VOROFACET: JUNCTION-AWARE MESH EXTRACTION FROM NEURAL UNSIGNED DISTANCE FIELDS

Anonymous authors
Paper under double-blind review

ABSTRACT

Unsigned distance fields (UDFs) can represent surfaces of arbitrary topology, but recovering mesh connectivity from neural UDFs remains challenging, especially at non-manifold junctions where multiple sheets meet. To this end, we propose **VoroFacet**, a novel pipeline that extracts single-layer meshes from neural UDFs and reconstructs junction connectivity accurately through Voronoi facets. Specifically, we first use a near-surface double-layered proxy derived from the UDF as a local spatial partition for facet selection, where the facets belong to the Voronoi diagram of the sites generated by an adaptive ball cover of the proxy. We then estimate junctions from the mesh of this first extraction and place balls along them, so that in a re-extraction process, the resulting cospherical sites make incident sheets share Voronoi vertices and edges at these junctions. Finally, we propose a three-stage optimization to refine the geometry of the extracted mesh. Experiments demonstrate superior reconstruction quality over state-of-the-art methods, particularly at non-manifold junctions. Interactive 3D comparisons are available at https://anonymous2180.github.io/Vorofacet-project/. The source code will be publicly available.

1 INTRODUCTION

With the rapid development of 3D vision, neural implicit fields have attracted increasing attention as a representation of geometry. The most widely used are binary occupancy fields (BOFs) (Mescheder et al., 2019) and signed distance fields (SDFs) (Park et al., 2019), but both rely on an inside/outside classification and can therefore represent only watertight surfaces. Unsigned distance fields (UDFs) (Chibane et al., 2020) avoid this by assigning to each spatial location only its distance to the nearest surface point, and can therefore represent surfaces of arbitrary topology, including open boundaries, non-orientable surfaces, and non-manifold junctions. However, converting neural UDFs into meshes that faithfully represent the underlying surface remains difficult, which limits their use in downstream applications that require explicit geometry.

The first difficulty is that a UDF is non-negative, so field values alone cannot indicate surface crossings; classical sign-based isosurfacing methods (Lorensen & Cline, 1998; Ju et al., 2002) do not apply. The second is that the ideal UDF is not differentiable on the surface, nor on the medial axis, where the nearest point on the surface is not unique. Since neural UDFs are typically smooth, they poorly approximate the ideal field in the narrow regions around them, where the distance values are biased and the gradients are substantially less reliable (Zhang et al., 2023).

Despite these difficulties, existing methods extract meshes from neural UDFs with remarkable geometric accuracy, but their connectivity is often incorrect at non-manifold junctions, where three or more sheets meet. Since the field is unsigned, most methods (Guillard et al., 2022; Zhou et al., 2024; Ren et al., 2023; Stella et al., 2025; Chen et al., 2025a; Zhang et al., 2023; Chen et al., 2022; Bao et al., 2026) infer connectivity from local surface evidence, such as gradients, using handcrafted rules or learned networks. At such a junction, these sheets create additional medial axes that degrade this evidence, so these methods tend to leave holes or spurious connections. By contrast, offset-based methods (Hou et al., 2023; Chen et al., 2025b) deform the more reliable positive level set of the neural UDF toward the surface, but the near-surface double-layered approximation they produce leaves a gap at a junction instead of a shared connection.

Different from existing methods, we introduce **VoroFacet**, which reconstructs the target surface as a selected set of facets in an ambient Voronoi diagram. Our starting observation is that the near-surface double-layered approximation of DCUDF2 (Chen et al., 2025b) can serve as a useful proxy even when its connectivity is unsuitable for the desired output. Instead of transferring the proxy's vertices or faces, we use it only to guide the selection of facets from the Voronoi diagram, so that the output connectivity is constructed anew rather than inherited. Since Voronoi cells partition the space without gaps or overlaps, the selected facets meet along shared Voronoi edges: a smooth sheet is extracted as a single layer, and sheets meeting at a junction are joined along common edges rather than left with holes, gaps, or spurious connections. Moreover, as connectivity follows from this spatial partition rather than from gradient cues or global inside/outside labels, it is less sensitive to the approximation errors of neural UDFs and imposes no restriction on the topology of the target, handling open boundaries, non-orientable surfaces, and non-manifold junctions within a single pipeline.

Our main contributions are summarized as follows.

- We propose VoroFacet, which formulates mesh extraction from neural UDFs as facet selection in a Voronoi diagram guided by a near-surface double-layered proxy, producing single-layer meshes for surfaces of arbitrary topology.
- We design a junction-aware re-extraction strategy that estimates non-manifold junctions and enforces shared connectivity among the sheets meeting there, recovering junctions that existing methods leave as holes, spurious connections, or gaps.
- We develop an algorithm to extract junction curves from CAD meshes to quantitatively evaluate junction reconstruction. Extensive experiments on diverse datasets show that VoroFacet achieves state-of-the-art junction reconstruction while maintaining competitive surface accuracy.

2 Related Work

Existing UDF mesh extraction methods differ mainly in how they determine mesh connectivity, that is, how surface points are connected into faces.

Connectivity from Inferred Pseudo-signs. Pseudo-sign methods infer from local values and gradients, by handcrafted rules or learned classifiers, which corners of each grid cell lie on opposite sides of the surface, and let a Marching Cubes lookup table determine the connectivity within the cell (Guillard et al., 2022; Zhou et al., 2024; Ren et al., 2023; Stella et al., 2024; 2025). MIND (Chen et al., 2025a) assembles such local two-sided fields into a global multi-label field to extract non-manifold junctions. However, at a junction the corners of a cell cannot be split into two sides, so the local two-sided field is ill-defined and the gradients behind it are unreliable, leaving holes or spurious connections; a global labeling is further undefined for non-orientable surfaces.

Connectivity from Cell Adjacency. Dual methods follow Dual Contouring (Ju et al., 2002). They place one representative point in each cell that contains the surface, and cell adjacency proposes candidate faces between neighboring cells, which handcrafted rules or learned networks accept or reject using local surface evidence such as values or gradients (Zhang et al., 2023; Chen et al., 2022; Bao et al., 2026). Since such decisions are heuristic and rely on local surface evidence, these methods carry no topological guarantee and can be wrong where that evidence is unreliable, as at a junction, so the output may contain holes or spurious connections. Other variants instead base the decision on an explicit construction from distances or gradients of an ideal UDF (Kohlbrenner & Alexa, 2026; Wang et al., 2026a;b), which a neural UDF only approximates.

Connectivity from Offset Surfaces. Offset-based methods decide connectivity away from the zero level set: they extract a small positive level set as a manifold double cover of the target, deform it toward the surface, and cut a single layer out of the cover (Hou et al., 2023; Chen et al., 2025b). Gradients thus guide a deformation rather than decide connectivity, so these methods are less sensitive to gradient errors. Their result, however, inherits the connectivity of this cover: for non-manifold or non-orientable targets no valid single layer exists, so the deformed cover is returned as duplicated geometry. The layers bend around a junction and leave a gap instead of a shared connection.

Our Positioning. In summary, when applied to neural UDFs, existing methods either derive connectivity from surface evidence that degrades at a junction, or inherit connectivity from an offset surface whose layers leave a gap instead of a shared connection. OffsetAxis (Huang et al., 2026),

the most recent of the latter, reconstructs the medial axis of the offset volume but connects it heuristically, without any topological guarantee. Our VoroFacet keeps the robust double cover but selects facets from an ambient Voronoi diagram with the proxy, yielding a single layer in which sheets meeting at a junction share Voronoi edges by construction rather than by heuristic connection.

3 PRELIMINARIES

VoroFacet selects facets from the Voronoi diagram of ball-induced sites, which we recall below.

Voronoi Diagrams. Let $\mathcal{P} \subset \mathbb{R}^3$ be a finite set of sites. For a site $\mathbf{p}_a \in \mathcal{P}$, its Voronoi cell is

$$V_a = \left\{ \mathbf{x} \in \mathbb{R}^3 : \|\mathbf{x} - \mathbf{p}_a\| \leq \|\mathbf{x} - \mathbf{p}_b\|, \quad \forall \mathbf{p}_b \in \mathcal{P} \setminus \{\mathbf{p}_a\} \right\}. \tag{1}$$

Each inequality defines the half-space of points closer to $\mathbf{p}_a$ than to $\mathbf{p}_b$, bounded by the bisector plane of $\mathbf{p}_a$ and $\mathbf{p}_b$; thus V_a is the intersection of these half-spaces, hence convex. The collection of all such cells is the Voronoi diagram of $\mathcal{P}$. A Voronoi facet shared by V_a and V_b is the portion of the bisector plane of $\mathbf{p}_a$ and $\mathbf{p}_b$ clipped by the other half-spaces; we denote it by F_{ab}.

Ball-induced Voronoi Sites. VoroCrust (Abdelkader et al., 2020) can convert *mesh surface* into a volumetric Voronoi mesh, as shown in Fig. 1. Since arbitrarily placing the sites near the surface does not guarantee that the boundary of the mesh is aligned with the input surface (Fig. 1, middle), it generates sites from a ball cover with deliberate cosphericity (Fig. 1, right). The surface is first partitioned into smooth patches separated by sharp features and non-manifold junctions. Then, surface balls are centered within patches and kept disjoint from those in other patches, whereas feature balls are centered on patch boundaries and connect the surface balls of adjacent patches.

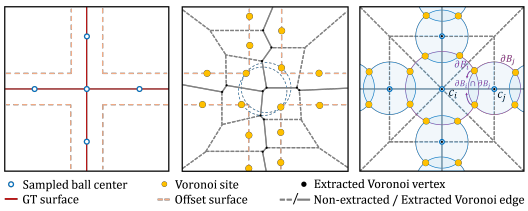


Figure 1: Site placement controls the connectivity of the Voronoi edges at a non-manifold junction, shown here in 2D. Left: the input surface. Middle: arbitrarily placed sites, which leave the junction without a shared vertex and add spurious edges in smooth regions. Right: ball-induced sites, which make the edges incident to the junction share a vertex.

For the resulting ball cover $\mathcal{B} = \{B_i = B(\mathbf{c}_i, r_i)\}$ with union $U(\mathcal{B}) = \bigcup_i B_i$, the Voronoi sites are the exposed triple-sphere intersections of the balls in the cover:

$$\mathcal{P}(\mathcal{B}) = \{\mathbf{p} \in \partial B_i \cap \partial B_j \cap \partial B_k : i < j < k,\ \mathbf{p} \notin \mathrm{int}(B_\ell), \quad \forall B_\ell \in \mathcal{B}\}. \tag{2}$$

The sites on ∂B_i are the closest sites to $\mathbf{c}_i$, so their cells share $\mathbf{c}_i$. If four of these sites are noncoplanar, $\mathbf{c}_i$ is a Voronoi vertex. Similarly, the sites on the intersection circle $\partial B_i \cap \partial B_j$ are the nearest sites to both centers, so their cells contain both centers and, by convexity, the segment $\overline{\mathbf{c}_i\mathbf{c}_j}$. If the circle contains at least three sites and both centers are Voronoi vertices, this segment is a shared Voronoi edge. Thus, the ball construction suppresses spurious facets in smooth regions and enforces shared Voronoi vertices or edges along sharp features and junctions, as visualized in Fig. 1.

It is worth noting that VoroCrust relies on an explicit input mesh both to build the ball cover and to locate the sharp features and non-manifold junctions, whereas a neural UDF provides no explicit surface.

4 PROPOSED METHOD

Given a neural UDF f_θ approximating the distance to an unknown target surface $\bar{\mathcal{S}}$, our goal is to recover a single-layer triangular mesh $\mathcal{M}$, with emphasis on the connectivity of surface sheets at non-manifold junctions. Since the field provides no explicit surface, we start from the near-surface double-layered proxy $\mathcal{D}$ of DCUDF2 (Chen et al., 2025b), which is geometrically faithful but topologically incorrect; every sheet is duplicated, and every junction is buried in the duplicate-layer gap. We therefore treat $\mathcal{D}$ only as a geometric prior and rebuild connectivity from the ball-induced Voronoi sites.

Our pipeline, namely VoroFacet, proceeds in two passes, as illustrated in Fig. 2. In the first pass, we construct an adaptive ball cover of $\mathcal{D}$, with balls sized to fill the duplicate-layer gap. From this

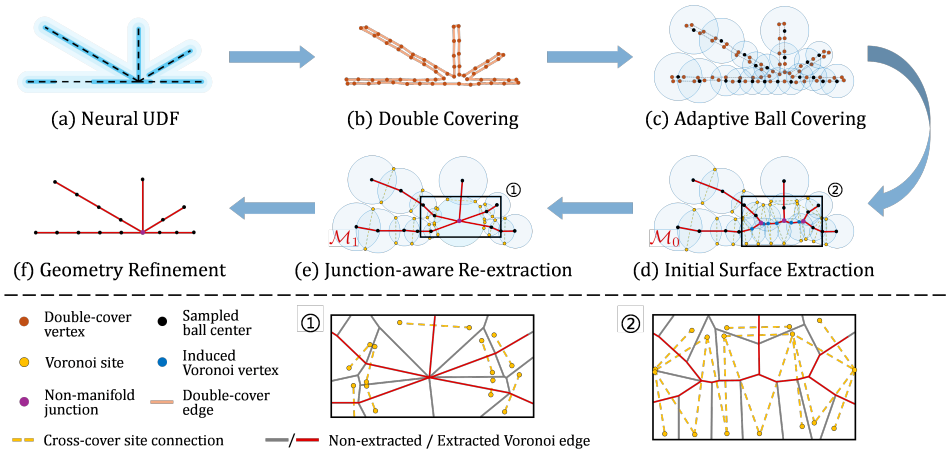


Figure 2: Illustration of the algorithmic pipeline of our VoroFacet on a 2D toy model with boundaries and non-manifold junctions.

cover, Voronoi sites are generated using the exposed triple-sphere intersection rule, and a geometric facet-selection rule is applied to produce a single-layer intermediate mesh $\mathcal{M}_0$. In the second pass, candidate junction curves of $\mathcal{M}_0$ are consolidated into a junction graph $\mathcal{G}$; junction balls are then placed along $\mathcal{G}$, and a finer manifold ball cover is added within surface patches. From the refined cover, sites are again generated via the same rule, and the resulting sites around junctions are deliberately cospherical, inducing shared Voronoi vertices or edges. Applying the same facet-selection rule yields $\mathcal{M}_1$. Thus, the two passes realize geometric observation, junction estimation, and connectivity construction, while the subsequent constrained remeshing and three-stage refinement carry out geometric refinement.

4.1 Initial Surface Extraction

Adaptive Ball Cover. Given the double-layered proxy $\mathcal{D}$, our aim is to construct a ball cover that fills the duplicate-layer gap, especially the wider gap near junctions, while keeping distinct sheets separate. We first set a baseline radius R^0_{gap}, the minimum admissible radius that ensures the desired gap coverage. However, using this radius uniformly would cover flat regions with many unnecessarily small balls, making the cover inefficient. We therefore initialize every ball with a larger radius R_{init} and shrink it only where it would bridge a gap between distinct sheets, never below R^0_{gap}.

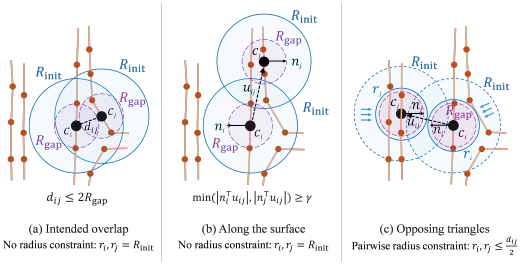


Figure 3: Adaptive ball cover construction. The balls are left unconstrained when the space between the two triangles has to be covered, which holds both for a pair closer than $2R^0_{\text{gap}}$ (left) and for a pair that extends along the surface (center). For an opposing pair, that space has to remain uncovered, so the radii are capped at half their distance (right).

VoroCrust prevents such bridging through co-smoothness, which allows balls to overlap only within smoothly connected surface patches. Unfortunately, the duplicated layers of $\mathcal{D}$ are disconnected, making this connectivity-based criterion inapplicable here. We instead introduce an explicit opposing-triangle test to prevent balls from bridging gaps between distinct sheets.

Specifically, for each triangle T_i of $\mathcal{D}$, we place a candidate ball $B_i = B(\mathbf{c}_i, r_i)$ centered at its centroid $\mathbf{c}_i$. For two nonadjacent triangles T_i and T_j, define $d_{ij} = |\mathbf{c}_i - \mathbf{c}_j|$ and $\mathbf{u}_{ij} = \frac{\mathbf{c}_i - \mathbf{c}_j}{d_{ij}}$. We declare T_i and T_j *opposing* when

$$2R^0_{\text{gap}} < d_{ij} < 2R_{\text{init}} \quad \text{and} \quad \min\left(|\mathbf{n}_i^\top \mathbf{u}_{ij}|, |\mathbf{n}_j^\top \mathbf{u}_{ij}|\right) \geq \gamma. \tag{3}$$

Since all radii start at R_{init}, balls may overlap. Overlap is allowed when the gap between the two triangles has to be filled, and when they extend along the surface rather than facing each other (Fig. 3, left and center). Only opposing pairs, which face each other across a wider separation, are constrained to $r_i, r_j \leq \frac{d_{ij}}{2}$, so that their balls do not bridge the gap between distinct sheets (Fig. 3, right).

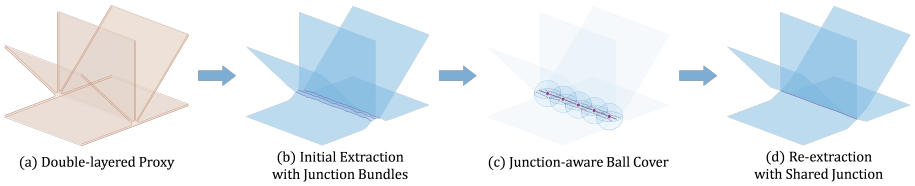


Figure 4: Illustration of junction-aware re-extraction. (a) The double-layered proxy leaves a gap around the non-manifold junction. (b) Initial extraction produces a bundle of nearby junction curves. (c) The curves are consolidated into a junction graph, along which junction balls are placed. (d) Re-extraction joins the incident sheets along a single shared junction curve.

To ensure coverage, let ρ_i denote the radius required for a ball centered at $\mathbf{c}_i$ to cover its seed triangle T_i with a small margin. If $r_i < \rho_i$, we split T_i at its edge midpoints into four child triangles and recompute their balls; the required coverage radius is halved after each subdivision. Since $r_i \geq R^0_{\text{gap}}$, each triangle is subdivided at most $O\left(\log\left(1 + \rho_i / R^0_{\text{gap}}\right)\right)$ times. We then greedily select balls in decreasing order of radius, skipping those whose seed triangles are already covered, yielding $\mathcal{B}_0$ with $\mathcal{D} \subseteq U(\mathcal{B}_0)$.

Surface Extraction. For a ball cover $\mathcal{B}$ of $\mathcal{D}$, we construct the Voronoi diagram of the exposed sites via Eq. (2) and select each bounded facet F_{ab} that satisfies

$$\text{vert}(F_{ab}) \subseteq U(\mathcal{B}) \quad \text{and} \quad \overline{\mathbf{p}_a\mathbf{p}_b} \cap \mathcal{D} \neq \emptyset. \tag{4}$$

This purely local test requires no global inside/outside labels and thus applies to open surfaces and non-manifold junctions. Applying this selection to $\mathcal{B}_0$ yields the initial single-layer mesh $\mathcal{M}_0$, whose non-manifold edges form candidate junction curves. Since several nearby curves may correspond to the same junction, we consolidate them in the second pass.

4.2 Junction-Aware Re-extraction

Junction Graph Estimation. Edges of $\mathcal{M}_0$ incident to more than two facets form candidate junction curves. Since the first pass lacks junction positions, a junction on $\bar{\mathcal{S}}$ where four or more sheets meet may appear in $\mathcal{M}_0$ as a bundle of nearby but disjoint curves (Fig. 4b). We densely sample these curves, then greedily select cluster seeds with spacing δ_G larger than the typical bundle width, so that seeds align along the curves rather than across them. Assigning each sample to its nearest seed and connecting clusters that contain adjacent samples gives the junction graph $\mathcal{G}$, with cluster centroids as vertices (Fig. 4c).

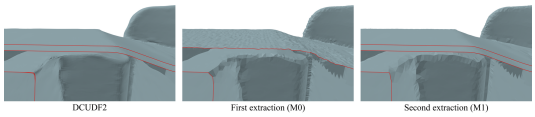


Figure 5: Two-pass extraction of closely spaced sheets. The first extraction (middle) merges two nearby sheets present in the DCUDF2 proxy (left). Re-extraction (right) restores their separation using a smaller R^1_{gap} and the adaptive cover in Fig. 3.

Junction-aware Cover and Re-extraction. We then place overlapping junction balls $\mathcal{B}_G$ of radius R_G along $\mathcal{G}$, sized to fill the junction gaps and isolate adjacent sheets. Let $U_G = U(\mathcal{B}_G)$. We clip $\mathcal{D}$ to the region outside U_G, cutting triangles that cross the junction balls. The remaining surface has only thinner duplicate-layer gaps. We therefore repeat the cover-and-extraction step of Section 4.1 on the clipped triangles, with only two changes: the balls are placed on the remaining triangles outside U_G, and the minimum admissible radius is reduced from R^0_{gap} to R^1_{gap}. This lower threshold marks more pairs as opposing, thereby restoring gaps between distinct sheets that are wider than $2R^1_{\text{gap}}$ but were bridged by the first cover (Fig. 5).

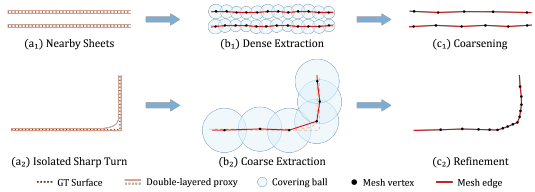


Figure 6: 2D illustration of geometry-adaptive remeshing. Top: two closely spaced flat sheets lead to an unnecessarily dense initial extraction, which remeshing coarsens. Bottom: an isolated sharp turn is extracted too coarsely, so remeshing adds vertices around it.

We then generate exposed sites from the refined cover $\mathcal{B}_1$ via Eq. (2) and apply the facet selection, i.e., Eq. (4), to obtain $\mathcal{M}_1$, a single-layer, self-intersection-free mesh whose non-manifold incidence is dictated by the junction balls (Fig. 4d).

4.3 Remeshing and Geometry Refinement

Geometry-adaptive Remeshing. The triangle density of $\mathcal{M}_1$ follows inter-sheet spacing rather than local geometry, over-sampling closely spaced flat sheets while leaving detailed regions coarse; this limits the accuracy of subsequent fitting (Fig. 6b). We therefore replace it with a geometry-adaptive triangulation. Since isotropic remeshing requires manifold inputs, we cut $\mathcal{M}_1$ along its non-manifold curves and remesh each sheet independently, resampling non-manifold curves and open boundaries at the target resolution as constraints. Using CGAL (Pouget & Cazals, 2024; Coeurjolly et al., 2024), we estimate curvature, derive target edge lengths, and perform constrained isotropic remeshing (Fig. 6c); the sheets are then reassembled through their shared constrained curves.

Three-stage Geometry Refinement. Remeshing improves sampling but not vertex positions: non-manifold vertices remain at the cluster centroids of the candidate junction curves in $\mathcal{M}_0$. These centroids are derived from point clusters rather than from the intersections of the incident sheet extensions, so they may deviate substantially from the true junction. Denser sampling near junctions also places adjacent manifold vertices close to these incorrect positions. Jointly minimizing the UDF and Laplacian would still pull each non-manifold vertex toward a local centroid of its neighbors, not the true junction. We therefore decouple refinement into three stages: (I) fit the manifold sheets away from junctions; (II) reposition the non-manifold vertices from the fitted sheets; (III) refine the manifold surface with the non-manifold vertices fixed (Fig. 7).

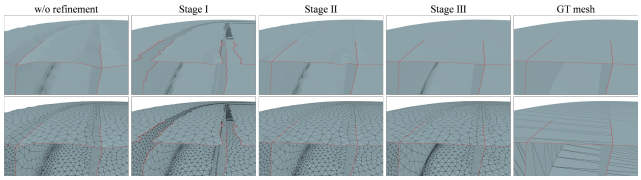


Figure 7: Visual effect of the three-stage geometry refinement at a non-manifold junction. From left to right: before refinement; Stage I, fitting manifold sheets; Stage II, repositioning non-manifold vertices; Stage III, refining the manifold surface with non-manifold vertices fixed; and the ground truth. Shaded and wireframe views are shown in the top and bottom rows, respectively.

To stabilize the fitting and keep the vertices uniformly distributed, we use the Laplacian $(\Delta \mathbf{x})_v = \frac{1}{|\mathcal{N}(v)|}\sum_{u\in\mathcal{N}(v)}(\mathbf{x}_u - \mathbf{x}_v)$, where $\mathcal{N}(v)$ is the one-ring neighbor set for manifold vertices and the adjacent curve-neighbor set for non-manifold vertices, the latter denoted Δ_{curve}.

Stage I: Manifold-sheet fitting. First, we remove the non-manifold vertices, denoted $\mathcal{V}_N$, and their incident faces, then optimize the remaining manifold sheets $\mathcal{M}^- = (\mathcal{V}^-, \mathcal{E}^-, \mathcal{F}^-)$ by

$$E_{\text{sheet}} = \lambda_u \sum_{v\in\mathcal{V}^-} \left|f_\theta(\mathbf{x}_v)\right| + \lambda_e \sum_{\overline{uv}\in\mathcal{E}^-} \left(\frac{\|\mathbf{x}_u - \mathbf{x}_v\| - \ell^0_{uv}}{\ell^0_{uv}}\right)^2, \tag{5}$$

where ℓ^0_{uv} is the post-remeshing edge length. The UDF term pulls the sheets onto $\bar{\mathcal{S}}$, while the edge-length term prevents neighboring vertices from collapsing into the same local UDF minimum.

Stage II: Non-manifold vertex fitting. Next, we restore the non-manifold vertices $\mathcal{V}_N$ and their incident faces, keeping all manifold vertices fixed. For each $v \in \mathcal{V}_N$, we collect a set $\mathcal{F}^q_v$ of nearby fitted faces from each incident sheet. Assuming local planarity near the junction, each face $f \in \mathcal{F}^q_v$ defines a plane through its centroid $\mathbf{c}_f$ with normal $\mathbf{n}_f$, and the junction vertex should lie on all such planes. This leads to

$$E_{\text{nm}} = \frac{\lambda_q}{N_q}\sum_{v\in\mathcal{V}_N}\sum_{f\in\mathcal{F}^q_v}\left[\mathbf{n}_f^\top(\mathbf{x}_v - \mathbf{c}_f)\right]^2 + \frac{\lambda_c}{|\mathcal{V}_N|}\sum_{v\in\mathcal{V}_N}\|\Delta_{\text{curve}}\mathbf{x}_v\|^2 + \frac{\lambda_a}{|\mathcal{V}_N|}\sum_{v\in\mathcal{V}_N}\|\mathbf{x}_v - \mathbf{x}^0_v\|^2, \tag{6}$$

with $N_q = \sum_{v\in\mathcal{V}_N}|\mathcal{F}^q_v|$ and $\mathbf{x}^0_v$ the position of v before this stage. The curve Laplacian smooths the junction curve, and the weak anchor limits drift where the plane constraints are degenerate.

Stage III: Manifold refinement. Finally, with the non-manifold and boundary vertices fixed, we refine the manifold surface following DCUDF2 (Chen et al., 2025b). Manifold vertices within three

graph rings of the non-manifold curves, denoted $\mathcal{V}_B$, receive only Laplacian regularization, whereas the remaining vertices $\mathcal{V}_F$ and the faces $\mathcal{F}_F$ whose vertices all lie in $\mathcal{V}_F$ also minimize their UDF values:

$$E_{\text{ref}} = \lambda_l \sum_{v \in \mathcal{V}_B \cup \mathcal{V}_F} \omega_v \|\Delta \mathbf{x}_v\|^2 + \lambda_u \left(\sum_{v \in \mathcal{V}_F} \alpha_v \left| f_\theta(\mathbf{x}_v) \right| + \sum_{f \in \mathcal{F}_F} \alpha_f \left| f_\theta(\mathbf{c}_f) \right| \right), \tag{7}$$

where ω_v is the Laplacian weight and α_v, α_f are per-vertex and per-face weights. The output thus retains the connectivity of $\mathcal{M}_1$ while aligning both manifold and non-manifold vertices with $\bar{\mathcal{S}}$.

5 Experiments

5.1 Experiment Settings

Baselines. We compare nine UDF mesh extractors from the three families discussed in Sec. 2. MeshUDF (Guillard et al., 2022), CAP-UDF (Zhou et al., 2024), GeoUDF (Ren et al., 2023), HR-UDF (Stella et al., 2025), and MIND (Chen et al., 2025a) infer grid-corner signs or labels and use Marching Cubes lookup tables to determine face connectivity within each cell. DualMesh-UDF (Zhang et al., 2023) and DCx (Bao et al., 2026) place one vertex in each cell containing the surface and connect vertices in neighboring cells to form faces. DCUDF (Hou et al., 2023) and DCUDF2 (Chen et al., 2025b) establish connectivity by applying Marching Cubes to a positive UDF level set, then deform the resulting mesh toward the target surface.

Implementation Details. All experiments run on an AMD EPYC 7542 system with a single NVIDIA RTX A6000 GPU (48 GB). We optimize vertex positions with VectorAdam (Ling et al., 2022) and use a default extraction resolution of 256^3. We obtain each UDF using a state-of-the-art learning method (Yang et al., 2025), trained on 100k unoriented surface points per object or 1,000 points/m^2 for scenes. All extractors query the same UDF; for CAP-UDF and GeoUDF, we use only their extraction modules. DCUDF and DCUDF2 apply minimum-cut post-processing only in the manifold-object experiments. Before evaluation, we normalize each reference mesh to a unit box and apply the same transformation to its reconstruction.

Hyperparameters. All length parameters are specified in normalized neural-field coordinates. We initialize $\mathcal{D}$ using the corresponding DCUDF2 baseline configuration, without minimum-cut separation. We set $(R_{\text{init}}, R^0_{\text{gap}}, R^1_{\text{gap}}, R_G) = (0.01, 0.0075, 0.003, 0.0125)$, $\gamma = 0.5$, and $\delta_G = 0.015$. Adaptive remeshing uses target edge lengths in $[0.003, 0.01]$. The three refinement stages use $(200, 200, 120)$ iterations and a base learning rate of 5×10^{-4}. The UDF and edge-length weights in Stage I are $(1000, 1000)$; the QEF, curve-Laplacian, and anchoring weights in Stage II are $(1000, 1600, 100)$. In Stage III, both UDF terms have unit weight, and both Laplacian terms have weight 2000. These extraction and refinement parameters are shared across datasets.

Evaluation Metrics. For surface reconstruction we report Chamfer distance (CD), F1-score (threshold 0.004), and normal consistency (NC) over 100k area-uniform samples per mesh, plus the number of connected components (Comp). We evaluate non-manifold (NM) junction reconstruction using NM-CD, NM-F1, valence-aware NM-F1 (VA-NM-F1), and NM-LR. Unlike NM-F1, which evaluates junction proximity, VA-NM-F1 additionally requires agreement in the number of incident sheets. Details of ground-truth junction extraction and metric computation are provided in Appendix A. CD and NM-CD are scaled by 10^3, and F1, NC, NM-F1, and VA-NM-F1 are percentages. We report mean $\pm$ standard deviation across objects or scenes.

5.2 Results and Comparisons

To assess VoroFacet across diverse geometries, we evaluate non-manifold objects, complex interior structures, 3D scenes, manifold objects, and open-boundary objects, with object subsets randomly sampled from their respective datasets. Quantitative results are reported in Tables 1–3 and Appendix D (Tables 10 and 11).

Evaluation on Non-manifold Objects. We evaluate on 50 shapes from the non-manifold ABC dataset introduced by SALS (Ren & Hou, 2025), constructed by intersecting pairs of ABC shapes (Koch et al., 2019). VoroFacet achieves the best surface CD and F1, with a larger advantage

Table 1: Quantitative results on **Non-manifold Objects**. **Bold** / underline: best / second-best across all methods. —: no explicit NM edges on any model.

Method	CD ↓	F1 ↑	NC ↑	NM-CD ↓	NM-F1 ↑	VA-NM-F1 ↑	NM-LR → 1	Comp	Time (s)
MeshUDF [ECCV'22]	5.727±0.894	78.97±8.99	97.12±1.77	69.58±77.90	26.26±13.59	7.66±5.39	**0.151**±0.217	241.9±318.4	9.124±1.666
CAP-UDF [TPAMI'24]	5.939±0.734	76.26±8.31	96.00±2.31	—	—	—	—	4157±3384	174.7±19.8
GeoUDF [ICCV'23]	5.580±0.803	80.26±8.95	96.44±3.28	113.9±51.8	21.75±11.37	4.55±3.08	98.772±665.577	346.5±507.0	6.859±0.647
MIND [NeurIPS'25]	5.959±0.752	76.70±7.41	95.14±4.33	45.81±46.73	38.62±19.13	2.33±4.03	3.526±12.001	76.76±113.25	200.6±222.3
HR-UDF [NeurIPS'25]	5.580±0.818	80.30±9.08	97.17±1.31	—	—	—	—	5421±5658	102.9±23.0
DualMesh-UDF [ICCV'23]	5.632±0.814	79.77±9.05	96.75±1.52	130.3±50.6	8.715±3.960	3.82±2.04	1524.460±8869.958	17.06±36.54	12.98±3.01
DCx [TOG'26]	5.645±0.751	79.02±8.07	96.92±2.90	46.44±50.70	53.19±19.26	6.22±8.64	6.629±11.179	8.52±10.34	75.59±38.28
DCUDF [TOG'23]	5.563±0.826	80.18±8.89	**98.07**±0.77	—	—	—	—	20.42±36.30	178.2±56.9
DCUDF2 [TVCG'25]	5.562±0.824	80.24±8.88	98.05±0.76	—	—	—	—	20.42±36.30	153.0±48.0
VoroFacet (Ours)	**5.549**±0.831	**80.46**±8.98	98.02±0.85	**37.74**±48.06	**65.80**±21.37	**37.07**±17.23	2.167±1.741	1.380±1.086	418.0±123.6

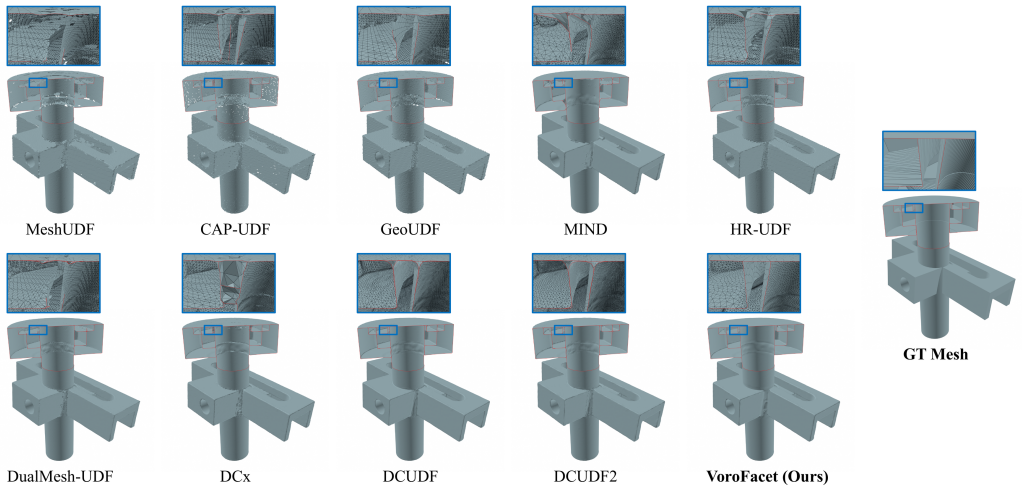


Figure 8: Visual comparison on a **non-manifold** object. Close-ups highlight junction connectivity: our VoroFacet joins incident sheets along shared junction curves, while methods under comparison leave holes, introduce spurious connections, or retain duplicated layers that leave gaps at junctions. *Zoom in for details.*

in junction reconstruction (Table 1). Its NM-F1 exceeds the strongest baseline by 12.61 percentage points, while its VA-NM-F1 reaches 37.07, compared with at most 7.66 among the baselines. This indicates better recovery of junctions with the correct incident-sheet counts, beyond geometric proximity alone. Fig. 8 shows that DCUDF and DCUDF2 bend separate manifold sheets around the junction, forming a gap rather than a shared intersection. The other baselines exhibit holes or spurious connections near junctions. VoroFacet instead joins incident sheets along shared junction curves while maintaining high geometric accuracy.

Evaluation on Complex Interior Structures. ShapeNet-Car models (Chang et al., 2015) contain many thin sheets and non-manifold junctions in their interiors; we evaluate on 20 models, retaining all interior faces. VoroFacet keeps competitive surface accuracy while achieving the best NM-CD, NM-F1, and VA-NM-F1, with an NM-LR of 0.982 (Table 3). Its VA-NM-F1 reaches 18.91, compared with 15.09 for the strongest baseline, indicating better recovery of junctions with the correct incident-sheet counts in these complex interiors. Fig. 9 shows that VoroFacet resolves closely spaced sheets and their junctions, whereas the baselines merge distinct sheets or introduce holes, spurious connections, or duplicated layers.

Table 2: Quantitative results on **Manifold Objects**. **Bold** / underline: best / second-best.

Method	CD ↓	F1 ↑	NC ↑	Comp	Time (s)
MeshUDF [ECCV'22]	4.608±1.128	89.26±9.33	97.89±2.21	104.88±305.40	7.67±6.10
CAP-UDF [TPAMI'24]	4.615±1.043	89.22±8.94	97.63±2.56	2415.96±3506.15	167.62±60.73
GeoUDF [ICCV'23]	4.518±1.012	**90.34**±8.68	97.69±2.26	65.4±156.6	6.39±2.27
MIND [NeurIPS'25]	4.572±1.032	89.90±8.81	97.81±2.42	6.78±25.44	121.17±395.95
HR-UDF [NeurIPS'25]	4.530±1.019	90.24±8.79	97.77±2.34	3190.04±2813.09	78.22±22.98
DualMesh-UDF [ICCV'23]	4.569±1.040	89.57±8.95	97.72±2.25	1.86±3.28	9.45±9.10
DCx [TOG'26]	4.539±0.996	90.00±8.50	97.95±2.32	6.52±14.89	73.54±281.25
DCUDF [TOG'23]	4.529±1.032	90.08±8.71	**98.36**±1.67	1.08±0.27	303.04±830.70
DCUDF2 [TVCG'25]	4.528±1.031	90.15±8.73	98.32±1.68	1.06±0.24	216.11±441.05
VoroFacet (Ours)	**4.515**±1.012	90.33±8.65	98.24±1.88	1.28±0.95	306.59±283.42

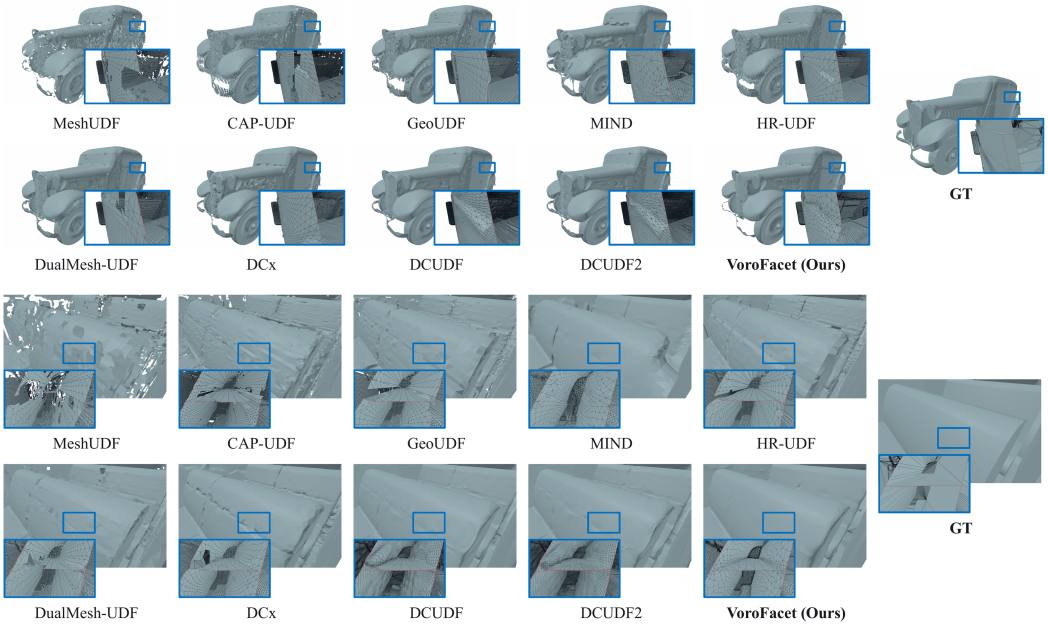


Figure 9: Visual comparison on a **ShapeNet-Car** model, showing its exterior (top) and interior (bottom). Close-ups show that VoroFacet keeps nearby but distinct sheets separate and joins incident sheets at junctions, whereas methods under comparison merge distinct sheets or fail to recover the correct junction connectivity. *Zoom in for details.*

Table 3: Quantitative results on **ShapeNet-Car**. **Bold** / underline: best / second-best across all methods. —: no explicit NM edges on any model.

Method	CD ↓	F1 ↑	NC ↑	NM-CD ↓	NM-F1 ↑	VA-NM-F1 ↑	NM-LR → 1	Comp	Time (s)
MeshUDF [ECCV'22]	6.073±0.968	76.51±8.62	91.71±3.22	52.19±15.75	6.354±2.188	2.84±1.12	0.021±0.008	1154±994	12.60±6.83
CAP-UDF [TPAMI'24]	5.784±0.791	78.55±8.54	89.15±3.40	—	—	—	—	13342±8304	267.5±53.1
GeoUDF [ICCV'23]	5.299±0.701	83.75±7.23	91.00±2.96	37.86±8.79	22.02±4.48	11.20±3.30	0.539±0.256	792.5±626.7	8.131±2.477
MIND [NeurIPS'25]	6.393±0.813	75.52±7.37	87.17±3.42	59.51±17.64	8.881±5.194	4.78±4.29	0.323±0.171	126.1±74.5	299.6±206.0
HR-UDF [NeurIPS'25]	5.309±0.718	**84.15**±7.47	90.75±3.16	—	—	—	—	3432±1996	157.4±30.2
DualMesh-UDF [ICCV'23]	5.388±0.739	83.05±7.37	89.71±3.66	31.19±8.22	22.47±4.72	12.71±3.56	21.111±10.808	51.7±44.7	21.38±11.83
DCx [TOG'26]	5.605±0.701	81.30±6.88	88.64±3.70	25.47±9.06	30.80±7.19	15.09±7.22	1.696±0.678	52.75±44.60	361.7±321.5
DCUDF [TOG'23]	5.304±0.690	83.48±7.04	**92.27**±2.56	—	—	—	—	185.7±151.2	320.4±75.2
DCUDF2 [TVCG'25]	**5.291**±0.695	83.74±7.15	92.01±2.53	—	—	—	—	185.9±151.4	277.5±53.4
VoroFacet (Ours)	5.370±0.687	83.61±7.05	91.29±2.75	**21.79**±6.75	**36.67**±8.44	**18.91**±8.50	**0.982**±0.320	13.95±11.90	609.2±119.2

Evaluation on 3D Scenes. To assess generalization from individual objects to scenes, we evaluate the same five real scans from the 3D Scene dataset (Zhou & Koltun, 2013) used in CAP-UDF (Zhou et al., 2024). VoroFacet produces the fewest connected components while keeping competitive geometric accuracy, so each scene is recovered as a few coherent pieces rather than many fragments. We refer readers to Appendix D for the scene results.

Evaluation on Manifold Objects and Open Boundaries. We evaluate on 50 manifold shapes from FAMOUS (Erler et al., 2020) and Thingi32 (Maruani et al., 2023), and on 20 MGN models (Bhatnagar et al., 2019), whose garment surfaces have open boundaries. On both sets, VoroFacet is at or near the best surface accuracy and keeps the number of connected components close to one (Tables 2 and 11). Manifold, open-boundary, and non-manifold structures are therefore all reconstructed within a single pipeline, which produces accurate single-layer meshes without minimum-cut post-processing. We refer readers to Appendix D for the open-boundary results.

5.3 Ablation Studies

VoroFacet constructs the output connectivity through two extraction passes and subsequently refines the geometry. We evaluate these designs on 20 non-manifold objects with four variants: replacing

Table 4: Component ablations on 20 non-manifold objects.

Variant	CD ↓	F1 ↑	NC ↑	NM-CD ↓	NM-F1 ↑	VA-NM-F1 ↑	NM-LR → 1	Time (s)
Fixed minimum radii	**5.608**±0.764	79.78±8.50	98.15±0.83	28.166±28.714	66.43±14.11	35.96±15.14	1.771±0.758	1824.6±613.7
w/o junction-aware re-extraction	5.644±0.741	79.17±7.99	98.05±0.87	25.886±24.753	**75.67**±11.39	6.09±15.04	1.916±0.510	372.3±115.4
w/o remeshing	5.631±0.752	79.71±8.44	97.96±0.91	24.639±28.889	72.72±15.74	37.05±14.40	1.639±0.691	397.9±107.6
w/o three-stage refinement	5.630±0.758	79.72±8.52	98.01±0.89	25.855±28.647	63.44±16.40	36.39±15.39	1.696±0.698	387.0±112.6
Full method	5.609±0.767	**79.80**±8.51	**98.16**±0.82	**23.331**±27.554	74.79±14.04	**38.50**±15.88	**1.615**±0.643	428.2±115.2

the adaptive ball cover with fixed minimum radii, and removing the junction-aware re-extraction, the geometry-adaptive remeshing, or the three-stage refinement.

The three-stage refinement matters most for the junction geometry. Without it (Fig. 7), the non-manifold vertices stay at the cluster centroids, so the junction curves are kinked, deviate from the intersection of the incident sheets, and fold the adjacent surfaces; with it, each junction is a single curve on that intersection, matching the ground truth. We refer readers to Appendix C for the other variants.

Table 4 is consistent with these observations. The adaptive ball cover primarily improves efficiency: with fixed minimum radii, flat regions are covered by many unnecessarily small balls, which increases the runtime by $4.3\times$ without improving accuracy. Removing the re-extraction leaves nearby curves in $\mathcal{M}_0$ that represent the same junction unmerged, increasing NM-LR from 1.615 to 1.916. These curves can still fall within the matching threshold of the ground truth, giving this variant the highest NM-F1. However, VA-NM-F1 drops from 38.50 to 6.09, revealing incorrect incident-sheet counts despite geometric proximity. Removing the remeshing slightly degrades the overall accuracy, since the triangle density of $\mathcal{M}_1$ follows the inter-sheet spacing rather than the local geometry; its effect is more apparent around sharp features than in the averaged metrics. Removing the three-stage refinement causes the largest drop in NM-F1, whereas the surface metrics decrease only marginally. The extracted sheets already lie close to the target surface, while the non-manifold vertices stay inaccurate until Stage II. Appendix F breaks down the runtime of the full method, and Appendix E further analyzes the sensitivity to the main parameters and the extraction resolution. Additional robustness experiments and failure cases are reported in Appendix B.

6 Conclusion and Discussion

We have presented VoroFacet, a novel pipeline that extracts single-layer meshes from neural UDFs while handling complex structures including non-manifold junctions. VoroFacet treats a near-surface double-layered proxy of the target as a geometric prior and derives the output connectivity from the Voronoi diagram of the sites generated by an adaptive ball cover of that proxy, producing a single layer rather than duplicated geometry. Then a junction-aware re-extraction joins the sheets that meet at a junction, and remeshing and a three-stage refinement refine the geometry of the resulting mesh. Experiments on complex shapes show that VoroFacet recovers junctions that existing methods leave as holes, spurious connections, or gaps, while maintaining competitive surface accuracy. More broadly, these results suggest that a geometrically faithful proxy can guide surface extraction even when its topology is incorrect, and that the connectivity of the incident sheets can instead be constructed from a Voronoi partition.

Three limitations remain and suggest directions for future work. First, the ball radii are determined by the gaps of the double-layered proxy, i.e., the duplicate-layer gap on smooth sheets and the wider gap at junctions. If two thin sheets are separated by a distance comparable to or smaller than these gaps, the balls cover the space between them, so no sites are generated there and the two sheets are extracted as a single layer. A double cover that stays closer to the surface, or a more robust spatial partition in its place, would allow smaller ball radii and thus support thinner structures. Second, our method is not designed to preserve sharp features, since the Laplacian-based refinement smooths the normal discontinuity across a sharp edge. Such an edge could instead be constructed by feature balls placed along it, as junction balls do at non-manifold junctions, but this requires detecting the sharp edges from the UDF and preserving them during remeshing and refinement. Third, both extraction passes yield self-intersection-free meshes, but the three-stage refinement can introduce a few self-intersections. It remains open whether constrained or intersection-aware optimization can keep the output self-intersection-free without losing geometric accuracy.

AI USE STATEMENT

In preparing this manuscript, we used OpenAI's ChatGPT-5.6 solely to assist with improving the clarity and quality of the writing. Generative AI tools were not used to generate, modify, or preprocess experimental data. All AI-assisted text was critically reviewed, verified, and revised by the authors as appropriate. The authors take full responsibility for the final content of the manuscript, including its text, technical claims, and mathematical results.

ETHICS STATEMENT

This work does not involve human participants, personal or sensitive data, or automated decisions concerning individuals. We are not aware of ethical risks specific to the proposed method beyond those generally associated with mesh extraction and geometric processing.

REFERENCES

Ahmed Abdelkader, Chandrajit L Bajaj, Mohamed S Ebeida, Ahmed H Mahmoud, Scott A Mitchell, John D Owens, and Ahmad A Rushdi. Vorocrust: Voronoi meshing without clipping. *ACM Transactions on Graphics (TOG)*, 39(3):1–16, 2020.

Qingchao Bao, Xuhui Chen, Jingpeng Yin, Fei Hou, Wencheng Wang, Hong Qin, and Ying He. Dual contouring over expanded cubes (dcx) for zero-level set extraction from neural unsigned distance functions. *ACM Transactions on Graphics (TOG)*, 45(4):1–20, 2026.

Bharat Lal Bhatnagar, Garvita Tiwari, Christian Theobalt, and Gerard Pons-Moll. Multi-garment net: Learning to dress 3d people from images. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 5420–5430, 2019.

Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.

Xuhui Chen, Fei Hou, Wencheng Wang, Hong Qin, and Ying He. Mind: Material interface generation from udfs for non-manifold surface reconstruction. *Advances in Neural Information Processing Systems*, 38:120283–120306, 2025a.

Xuhui Chen, Fugang Yu, Fei Hou, Wencheng Wang, Zhebin Zhang, and Ying He. Dcudf2: Improving efficiency and accuracy in extracting zero level sets from unsigned distance fields. *IEEE Transactions on Visualization and Computer Graphics*, 31(10):9052–9065, 2025b.

Zhiqin Chen, Andrea Tagliasacchi, Thomas Funkhouser, and Hao Zhang. Neural dual contouring. *ACM Transactions on Graphics (TOG)*, 41(4):1–13, 2022.

Julian Chibane, Gerard Pons-Moll, et al. Neural unsigned distance fields for implicit function learning. *Advances in Neural Information Processing Systems*, 33:21638–21652, 2020.

David Coeurjolly, Jaques-Olivier Lachaud, Konstantinos Katrioplas, Sébastien Loriot, Ivan Paden, Mael Rouxel-Labbé, Hossam Saeed, Jane Tournois, and Ilker O. Yaz. Polygon mesh processing. In *CGAL User and Reference Manual*. CGAL Editorial Board, 6.0.1 edition, 2024. URL https://doc.cgal.org/6.0.1/Manual/packages.html#PkgPolygonMeshProcessing.

Philipp Erler, Paul Guerrero, Stefan Ohrhallinger, Niloy J Mitra, and Michael Wimmer. Points2surf learning implicit surfaces from point clouds. In *European conference on computer vision*, pp. 108–124. Springer, 2020.

Benoit Guillard, Federico Stella, and Pascal Fua. Meshudf: Fast and differentiable meshing of unsigned distance field networks. In *European conference on computer vision*, pp. 576–592. Springer, 2022.

Fei Hou, Xuhui Chen, Wencheng Wang, Hong Qin, and Ying He. Robust zero level-set extraction from unsigned distance fields based on double covering. *ACM Transactions on Graphics (TOG)*, 42(6):1–15, 2023.

Qijia Huang, Pierre Kraemer, and Dominique Bechmann. Offsetaxis: Udf mesh reconstruction via offset-volume medial axis extraction. *arXiv preprint arXiv:2605.15369*, 2026.

Tao Ju, Frank Losasso, Scott Schaefer, and Joe Warren. Dual contouring of hermite data. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pp. 339–346, 2002.

Sebastian Koch, Albert Matveev, Zhongshi Jiang, Francis Williams, Alexey Artemov, Evgeny Burnaev, Marc Alexa, Denis Zorin, and Daniele Panozzo. Abc: A big cad model dataset for geometric deep learning. In *2019 IEEE/CVF conference on computer vision and pattern Recognition (CVPR)*, pp. 9593–9603. Ieee, 2019.

M Kohlbrenner and M Alexa. Meshing unsigned distance fields with regular triangulations. In *Computer Graphics Forum*, pp. e70524. Wiley Online Library, 2026.

Selena Zihan Ling, Nicholas Sharp, and Alec Jacobson. Vectoradam for rotation equivariant geometry optimization. *Advances in Neural Information Processing Systems*, 35:4111–4122, 2022.

William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Seminal graphics: pioneering efforts that shaped the field*, pp. 347–353. 1998.

Nissim Maruani, Roman Klokov, Maks Ovsjanikov, Pierre Alliez, and Mathieu Desbrun. Voromesh: Learning watertight surface meshes with voronoi diagrams. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 14519–14528. IEEE, 2023.

Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4455–4465. IEEE, 2019.

Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *2019 IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, pp. 165–174. IEEE, 2019.

Marc Pouget and Frédéric Cazals. Estimation of local differential properties of point-sampled surfaces. In *CGAL User and Reference Manual*. CGAL Editorial Board, 6.0.1 edition, 2024. URL https://doc.cgal.org/6.0.1/Manual/packages.html#PkgJetFitting3.

Siyu Ren and Junhui Hou. Shape as line segments: Accurate and flexible implicit surface representation. In *International Conference on Learning Representations*, volume 2025, pp. 82842–82862, 2025.

Siyu Ren, Junhui Hou, Xiaodong Chen, Ying He, and Wenping Wang. Geoudf: Surface reconstruction from 3d point clouds via geometry-guided distance representation. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 14168–14178. IEEE, 2023.

Federico Stella, Nicolas Talabot, Hieu Le, and Pascal Fua. Neural surface detection for unsigned distance fields. In *European Conference on Computer Vision*, pp. 394–409. Springer, 2024.

Federico Stella, Nicolas Talabot, Hieu Le, and Pascal Fua. High resolution udf meshing via iterative networks. In *Advances in Neural Information Processing Systems*, pp. 170903–170933, 2025. doi: 10.52202/085713-5692.

Ningna Wang, Xiana Carrera, Christopher Batty, Oded Stein, and Silvia Sellán. Spudd: Superpower contouring of unsigned distance data. *arXiv preprint arXiv:2604.19568*, 2026a.

Ningna Wang, Zilong Wang, Xiana Carrera, Xiaohu Guo, and Silvia Sellán. Voroudf: Meshing unsigned distance fields with voronoi optimization. *arXiv preprint arXiv:2602.02907*, 2026b.

Chuanxiang Yang, Yuanfeng Zhou, Guangshun Wei, Long Ma, Junhui Hou, Yuan Liu, and Wenping Wang. Monge-ampere regularization for learning arbitrary shapes from point clouds. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(8):6809–6822, 2025.

Congyi Zhang, Guying Lin, Lei Yang, Xin Li, Taku Komura, Scott Schaefer, John Keyser, and Wenping Wang. Surface extraction from neural unsigned distance fields. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 22474–22483. IEEE, 2023.

Junsheng Zhou, Baorui Ma, Shujuan Li, Yu-Shen Liu, Yi Fang, and Zhizhong Han. Cap-udf: Learning unsigned distance functions progressively from raw point clouds with consistency-aware field optimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12):7475–7492, 2024.

Qian-Yi Zhou and Vladlen Koltun. Dense scene reconstruction with points of interest. *ACM Transactions on Graphics (ToG)*, 32(4):1–8, 2013.

A GROUND-TRUTH JUNCTION EXTRACTION AND EVALUATION

Ground-truth junction extraction. The CAD meshes used in our experiments originate from public online platforms and lack explicit junction annotations. Since sheets may intersect or nearly touch without sharing mesh edges, we identify GT junctions using both mesh connectivity and local geometry.

We preserve the input vertex indices and sample all mesh edges and exact triangle-intersection segments. Samples on edges incident to three or more triangles are retained directly; the other candidates are classified by their local surface directions. At each sample, the source segment defines the local axis. We collect triangles within $\varepsilon = 10^{-5}S$, where S is the GT bounding-box longest side, whose planes are within 5° of parallel to this axis. Each triangle contributes directed extensions perpendicular to the axis: one toward its interior if the sample is within ε of a triangle edge, and two opposite directions otherwise. We group these directions using a 3° angular tolerance and retain samples with at least three direction classes. The class count gives the local number of incident sheets. We greedily subsample the retained detector points to 10k samples, using a minimum Euclidean separation adjusted by bisection. Table 5 reports the sensitivity of the junction metrics to the GT extraction thresholds on Non-manifold Objects and ShapeNet-Car. VoroFacet achieves the highest VA-NM-F1 under every tested GT extraction setting on both datasets, indicating that its advantage is not specific to the default thresholds.

Junction evaluation. Predicted junctions consist only of mesh edges incident to at least three faces; each nonempty junction set is sampled at 10k points uniformly by arc length. For objects with nonempty GT junction references, let $\mathcal{P}$ and $\mathcal{G}$ denote the predicted and GT sample sets after normalization by S, with $N_P = |\mathcal{P}|$ and $N_G = |\mathcal{G}|$. Define $d(x, \mathcal{Y}) = \min_{y\in\mathcal{Y}} \|x - y\|_2$ and use $\tau = 0.004$.

NM-CD and NM-F1. NM-CD measures bidirectional junction distance, reported at a scale of 10^3:

$$\text{NM-CD} = 10^3 \left[\frac{1}{N_P} \sum_{p\in\mathcal{P}} d(p, \mathcal{G}) + \frac{1}{N_G} \sum_{g\in\mathcal{G}} d(g, \mathcal{P}) \right]. \tag{8}$$

NM-F1 is the harmonic mean of precision and recall at threshold τ:

$$P = \frac{1}{N_P} \sum_{p\in\mathcal{P}} \mathbf{1}[d(p, \mathcal{G}) < \tau], \qquad R = \frac{1}{N_G} \sum_{g\in\mathcal{G}} \mathbf{1}[d(g, \mathcal{P}) < \tau],$$
$$\text{NM-F1} = 100\, \frac{2PR}{P + R}. \tag{9}$$

VA-NM-F1. VA-NM-F1 measures junction reconstruction accuracy in both location and valence (the number of incident surface sheets). It extends NM-F1 by requiring matched samples to have the same valence. GT valences are given by the direction counts above; predicted valences are given by the numbers of faces incident to the sampled edges. Let $\mathcal{P}_k$ and $\mathcal{G}_k$ contain samples with valence k. Then

$$P_{\text{VA}} = \frac{1}{N_P} \sum_{k\geq 3} \sum_{p\in\mathcal{P}_k} \mathbf{1}[d(p, \mathcal{G}_k) < \tau],$$
$$R_{\text{VA}} = \frac{1}{N_G} \sum_{k\geq 3} \sum_{g\in\mathcal{G}_k} \mathbf{1}[d(g, \mathcal{P}_k) < \tau], \tag{10}$$
$$\text{VA-NM-F1} = 100\, \frac{2P_{\text{VA}} R_{\text{VA}}}{P_{\text{VA}} + R_{\text{VA}}}.$$

To apply the same matching rule to all valence classes, we define $d(x, \emptyset) = +\infty$. Both F1 scores are zero when their denominators vanish or the prediction is empty; NM-CD is averaged over objects with nonempty predicted junction sets.

NM-LR. NM-LR measures junction redundancy using the predicted-to-GT total length ratio:

$$\text{NM-LR} = \frac{L_P}{L_G}, \tag{11}$$

where L_P sums the lengths of predicted junction edges and L_G is recovered from accepted GT segment intervals with collinear overlaps merged. The ideal value is 1; values above or below 1 indicate excessive or insufficient total junction length, respectively. All metrics are computed per object before reporting their mean and standard deviation.

Table 5: Sensitivity to ground-truth junction extraction thresholds on **Non-manifold Objects** and **ShapeNet-Car**. **Bold**: best mean within each dataset, metric, and setting.

Metric	Method	Default	Distance tolerance		Direction merging		Axis–plane alignment	
			$\varepsilon_0/2$	$2\varepsilon_0$	1°	5°	3°	7°
(a) Non-manifold Objects								
NM-CD ↓	MeshUDF	69.530±77.84	71.145±86.00	69.257±71.99	80.851±81.04	67.502±77.82	69.522±77.84	69.542±77.83
	GeoUDF	113.928±51.81	118.983±57.77	105.833±47.64	111.445±48.19	114.025±51.82	113.935±51.81	113.929±51.81
	MIND	45.954±47.12	45.597±47.17	46.216±45.87	56.867±54.86	44.032±47.00	45.931±47.13	45.954±47.12
	DualMesh-UDF	130.277±50.62	136.777±56.59	119.992±47.35	123.129±50.68	131.023±50.01	130.306±50.59	130.278±50.62
	DCx	46.424±50.80	47.932±55.84	42.080±40.53	57.558±58.13	45.039±51.11	46.413±50.80	46.430±50.80
	VoroFacet	**37.699**±48.19	**39.169**±52.92	**35.799**±44.63	**49.785**±56.92	**36.228**±48.26	**37.689**±48.19	**37.710**±48.19
NM-F1 ↑	MeshUDF	26.25±13.60	26.61±13.49	25.77±13.63	23.15±14.02	26.46±13.60	26.26±13.59	26.25±13.60
	GeoUDF	21.74±11.36	21.78±11.45	21.66±11.28	20.83±11.29	21.81±11.40	21.74±11.37	21.74±11.36
	MIND	38.61±19.14	38.89±19.19	38.18±19.05	35.66±19.42	38.84±19.18	38.62±19.14	38.61±19.14
	DualMesh-UDF	8.71±3.96	8.62±3.96	8.90±4.00	9.19±3.87	8.68±3.96	8.71±3.96	8.72±3.96
	DCx	53.19±19.26	53.56±19.29	52.69±19.40	48.91±20.73	53.47±19.31	53.19±19.26	53.19±19.26
	VoroFacet	**65.80**±21.37	**66.23**±21.31	**65.15**±21.55	**59.32**±24.46	**66.14**±21.33	**65.80**±21.37	**65.80**±21.36
VA-NM-F1 ↑	MeshUDF	7.67±5.38	7.70±5.35	7.62±5.43	6.73±5.40	7.72±5.39	7.69±5.39	7.68±5.38
	GeoUDF	4.55±3.07	4.41±2.99	4.77±3.09	4.31±2.53	4.55±3.06	4.55±3.06	4.55±3.07
	MIND	2.33±4.03	2.22±4.02	2.68±4.04	2.46±3.93	2.37±4.04	2.35±4.03	2.32±4.02
	DualMesh-UDF	3.82±2.04	3.74±2.01	3.98±2.13	4.25±2.23	3.79±2.03	3.82±2.04	3.82±2.04
	DCx	6.22±8.64	5.90±8.47	6.90±8.76	5.38±7.36	6.27±8.61	6.27±8.63	6.21±8.62
	VoroFacet	**37.07**±17.24	**36.88**±17.07	**37.28**±17.40	**32.50**±17.76	**37.26**±17.27	**37.09**±17.23	**37.07**±17.23
NM-LR → 1	MeshUDF	**0.151**±0.217	**0.202**±0.553	**0.122**±0.073	0.132±0.221	**0.153**±0.217	**0.151**±0.217	**0.151**±0.217
	GeoUDF	98.772±666	245±1688	22.312±131	97.849±666	98.870±666	98.773±666	98.770±666
	MIND	3.526±12.001	6.251±30.778	2.071±2.260	3.221±12.056	3.561±11.996	3.527±12.001	3.525±12.001
	DualMesh-UDF	1524±8870	3486±22537	493±1730	1463±8879	1528±8869	1525±8870	1524±8870
	DCx	6.629±11.179	8.618±21.955	5.448±8.063	5.266±10.861	6.689±11.168	6.631±11.185	6.627±11.180
	VoroFacet	2.167±1.741	2.615±4.288	1.885±0.845	**1.681**±1.762	2.196±1.734	2.168±1.741	2.167±1.742
(b) ShapeNet-Car								
NM-CD ↓	MeshUDF	52.187±15.75	51.931±15.75	52.324±15.75	52.167±15.60	52.131±15.74	52.160±15.73	52.138±15.65
	GeoUDF	37.861±8.79	38.418±9.47	37.276±8.27	37.464±8.71	37.897±8.79	37.950±8.86	37.846±8.79
	MIND	59.505±17.64	59.722±17.55	59.343±17.63	59.556±17.59	59.487±17.63	59.532±17.63	59.477±17.60
	DualMesh-UDF	31.190±8.22	31.812±8.62	30.363±7.68	30.785±7.97	31.271±8.20	31.269±8.27	31.168±8.21
	DCx	25.470±9.06	25.655±9.13	25.380±9.03	25.528±9.20	25.405±9.04	25.473±9.07	25.435±9.04
	VoroFacet	**21.793**±6.75	**21.954**±6.88	**21.662**±6.90	**21.831**±6.78	**21.729**±6.77	**21.796**±6.75	**21.768**±6.75
NM-F1 ↑	MeshUDF	6.35±2.19	6.45±2.22	6.10±2.20	6.25±2.19	6.36±2.20	6.32±2.17	6.32±2.20
	GeoUDF	22.02±4.48	21.98±4.41	21.89±4.61	22.01±4.46	22.05±4.48	22.02±4.49	22.04±4.49
	MIND	8.88±5.19	8.87±5.29	8.78±5.10	8.82±5.19	8.88±5.22	8.86±5.19	8.91±5.23
	DualMesh-UDF	22.47±4.72	22.28±4.75	22.68±4.70	22.49±4.71	22.42±4.73	22.46±4.71	22.47±4.70
	DCx	30.80±7.19	30.71±7.48	30.90±7.23	30.74±7.21	30.86±7.31	30.79±7.28	30.79±7.21
	VoroFacet	**36.67**±8.44	**36.67**±8.57	**36.73**±8.38	**36.62**±8.35	**36.75**±8.52	**36.71**±8.43	**36.67**±8.47
VA-NM-F1 ↑	MeshUDF	2.84±1.12	2.87±1.15	2.83±1.12	2.79±1.11	2.85±1.15	2.83±1.13	2.84±1.13
	GeoUDF	11.20±3.30	10.94±3.36	11.58±3.29	11.24±3.31	11.32±3.28	11.19±3.30	11.21±3.31
	MIND	4.78±4.29	4.63±4.35	4.85±4.15	4.72±4.26	4.75±4.30	4.76±4.28	4.80±4.30
	DualMesh-UDF	12.71±3.56	12.49±3.57	13.16±3.39	12.74±3.53	12.73±3.52	12.75±3.54	12.72±3.54
	DCx	15.09±7.22	14.75±7.29	16.12±7.24	15.05±7.18	15.23±7.29	15.13±7.26	15.11±7.21
	VoroFacet	**18.91**±8.50	**18.52**±8.61	**19.66**±8.31	**18.83**±8.42	**19.11**±8.62	**18.96**±8.48	**18.93**±8.48
NM-LR → 1	MeshUDF	0.021±0.008	0.022±0.008	0.019±0.008	0.020±0.008	0.021±0.008	0.021±0.008	0.021±0.008
	GeoUDF	0.539±0.256	0.561±0.265	0.487±0.232	0.526±0.250	0.545±0.256	0.540±0.256	0.539±0.256
	MIND	0.323±0.171	0.337±0.178	0.292±0.159	0.315±0.167	0.327±0.172	0.323±0.171	0.323±0.171
	DualMesh-UDF	21.111±10.808	22.001±11.250	18.907±9.604	20.607±10.564	21.338±10.799	21.144±10.814	21.089±10.800
	DCx	1.696±0.678	1.764±0.700	1.527±0.643	1.660±0.676	1.715±0.680	1.699±0.678	1.694±0.677
	VoroFacet	**0.982**±0.320	**1.022**±0.330	**0.886**±0.312	**0.959**±0.320	**0.993**±0.321	**0.983**±0.320	**0.981**±0.320

Table 6: VoroFacet with different proxy constructors and identical downstream parameters. CD is scaled by 10^3; F1 scores are percentages.

Proxy	CD ↓	F1 ↑	NM-F1 ↑
DCUDF	5.563±0.826	80.32±8.94	62.15±23.16
DCUDF2	5.549±0.831	80.46±8.98	65.80±21.37

B Robustness and Failure Cases

Closely spaced surfaces. Fig. 10 shows one example each from ShapeNet-Car (top) and Non-manifold Objects (bottom). In the upper example, the gap between distinct sheets in the DCUDF2 proxy exceeds the spacing between the duplicate layers representing a single sheet. Like the double-covering methods, VoroFacet keeps the two target sheets separate, whereas the other baselines either merge them into a single layer or produce irregular geometry. In the lower example, the inter-sheet gap is comparable to the duplicate-layer spacing, so the ball cover spans this gap and merges the two target sheets into a single layer.

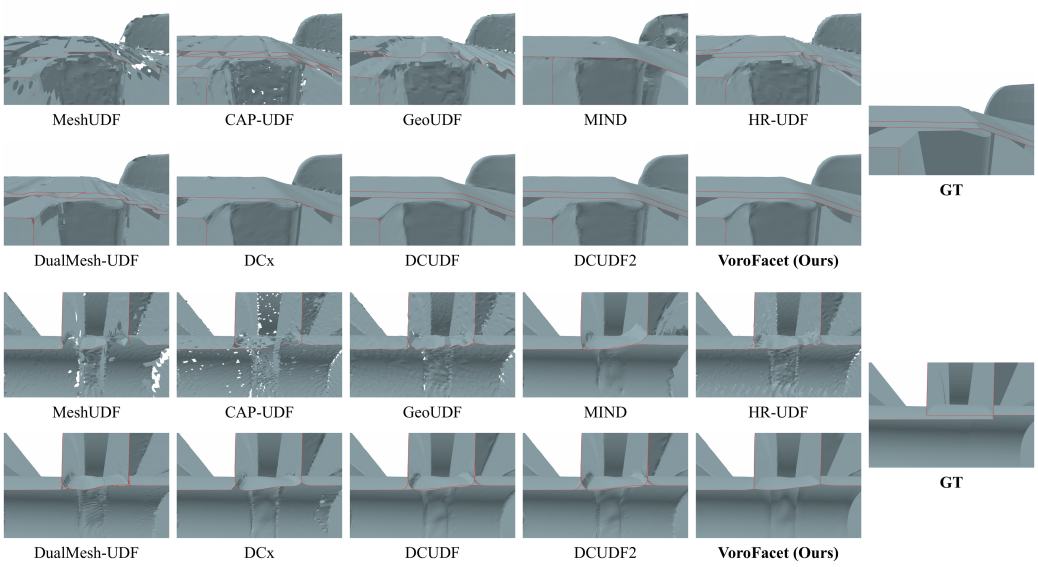


Figure 10: Visual comparison on closely spaced surfaces.

Choice of proxy. We replace DCUDF2 with DCUDF for proxy construction on the 50 Non-manifold Objects, keeping the learned UDFs and all downstream parameters unchanged. Surface accuracy remains comparable (Table 6), while NM-F1 reaches 62.15, exceeding the strongest baseline, DCx (53.19). Figure 11 illustrates successful junction recovery using either proxy. These results demonstrate that VoroFacet does not require DCUDF2 specifically, retaining strong reconstruction performance with an alternative proxy constructor.

High-Valence Junctions. We construct three synthetic junctions with 3, 6, or 9 rectangular sheets sharing a common edge and arranged at equal angles around it. For each configuration, we learn a UDF from 100k unoriented surface points using the same method (Yang et al., 2025) as in the main experiments; all extractors query this UDF. With the default parameters, VoroFacet achieves VA-NM-F1 scores of 99.39% and 97.01% for the 3- and 6-sheet junctions, respectively (Table 7). In contrast, DCx obtains zero VA-NM-F1 for the 6-sheet case despite an NM-F1 of 99.10%, indicating junctions close to the ground truth but with incorrect sheet incidence.

In the 9-sheet case, the DCUDF2 proxy leaves a wide junction gap where adjacent sheets meet at acute angles (Fig. 12). If the cover balls are too small to span this gap, Voronoi sites can be generated inside it. Facets selected between these sites and those outside then retain the gap rather

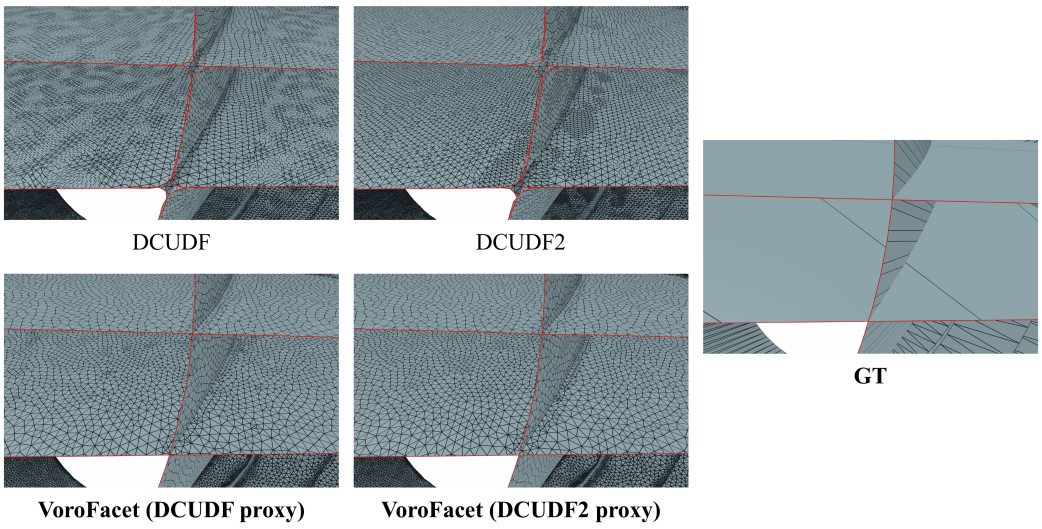


Figure 11: Effect of proxy choice. Top: DCUDF and DCUDF2 reconstructions. Bottom: VoroFacet results using the respective proxies with identical downstream parameters. Ground truth is shown on the right.

Table 7: Junction reconstruction on synthetic equal-angle junctions with 3, 6, and 9 incident sheets. NM-CD is scaled by 10^3; both F1 scores are percentages. The $2\times$ setting doubles $(R_{\text{init}}, R^0_{\text{gap}}, R^1_{\text{gap}}, R_G, \delta_G)$. — denotes an undefined NM-CD for an empty predicted junction set.

Method	3 sheets				6 sheets				9 sheets			
	NM-CD ↓	NM-F1 ↑	VA-NM-F1 ↑	NM-LR → 1	NM-CD ↓	NM-F1 ↑	VA-NM-F1 ↑	NM-LR → 1	NM-CD ↓	NM-F1 ↑	VA-NM-F1 ↑	NM-LR → 1
MeshUDF	—	0.00	0.00	0.000	72.549	0.00	0.00	0.035	63.784	18.98	0.00	0.088
GeoUDF	18.234	96.55	96.55	1.068	116.981	38.71	0.00	0.410	148.565	29.34	0.00	3.701
MIND	51.807	86.93	86.93	1.270	42.426	21.79	0.00	6.608	58.086	1.91	0.00	10.304
DualMesh-UDF	198.206	5.92	3.79	147.333	196.396	6.89	0.00	302.469	204.838	3.36	0.00	459.370
DCx	1.071	99.33	99.33	1.046	2.525	99.10	0.00	4.153	5.202	71.51	0.00	11.874
VoroFacet	0.230	99.39	99.39	0.982	3.520	99.12	97.01	1.035	11.852	1.33	0.00	8.289
VoroFacet ($2\times$ scales)	0.301	99.25	99.25	0.977	0.389	99.84	96.74	1.015	2.135	95.87	23.23	2.938

than form a common junction. Simply doubling the five extraction scales listed in Table 7 increases ball coverage of this gap and substantially improves reconstruction: NM-F1 increases from 1.33% to 95.87%, and VA-NM-F1 from zero to 23.23%, with the UDF and all other parameters unchanged.

Sharp features. We construct two rectangular sheets sharing an edge with a 30° opening angle and train a UDF using the same protocol as above. Fig. 13 compares reconstructions from this shared UDF. DCUDF and DCUDF2 retain double layers with extensive self-intersections, while several other baselines leave gaps or introduce numerous non-manifold edges. VoroFacet extracts a single-layer surface with only a small number of non-manifold edges near the crease. These can arise from residual separation between the proxy layers.

Self-intersections. On the 50 Non-manifold Objects, DCUDF and DCUDF2 exhibit extensive self-intersections in their double-layer outputs, affecting 68.352% and 43.947% of triangles on average, respectively (Table 8). VoroFacet constructs its surface by selecting Voronoi facets, which does not introduce self-intersections. Its subsequent remeshing and geometric refinement can introduce intersections, though Stage III reduces the average percentage from 0.411% to 0.064%.

Manifold objects. Table 9 reports non-manifold edges and self-intersections on the 50 Manifold Objects. VoroFacet produces non-manifold edges in 46 outputs, with incident triangles accounting for 1.321% of all triangles on average. Self-intersections occur in 43 outputs but involve only 0.070% of triangles on average. These reconstructions achieve the lowest CD and second-highest F1 on this set (Table 2).

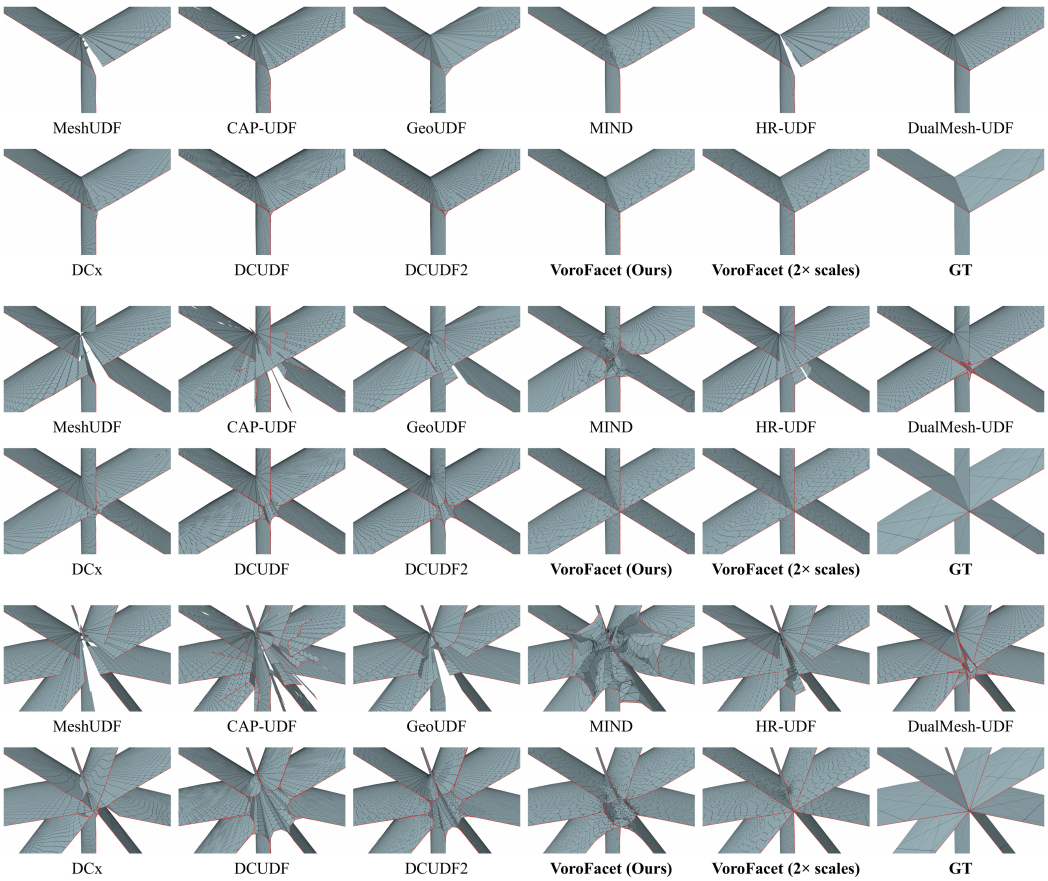


Figure 12: Visual comparison on synthetic junctions with 3, 6, and 9 incident sheets (top, middle, and bottom groups).

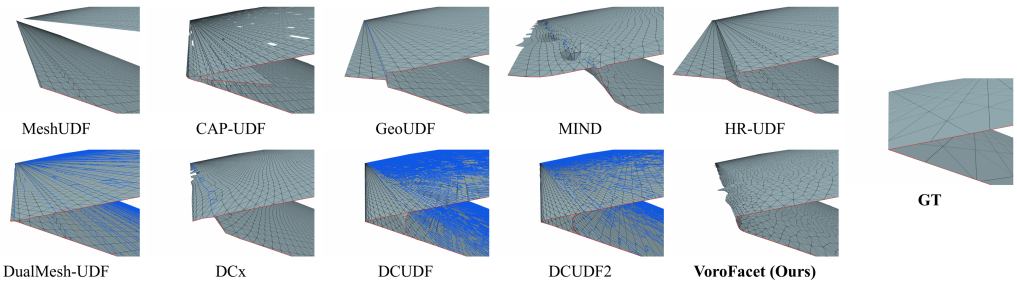


Figure 13: Visual comparison on a 30° sharp crease. Blue lines mark non-manifold edges or self-intersections.

C Additional Qualitative Results

Figures 14 and 15 show the visual effect of the two components. Without the junction-aware re-extraction, a junction is split into two nearby curves, and the faces between them are distorted; with it, the incident sheets meet along a single shared curve. Without the geometry-adaptive remeshing, the triangles around a sharp feature are coarse and irregular, and the fitted surface is jagged; with it, the sampling follows the local geometry and the feature is recovered. These observations are consistent with the quantitative results in Table 4.

Table 8: Self-intersecting triangle percentages on 50 Non-manifold Objects, reported as mean ± standard deviation and rounded to three decimal places. Stage I is evaluated with the temporarily removed faces restored.

	Baselines								VoroFacet				
Method	MeshUDF	CAP-UDF	GeoUDF	MIND	HR-UDF	DualMesh-UDF	DCx	DCUDF	DCUDF2	Extraction	Remeshing	Stage I	Stage II	Stage III
SI (%) ↓	0.005±0.005	0.000±0.000	0.789±1.052	3.451±3.528	0.000±0.000	17.941±4.530	0.276±0.738	68.352±10.895	43.947±7.754	0.000±0.000	0.001±0.001	0.228±0.523	0.411±0.554	0.064±0.053

Table 9: Non-manifold edges and self-intersections on **Manifold Objects**. Affected triangle percentages are reported as mean ± standard deviation across the 50 objects.

Method	Non-manifold edges		Self-intersections	
	Models	Triangles (%)	Models	Triangles (%)
MeshUDF	29/50	0.014±0.047	6/50	0.002±0.006
CAP-UDF	0/50	0.000±0.000	0/50	0.000±0.000
GeoUDF	50/50	2.088±2.907	50/50	0.336±0.873
MIND	32/50	0.431±1.313	38/50	0.585±2.723
HR-UDF	0/50	0.000±0.000	0/50	0.000±0.000
DualMesh-UDF	50/50	71.527±8.822	50/50	35.914±8.628
DCx	50/50	0.551±1.450	46/50	0.040±0.080
DCUDF	0/50	0.000±0.000	37/50	2.233±10.710
DCUDF2	0/50	0.000±0.000	40/50	1.880±8.698
VoroFacet (Ours)	46/50	1.321±2.142	43/50	0.070±0.103

D Additional Quantitative Results

This section completes the comparisons of Sec. 5.2. Table 10 reports the 3D scenes and Table 11 the open-boundary objects.

Table 10: Quantitative results on **3D Scenes**. **Bold** / underline: best / second-best.

Method	CD ↓	F1 ↑	NC ↑	Comp	Time (s)
MeshUDF [ECCV'22]	3.866±1.056	92.88±6.45	**94.32**±1.73	455±461	9.4±1.9
CAP-UDF [TPAMI'24]	3.638±0.862	95.87±4.60	91.91±3.33	4316±3618	151.1±20.9
GeoUDF [ICCV'23]	3.547±0.794	96.44±3.90	93.20±2.22	450.6±347.3	9.2±1.9
MIND [NeurIPS'25]	3.848±0.890	94.14±5.15	92.72±2.68	19.6±22.1	103.4±72.9
HR-UDF [NeurIPS'25]	3.575±0.818	96.72±3.96	92.84±2.41	1963±1520	87.4±16.4
DualMesh-UDF [ICCV'23]	3.612±0.812	96.30±4.14	92.05±2.96	9.2±10.5	11.0±3.0
DCx [TOG'26]	3.835±0.935	94.01±5.02	92.97±2.86	56.8±51.0	35.7±22.7
DCUDF [TOG'23]	**3.510**±0.776	96.66±3.68	94.18±1.94	24.2±20.8	102.0±45.2
DCUDF2 [TVCG'25]	3.524±0.773	**96.76**±3.58	94.09±1.97	23.4±21.2	91.0±37.2
VoroFacet (Ours)	3.664±0.816	95.86±3.71	93.40±2.58	6.4±5.0	199.9±93.3

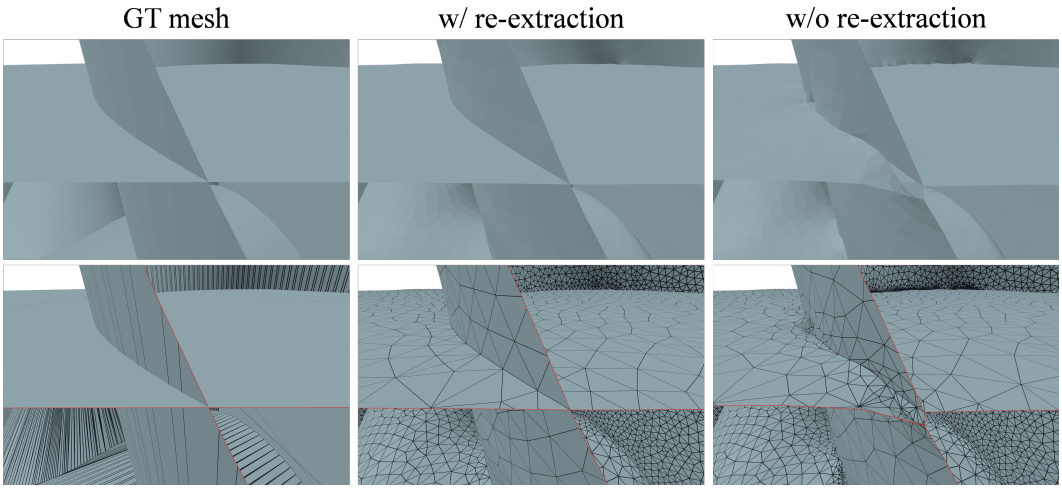


Figure 14: Visual effect of junction-aware re-extraction.

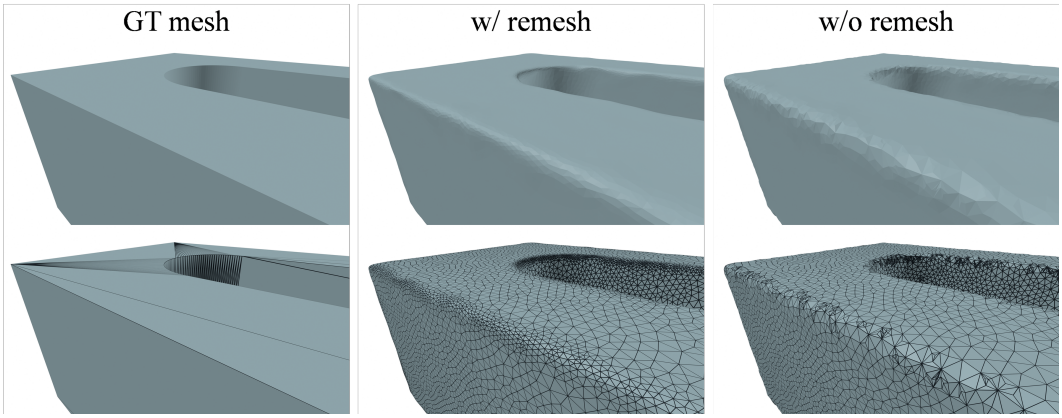


Figure 15: Visual effect of geometry-adaptive remeshing.

E Sensitivity Analysis

Tables 12 to 16 vary one parameter at a time on the same 20 non-manifold objects used in Sec. 5.3. VoroFacet is stable across the tested ranges except at 128^3, where the proxy follows the surface less closely and the wider duplicate-layer gap admits spurious sites (Fig. 16), and at the two smallest gap thresholds, which fail in the same way. These settings reduce NM-F1 from 74.79 to 20.44, 71.75, and 70.27, respectively. The remaining parameters change NM-F1 by less than 1.5 points. VA-NM-F1 shows a larger variation with R_G, reflecting changes in sheet incidence despite similar junction proximity. Increasing the resolution to 512^3 improves NM-F1 only slightly, so we retain 256^3 and the default values listed in Sec. 5.1.

F Runtime Efficiency

VoroFacet takes 418 seconds per non-manifold object on average (Table 1). Constructing $\mathcal{D}$ with DCUDF2 accounts for 41.2% of the runtime; the two passes of ball covering and surface extraction, including junction estimation, account for 33.9%, remeshing for 14.5%, and three-stage refinement for 8.8%. The remaining 1.6% is spent on final checks and other overhead.

Our implementation uses GPU acceleration for proxy optimization and three-stage refinement, CPU multithreading for ball-radius constraints, independent site construction and facet tests, and con-

Table 11: Quantitative results on **Open-boundary Objects**. **Bold** / underline: best / second-best.

Method	CD ↓	F1 ↑	NC ↑	Comp	Time (s)
MeshUDF [ECCV'22]	2.830±0.509	99.22±0.58	98.93±0.65	12.65±17.89	5.46±0.29
CAP-UDF [TPAMI'24]	2.943±0.524	98.53±0.72	98.47±1.02	1056.35±497.79	155.75±15.07
GeoUDF [ICCV'23]	2.889±0.509	99.03±0.56	98.58±0.73	28.9±20.9	5.76±0.11
MIND [NeurIPS'25]	2.942±0.506	98.58±0.59	98.66±0.90	2.05±1.93	47.64±4.97
HR-UDF [NeurIPS'25]	2.938±0.503	98.86±0.52	98.68±0.71	991.65±305.78	64.40±9.03
DualMesh-UDF [ICCV'23]	2.860±0.516	99.29±0.65	98.73±0.76	1.3±0.7	5.54±0.76
DCx [TOG'26]	2.869±0.509	98.94±0.56	98.69±0.90	4.1±3.0	16.51±3.70
DCUDF [TOG'23]	**2.800**±0.514	**99.48**±0.49	**99.13**±0.51	3.3±2.7	122.21±25.22
DCUDF2 [TVCG'25]	2.809±0.515	99.43±0.50	99.09±0.52	3.3±2.7	101.70±19.69
VoroFacet (Ours)	2.843±0.506	99.19±0.45	98.93±0.66	1.05±0.22	175.04±35.36

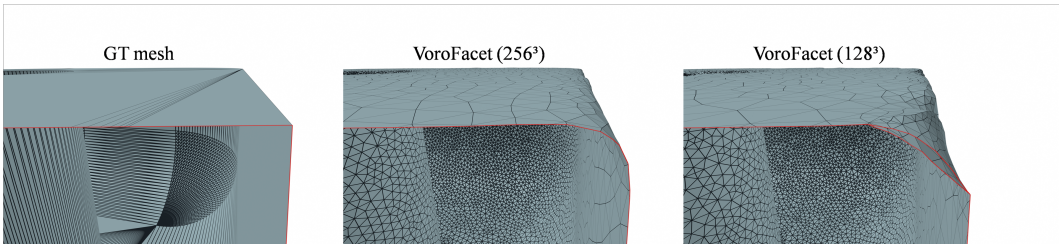


Figure 16: Visual effect of the extraction resolution near a sharp feature.

current remeshing across sheets. Greedy ball selection and dependent connectivity updates within each sheet's remeshing remain sequential in the current implementation. Proxy construction and the two extraction passes together account for approximately 75% of the runtime, making them the main targets for acceleration. Future directions include faster, geometrically accurate proxy construction tailored to VoroFacet and incremental second-pass extraction that reuses unaffected parts of the first-pass ball cover and Voronoi structure.

Table 12: Effect of extraction resolution on 20 non-manifold objects.

Resolution	CD ↓	F1 ↑	NC ↑	NM-CD ↓	NM-F1 ↑	VA-NM-F1 ↑	NM-LR → 1
128^3	5.722±0.767	78.68±8.58	96.93±1.60	60.303±29.803	20.44±15.37	6.26±7.44	6.158±2.860
256^3 (default)	5.609±0.767	79.80±8.51	98.16±0.82	**23.331**±27.554	74.79±14.04	38.50±15.88	1.615±0.643
512^3	**5.602**±0.761	**79.86**±8.46	**98.21**±0.78	25.356±26.771	**78.28**±12.01	**40.10**±19.61	**1.417**±0.516

Table 13: Effect of the first-pass gap threshold R^0_{gap}.

R^0_{gap}	CD ↓	F1 ↑	NC ↑	NM-CD ↓	NM-F1 ↑	VA-NM-F1 ↑	NM-LR → 1
0.006	**5.608**±0.761	**79.83**±8.44	98.15±0.81	26.181±28.483	71.75±13.42	38.44±16.65	1.766±0.675
0.0075 (default)	5.609±0.767	79.80±8.51	**98.16**±0.82	**23.331**±27.554	**74.79**±14.04	**38.50**±15.88	**1.615**±0.643
0.009	5.613±0.765	79.81±8.49	**98.16**±0.82	24.001±28.190	74.58±14.07	35.88±15.53	1.630±0.658

Table 14: Effect of the second-pass gap threshold R^1_{gap}.

R^1_{gap}	CD ↓	F1 ↑	NC ↑	NM-CD ↓	NM-F1 ↑	VA-NM-F1 ↑	NM-LR → 1
0.002	5.611±0.763	79.77±8.48	98.14±0.86	29.053±34.653	70.27±17.69	35.32±15.89	2.008±1.026
0.003 (default)	**5.609**±0.767	**79.80**±8.51	98.16±0.82	23.331±27.554	74.79±14.04	38.50±15.88	1.615±0.643
0.004	**5.609**±0.767	79.79±8.51	**98.17**±0.81	**21.836**±23.236	**75.74**±12.69	**39.58**±15.94	**1.547**±0.520

Table 15: Effect of the junction-ball radius R_G.

R_G	CD ↓	F1 ↑	NC ↑	NM-CD ↓	NM-F1 ↑	VA-NM-F1 ↑	NM-LR → 1
0.01	**5.609**±0.763	79.81±8.47	98.13±0.88	23.596±26.684	74.52±14.49	34.34±14.29	1.838±0.693
0.0125 (default)	**5.609**±0.767	79.80±8.51	98.16±0.82	**23.331**±27.554	**74.79**±14.04	38.50±15.88	1.615±0.643
0.015	5.610±0.764	**79.82**±8.48	**98.18**±0.80	23.742±28.468	74.64±13.72	**39.96**±16.77	**1.507**±0.633

Table 16: Effect of the junction clustering spacing δ_G.

δ_G	CD ↓	F1 ↑	NC ↑	NM-CD ↓	NM-F1 ↑	VA-NM-F1 ↑	NM-LR → 1
0.01	**5.608**±0.764	**79.83**±8.51	**98.17**±0.81	23.307±27.717	**75.32**±13.76	38.75±16.57	1.643±0.608
0.015 (default)	5.609±0.767	79.80±8.51	98.16±0.82	23.331±27.554	74.79±14.04	38.50±15.88	**1.615**±0.643
0.02	5.613±0.766	79.76±8.52	98.16±0.82	**23.016**±26.668	73.85±14.62	**39.08**±15.92	1.685±0.730