Eurographics Symposium on Geometry Processing (2026)
M. Skouras and M. Livesu (Guest Editors)

Neural Field-Based Sequence Planning for Additive-Subtractive Hybrid Manufacturing

Shudan Guo[1], Fanchao Zhong[1], Liyuan Wang[1] and Haisen Zhao[1]†

[1]Shandong University,Qingdao,China

Abstract

*Additive-Subtractive Hybrid Manufacturing (ASHM) combines additive geometric freedom with subtractive precision. However, sequence planning remains a major computational bottleneck. Existing discrete graph-based methods like VASCO [ZZL*23] frequently suffer from combinatorial memory explosion (O.O.M.) on complex models. We introduce a continuous neural field approach. Our framework utilizes neural networks to implicitly encode manufacturing constraints into a differentiable dual-field representation: an Accessibility Field predicting tool orientations and a Priority Field capturing global occlusion. Fusing these into a unified score, we formulate cutting plane selection as a maximal-margin Linear Support Vector Machine (SVM) optimization. To ensure strict manufacturability, we introduce BlockChecker for robust kinematic validation. Experiments demonstrate our framework successfully processes high-genus models where traditional methods fail, maintaining highly scalable, near-constant computation times.*

CCS Concepts

• ***Computing methodologies*** → ***Shape modeling;*** *Graphics systems and interfaces;*

1. Introduction

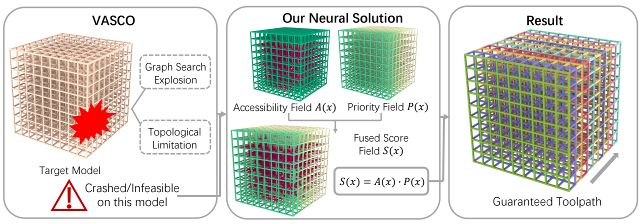

Figure 1: *Our continuous neural framework avoids the O.O.M. crashes of discrete graph searches (e.g., VASCO), guaranteeing manufacturability.*

Additive-Subtractive Hybrid Manufacturing (ASHM) integrates the geometric flexibility of 3D printing (AM) with the high precision of multi-axis milling (SM) on a unified platform [CLH*25]. Unlocking its full potential requires advanced sequence planning to decompose a complex model into a sequence of hybrid-fabricable sub-blocks $\mathscr{B} = \{B_1, B_2, ..., B_k\}$. Because subtractive tools are limited by line-of-sight and physical spindle collisions, the manufacturing process must alternate between adding a block and immediately machining its exposed critical surfaces before they are occluded by subsequent layers.

State-of-the-art frameworks like VASCO [ZZL*23] rely on discrete slab-based graph searches. While pioneering, they suffer from combinatorial explosion on high-resolution meshes, and explicit rule-based checks frequently fail to capture continuous accessibility transitions in high-genus lattice structures, leading to Out-Of-Memory (O.O.M.) crashes. To mitigate these scalability limits, we present a continuous, data-driven alternative for 3+2 axis ASHM process planning. By mapping multi-modal manufacturing constraints into continuous implicit fields over the mesh domain, we transform a intractable combinatorial optimization into a convex margin-maximization problem.

2. Problem Setup and Dual Neural Fields

We assume a multi-axis platform with a multi-stage ball-end cutter and a fixed build direction d_{am}. The AM phase deposits a near-net-shape stock, which SM mills to the exact boundaries of B_i, satisfying overhang (C_{AM}) and accessibility (C_{SM}) constraints.

Continuous Neural Field Representation: In discrete frameworks, ensuring C_{SM} requires explicitly computing dense visibility matrices via ray-casting, inevitably triggering memory explosion. To circumvent this, we reformulate the problem from discrete graph traversal to continuous implicit field analysis, amortizing physical ray-casting overhead by encoding constraints into differentiable scalar fields over the mesh domain.

† Corresponding author.

Using an Octree-based CNN (O-CNN) [WLG*17], we transform abstract constraints into spatial probability distributions:

- **Accessibility Field** $A(x)$: Evaluates local kinematic machinability [ZWW*25]. Trained on simulated multi-stage cutter envelopes Θ_{cutter}, $A(x)$ regresses a continuous probability distribution mapping surface vertices to their collision-free orientation spaces. Learning this field replaces computationally prohibitive real-time Boolean intersections with rapid forward-pass inferences.
- **Priority Field** $P(x)$: Provides the global temporal context. Because $A(x)$ is evaluated in local isolation, we supervise $P(x)$ using a geometric look-ahead heuristic to guide sequence planning: a surface region receives a high priority label if its immediate removal maximizes the "unlocking" of previously occluded volumes. We bootstrap this dataset using accelerated discrete rollouts to capture global occlusion boundaries.

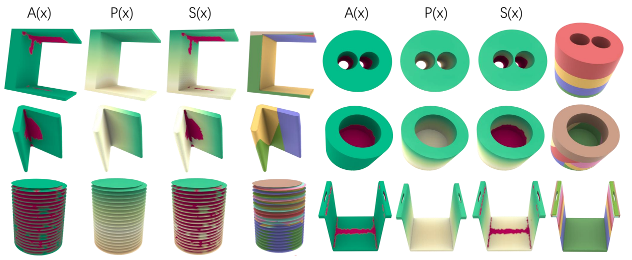

Figure 2: *Visualization of the proposed dual fields ($A(x)$, $P(x)$) and the fused manufacturability score $S(x)$ over complex meshes.*

3. SVM-Based Partitioning and Geometric Verification

We form an integrated score field $S(x) = A(x) \cdot P(x)$, heavily penalizing regions where either metric approaches zero to isolate optimal candidates.

We sample mesh vertices into high-priority bottlenecks $\mathscr{V}^+$ ($S(v_i) \geq \tau_{high}$) and low-feasibility regions $\mathscr{V}^-$ ($S(v_i) \leq \tau_{low}$). A linear Support Vector Machine (SVM) computes the maximal-margin separating normal n^* between these classes:

$$\min_{n,t,\xi} \frac{1}{2}\|n\|^2 + C\sum_i \xi_i \quad \text{s.t.} \quad y_i(n^\top v_i - t) \geq 1 - \xi_i \tag{1}$$

The normal n^* aligns with constraint boundaries. We fix the final translation t^* by finding the first local extremum of a feasibility response function $f(t) = \mathrm{Vol}(H(n^*, t) \cap \mathscr{M})$ that yields a valid component.

A deterministic validator, **BlockChecker**, uses Golden-Spiral sampling and BVH collision checks to rigorously verify B_i. If manufacturability falls below $\Gamma = 95\%$, a backtracking loop adjusts t^* until a physically sound cut is confirmed.

4. Preliminary Results and Discussion

We evaluated our framework across a dataset of 50+ complex models using an Intel i9 workstation and an NVIDIA RTX 4090 GPU.

Superior Scalability and Robustness: As shown in Table 1, our continuous field-based approach circumvents the combinatorial memory explosion inherent to explicit discrete graph searches. For complex, high-genus topologies like the Lattice model, where the state-of-the-art VASCO algorithm crashes due to Out-Of-Memory (O.O.M.) errors, our framework maintains a bounded, lightweight footprint (< 2.0 GB) and successfully extracts a valid manufacturing sequence. Furthermore, it demonstrates exceptional scalability. Compared to VASCO, our framework achieves an aggregate **43.7% reduction in average computation time** across the entire dataset, delivering a stable, near-constant average optimization time per block ($T_{avg} \approx 7-25$s) that remains robust regardless of geometric complexity.

Table 1: *Representative quantitative results.*

Model	Vertices	VASCO (N_b/Time)	Ours (N_b/Time)
Kitten	20k	2 / 102.0s	**2 / 24.07s**
Fertility	7k	N.S. (Failed)	**40** / 271.7s
Lattice	23k	O.O.M. (Crashed)	**140** / 1567.0s
TPMS	30k	**47** / 804.0s	161 / 1147.0s

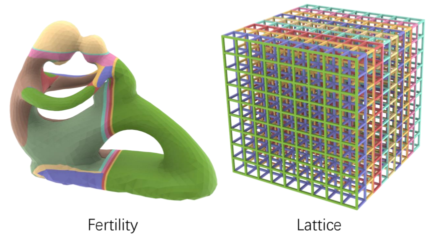

Figure 3: *Visual validation.*

Limitations and Future Work: Our pipeline achieves a 66.7% overall sequence planning success rate across the extensive benchmark dataset. Occasional failures in the remaining highly entangled models primarily stem from localized false-positive neural predictions, which can temporarily misguide the SVM hyperplane and exhaust the verification loop. Furthermore, to strictly guarantee physical manufacturability where graph methods fail, our current reliance on iterative single-planar cuts inherently produces a higher total block count for extreme geometries (e.g., $N_b = 161$ for TPMS). Enhancing neural precision and extending the optimization to multi-plane or non-planar cutting manifolds represent compelling directions for future work to further compress the sequence length.

References

[CLH*25] Chen Y., Liu T., Huang Y., Wang W., Zhang T., Qian K., Shi Z., Wang C. C.: Can any model be fabricated? inverse operation based planning for hybrid additive-subtractive manufacturing. *ACM Transactions on Graphics (TOG) 44*, 6 (2025), 1–14. 1

[WLG*17] Wang P.-S., Liu Y., Guo Y.-X., Sun C.-Y., Tong X.: O-CNN: Octree-based convolutional neural networks for 3d shape analysis. *ACM Transactions on Graphics (TOG) 36*, 4 (2017), 1–11. 2

[ZWW*25] Zhong F., Wang Y., Wang P.-S., Lu L., Zhao H.: Deepmill: Neural accessibility learning for subtractive manufacturing. *arXiv preprint arXiv:2502.06093* (2025). 2

[ZZL*23] Zhong F., Zhao H., Li H., Yan X., Liu J., Chen B., Lu L.: Vasco: Volume and surface co-decomposition for hybrid manufacturing. *ACM Trans. Graph. 42*, 6 (2023), 1–17. 1