive-Subtractive Manufacturing

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

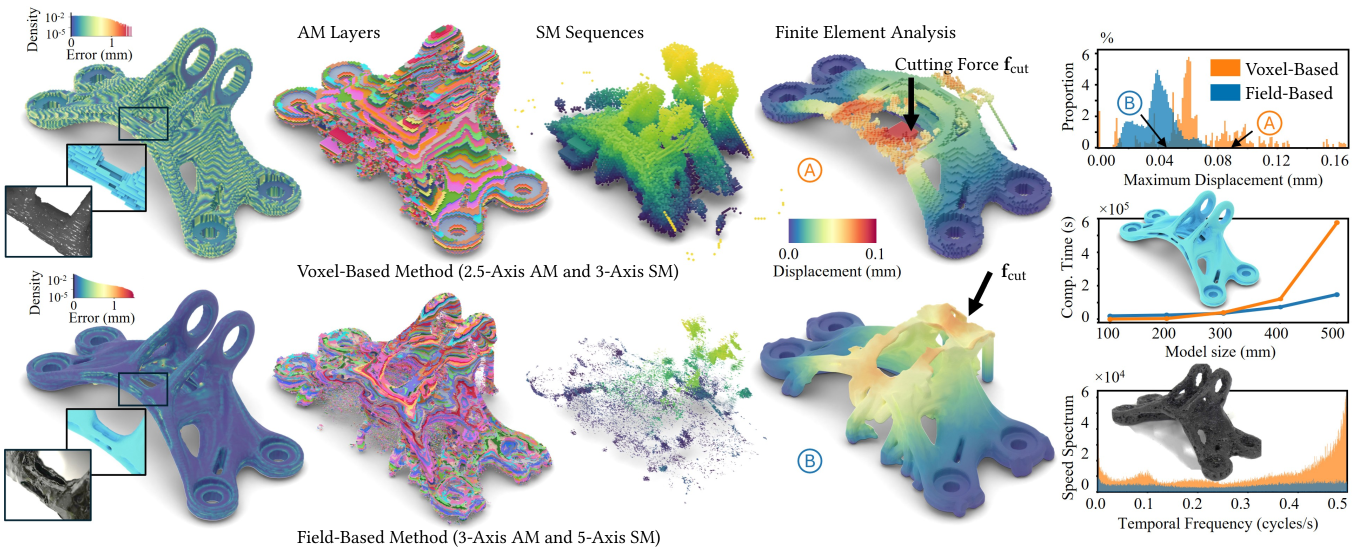

Fig. 1. By formulating hybrid manufacturing (HM) process planning as an optimization problem over continuous fields, our method reduces the estimated fabrication time by 68.3% on the GE-Bracket model and controls distortion in intermediate structures, as indicated by the displacements and the histograms obtained from finite element analysis. Owing to the continuous field representation, our method gives much smaller geometric approximation errors (see the color map in the left) and preserves small geometric features (see the zoom-views) while maintaining a computation time comparable to that of the voxel-based method [Chen et al. 2025a] in low-resolution. The frequency spectra of local tool-motion velocities further demonstrate the improved continuity of the operations generated by our approach. Moreover, the proposed formulation naturally extends to advanced HM platforms with 3-axis additive and 5-axis subtractive-manufacturing capabilities.

This paper presents a field-based optimization method for process planning in hybrid manufacturing, where the goal is to determine a feasible and efficient sequence of additive and subtractive operations for fabricating a target shape. Existing approaches rely on deterministic or discretized formulations, which lead to unoptimized fabrication time and make planning sensitive to voxel resolution. They also do not explicitly account for distortion in intermediate structures formed during manufacturing. To address these limitations, we represent manufacturing states using continuous fields, which scale naturally to models with large dimensions and fine geometric features while enabling smoother fabrication sequences. Based on this representation, we formulate process planning as a numerical optimization problem under multiple objectives, including final shape conformance, intermediate structural strength, manufacturing time, self-supporting and collision-free. Experimental results show that our method can generate distortion-aware process plans on a variety of models while substantially reducing fabrication time through up to a 72.6% reduction in the volume of extra support.

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

Additional Key Words and Phrases: Field Optimization, Scalable, Distortion-Aware, Process Planning, Hybrid Manufacturing

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

Authors' Contact Information: Yongxue Chen, The University of Manchester, Manchester, United Kingdom, yongxue.chen@manchester.ac.uk; Tao Liu, The University of Manchester, Manchester, United Kingdom, tao.liu@manchester.ac.uk; Aoran Lyu, The University of Manchester, Manchester, United Kingdom, aoran.lyu@postgrad.manchester.ac.uk; Yu Jiang, The University of Manchester, Manchester, United Kingdom, yu.jiang@manchester.ac.uk; Neelotpal Dutta, The University of Manchester, Manchester, United Kingdom, neelotpal.dutta@manchester.ac.uk; Charlie C.L. Wang, The University of Manchester, Manchester, United Kingdom, charlie.wang@manchester.ac.uk.

1 Introduction

Hybrid manufacturing (HM), which flexibly switches between additive manufacturing (AM) operations (e.g., filament-deposition-based 3D printing) and subtractive manufacturing (SM) operations (e.g., CNC milling), enables a broader design space for the generative design of smart structures that satisfy a variety of objectives, such as mechanical stiffness [Sigmund 2001], energy absorption [Snapp et al. 2024], cell proliferation [Wang et al. 2025], and sport biomechanics [Han et al. 2024]. In particular, HM makes it possible to fabricate models with highly complex geometries without compromising manufacturing constraints, a capability that cannot be achieved by AM alone – due to the need for additional support structures [Dumas et al. 2014] or by SM alone – due to the limit of tool accessibility [Mahdavi-Amiri et al. 2020; Zhong et al. 2025a]. Chen et al. [2025a] showed that any model with a connected shape can be fabricated if an unlimited number of switches between AM and SM operations and arbitrarily small SM tools are allowed. However, their method is a deterministic algorithm based on voxel discretization, leaving substantial room for improving the generated HM sequences.

1.1 Challenges

When generating operation sequences in HM process planning, existing methods face several, if not all, of the following challenges.

- **Loss of Shape-Conformance:** The sequence of AM/SM operations and the switches between them are typically generated by deterministic algorithms, which can easily drive the search into a 'dead end' and thus fail to completely fabricate the target shape. Moreover, voxel-based methods such as [Chen et al. 2025a] may not capture small geometric features faithfully in their discrete representation, which can further lead to the problem of shape fidelity.
- **Poor Scalability:** While most existing methods do not guarantee completeness, the voxel-based method of [Chen et al. 2025a] theoretically proves it. However, its voxel-based nature leads to poor scalability; see the computing time comparison in Fig. 1 as the model size increases. When the model size increases while the voxel width is kept fixed, the computation time of the voxel-based method grows to more than $\times 3.93$ that of our method.[1]
- **Lack of Optimization:** For a given target shape, there may exist many feasible AM/SM routines for fabrication. However, determining a routine that optimizes total manufacturing time remains an open problem. For example the GE-Bracket model shown in Fig. 1, the method of [Chen et al. 2025a] produces a sequence whose estimated fabrication time is 139,770 seconds, while the optimized sequence generated by our method results in a much shorter fabrication time of 44,352 seconds.
- **Uncontrolled Distortion:** Intermediate structures may deform under gravity and the cutting forces introduced during SM operations. As shown in Fig. 1, a sequence determined without considering these effects can lead to large distortion in intermediate structures. This is evident from the comparison of displacements computed by Finite Element Analysis (FEA) for the intermediate structures.
- **High Discontinuity:** Continuity in operations is desirable, since discontinuous toolpaths and cutter poses can lead to difficulties in control and degraded fabrication quality. Discrete-representation based methods such as [Chen et al. 2025a] usually perform poorly in this regard compared with the continuous field-based computation proposed in this paper (see the comparison of the frequency spectra of local motion velocities for different approaches shown in Fig. 1).

All these challenges will be effectively addressed by the field-based optimization method introduced in this paper.

1.2 Our Method

We parameterize the HM process-planning problem using three groups of fields, each defined implicitly over a position $\mathbf{x} \in \mathbb{R}^3$ within the computational domain Ω, to encode the information required by different aspects of the planning process. Specifically, these fields are defined for:

(1) **Manufacturing Sequence**, represented by two scalar fields $t_{\text{AM}}(\mathbf{x}) \in [0,1]$ and $t_{\text{SM}}(\mathbf{x}) \in [0,1]$, which indicate the relative time of the AM and SM operations at a query point $\mathbf{x}$, respectively;
(2) **Operation Occupancy**, represented by two scalar fields $p_{\text{AM}}(\mathbf{x}) \in [0,1]$ and $p_{\text{SM}}(\mathbf{x}) \in [0,1]$, which indicate the probability of conducting an AM or SM operation at point $\mathbf{x}$, regardless of when the operation is executed;
(3) **Cutting Configuration**, represented by two vector fields $\mathbf{v}_{\text{tl}}(\mathbf{x}) \in \mathbb{S}^2$ and $\mathbf{v}_{\text{bc}}(\mathbf{x}) \in \mathbb{S}^2$, where $\mathbf{v}_{\text{tl}}(\mathbf{x})$ defines the tool orientation in the world coordinate system, and $\mathbf{v}_{\text{bc}}(\mathbf{x})$ specifies the relative position of the contact point on the spherical surface of a ball-end cutter.

These fields are encoded by five neural networks (NNs), while structural relationships among them (e.g., $t_{\text{AM}}(\mathbf{x}) < t_{\text{SM}}(\mathbf{x})$) are enforced within the formulation. Manufacturing objectives and constraints are formulated as corresponding loss terms. The coefficients of NNs are then determined through a self-learning computational pipeline using backpropagation.

The technical contributions of this work are summarized as follows:

- We introduce a continuous field-based representation for process planning in hybrid additive-subtractive manufacturing, which scales to models with large dimensions and fine geometric features, while also enabling smoother manufacturing sequences.
- We propose an optimization-based formulation for computing manufacturing sequences under multiple objectives, including shape conformance, fabrication time, intermediate structural strength, self-supporting and collision-free.
- We develop a strategy to explicitly account for the mechanical strength of intermediate structures formed during manufacturing, allowing the generated plans to effectively reduce structural distortion caused by machining forces and gravity, while also avoiding the formation of 'floating' regions.

To the best of our knowledge, this is the first field-optimization based HM process planning approach that jointly reduces the total fabrication time and the distortion in intermediate structures. The source code of our implementation in Python can be accessed at: https://github.com/Yongxue-Chen/hybManuFieldOpt.

[1] The sampling density used in our method is kept unchanged to ensure a fair comparison and to maintain sufficient detail for capturing small features.

2 Related Work

2.1 Process Planning for Hybrid Manufacturing

Existing HM process planning methods can be broadly classified into surface-quality-driven and structure-driven approaches. Early HM studies mainly followed surface-quality-driven strategies, where AM is used to produce a near-net shape and SM is subsequently applied to improve dimensional accuracy and surface quality [Flynn et al. 2016]. Many existing methods rely on manufacturing features to plan operations [He et al. 2023; Zheng et al. 2020], which works well for regular geometries but is less reliable for freeform models (e.g., those generated by topology optimization [Aage et al. 2017]) because of the challenging feature recognition [Hu et al. 2014; Zhang et al. 2024]. To avoid explicit feature extraction, decomposition-based strategies have also been adopted. Chen et al. [2018] optimized the build-up height of each AM/SM alternation to reduce process switches, which was later extended to cover quasi-rotational parts [Chen et al. 2020]. More recently, VASCO [Zhong et al. 2023] introduced a volume-and-surface co-decomposition method that jointly considers AM's self-supporting and SM's accessibility requirements. However, these methods cannot guarantee the shape completeness of resultant models.

Structure-driven methods aim to address the issues of shape completeness, i.e., whether complex geometry can be completely fabricated by alternating AM and SM operations. AM and SM primitives were introduced in [Behandish et al. 2018] for automated HM planning and later extended to arbitrary part and tool shapes by considering self-supporting requirements and milling accessibility [Harabin and Behandish 2022]. Chen et al. [2025a] recently proposed a voxel-based planning method for HM that theoretically guarantees the completeness for arbitrary shapes. Besides, HM has also been explored for remanufacturing, where the geometry of an existing part is partially reused to form a new shape [Zhong et al. 2025b]. Nevertheless, existing approaches do not explicitly optimize fabrication efficiency. Voxel-based methods are also less scalable, becoming time-consuming when high resolutions are used for large models or fine geometric details. In addition, distortions in intermediate structures are not considered in existing HM planning methods.

2.2 Manufacturability-Aware Computational Design

Computational design has been explored to support 3D model creation [Schulz et al. 2014]. Among computational design methods, topology optimization (TO) has been widely used to generate lightweight, high-performance structures [Sigmund 2001; Zhu et al. 2017]. However, the optimized geometries often contain features that are difficult to fabricate using AM alone – such as large overhangs, or SM alone – such as deep inaccessible cavities. To address this issue, manufacturing constraints have been incorporated into topology optimization, including self-supporting constraints for AM [Langelaar 2016; Wang et al. 2023] and accessibility constraints for SM [Langelaar 2019; Lee et al. 2022]. Topology optimization has also been extended to HM settings [Liu and To 2017], including formulations that consider removable supports and tool accessibility [Mirzendehdel et al. 2022; Xu et al. 2024]. While these methods improve manufacturability by constraining the design space, they may also exclude structures with higher performance. In contrast, HM plan with the guarantee of shape completeness can fabricate structures with arbitrary connected shapes, allowing manufacturability requirements to be relaxed during design.

Beyond static constraints of manufacturability, the influence of self-weight of intermediate structures during fabrication has been considered in the formulation of topology optimization (e.g., [Allaire et al. 2017; Lan and Panetta 2022; Wang et al. 2020; Yang et al. 2025]). Inspired by these works, our formulation explicitly minimizes the distortion of intermediate structures at different timestamps during hybrid manufacturing, which significantly enhances the robustness of physical fabrication. In addition to ensuring the feasibility of fabricating the target shape, our method optimizes the HM process to reduce total fabrication time while considering other manufacturing constraints.

2.3 Neural Fields for Design and Manufacturing

In recent years, continuous fields have been increasingly employed in manufacturing process planning, as they provide compact and differentiable representations for encoding fabrication-related information, such as layer order, deposition direction, tool orientation, and material distribution. Scalar fields have been used in multi-axis AM to generate curved layers for support-free volume printing [Dai et al. 2018]. Field-based formulations have also been developed to control curved layers' orientations via deformation for improved surface quality [Etienne et al. 2019], to control filament alignment for anisotropic strength [Fang et al. 2020], and to jointly satisfy multiple manufacturing objectives [Zhang et al. 2022]. Implicit neural fields have been used to optimize curved layers for multi-axis AM [Liu et al. 2024], benefiting from the efficient optimization and automatic differentiation capabilities of neural-network-based computational pipelines. More recently, implicit neural fields have been further applied to collision-free multi-axis printing [Dutta et al. 2026; Qu et al. 2025].

Beyond manufacturing process planning, neural implicit fields have also been employed to generate optimized density fields for structural optimization [Chandrasekhar and Suresh 2021; Doosti et al. 2021; Han et al. 2024; Liu et al. 2026; Zehnder et al. 2021]. These methods similarly benefit from the automatic differentiation provided by neural network training pipelines. Recent developments have explored the co-optimization of design and manufacturing objectives in end-to-end computational frameworks to achieve manufacturable designs with better mechanical performance, such as fiber-reinforced composites [Liu et al. 2025]. In contrast, our method leverages the flexibility of HM to fabricate a broader range of shapes, thereby providing a larger space for optimized structural design.

3 Representation and Framework Overview

3.1 Preliminaries

Without loss of generality, for a given target model M, the computational domain is defined as a rectangular region slightly larger than the bounding box of M, denoted by $\mathcal{D}_M \subset \mathbb{R}^3$ (see Fig. 2(a)). The implicit fields used to determine operation sequences are defined and evaluated over $\mathcal{D}_M$.

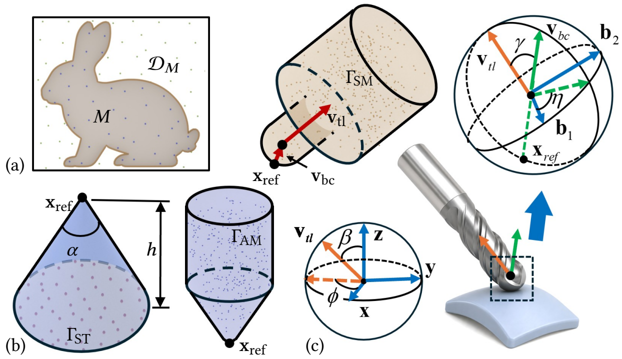

Fig. 2. Preliminary illustration about (a) the computation domain, (b) the self-support evaluation region Γ_{ST} and the printer head's solid region Γ_{AM}, and (c) the cutting configuration and the cutter's solid region Γ_{SM}.

The AM operation is defined as layer-based material deposition while allowing 3-axis motion of the printer head. Material is deposited at the tip of the printer head, denoted by $\mathbf{x}_{\text{ref}}$, as illustrated in Fig. 2(b). In contrast, the SM operation is defined as material removal by positioning the cutter at $\mathbf{x}_{\text{ref}}$ according to a cutting configuration specified by two vectors, $\mathbf{v}_{\text{bc}}$ and $\mathbf{v}_{\text{tl}}$ (see Fig. 2(c)). Our computational framework allows 5-axis motion for SM operations.

The sequence of AM and SM operations is represented by two scalar fields, $t_{\text{AM}}(\mathbf{x})$ and $t_{\text{SM}}(\mathbf{x})$, which indicate the order of material processing – i.e., deposition by AM and removal by SM – at different locations $\mathbf{x}$ in the computational domain $\mathcal{D}_M$. These values can be interpreted as virtual time: they do not represent physical time, but rather the relative ordering of operations. Moreover, whether an operation is actually applied is determined by two occupancy fields. Specifically, AM is applied when $p_{\text{AM}}(\mathbf{x}) > 0.5$, and SM is applied when both $p_{\text{AM}}(\mathbf{x}) > 0.5$ and $p_{\text{SM}}(\mathbf{x}) > 0.5$.

Several regions of interest (RoIs) are defined for AM and SM. First, a self-support evaluation region, Γ_{ST}, is defined as the conical volume with height h below the reference point $\mathbf{x}_{\text{ref}}$ for AM. Its apex angle α is determined by the self-supporting angle, which depends on the material (see the left of Fig. 2(b)). Second, the RoI for collision detection of the printer head, Γ_{AM}, is defined as the solid region of the printer head above $\mathbf{x}_{\text{ref}}$ for AM (see the right of Fig. 2(b)). Similarly, the RoI for collision detection of the cutter, Γ_{SM}, is defined as the solid region of the cutter located at $\mathbf{x}_{\text{ref}}$ for SM. Note that the pose of Γ_{SM} in Euclidean space is determined by the cutting configuration, i.e., $\mathbf{v}_{\text{tl}}$ and $\mathbf{v}_{\text{bc}}$ as explained in Fig. 2(c).

3.2 Field Representation

We next introduce the NN-based field representation for HM process planning. To determine whether a point $\mathbf{x}$ lies inside the target model M, we represent M using a *signed distance field* (SDF) encoded by a neural network, denoted by $\mathcal{N}_M(\mathbf{x})$, which can be pre-trained. Specifically, $\mathbf{x} \in M$ if $\mathcal{N}_M(\mathbf{x}) \leq 0$, and $\mathbf{x} \notin M$ otherwise.

3.2.1 Manufacturing Sequence Fields. To establish the global chronological order of operations, we define two continuous scalar fields, $t_{\text{AM}}(\mathbf{x})$ and $t_{\text{SM}}(\mathbf{x})$, which assign relative activation timestamps to each point within the domain $\mathcal{D}_M$ and thereby determine the execution order of AM and SM operations. Specifically, these fields are represented by two implicit neural networks, $\mathcal{N}_{t,\text{AM}}(\mathbf{x}) \in [0,1]$ and $\mathcal{N}_{\alpha,\text{SM}}(\mathbf{x}) \in [0,1]$, as

$$t_{\text{AM}}(\mathbf{x}) = \mathcal{N}_{t,\text{AM}}(\mathbf{x}), \tag{1}$$

$$t_{\text{SM}}(\mathbf{x}) = t_{\text{AM}}(\mathbf{x}) + \mathcal{N}_{\alpha,\text{SM}}(\mathbf{x}) \cdot (1 - t_{\text{AM}}(\mathbf{x})). \tag{2}$$

Note that $\mathcal{N}_{\alpha,\text{SM}}(\mathbf{x})$ encodes the relative temporal offset of the SM operation w.r.t. the AM operation, rather than directly representing the SM timestamp itself. This design eliminates invalid regions of the search space and inherently enforces the physical causality that AM must precede SM at every point, i.e., $t_{\text{SM}}(\mathbf{x}) > t_{\text{AM}}(\mathbf{x})$.

In the current implementation, we define only one field for AM and one field for SM to maintain computational efficiency, although using multiple fields for each operation type could provide greater optimization flexibility. Based on the theoretical analysis of HM process planning in Sec. 3.2 of [Chen et al. 2025a], applying either an AM operation or an AM operation followed by an SM operation at each local region is sufficient to realize arbitrary geometries.

3.2.2 Operation Occupancy Fields. The decision to execute an operation at a location $\mathbf{x}$ is determined by the operation-occupancy fields p_{AM} and p_{SM}. An AM operation is performed at point $\mathbf{x}$ if $p_{\text{AM}}(\mathbf{x}) > 0.5$, and omitted otherwise. An SM operation is performed at $\mathbf{x}$ if both $p_{\text{AM}}(\mathbf{x}) > 0.5$ and $p_{\text{SM}}(\mathbf{x}) > 0.5$. Both fields are represented with the help of implicit NNs. Specifically, these probability fields are computed from logit fields through the Sigmoid function $\sigma(\cdot)$, where the fields are parameterized by the neural networks $\mathcal{N}_{l,\text{AM}}$ and $\mathcal{N}_{l,\text{SM}}$ as:

$$p_{\text{AM}}(\mathbf{x}) = \max\left(1 - \sigma(\mathcal{N}_M(\mathbf{x})), \sigma(\mathcal{N}_{l,\text{AM}}(\mathbf{x}))\right), \tag{3}$$

$$p_{\text{SM}}(\mathbf{x}) = \min\left(\sigma(\mathcal{N}_M(\mathbf{x})), \sigma(\mathcal{N}_{l,\text{SM}}(\mathbf{x}))\right). \tag{4}$$

Here, the soft gate $\sigma(\mathcal{N}_M(\mathbf{x}))$ ensures that $p_{\text{AM}}(\mathbf{x}) \approx 1$ and $p_{\text{SM}}(\mathbf{x}) \approx 0$ for all points $\mathbf{x} \in M$. This is essential for ensuring that the final part remains solid and conforms strictly to the intended geometry after the full operational sequence. In this paper, differentiable approximations of the maximum / minimum evaluations are obtained using the integral log-sum-exp formulation and its sign-reversed counterpart, respectively.

3.2.3 Cutter Configuration Field. The cutter configuration at a given point $\mathbf{x}$, treated as the contact point for ball-end milling, is determined by the tool orientation and the relative position of the sphere center with respect to the contact point (see Fig. 2(c)). This information is parameterized by a neural network with four-dimensional output, denoted by $[\beta, \phi, \gamma, \eta] = \mathcal{N}_{\text{CC}}(\mathbf{x})$. Specifically, (β, ϕ) defines the tool orientation on the Gaussian sphere within a feasible region, where $\beta \in [0, \beta_{\max}]$ and $\phi \in [0, 2\pi]$. The corresponding tool-orientation vector is defined as

$$\mathbf{v}_{\text{tl}} = (\sin\beta\cos\phi,\ \sin\beta\sin\phi,\ \cos\beta). \tag{5}$$

Based on $\mathbf{v}_{\text{tl}}$, we define a local orthonormal frame $(\mathbf{b}_1, \mathbf{b}_2)$ centered at the sphere center, with both $\mathbf{b}_1 \perp \mathbf{v}_{\text{tl}}$ and $\mathbf{b}_2 \perp \mathbf{v}_{\text{tl}}$. The ball-center shifting vector is then defined as

$$\mathbf{v}_{\text{bc}} = \cos(\gamma)\,\mathbf{v}_{\text{tl}} + \sin(\gamma)\,(\cos(\eta)\mathbf{b}_1 + \sin(\eta)\mathbf{b}_2), \tag{6}$$

where $\gamma \in [0, \gamma_{\max}]$ and $\eta \in [0, 2\pi]$.

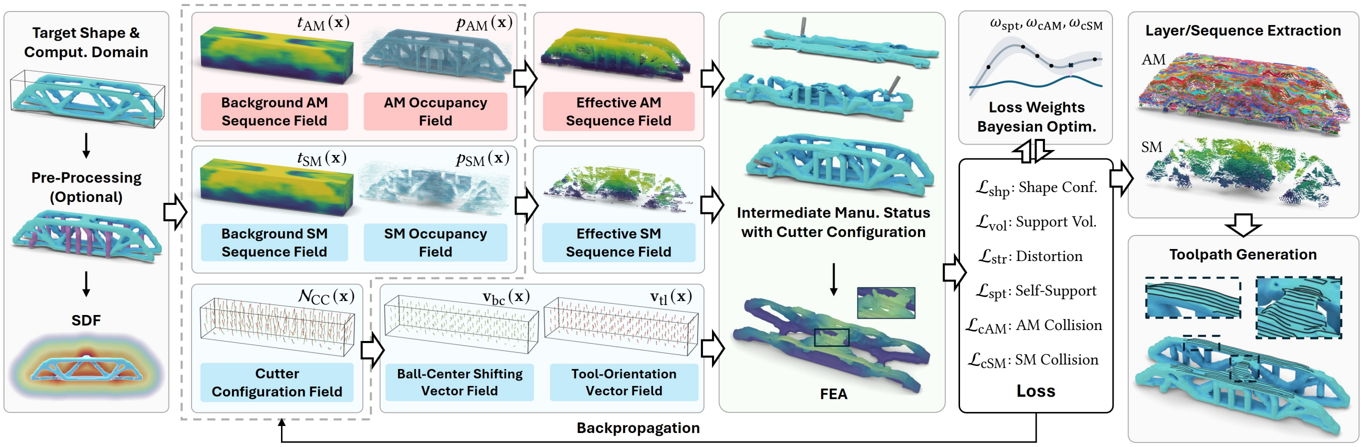

Fig. 3. Illustration of the computational pipeline of our field-optimization-based HM process planning approach. The sequence of AM/SM operations for an input model M is determined by five implicit fields: $(t_{AM}(\mathbf{x}), t_{SM}(\mathbf{x}), p_{AM}(\mathbf{x}), p_{SM}(\mathbf{x}), \mathcal{N}_{CC}(\mathbf{x}))$. The optimization of these fields is driven by minimizing loss terms associated with shape-conformance $\mathcal{L}_{shp}$, support volume $\mathcal{L}_{vol}$, structural distortion $\mathcal{L}_{str}$, self-supporting constraints $\mathcal{L}_{spt}$, and AM/SM collision avoidance, denoted by $\mathcal{L}_{cAM}$ and $\mathcal{L}_{cSM}$, respectively. After the implicit fields are determined, the AM layers and SM sequences are extracted and sorted in the post-processing step, followed by toolpath generation.

In summary, the coefficients of five NNs, $\mathcal{N}_{t,AM}$, $\mathcal{N}_{\alpha,SM}$, $\mathcal{N}_{l,AM}$, $\mathcal{N}_{l,SM}$, and $\mathcal{N}_{CC}$, collectively define the design-variable vector $\boldsymbol{\theta}$. The optimal design-variable vector is obtained by minimizing the total loss function $\mathcal{L}_{total}$ as will be defined in Sec. 4. That is

$$\boldsymbol{\theta}^* = \arg\min_{\boldsymbol{\theta}} \mathcal{L}_{total}\left(\mathcal{N}_{t,AM}, \mathcal{N}_{\alpha,SM}, \mathcal{N}_{l,AM}, \mathcal{N}_{l,SM}, \mathcal{N}_{CC}\right). \quad (7)$$

The HM operation sequence, consisting of AM and SM operations, is then determined by the five NNs with coefficients $\boldsymbol{\theta}^*$.

3.3 Computational Pipeline

The optimization is carried out via backpropagation, which updates $\boldsymbol{\theta}$ to minimize $\mathcal{L}_{total}$. The computational pipeline of our approach is summarized below (see also Fig. 3).

(1) Compute $\mathcal{N}_M(\mathbf{x})$ as SDF of the input model M;
(2) Initialize the five NNs: $\mathcal{N}_{t,AM}$, $\mathcal{N}_{\alpha,SM}$, $\mathcal{N}_{l,AM}$, $\mathcal{N}_{l,SM}$, and $\mathcal{N}_{CC}$;
(3) Apply a regular grid discretization to $\mathcal{D}_M$ for FEA to evaluate the mechanical strength of intermediate structures;
(4) Randomly sample points in the computational domain $\mathcal{D}_M$ and randomly sample timestamps in the time domain;
(5) Evaluate the loss terms (Sec. 4);
(6) Compute gradients $\partial\mathcal{L}_{total}/\partial\boldsymbol{\theta}$ via backpropagation;
(7) Update the network coefficients $\boldsymbol{\theta}$ using the gradients and the optimizer;
(8) Repeat Steps (4)–(7) until the optimization convergences;
(9) Extract the AM and SM layers adaptively according to the prescribed printing layer thickness and cutting depth (Sec. 5.4);
(10) Sort the extracted AM/SM layers according to their timestamps to determine the execution order of operations;
(11) Finally, generate toolpaths layer by layer.

After obtaining the toolpaths for AM/SM operations, they are sampled into nearly uniform waypoints, and inverse kinematics [Zhang et al. 2021] is applied to generate G-code for machine execution.

4 Optimization Formulation

The field-based optimization is driven by the loss terms defined according to multiple objectives, including final shape conformance, fabrication time, intermediate structural distortion, and manufacturability. They are formulated according to the aforementioned fields.

4.1 Shape Conformance

To ensure that no extraneous support structure remains at the end of manufacturing, we require every point $\mathbf{x}$ outside a given model M to satisfy at least one of the following conditions:

- $p_{AM}(\mathbf{x}) \to 0$, indicating that no material is accumulated by the AM operation;
- $p_{SM}(\mathbf{x}) \to 1$, indicating that the deposited material is removed by the SM operation.

Specifically, this can be achieved by enforcing $\sigma(\mathcal{N}_{l,AM}(\mathbf{x})) = 0$ or $\sigma(\mathcal{N}_{l,SM}(\mathbf{x})) = 1$ in the region outside M. We therefore minimize the following loss term:

$$\mathcal{L}_{shp} = \int_{\mathcal{D}_M} \sigma(\mathcal{N}_M(\mathbf{x}))\sigma(\mathcal{N}_{l,AM}(\mathbf{x}))(1-\sigma(\mathcal{N}_{l,SM}(\mathbf{x})))d\mathbf{x}. \quad (8)$$

Here, $\sigma(\mathcal{N}_M(\mathbf{x}))$ suppresses contributions from region inside M, so that the loss only accounts for regions outside M, where $\mathcal{N}_M(\mathbf{x}) > 0$.

4.2 Fabrication Time

The total fabrication time can be indirectly reduced by minimizing the total volume of regions generated by AM operations outside M. By definition, an AM operation is performed at position $\mathbf{x}$ when $p_{AM}(\mathbf{x}) > 0.5$. Therefore, the volume loss associated with AM-generated material outside M can be defined as

$$\mathcal{L}_{vol} = \int_{\mathcal{D}_M} \sigma(\mathcal{N}_M(\mathbf{x}))\max(\sigma(\mathcal{N}_{l,AM}(\mathbf{x})) - 0.5, 0)d\mathbf{x}. \quad (9)$$

Note that a simplified form of $p_{\text{AM}}(\mathbf{x})$ (Eq. (3)) is employed here as $\sigma(\mathcal{N}_M(\mathbf{x}))$ has been used to neglect the region inside M.

4.3 Distortion under Fabrication

The distortion of the intermediate structures during fabrication at different timestamps t is evaluated by applying FEA to discrete elements defined in $\mathcal{D}_M$. This analysis can also help detect physically infeasible 'floating' regions that are disconnected from the build plate, as such regions manifest as intermediate structures with extremely large displacements. In this subsection, we first define the density distribution of an intermediate structure in terms of t. After that, the equilibrium equation is formulated to compute the displacements of every element due to the gravity and the cutting forces. Lastly, the loss function of temporal distortion $\mathcal{L}_{\text{str}}$ is defined.

4.3.1 Spatiotemporal Density Representation. For a query point $\mathbf{x}$, its status as either *solid* or *void* at any time t is determined by the values of $p_{\text{AM}}(\mathbf{x})$, $t_{\text{AM}}(\mathbf{x})$, $p_{\text{SM}}(\mathbf{x})$, and $t_{\text{SM}}(\mathbf{x})$. The point is considered solid only when $p_{\text{AM}}(\mathbf{x}) \geq 0.5$ and $t_{\text{AM}}(\mathbf{x}) \leq t$ (i.e., AM operation has been applied at $\mathbf{x}$ before t), together with either $p_{\text{SM}}(\mathbf{x}) < 0.5$ (i.e., $\mathbf{x}$ is not subject to SM) or $t_{\text{SM}}(\mathbf{x}) > t$ (i.e., SM operation being applied after time t). Therefore, its status value $s(\mathbf{x}, t)$ can be computed as

$$\begin{aligned} s(\mathbf{x}, t) = \sigma(p_{\text{AM}}(\mathbf{x}) - 0.5)\sigma(t - t_{\text{AM}}(\mathbf{x})) \\ \max(\sigma(0.5 - p_{\text{SM}}(\mathbf{x})), \sigma(t_{\text{SM}}(\mathbf{x}) - t)), \end{aligned} \tag{10}$$

where $s(\mathbf{x}) = 1$ indicates *solid* and $s(\mathbf{x}) = 0$ for *void*. The continuous density function is then defined as:

$$\rho(\mathbf{x}, t) = s(\mathbf{x}, t). \tag{11}$$

For the computation of FEA, we discretize $\mathcal{D}_M$ into a uniform Cartesian grid with isotropic cubic elements. For each element e occupying a region Γ_e, the density $\rho_e(t)$ is computed as

$$\rho_e(t) = \frac{1}{V_e} \int_{\mathbf{x} \in \Gamma_e} \rho(\mathbf{x}, t) d\mathbf{x}, \tag{12}$$

with V_e denoting the volume of the element.

4.3.2 Forward Physical Simulation. Given the element density field at a timestamp t, we establish an equilibrium equation to compute the displacement field over all elements:

$$\mathbf{K}(\boldsymbol{\rho})\,\mathbf{u} = \mathbf{f}, \tag{13}$$

where $\mathbf{K}(\boldsymbol{\rho})$ is the system stiffness matrix and $\mathbf{f}$ is the vector of external forces.

The system stiffness matrix $\mathbf{K}(\boldsymbol{\rho})$ is determined by first computing the Young's modulus E_e of each element using the Solid Isotropic Material with Penalization (SIMP) model [Sigmund 2001]:

$$E_e = \rho_e^p E_0 + \epsilon,$$

where E_0 is the base Young's modulus, $p = 3$ is the penalization factor, and ϵ is a small constant introduced for numerical stability. The global stiffness matrix $\mathbf{K}$ is then assembled through finite element discretization according to the spatial distribution of E_e.

During AM operations, the external force vector is defined solely by gravity, i.e.,

$$\mathbf{f} = \mathbf{f}_{\text{grav}}(\boldsymbol{\rho}).$$

During an SM operation at the reference point $\mathbf{x}_{\text{ref}}$, an additional cutting force $\mathbf{f}_{\text{cut}}$ is applied in the direction opposite to $\mathbf{v}_{\text{bc}}$ as $\mathbf{f}_{\text{cut}} = -f_c \mathbf{v}_{\text{bc}}$, where f_c is a material- and machine-dependent coefficient obtained through physical calibration. The total external force vector for an SM operation is therefore

$$\mathbf{f} = \mathbf{f}_{\text{grav}}(\boldsymbol{\rho}) + \mathbf{f}_{\text{cut}}.$$

Without boundary conditions, Eq. (13) is singular and cannot be solved. To make the system well-posed, we impose zero displacement on all the elements attached to the build plate. We solve Eq. (13) by the conjugate gradient method to determine the element displacements $\mathbf{u}_e$, which are used to evaluate the structural-distortion loss introduced below.

Structural stability analysis is generally unnecessary for a standard planar-layer AM process. In our formulation, however, whether an AM or SM operation is performed at a given time t is jointly determined by the AM/SM sequence and the occupancy fields. Consequently, restricting stability analysis to SM operations is nontrivial in practice. Furthermore, in certain extreme geometries, such as an asymmetric V-shaped model, structural stability may still need to be considered even when only planar AM operations are performed.

4.3.3 Timestamp and Loss Formulation. Structural distortion must be controlled throughout the entire HM process. To evaluate the displacement distribution at different stages of fabrication, we sample a set of effective timestamps and store them in two sets, Υ_{AM} and Υ_{SM}, for AM and SM operations, respectively. These timestamp samples are obtained as follows:

- Randomly sample $2n$ points $\{\mathbf{x}_i\}$ from the computational domain $\mathcal{D}_M$, with n inside M and n outside M;
- For each point $\mathbf{x}_i$ with $p_{\text{AM}}(\mathbf{x}_i) > 0.5$, add $t_{\text{AM}}(\mathbf{x}_i)$ to Υ_{AM};
- For each point $\mathbf{x}_i$ with $p_{\text{AM}}(\mathbf{x}_i) > 0.5$ and $p_{\text{SM}}(\mathbf{x}_i) > 0.5$, add $t_{\text{SM}}(\mathbf{x}_i)$ to Υ_{SM}.

For timestamp samples in Υ_{AM} and Υ_{SM}, different external force vectors are used to compute the corresponding element displacements.

The overall distortion of intermediate structures is then minimized using the following loss:

$$\begin{aligned} \mathcal{L}_{\text{str}} = \frac{1}{N_\Upsilon} \sum_{t_k \in \Upsilon_{\text{AM}} \cup \Upsilon_{\text{SM}}} \Bigg(\sum_e \rho_e \left(\max(0, \|\mathbf{u}_e\| - \tau)\right)^2 \\ + \omega_{\text{reg}} \sum_e \rho_e (1 - \rho_e) \Bigg), \end{aligned} \tag{14}$$

where N_Υ denotes the total number of timestamp samples in Υ_{AM} and Υ_{SM}, and both ρ_e and $\mathbf{u}_e$ are functions of t_k. The first term penalizes elements whose displacement magnitude exceeds a threshold τ. The second term is a SIMP-style regularizer, which encourages binary density values and thereby improves the physical interpretability and accuracy of the structural evaluation. Note that this loss not only reduces structural distortion, but also prevents the formation of 'floating' regions (see Fig. 10 for an example), thereby ensuring the connectivity of all intermediate structures. $\tau = 0.05$ and $\omega_{\text{reg}} = 0.001$ are determined by experiments and employed for all examples. In our implementation, we set $n = 15$ to evaluate $\mathcal{L}_{\text{str}}$, and the timestamps are re-sampled at each iteration.

4.4 Manufacturability

Three loss terms are defined according to manufacturing constraints, including self-supporting and collision-free. All are formulated by checking the existence of points inside the relevant RoIs (Fig. 2) at a certain timestamp t. Details are presented below.

4.4.1 Self-Supporting. When material is deposited at a point $\mathbf{x}$, the corresponding AM operation is considered self-supported if there exists at least one solid point within the RoI for self-support, $\Gamma_{\mathrm{ST}}(\mathbf{x})$. This region is defined in Sec. 3.1. For a query point $\mathbf{x}_q \in \Gamma_{\mathrm{ST}}(\mathbf{x})$, its status at timestamp t can be evaluated using the function $s(\mathbf{x}_q, t)$ defined in Eq. (10). For any point $\mathbf{x}$ with $p_{\mathrm{AM}}(\mathbf{x}) > 0.5$, the corresponding AM operation time is given by $t_{\mathrm{AM}}(\mathbf{x})$. Therefore, its self-supporting status can then be evaluated as

$$\max_{\mathbf{x}_q \in \Gamma_{\mathrm{ST}}(\mathbf{x})} s\left(\mathbf{x}_q, t_{\mathrm{AM}}(\mathbf{x})\right), \tag{15}$$

which returns *one* if the AM operation is self-supported and *zero* otherwise. A differentiable approximation of this maximum evaluation can be obtained using the integral log-sum-exp formulation as a soft maximum.

Based on this formulation, the self-supporting loss for the HM process is defined by integrating over the computational domain:

$$\mathcal{L}_{\mathrm{spt}} = \int_{\mathcal{D}_M} \sigma(p_{\mathrm{AM}}(\mathbf{x}) - 0.5) \left(1 - \max_{\mathbf{x}_q \in \Gamma_{\mathrm{ST}}(\mathbf{x})} s\left(\mathbf{x}_q, t_{\mathrm{AM}}(\mathbf{x})\right)\right) d\mathbf{x}, \tag{16}$$

where the sigmoid term suppresses contributions from points where no AM operation is performed. Unlike $\mathcal{L}_{\mathrm{str}}$ that controls the distortion of an already fabricated intermediate structure (i.e., the post-deposition state), this loss $\mathcal{L}_{\mathrm{spt}}$ is used to guide the pre-deposition decision of where to print.

4.4.2 Collision-Free. The collision-free feasibility of both AM and SM operations is evaluated using a strategy similar to the self-supporting status defined in Eq. (15).

For an AM operation performed at point $\mathbf{x}$, whether the operation is collision-free depends on whether any solid material already exists within its collision-checking RoI, $\Gamma_{\mathrm{AM}}(\mathbf{x})$. This status can be evaluated as

$$\max_{\mathbf{x}_q \in \Gamma_{\mathrm{AM}}(\mathbf{x})} s\left(\mathbf{x}_q, t_{\mathrm{AM}}(\mathbf{x})\right), \tag{17}$$

which returns *one* if a collision occurs and *zero* otherwise. Accordingly, the AM collision-free loss can be defined as

$$\mathcal{L}_{\mathrm{cAM}} = \int_{\mathcal{D}_M} \sigma(p_{\mathrm{AM}}(\mathbf{x}) - 0.5) \max_{\mathbf{x}_q \in \Gamma_{\mathrm{AM}}(\mathbf{x})} s\left(\mathbf{x}_q, t_{\mathrm{AM}}(\mathbf{x})\right) d\mathbf{x}, \tag{18}$$

where the sigmoid term ensures that only points with sufficiently high probability of AM operation contribute to the loss.

For an SM operation performed at point $\mathbf{x}$, the pose of the SM tool is determined by the cutter configuration field $\mathcal{N}_{\mathrm{CC}}(\mathbf{x})$ – therefore also its collision-checking RoI as $\Gamma_{\mathrm{SM}}(\mathbf{x}, \mathcal{N}_{\mathrm{CC}}(\mathbf{x}))$. The status of whether the operation is collision-free can then be evaluated as

$$c_{\mathrm{SM}}(\mathbf{x}) = \max_{\mathbf{x}_q \in \Gamma_{\mathrm{SM}}(\mathbf{x}, \mathcal{N}_{\mathrm{CC}})} s\left(\mathbf{x}_q, t_{\mathrm{SM}}(\mathbf{x})\right), \tag{19}$$

which returns *zero* for collision-free. Using this formula, we define the SM collision-free loss as

$$\mathcal{L}_{\mathrm{cSM}} = \int_{\mathcal{D}_M} \sigma(p_{\mathrm{AM}}(\mathbf{x}) - 0.5)\sigma(p_{\mathrm{SM}}(\mathbf{x}) - 0.5)c_{\mathrm{SM}}(\mathbf{x})d\mathbf{x}, \tag{20}$$

where the sigmoid terms help select only high probability points for SM to contribute to the loss evaluation.

For SM operations, the query point $\mathbf{x}_q$ may fall within a region below the build plate. In such cases, we assign $s(\mathbf{x}_q) = 1$, treating the region as occupied, to prevent collisions between the SM tool and the build plate. This strategy can be further extended to account for other environmental obstacles for both AM and SM operations.

4.5 Total Loss

Combining all the loss terms discussed above yields the total loss used to drive the optimization for HM process planning:

$$\begin{aligned} \mathcal{L}_{\mathrm{total}} = & \mathcal{L}_{\mathrm{shp}} + \mathcal{L}_{\mathrm{str}} + 1.5\mathcal{L}_{\mathrm{vol}} \\ & + \omega_{\mathrm{spt}}\mathcal{L}_{\mathrm{spt}} + \omega_{\mathrm{cAM}}\mathcal{L}_{\mathrm{cAM}} + \omega_{\mathrm{cSM}}\mathcal{L}_{\mathrm{cSM}}. \end{aligned} \tag{21}$$

Here, the weight 1.5 is chosen experimentally to balance the soft-objective losses associated with shape conformance, support volume for fabrication time, and distortion of intermediate structures. For the manufacturing constraints, the weights ω_{spt}, ω_{cAM} and ω_{cSM} are dynamically adjusted during optimization to ensure the manufacturability of the resulting sequence of AM / SM operations. Details are provided in Sec. 5.5.

5 Implementation Details

5.1 Network Architecture

We define the neural networks $\mathcal{N}_{t,\mathrm{AM}}$, $\mathcal{N}_{\alpha,\mathrm{SM}}$, $\mathcal{N}_{l,\mathrm{AM}}$, $\mathcal{N}_{l,\mathrm{SM}}$, and $\mathcal{N}_{\mathrm{CC}}$ based on the Instant-NGP framework [Müller et al. 2022]. Instant-NGP is adopted for its high *computational efficiency* and strong *representational capacity*, particularly for capturing high-frequency details and quantities with significant spatial variation. Specifically, each neural field consists of a multi-resolution hash-grid encoder followed by a lightweight MLP decoder. For all fields, the hash-grid encoder uses 16 resolution levels, with 2 feature channels per level.

- The two manufacturing sequence fields, $\mathcal{N}_{t,\mathrm{AM}}$ and $\mathcal{N}_{\alpha,\mathrm{SM}}$, use a hash table size of 2^{16}, with a base resolution of 16 and a maximum resolution of 512.
- The two operation occupancy fields, $\mathcal{N}_{l,\mathrm{AM}}$ and $\mathcal{N}_{l,\mathrm{SM}}$, use a hash table size of 2^{16}, with a base resolution of 16 and a maximum resolution of 2048.
- The cutter configuration field, $\mathcal{N}_{\mathrm{CC}}$, uses a hash table size of 2^{19}, with resolutions ranging from 16 to 4096.

The 32-dimensional feature vector output by each hash-grid encoder is then fed into the corresponding MLP decoder. All five MLPs share the same architecture, consisting of two hidden layers with 64 neurons each and ReLU activation functions. To ensure that the outputs of $\mathcal{N}_{t,\mathrm{AM}}$ and $\mathcal{N}_{\alpha,\mathrm{SM}}$ remain within the physically valid range $[0, 1]$, we apply a sigmoid activation function at the final layer of these two networks. Moreover, the SDF of each input model M is represented by a SIREN-based network $\mathcal{N}_M$ consisting of five layers with 256 neurons per layer (ref. [Sitzmann et al. 2020]).

5.2 Network Initialization

A good initialization is crucial for achieving high-quality optimization results and fast convergence. In this work, we initialize the networks using the output of a voxel-based method with a grid resolution of 100 voxels along its longest dimension, which preserves the scalability of our framework. Specifically, we first employ the method from [Chen et al. 2025a] to generate additive and subtractive sequences for all voxels. These sequences are then normalized to the temporal domain $[0, 1]$, and the corresponding AM and SM timestamps are assigned to the spatial coordinates of each voxel center. The neural networks are then initialized sequentially as follows:

(1) **AM Sequence Network:** We train $\mathcal{N}_{t,\mathrm{AM}}$ using the AM timestamps at voxel centers as ground-truth values, optimized with a mean squared error (MSE) loss.
(2) **SM Offset Network:** Based on the pre-trained $\mathcal{N}_{t,\mathrm{AM}}$, we train $\mathcal{N}_{\alpha,\mathrm{SM}}$ using the SM timestamps as supervision to satisfy Eq. (2), also optimized with an MSE loss.
(3) **Occupancy Networks:** The two operation occupancy networks, $\mathcal{N}_{l,\mathrm{AM}}$ and $\mathcal{N}_{l,\mathrm{SM}}$, are trained at the voxel centers according to the type of AM/SM operations obtained from the voxel-based solution.
(4) **Cutter Configuration Network:** The voxel-based method provides a tool-axis direction for each SM operation. These directions are used as supervised samples to determine the corresponding $\mathbf{v}_{\mathrm{tl}}$ and $\mathbf{v}_{\mathrm{bc}}$, and are then used to train the cutter configuration network $\mathcal{N}_{\mathrm{CC}}$.

The SDF network $\mathcal{N}_M$ is trained using points randomly sampled in the computational domain $\mathcal{D}_M$, with signed distance values evaluated by the PQP library [Larsen et al. 1999].

5.3 Optional Pre-Processing: Support Generation

To accelerate optimization convergence and improve workpiece stability during manufacturing, we introduce an optional pre-processing step that generates supports removable by the final SM operations. In this step, auxiliary support structures are generated for the input model M using a rapidly exploring random tree (RRT)-based method [LaValle 1998]. Specifically, points are sampled from surface regions of M with large overhangs, and support branches are grown downward from these points subject to a prescribed self-supporting angle. These branches are then connected into skeletons, from which convolution surfaces are generated to form the solid support structures, as visualized in purple in Figs. 3 and 7(a). In addition, the generated supports, represented by the sampled skeletons, are required to remain accessible to SM operations after manufacturing is completed, ensuring that they can be removed in the final stage. These supports are combined with the target solid to compute the SDF used as the input model for our optimization pipeline.

5.4 Post-Processing: Layers and Toolpaths

To generate the final sequence for hybrid manufacturing, we first adaptively extract the AM and SM layers according to the prescribed printing layer thickness d_{AM} and cutting depth d_{SM}. Specifically, the optimized NNs provide the manufacturing sequence fields $t_{\mathrm{AM}}(\mathbf{x})$ and $t_{\mathrm{SM}}(\mathbf{x})$, together with the operation occupancy fields $p_{\mathrm{AM}}(\mathbf{x})$ and $p_{\mathrm{SM}}(\mathbf{x})$. Given a timestamp t_i, an AM isosurface $\mathcal{S}_i^{\mathrm{AM}}$ can be defined as the set of points satisfying $t_{\mathrm{AM}}(\mathbf{x}) = t_i$. Similarly, an SM isosurface $\mathcal{S}_j^{\mathrm{SM}}$ is obtained from the set of points satisfying $t_{\mathrm{SM}}(\mathbf{x}) = t_j$. These isosurfaces are used as layers for generating AM and SM toolpaths.

The AM and SM layers are extracted through the following steps:

(1) Starting from $t_0 = \epsilon$, we adaptively determine t_{i+1} such that the distance between $\mathcal{S}_i^{\mathrm{AM}}$ and $\mathcal{S}_{i+1}^{\mathrm{AM}}$ lies within the range of $d_{\mathrm{AM}} \pm 0.1 d_{\mathrm{AM}}$.
(2) Starting from $t_0 = 2.0\epsilon$, we adaptively determine t_{j+1} such that the distance between $\mathcal{S}_j^{\mathrm{SM}}$ and $\mathcal{S}_{j+1}^{\mathrm{SM}}$ lies within the range of $d_{\mathrm{SM}} \pm 0.1 d_{\mathrm{SM}}$.
(3) Each isosurface, either AM or SM, is represented as a triangular mesh extracted using the Marching Cubes algorithm [Lorensen and Cline 1987]. The extracted mesh is then trimmed by the corresponding implicit occupancy condition, i.e., $p_{\mathrm{AM}}(\mathbf{x}) > 0.5$ for AM layers, and $p_{\mathrm{AM}}(\mathbf{x}) > 0.5$ and $p_{\mathrm{SM}}(\mathbf{x}) > 0.5$ for SM layers.

After these steps, all AM and SM layers are jointly sorted according to their timestamps to determine the execution order of operations. In our implementation, $\epsilon = 0.001$ is chosen experimentally.

Finally, toolpaths are generated layer by layer. For ease of implementation, we employ zig-zag toolpaths to fill each layer. More sophisticated filling schemes, such as Fermat spirals [Zhao et al. 2016] or medial-axis-based filling [Hornus et al. 2020; Kuipers et al. 2020], can also be adopted. Waypoints are uniformly sampled along the toolpaths, and each SM waypoint is assigned an initial pose determined by the cutter configuration field $\mathcal{N}_{\mathrm{CC}}$. The SM tool pose is further adjusted during inverse kinematics solving to ensure collision-free execution and improve kinematic optimality, following the method of [Chen et al. 2025b].

5.5 Weight Tuning via Bayesian Optimization

To effectively optimize the HM process, appropriate values must be selected for the weights ω_{spt}, ω_{cAM}, and ω_{cSM}, which balance the contributions of the constraint-related terms in the total loss defined in Eq. (21). These weights correspond to the self-supporting loss $\mathcal{L}_{\mathrm{spt}}$, the AM collision-free loss $\mathcal{L}_{\mathrm{cAM}}$, and the SM collision-free loss $\mathcal{L}_{\mathrm{cSM}}$, respectively. Since these terms are intended to enforce manufacturability requirements as hard constraints, the optimal weight values are model-dependent. Therefore, in our implementation, Bayesian optimization (BO) is employed to determine suitable weights for each model by minimizing an objective that encourages all three constraint-related losses to approach zero.

Specifically, we use Optuna's multivariate tree-structured Parzen estimator (TPE) sampler [Akiba et al. 2019; Bergstra et al. 2011] to optimize the weights assigned to the different loss terms, with the outer optimization limited to 20 trials. In each trial, the loss weights are fixed and passed to the inner learning loop, in which the neural fields are optimized using the total loss defined in Eq. (21). After the inner optimization is completed, the outer BO loop evaluates the resulting solution and updates the loss weights for the next trial.

5.6 Differentiability and Gradients

Among all the loss terms defined in Eq. (21), the sensitivity $d\mathcal{L}_{\mathrm{str}}/d\boldsymbol{\rho}$ is less straightforward to compute while all other loss terms are directly formulated w.r.t. the five NNs therefore differentiable.

Take the derivative of Eq. (13) with respect to $\boldsymbol{\rho}$, we can get

$$\frac{d\mathbf{u}}{d\boldsymbol{\rho}} = \mathbf{K}^{-1}\left(\frac{d\mathbf{f}}{d\boldsymbol{\rho}} - \frac{d\mathbf{K}}{d\boldsymbol{\rho}}\mathbf{u}\right). \tag{22}$$

Table 1. Computational statistics of our approach and comparison with the results obtained by the voxel-based method

Model	Fig.	Genus	Computational Domain Dimension (unit: mm)	Pre-proc. Time (sec.)	Computing Time (sec.)	Post-proc. Time (sec.)	Support Ratio†		# Jumps		Fab. Time‡ (sec.)	
							Ours	Voxel*	Ours	Voxel*	Ours	Voxel*
GE-Bracket	1	20	$150.00 \times 90.93 \times 52.59$	769.90	9,282.20	2,833.90	28.39%	43.02%	4,934	53,958	44,351.93	139,770.48
MBB Beam	3	17	$100.00 \times 21.77 \times 21.77$	5.49	9,408.57	643.56	41.23%	35.56%	1,169	3,533	9,660.46	10,466.06
Femur	4	98	$85.86 \times 43.01 \times 100.00$	368.15	15,923.60	1,131.89	13.16%	42.53%	6,683	61,934	65,852.14	165,814.95
Cantilever Bracket	5	6	$100.00 \times 32.30 \times 40.00$	51.26	12,476.60	505.96	5.70%	11.46%	2,773	12,660	14,617.38	41,673.57
Spiral Ellipsoid	6	0	$100.00 \times 21.77 \times 21.77$	2,543.80	43,737.40	2,292.18	36.17%	131.98%	6,049	174,994	29,416.53	412,904.29
Fertility	7	4	$100.00 \times 38.19 \times 74.48$	268.77	14,166.20	1,060.06	17.51%	17.80%	3,678	24,564	20,094.98	76,145.38

* For the results generated by the voxel-based method [Chen et al. 2025a], the voxel width as 1 mm is employed for all examples.
† The support ratio reports the volume of additional support structures as a percentage of the target model volume.
‡ The fabrication time is estimated through simulation with 2.0 sec. assigned for each jump and 10.0 sec. for each AM / SM switch.

According to Eq. (14), the derivative of $\mathcal{L}_{str}(\boldsymbol{\rho}, \mathbf{u}(\boldsymbol{\rho}))$ with respect to $\boldsymbol{\rho}$ is

$$\frac{d\mathcal{L}_{\text{str}}}{d\boldsymbol{\rho}} = \frac{\partial \mathcal{L}_{\text{str}}}{\partial \boldsymbol{\rho}} + \frac{\partial \mathcal{L}_{\text{str}}}{\partial \mathbf{u}} \frac{d\mathbf{u}}{d\boldsymbol{\rho}}. \tag{23}$$

Substituting Eq. (22) into Eq. (23), we get

$$\frac{d\mathcal{L}_{\text{str}}}{d\boldsymbol{\rho}} = \frac{\partial \mathcal{L}_{\text{str}}}{\partial \boldsymbol{\rho}} + \left(\frac{\partial \mathcal{L}_{\text{str}}}{\partial \mathbf{u}} \mathbf{K}^{-1}\right) \left(\frac{\partial \mathbf{f}}{\partial \boldsymbol{\rho}} - \frac{\partial \mathbf{K}}{\partial \boldsymbol{\rho}} \mathbf{u}\right), \tag{24}$$

where the term $(\partial \mathcal{L}_{\text{str}}/\partial \mathbf{u})\mathbf{K}^{-1}$ can be obtained by solving a linear system

$$\mathbf{K}^T \boldsymbol{\lambda} = \left(\frac{\partial \mathcal{L}_{\text{str}}}{\partial \mathbf{u}}\right)^T \tag{25}$$

without explicitly computing $\mathbf{K}^{-1}$. The solution gives

$$\boldsymbol{\lambda}^T = (\partial \mathcal{L}_{\text{str}}/\partial \mathbf{u})\mathbf{K}^{-1}. \tag{26}$$

5.7 Sampling Density and FEA Resolution

The spatial integrals in our loss functions are evaluated using Monte Carlo integration. At each optimization iteration, points are stochastically sampled within the computational domain $\mathcal{D}_M$. To ensure consistency across different examples, the total number of sample points is determined according to a fixed sampling density (e.g., on average 1.25 points / mm^3). After each optimization iteration, the sample points are regenerated using the same density, which improves the likelihood of capturing small geometric features for sampling-based loss evaluation. For the RoIs used in support and collision evaluation, we randomly sample k points within each RoI, where $k = 10$ is chosen empirically in our implementation.

In our FEA implementation, the domain $\mathcal{D}_M$ is discretized using a uniform Cartesian grid of voxel elements. The element size is selected such that the longest dimension of the domain is represented by 64 elements, providing a balance between analysis accuracy and computational cost. This resolution is used for all examples.

6 Results and Discussion

6.1 Computational Experiment

Our computational pipeline is implemented in Python. We use PyTorch with the AdamW optimizer for neural network training. All computational experiments are performed on a desktop PC running Ubuntu 24.04 LTS, equipped with an AMD EPYC 9655P CPU with 96 cores, 256 GB RAM, and an NVIDIA RTX 5090 GPU with 32 GB VRAM. All examples are computed using an SM tool with a diameter of 2mm and a length of 10mm, with the spindle diameter set to 100mm. For AM, the printer head is approximated by an envelope cone with a semi-vertex angle of 45° and a height of 50mm. Our approach has been tested on a variety of examples and compared with other methods. The corresponding computational statistics are reported in Table 1.

6.1.1 Examples and Comparisons. The first example is the GE-Bracket model obtained from topology optimization [Sigmund 2001], as shown in Fig. 1. As illustrated by the zoomed-in views in Fig. 1, our method better preserves small geometric features. In addition, the toolpaths generated by our method are significantly smoother than those produced by the prior voxel-based method [Chen et al. 2025a]. This improvement can be qualitatively evaluated through frequency analysis of the tool velocity during manufacturing, for both AM and SM operations, as shown by the spectrum in the bottom-right corner of Fig. 1. We also count the number of 'jumps' and summarize the results in Table 1; the number is reduced from 53.9k to 4.9k. Moreover, an important factor affecting manufacturing efficiency in HM is the volume of additional support structures required to fabricate the target shape. For the GE-Bracket model, the percentage of support volume w.r.t. the target model's volume is reduced from 43.02% to 28.39%. Compared with the voxel-based result, our method achieves a 68% reduction in estimated fabrication time, from 38.8 hours to 12.3 hours. This improvement benefits from the significant reductions in both the volume of additional support structures and the number of 'jumps'.

The second example is the MBB Beam generated by topology optimization, as shown in Fig. 3. Unlike TO approaches that impose self-supporting constraints, e.g., [Langelaar 2016], unconstrained TO can explore a larger design space [Sigmund 2001], potentially yielding designs with better mechanical performance. This improvement will be further validated through physical experiments presented in Sec. 6.2. The corresponding computational statistics are reported in Table 1. This example is particularly interesting because the support-volume ratio generated by our method as 41.23% is slightly higher than that produced by the voxel-based method as 35.56%. Our tests indicate that this is mainly because our method is distortion-aware and therefore generates more additional support structures to improve structural strength during fabrication. An ablation study of the structural distortion loss $\mathcal{L}_{\text{str}}$ is presented in the next subsection and in Fig. 11. Without the structural distortion loss, the resulting process generates supports with a volume of only 35.44% of the target model.

The third example is a femur model with high genus, whose complex geometry poses significant challenges for either SM-only

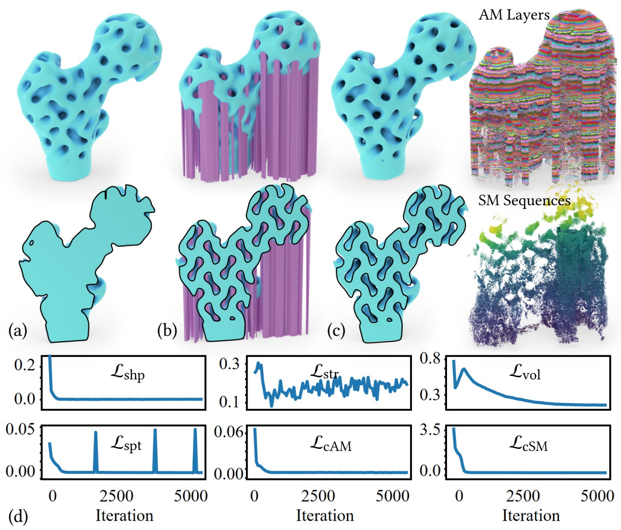

Fig. 4. Fabrication of a femur model with complex geometry (i.e., genus 98). (a) Using only SM operations cannot form the internal porous structures. (b) When using only AM operations, the internal support structures cannot be removed after fabrication. The results of (a) and (b) are both generated by Autodesk Fusion. (c) The exact target model that can be fabricated from the interlaced AM and SM operations generated by our approach with the AM and SM layers shown on the right. (d) The convergence curves of different loss terms.

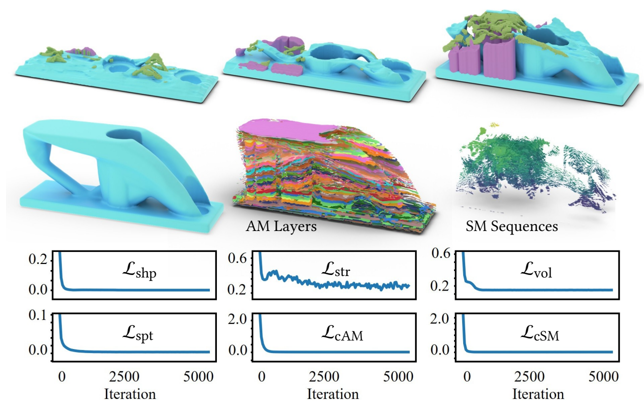

Fig. 5. Results for the Cantilever Bracket example, including the generated manufacturing process and the corresponding loss curves. The target model can be accurately fabricated, and the optimization shows stable convergence.

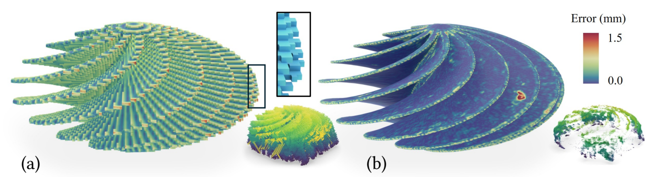

Fig. 6. Results for the Spiral Ellipsoid example using (a) the voxel-based method [Chen et al. 2025a] and (b) our field-optimization-based approach. The color maps show the corresponding shape approximation errors, while the bottom-right images visualize the resulting SM sequences applied to the temporary support structures.

or AM-only fabrication. As shown in Fig. 4, the internal porous structures are difficult to fabricate using purely SM operations, while the additional supports required for pure AM fabrication are difficult to remove afterward. In contrast, the HM process generated by our approach can accurately fabricate the input model through interlaced AM and SM operations. The process plan generated by our method completes this model through 263 switches between AM and SM operations. The convergence curves of the different loss terms are also reported in Fig. 4(d). All loss terms converge well except for the structural distortion loss $\mathcal{L}_{\mathrm{str}}$, which exhibits oscillations but remains at a very small value because the model itself has relatively high stiffness.

The fourth example is a Cantilever Bracket, which is also obtained through topology optimization. The manufacturing process and the corresponding AM layers and SM sequences are shown in Fig. 5. For all examples presented in this paper, the removable supports added during the pre-processing step are displayed in purple. They will be removed by SM operations after the target shape is produced. The temporary supports generated by our optimization approach are displayed in green. The convergence curves of the different loss terms are also shown at the bottom of Fig. 5. Compared with the voxel-based method, our approach generates optimized sequences with substantially smaller volumes of additional support in the Femur and the Cantilever Bracket examples, as indicated by the support-ratio statistics reported in Table 1. As a result, the manufacturing time is reduced by 60.3% and 64.9%, respectively.

Although the capability of our method to preserve fine geometric details has been illustrated in the zoomed-in views of Fig. 1, we further consider a more challenging example containing thin plate-like features throughout the model, namely the Spiral Ellipsoid shown in Fig. 6. Due to the geometric complexity of this model, voxel-based approaches such as [Chen et al. 2025a] tend to introduce significant staircase artifacts. Furthermore, because the features are extremely thin, many voxels are connected only through weak and unstable edge contacts, as shown in the zoomed-in view. The geometric approximation errors of both representations are visualized using color maps, from which it is clear that the voxel-based representation suffers from much larger aliasing errors. In addition, owing to the optimization-based nature of our method, the resulting HM process requires significantly less temporary support. In particular, the support ratio is reduced from 131.98% for the voxel-based method to 36.17% for our method. The corresponding SM sequences to remove support volumes are also illustrated at the bottom of Fig. 6.

The sixth example tested in our experiments is the Fertility model, as shown in Fig. 7. In this example, we compare our method with the recently published Waste-to-Value approach [Zhong et al. 2025b], although its original purpose is not HM process planning from scratch. To apply their method, we use the bounding box and the

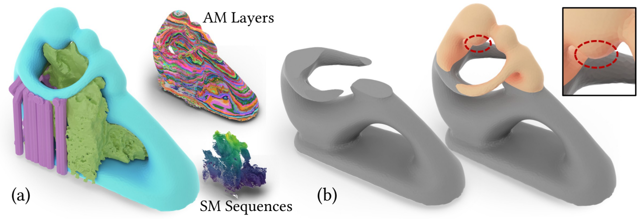

Fig. 7. Comparison of results on the Fertility model generated by (a) our method and (b) the Waste-to-Value approach [Zhong et al. 2025b]. In (a), supports added during the steps of optimization and pre-processing are shown in green and purple respectively. In (b), the gray region denotes the material retained after SM from the input bounding-box stock, while the orange region indicates the regions to be 3D printed; Large overhangs appear in the printed regions, as highlighted by the red dashed circles.

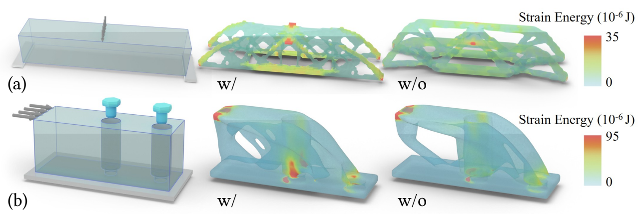

Fig. 8. Comparison of topology optimization results for (a) the MBB Beam example and (b) the Cantilever Bracket example. For each example, results with (w/) and without (w/o) the self-support constraint are computed under the same target volume fraction: 10% for the MBB Beam and 35% for the Cantilever Bracket. The strain energy distributions are visualized as color maps.

target Fertility model as the required input models for Waste-to-Value. The resultant process is shown in Fig. 7(b). The Waste-to-Value approach may still require additional support structures to complete the fabrication of the orange region, as highlighted by the red dash circles, which are not automatically generated and may be difficult to remove. In contrast, our method automatically generates a complete HM process plan for fabricating the target model.

To demonstrate the advantage of HM over AM-only fabrication, we conduct self-support-constrained topology optimization [Langelaar 2016] on the MBB Beam and Cantilever Bracket examples. For each example, results with and without self-support constraint are obtained in topology optimization subject to the same target volume fraction, and therefore have the same weights. The results are shown in Fig. 8. As indicated by the strain energy distributions, removing the self-support constraint substantially reduces the regions with high strain energy. Quantitatively, the total strain energy is reduced from 14.25 to 10.68 for the MBB Beam, and from 52.96 to 36.30 for the Cantilever Bracket (unit: 10^{-3} J). The comparisons will be further verified by mechanical tests taken on physically fabricated specimens (Sec. 6.2).

We further evaluate our method on the Coral model used in VASCO [Zhong et al. 2023]. VASCO decomposes the model into three support-free, hybrid-fabricable blocks with different printing directions, with SM used primarily for surface finishing. In contrast,

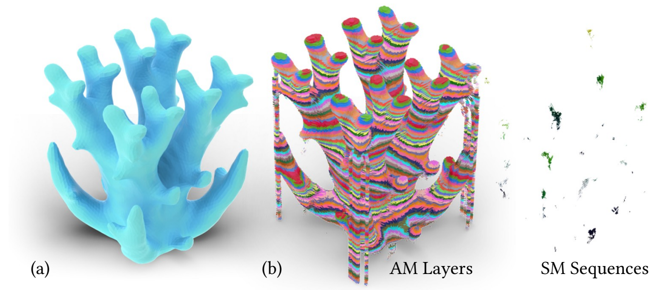

Fig. 9. Comparison on the Coral model: VASCO [Zhong et al. 2023] achieves support-free fabrication by changing the printing direction, whereas our method requires only 4.74% additional intermediate support material while remaining compatible with conventional machines with a fixed printing direction: (a) the input Coral model, and (b) the process generated by our method, with AM layers and SM operations illustrated.

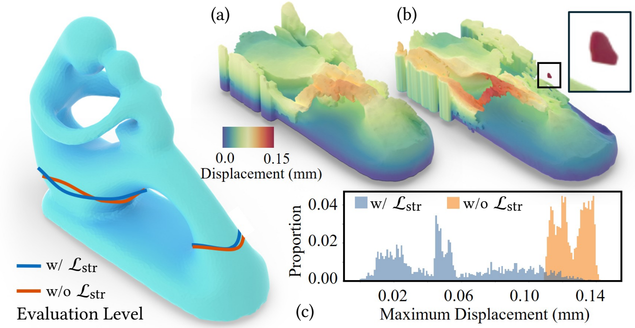

Fig. 10. Ablation study of the structural distortion loss $\mathcal{L}_{str}$ on the Fertility model. The distortion at the same fabrication height is compared between (a) the result optimized with (w/) $\mathcal{L}_{str}$ and (b) the result optimized without (w/o) $\mathcal{L}_{str}$, where a floating component is formed, as highlighted in the zoomed-in view of (b). (c) Histogram comparison of the maximum displacements over all intermediate structures.

our method integrates AM and SM operations to fabricate complex geometric features. For the same target geometry, our method generates a complete, self-supporting, and collision-free process plan while requiring additional intermediate support material amounting to only 4.74% of the target model volume (see Fig. 9). VASCO requires an AM system capable of accommodating changes in printing direction. Differently, all deposition directions in our process are vertically upward, making the method compatible with conventional three-axis AM systems. Moreover, VASCO does not account for the mechanical stability of intermediate structures.

6.1.2 Ablation Study. We first validate the importance of the structural distortion loss $\mathcal{L}_{str}$ through an ablation study. As shown in Fig. 10, large distortion and even floating components are generated when $\mathcal{L}_{str}$ is disabled. The color maps visualize the deformation displacements at selected timestamps when the model is fabricated to similar heights. The overall process behavior can be further observed from the histogram of maximum displacements over all intermediate structures, as shown in Fig. 10(c). A similar histogram

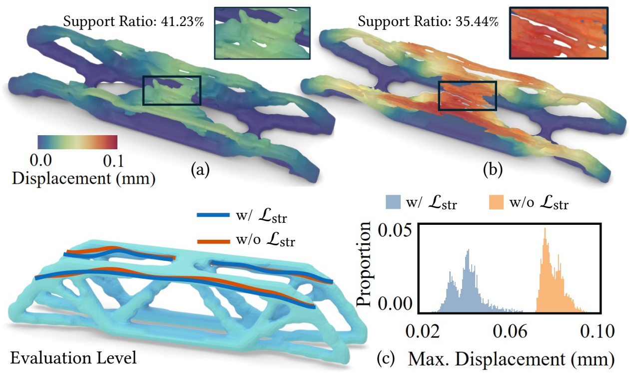

Fig. 11. Ablation study of the structural distortion loss $\mathcal{L}_{\text{str}}$ on the MBB Beam model – see the distortion at the same fabrication height for the results (a) with vs. (b) without $\mathcal{L}_{\text{str}}$, where the histogram comparison of the maximum displacements over all intermediate structures is given in (c).

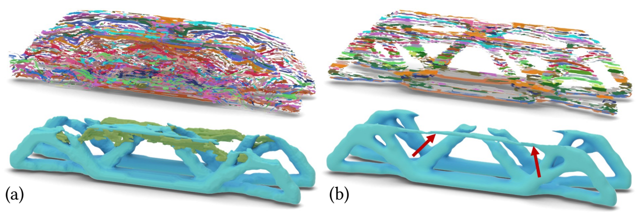

Fig. 12. Ablation study of the self-supporting loss $\mathcal{L}_{\text{spt}}$ on the MBB Beam example. The HM sequences generated (a) with $\mathcal{L}_{\text{spt}}$ and (b) without $\mathcal{L}_{\text{spt}}$ are compared. When the self-supporting loss is disabled, large overhangs are formed, as highlighted by the arrows in (b).

distribution is observed on another model, the MBB Beam, under the same ablation setting, as shown in Fig. 11. As discussed earlier, lower distortion is achieved by generating a more optimal operation-sequence together with support structures. Therefore, the result of MBB Beam optimized with $\mathcal{L}_{\text{str}}$ has the larger support ratio of 41.23%, i.e., a larger volume of supporting structures.

The second ablation study evaluates the self-supporting loss $\mathcal{L}_{\text{spt}}$. While the structural distortion loss $\mathcal{L}_{\text{str}}$ reinforces the mechanical stability of intermediate structures after they have been formed, the self-supporting loss $\mathcal{L}_{\text{spt}}$ prevents regions with large overhangs from being fabricated by AM operations. This study is conducted on the MBB Beam model, and the results are presented in Fig. 12. It can be clearly observed that, when $\mathcal{L}_{\text{spt}}$ is disabled, layers with large overhangs are generated in the HM sequence. Such overhangs can lead to material failure during the deposition process and should therefore be avoided.

We further evaluate the effectiveness of the shape-conformance loss $\mathcal{L}_{\text{shp}}$ and the support-volume loss $\mathcal{L}_{\text{vol}}$ on the Femur model. When $\mathcal{L}_{\text{shp}}$ is disabled, the residual material outside the target shape after fabrication increases from 0.70% to 9.71% of the target volume (see Fig. 13(a)), demonstrating its importance in ensuring shape conformance. When $\mathcal{L}_{\text{vol}}$ is disabled, the additional support generated during fabrication increases substantially from 13.16% to 224.77% of

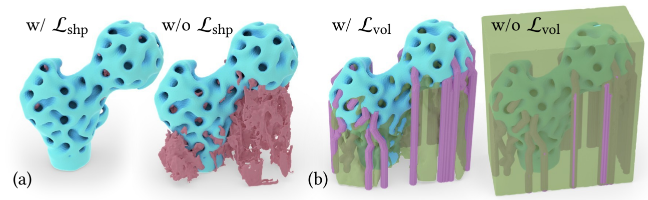

Fig. 13. Ablation study to demonstrate the effectiveness of the shape-conformance loss $\mathcal{L}_{\text{shp}}$ and the support-volume loss $\mathcal{L}_{\text{vol}}$ on the Femur model: (a) the residual material, shown in muted red, increases from 0.70% (with $\mathcal{L}_{\text{shp}}$) to 9.71% (without $\mathcal{L}_{\text{shp}}$) of the target volume, and (b) the additional support, shown in green, increases from 13.16% (with $\mathcal{L}_{\text{vol}}$) to 224.77% (without $\mathcal{L}_{\text{vol}}$) of the target volume.

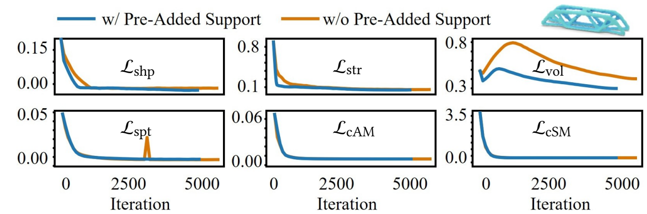

Fig. 14. Convergence curves of different loss terms for the MBB Beam example with (w/) and without (w/o) pre-added support, demonstrating that the computation converges faster with the aid of pre-added support.

the target volume, confirming that this loss is essential for suppressing excessive support generation. Without explicitly constraining the volume of additional support material, the optimizer tends to produce overly conservative process sequences that minimize the self-support and fabrication-distortion objectives. As a result, the additional support may expand to cover the entire computational domain (see Fig. 13(b)).

Another study investigates the effectiveness of the supports added during the pre-processing step. As shown by the curves in Fig. 14 for the MBB Beam example, the computation converges faster with the aid of the pre-added support structures. Moreover, a better solution, with less additional support volume and therefore reduced machining time, is obtained when the pre-processing step to add supports is applied. Note that the computational cost of pre-processing is mainly dominated by the network initialization step described in Sec. 5.2, whereas support generation is completed within one minute for all examples tested in this paper.

6.1.3 Scalability. In Fig. 1, the scalability of our method in terms of computational time is demonstrated on the GE-Bracket model by increasing the model dimensions and comparing the results with those of the voxel-based method. Here, we report additional aspects of this scalability study, including the support ratio, geometric error (measured by both mean and maximum error), number of jumps, and manufacturing time. For conducting a fair comparison, all these experiments were conducted using the same milling cutter with a diameter of 1mm and a length of 10mm. The results are shown in Fig. 15. It can be observed that our field-optimization-based method

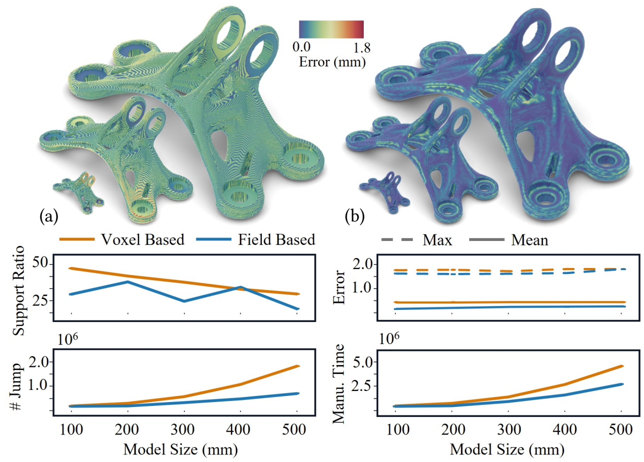

Fig. 15. Scalability study – comparison of (a) voxel-based method [Chen et al. 2025a] and (b) our field-based method for the GE-Bracket model in different dimensions.

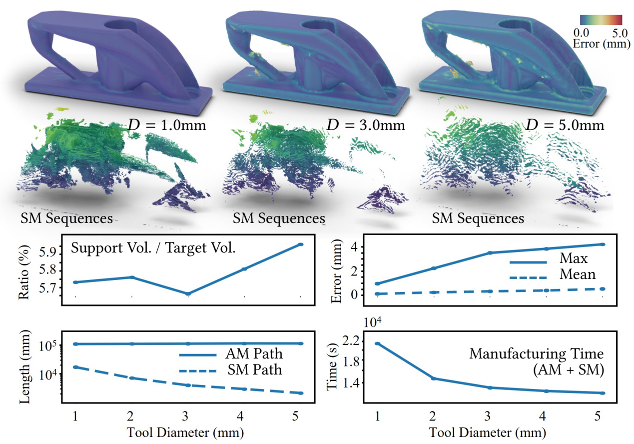

Fig. 16. Scalability tests by using SM tools with different diameter D while keeping the tool length as 10mm. The top row shows the geometric errors on the resultant model fabricated by different tool diameters.

outperforms the voxel-based method in all aspects, and the performance gaps in operation continuity and manufacturing time become larger as the model dimensions increase. This trend is also consistent with the computational-time results discussed earlier.

A distinctive advantage of our method over the voxel-based method is that the tool geometry is no longer constrained by the voxel resolution. To demonstrate this, we evaluate the Cantilever Bracket model using milling tools with different diameters, as shown in Fig. 16. In our experiments, the cutting depth and cutting width for SM operations are set proportional to the tool diameter, e.g., 20% of the diameter. The results are obtained by increasing the tool diameter while keeping the tool length fixed. As the tool diameter increases, additional support structures are required – i.e., the support ratio increases because the accessibility of the SM tool to deep

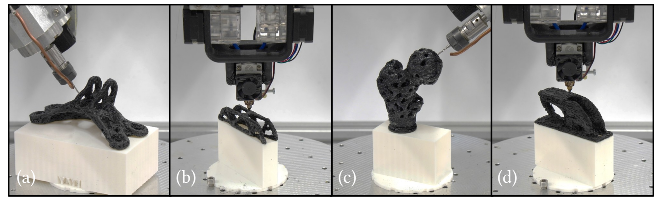

Fig. 17. Physical models fabricated using a desktop HM machine, with the corresponding fabrication times shown in parentheses: (a) GE-Bracket (16 hours), (b) MBB Beam (5 hours), (c) Femur (23 hours), and (d) Cantilever Bracket (6 hours).

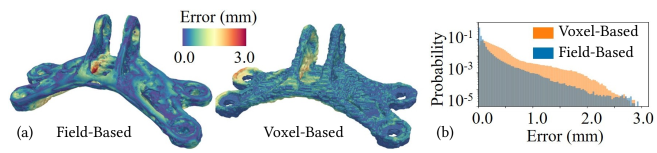

Fig. 18. Comparison of the surface quality of models fabricated using our field-based method and the voxel-based method [Chen et al. 2025a]: (a) color maps of geometric errors computed from scanned point clouds and (b) histograms showing the corresponding error distributions.

cavities becomes more limited. However, since the cutting depth and cutting width also increase with the tool diameter, larger SM tools reduce the total toolpath length and thus the machining time. This improvement in efficiency comes at the cost of reduced geometric accuracy. Overall, this capability provides users with greater flexibility in balancing manufacturing efficiency and part accuracy.

6.2 Physical Experiment

6.2.1 Hardware and Fabrication. We validate the HM sequences generated by our field-based optimization pipeline using a hybrid machine modified from a 5-axis CNC machine (Type: 5AxisMaker 5XM600XL). The hardware follows the setup proposed in [Chen et al. 2025a], but with minor variations. Specifically, all AM operations are performed using a filament-based extruder equipped with a 0.8 mm diameter nozzle, operating at an extrusion rate of 800 mm/min. For SM operations, a ball-end mill with a 2.0 mm diameter cutter and a 10.0 mm tool length is used. The cutting parameters are set to a spindle speed of 7,000 RPM and a feed rate of 650 mm/min.

We physically fabricate several examples using our hybrid machine. For AM operations, we use PLA filament reinforced with chopped carbon fibers due to its improved high-temperature resistance and machinability. As shown in Fig. 17, models with complex geometry can be successfully fabricated by following the HM sequences generated by our approach. The actual physical fabrication time for each model is affected by many non-computational factors.

Although surface quality is not directly optimized in our objective, we quantitatively evaluate the surface quality of the fabricated models using 3D scanning. Specifically, we compare the GE-bracket model fabricated using the HM process plan generated by [Chen et al. 2025a] and the model fabricated using our method. The scanned geometries are compared against the target CAD model with the

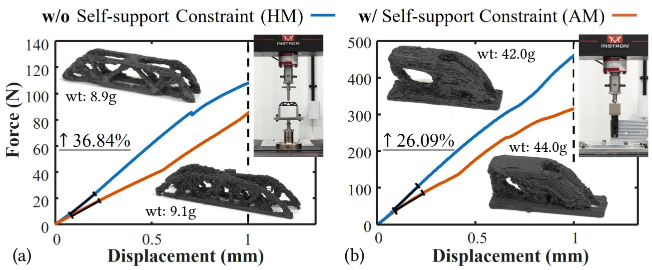

Fig. 19. Force-displacement curves obtained from mechanical tests of (a) the MBB Beam and (b) the Cantilever Bracket for structures generated by TO with (w/) and without (w/o) self-support constraints. The weights of all specimens are also reported.

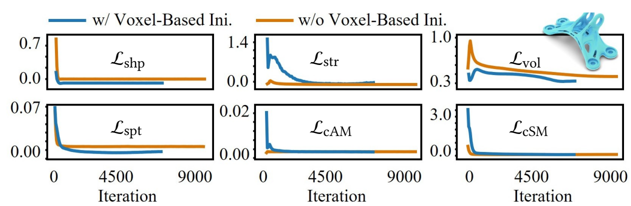

Fig. 20. Comparison of results for the GE-Bracket example with (w/) and without (w/o) the voxel-based initial guess; in the latter case, the height-based sequence is used to initialize the network. The self-supporting loss $\mathcal{L}_{\mathrm{spt}}$ does not converge well under this simple height-based initialization.

color maps of geometric errors shown in Fig. 18(a). The quantitative results demonstrate the advantage of our method. For the GE-bracket model with dimensions $150.00 \times 90.93 \times 52.59$ mm^3, the surface area with errors greater than 1 mm decreases from 6.59% (for the voxel-based method [Chen et al. 2025a]) to 1.22% (for our field-based method). The surface area with errors greater than 2 mm decreases from 0.44% (for the voxel-based method) to 0.12% (for our method). A histogram comparison of geometric errors has been given in Fig. 18(b).

6.2.2 Mechanical Tests. We physically fabricate the MBB Beam and Cantilever Bracket results presented in Fig. 8 and compare their mechanical stiffness through experimental testing. For the designs generated by TO with the self-support constraint, the specimens are fabricated using AM-only operations with planar layers. Three-point bending tests are performed for the MBB Beam example, while compression tests are performed for the Cantilever Bracket examples. The mechanical tests are limited to a displacement range of 0–1mm, which better matches the loading conditions used in TO. The resultant force–displacement curves are shown in Fig. 19, demonstrating the significant stiffness advantage of the models computed by TO without self-support constraints, which can only be fabricated using HM: stiffness was 36.84% higher for the MBB Beam and 26.09% higher for the Cantilever Bracket.

6.3 Limitations

The field-based optimization approach introduced in this paper performs well on a variety of examples, as demonstrated above.

Table 2. Effectiveness of initialization in different voxel resolutions.

Init. Res.	Time†	$\mathcal{L}_{\mathrm{shp}}$	$\mathcal{L}_{\mathrm{str}}$	$\mathcal{L}_{\mathrm{vol}}$	$\mathcal{L}_{\mathrm{spt}}$	$\mathcal{L}_{\mathrm{cAM}}$	$\mathcal{L}_{\mathrm{cSM}}$
$50 \times 29 \times 17$	0.43	0.0199	0.087	0.48	0.0037	0.0005	0.0028
$100 \times 59 \times 35$	25.45	0.0196	0.038	0.38	0.0035	0.0005	0.0037
$150 \times 89 \times 53$	367.98	0.0192	0.016	0.31	0.0030	0.0004	0.0034

†Computing time (sec.) to obtain the initialization by the voxel-based method [Chen et al. 2025a], where the experiments are taken on the GE-bracket model (Fig. 1).

However, like many numerical optimization approaches, its convergence rate depends strongly on a good initial guess. In this work, the voxel-based method presented in [Chen et al. 2025a] is employed to provide an initial guess for the HM processing plan. As shown in Fig. 20, replacing the meaningful voxel-based initialization with a simple height-based initialization leads to noticeably poorer convergence, especially for the self-supporting loss $\mathcal{L}_{\mathrm{spt}}$, demonstrating the importance of an appropriate initialization.

Although initialization quality affects the optimization, it only needs to provide the overall trend of the material processing flow; therefore, a high-resolution voxel initialization is not required. Our experiments on the GE-bracket model (Fig. 1) with multiple voxel resolutions show that the subsequent optimization is relatively insensitive to the voxel resolution in initializations – see statistics reported in Table 2. The generation time of the initial solution increased dramatically with resolution. However, from the resulting loss values, we observe that the three strictly constrained losses, $\mathcal{L}_{\mathrm{spt}}$, $\mathcal{L}_{\mathrm{cAM}}$, and $\mathcal{L}_{\mathrm{cSM}}$, all remain at very low levels. Increasing the resolution slightly reduces $\mathcal{L}_{\mathrm{shp}}$, $\mathcal{L}_{\mathrm{str}}$, and $\mathcal{L}_{\mathrm{vol}}$, but the improvements are not significant. Nevertheless, a very low resolution may fail to capture case-specific topological issues. As a trade-off, we use a resolution of 100 in all examples.

The second limitation of our approach is that the current formulation allows only one switch between AM and SM operations at each position. Although prior research [Chen et al. 2025a] has theoretically proven that this is sufficient to fabricate any shape, it is not necessarily optimal. In fact, we suspect that better solutions, such as those with lower distortion or reduced additional support volume, may exist if multiple switches between AM and SM operations are allowed at the same position. Moreover, such flexibility could enable the incorporation of additional optimization objectives, such as optimizing the orientation of fibre-reinforced filaments to improve mechanical strength [Liu et al. 2025]. We will consider this in our future work.

The last limitation is that, in our current implementation, the cutter configuration field is formulated to support a ball-end mill. However, this formulation can be readily extended to other types of milling cutters. This extension is relatively straightforward and will be considered in future work.

7 Conclusion

To address the limitations of existing HM process planning methods, including incomplete fabrication, poor scalability, lack of optimization, uncontrolled distortion, and highly discontinuous manufacturing sequences, we present a field-optimization based method for determining a feasible and efficient sequence of AM/SM operations to fabricate a target shape. Our method represents manufacturing states using continuous fields, which scale naturally to models with large dimensions and fine geometric features while enabling smoother manufacturing sequences. The field optimization problem is parameterized by five implicit neural networks, whose coefficients serve as the optimization variables. Loss functions are defined according to multiple objectives, including shape completeness, fabrication time (via the volume of extra support), distortion of intermediate structures, and manufacturability constraints (including self-supporting requirements and collision avoidance for both AM and SM operations). By optimizing these fields, our method determines the AM layers and SM sequences for HM process planning. Significant performance improvements over prior work are demonstrated on a variety of examples through both computational and physical experiments.

Acknowledgments

This project was supported by the Chair Professorship Fund at the University of Manchester (UoM) and the UK Engineering and Physical Sciences Research Council (EPSRC) Fellowship Grant (Ref: EP/X032213/1). The authors also gratefully acknowledge the technical support provided by 5AXISWORKS Ltd. and the 3D Printing Lab of UoM.

References

Niels Aage, Erik Andreassen, Boyan S. Lazarov, and Ole Sigmund. 2017. Giga-voxel computational morphogenesis for structural design. *Nature* 550, 7674 (2017), 84 – 86.

Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. 2019. Optuna: a next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '19)*. 2623–2631.

Grégoire Allaire, Charles Dapogny, Alexis Faure, and Georgios Michailidis. 2017. Shape optimization of a layer by layer mechanical constraint for additive manufacturing. *Comptes Rendus Mathematique* 355, 6 (2017), 699–717.

Morad Behandish, Saigopal Nelaturi, and Johan de Kleer. 2018. Automated process planning for hybrid manufacturing. *Computer-Aided Design* 102 (2018), 115 – 127.

James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. 2011. Algorithms for hyper-parameter optimization. In *Advances in Neural Information Processing Systems*, Vol. 24. 2546–2554.

Aaditya Chandrasekhar and Krishnan Suresh. 2021. TOuNN: topology optimization using neural networks. *Structural and Multidisciplinary Optimization* 63, 3 (2021), 1135–1149.

Li Chen, Tak Yu Lau, and Kai Tang. 2020. Manufacturability analysis and process planning for additive and subtractive hybrid manufacturing of quasi-rotational parts with columnar features. *Computer-Aided Design* 118 (2020), 102759.

Li Chen, Ke Xu, and Kai Tang. 2018. Optimized sequence planning for multi-axis hybrid machining of complex geometries. *Computers & Graphics* 70 (2018), 176 – 187.

Yongxue Chen, Tao Liu, Yuming Huang, Weiming Wang, Tianyu Zhang, Kun Qian, Zikang Shi, and Charlie C. L. Wang. 2025a. Can any model be fabricated? Inverse operation based planning for hybrid additive-subtractive manufacturing. *ACM Transactions on Graphics* 44, 6, Article 191 (2025).

Yongxue Chen, Tianyu Zhang, Yuming Huang, Tao Liu, and Charlie C. L. Wang. 2025b. Co-optimization of tool orientations, kinematic redundancy, and waypoint timing for robot-assisted manufacturing. *IEEE Transactions on Automation Science and Engineering* 22 (2025), 12102–12117.

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, Article 134 (2018).

Nikan Doosti, Julian Panetta, and Vahid Babaei. 2021. Topology Optimization via Frequency Tuning of Neural Design Representations. In *Proceedings of the 6th Annual ACM Symposium on Computational Fabrication (SCF '21)*. Article 1.

Jérémie Dumas, Jean Hergel, and Sylvain Lefebvre. 2014. Bridging the gap: automated steady scaffoldings for 3D printing. *ACM Transactions on Graphics* 33, 4, Article 98 (2014).

Neelotpal Dutta, Tianyu Zhang, Tao Liu, Yongxue Chen, and Charlie C. L. Wang. 2026. Implicit neural field-based process planning for multi-axis manufacturing: direct control over collision avoidance and toolpath geometry. arXiv:2511.17578 [cs.RO]

Jimmy Etienne, Nicolas Ray, Daniele Panozzo, Samuel Hornus, Charlie C. L. Wang, Jonàs Martínez, Sara McMains, Marc Alexa, Brian Wyvill, and Sylvain Lefebvre. 2019. CurviSlicer: slightly curved slicing for 3-axis printers. *ACM Transactions on Graphics* 38, 4, Article 81 (2019).

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).

Joseph M. Flynn, Alborz Shokrani, Stephen T. Newman, and Vimal Dhokia. 2016. Hybrid additive and subtractive machine tools – research and industrial developments. *International Journal of Machine Tools and Manufacture* 101 (2016), 79–101.

Xingjian Han, Yu Jiang, Weiming Wang, Guoxin Fang, Simeon Gill, Zhiqiang Zhang, Shengfa Wang, Jun Saito, Deepak Kumar, Zhongxuan Luo, Emily Whiting, and Charlie C.L. Wang. 2024. Motion-driven neural optimizer for prophylactic braces made by distributed microstructures. In *SIGGRAPH Asia 2024 Conference Papers* (Tokyo, Japan) *(SA '24)*. Article 57.

George P. Harabin and Morad Behandish. 2022. Hybrid manufacturing process planning for arbitrary part and tool shapes. *Computer-Aided Design* 151 (2022), 103299.

Yu He, Jiacheng Wei, Yuyu He, Xiaoxu Rong, Wenfeng Guo, Fei Wang, Yang Wang, and Junyan Liu. 2023. A process strategy planning of additive-subtractive hybrid manufacturing based multi-dimensional manufacturability evaluation of geometry feature. *Journal of Manufacturing Systems* 67 (2023), 296 – 314.

Samuel Hornus, Tim Kuipers, Olivier Devillers, Monique Teillaud, Jonàs Martínez, Marc Glisse, Sylvain Lazard, and Sylvain Lefebvre. 2020. Variable-width contouring for additive manufacturing. *ACM Transactions on Graphics* 39, 4, Article 131 (2020).

Ruizhen Hu, Honghua Li, Hao Zhang, and Daniel Cohen-Or. 2014. Approximate pyramidal shape decomposition. *ACM Transactions on Graphics* 33, 6, Article 213 (2014).

Tim Kuipers, Eugeni L. Doubrovski, Jun Wu, and Charlie C.L. Wang. 2020. A framework for adaptive width control of dense contour-parallel toolpaths in fused deposition modeling. *Computer-Aided Design* 128, Article 102907 (2020).

Weixian Lan and Julian Panetta. 2022. Efficient Layer-by-Layer Simulation for Topology Optimization. In *Proceedings of the 7th Annual ACM Symposium on Computational Fabrication (SCF '22)*. Article 4.

Matthijs Langelaar. 2016. Topology optimization of 3D self-supporting structures for additive manufacturing. *Additive Manufacturing* 12 (2016), 60 – 70.

Matthijs Langelaar. 2019. Topology optimization for multi-axis machining. *Computer Methods in Applied Mechanics and Engineering* 351 (2019), 226–252.

Eric Larsen, Stefan Gottschalk, Ming C Lin, and Dinesh Manocha. 1999. *Fast proximity queries with swept sphere volumes*. Technical Report TR 99-018. Department of Computer Science, University of North Carolina.

Steven M. LaValle. 1998. *Rapidly-exploring random trees: a new tool for path planning*. Technical Report TR 98-11. Computer Science Department, Iowa State University.

Hak Yong Lee, Mu Zhu, and James K. Guest. 2022. Topology optimization considering multi-axis machining constraints using projection methods. *Computer Methods in Applied Mechanics and Engineering* 390 (2022), 114464.

Jikai Liu and Albert C. To. 2017. Topology optimization for hybrid additive-subtractive manufacturing. *Structural and Multidisciplinary Optimization* 55, 4 (2017), 1281 – 1299.

Tao Liu, Aoran Lyu, Yongxue Chen, Yu Jiang, Michael James Petty, and Charlie C. L. Wang. 2026. Co-Optimization of Structure and Manufacturable Semi-Continuous Layers for Laminated Composites. *ACM Transactions on Graphics* 45, 4, Article 146 (2026).

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, Article 85 (2024).

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, Article 128 (2025).

William E. Lorensen and Harvey E. Cline. 1987. Marching Cubes: a high resolution 3D surface construction algorithm. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '87)*. 163–169.

Ali Mahdavi-Amiri, Fenggen Yu, Haisen Zhao, Adriana Schulz, and Hao Zhang. 2020. VDAC: volume decompose-and-carve for subtractive manufacturing. *ACM Transactions on Graphics* 39, 6, Article 203 (2020).

Amir M. Mirzendehdel, Morad Behandish, and Saigopal Nelaturi. 2022. Topology optimization for manufacturing with accessible support structures. *Computer-Aided Design* 142 (2022), 103117.

Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. 2022. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics* 41, 4, Article 102 (2022).

Jiasheng Qu, Zhuo Huang, Dezhao Guo, Hailin Sun, Aoran Lyu, Chengkai Dai, Yeung Yam, and Guoxin Fang. 2025. INF-3DP: implicit neural fields for collision-free multi-axis 3D printing. *ACM Transactions on Graphics* 44, 6, Article 195 (2025).

Adriana Schulz, Ariel Shamir, David I. W. Levin, Pitchaya Sitthi-amorn, and Wojciech Matusik. 2014. Design and fabrication by example. *ACM Transactions on Graphics* 33, 4, Article 62 (2014).

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

Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. 2020. Implicit neural representations with periodic activation functions. In *Advances in Neural Information Processing Systems*, Vol. 33. 7462–7473.

Kelsey L. Snapp, Benjamin Verdier, Aldair E. Gongora, Samuel Silverman, Adedire D. Adesiji, Elise F. Morgan, Timothy J. Lawton, Emily Whiting, and Keith A. Brown. 2024. Superlative mechanical energy absorbing efficiency discovered through self-driving lab-human partnership. *Nature Communications* 15, 1, Article 4290 (2024).

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

Weiming Wang, Yanhao Hou, Renbo Su, Weiguang Wang, and Charlie C. L. Wang. 2025. Simultaneous topology optimization of differentiable and non-differentiable objectives via morphology learning: stiffness and cell growth on scaffold. *Advanced Intelligent Discovery* 1, 1 (2025), 2400015.

Weiming Wang, Dirk Munro, Charlie C.L. Wang, Fred van Keulen, and Jun Wu. 2020. Space-time topology optimization for additive manufacturing: concurrent optimization of structural layout and fabrication sequence. *Structural and Multidisciplinary Optimization* 61, 1 (2020), 1 – 18.

Shuzhi Xu, Jikai Liu, Kentaro Yaji, and Lin Lu. 2024. Topology optimization for hybrid additive-subtractive manufacturing incorporating dynamic process planning. *Computer Methods in Applied Mechanics and Engineering* 431 (2024), 117270.

Li Yang, Weiming Wang, Ye Ji, Chun-Gang Zhu, and Charlie C.L. Wang. 2025. Space-time isogeometric topology optimization with additive manufacturing constraints. *Computer Methods in Applied Mechanics and Engineering* 441, 117976 (2025), 1 – 28.

Jonas Zehnder, Yue Li, Stelian Coros, and Bernhard Thomaszewski. 2021. NTopo: mesh-free topology optimization using implicit neural representations. In *Proceedings of the 35th International Conference on Neural Information Processing Systems (NIPS '21)*. Article 793.

Shuming Zhang, Zhidong Guan, Hao Jiang, Xiaodong Wang, and Pingan Tan. 2024. BrepMFR: enhancing machining feature recognition in B-rep models through deep learning and domain adaptation. *Computer Aided Geometric Design* 111 (2024), 102318.

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, Article 277 (2022).

Haisen Zhao, Fanglin Gu, Qi-Xing Huang, Jorge Garcia, Yong Chen, Changhe Tu, Bedrich Benes, Hao Zhang, Daniel Cohen-Or, and Baoquan Chen. 2016. Connected Fermat spirals for layered fabrication. *ACM Transactions on Graphics* 35, 4, Article 100 (2016).

Yufan Zheng, Jikai Liu, and Rafiq Ahmad. 2020. A cost-driven process planning method for hybrid additive–subtractive remanufacturing. *Journal of Manufacturing Systems* 55 (2020), 248 – 263.

Fanchao Zhong, Yang Wang, Peng-Shuai Wang, Lin Lu, and Haisen Zhao. 2025a. DeepMill: neural accessibility learning for subtractive manufacturing. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '25)*. Article 159.

Fanchao Zhong, Zhenmin Zhang, Liyuan Wang, Xin Yan, Jikai Liu, Lin Lu, and Haisen Zhao. 2025b. Waste-to-Value: reutilized material maximization for additive and subtractive hybrid remanufacturing. *ACM Transactions on Graphics* 44, 6, Article 190 (2025).

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, Article 188 (2023).

Bo Zhu, Mélina Skouras, Desai Chen, and Wojciech Matusik. 2017. Two-scale topology optimization with microstructures. *ACM Transactions on Graphics* 36, 4, Article 120b (2017).