

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 REGIONUDF: REGION-AWARE UNSIGNED DISTANCE FIELDS FOR SURFACE RECONSTRUCTION FROM POINT CLOUDS

Anonymous authors

Paper under double-blind review

## ABSTRACT

Distance fields offer a powerful representation for continuous geometry, yet current learning-based neural unsigned distance fields (UDFs) remain limited in their ability to capture data patterns and generalize to real-world open surfaces. Point-Based methods mitigate grid quantization errors but current work often oversmooth local details, as query features are obtained solely through interpolation of point-wise features which are aggregated over large receptive fields. To address this, we propose a *discriminative region representation* that fuses narrow neighborhood features with broader contextual point-wise features, and a *primitive-based region representation* that decomposes the query region into triplet-defined primitives, **enabling the detailed encoding of local surface geometry and the clear distinction of multi-layer structures**. Building on these designs, we propose *RegionUDF*, a region-aware UDF framework that achieves state-of-the-art open-surface reconstruction on both object- and room-level scenes, with additional validation on watertight shapes. Extensive experiments on synthetic and real-world datasets demonstrate superior accuracy and robust cross-domain generalization. Our source code will be available at *[no-name-for-blind-review]*.

## 1 INTRODUCTION

Implicit geometric representations have become a predominant paradigm in 3D vision, computer graphics, and robotics (Park et al., 2019; Huang et al., 2022; Mescheder et al., 2019; Chibane et al., 2020b; Yang et al., 2023). Unlike explicit representations such as meshes or point clouds (Hanocka et al., 2020; Badki et al., 2020; Lin et al., 2020), implicit representations model surfaces as continuous isosurfaces in space, overcoming the resolution limits of discrete methods and enabling high-fidelity reconstruction. Among them, signed distance fields (SDFs) and unsigned distance fields (UDFs) are the most common. SDFs (Tang et al., 2021; Peng et al., 2020; Huang et al., 2023) encode signed distances to the nearest surface, where the sign indicates whether a point lies inside or outside a watertight object. UDFs (Ren et al., 2023; Fainstein et al., 2024; Ye et al., 2022), in contrast, record only distance magnitudes, allowing representation of more general, non-watertight shapes and making them particularly suitable for real-world surface reconstruction.

Recent advances have explored direct neural mappings from spatial coordinates to UDFs, either via scene-specific optimization (Zhou et al., 2022; 2023) or by defining new distance fields such as orthogonal Lu et al. (2024) and line-segment fields (Ren & Hou, 2025). Optimization-Based methods, however, are restricted to single scenes without addressing cross-scene generalization. In contrast, learning-based approaches (Chibane et al., 2020b; Ye et al., 2022; Ren et al., 2023) must construct query features from point-cloud inputs to generalize across scenes. Grid-Based trilinear interpolation remains common but introduces quantization errors and surface detail loss. Point-Based alternatives, though less studied for UDFs, have been explored for SDFs (Boulch & Marlet, 2022; Wang et al., 2023; Ranade et al., 2025). While they better preserve fine details, most rely only on point locations to refine interpolation, neglecting richer regional information; thus, neighborhood structure is underrepresented and point-wise features alone fail to capture fine-grained geometry.

Point-Wise features capture broad receptive fields but often oversmooth local details. In point-based methods, they encode regional characteristics over large spatial ranges, whereas narrow neighborhood

054 features we defined are confined to the immediate vicinity of a query point. These two levels of  
 055 representation, reflecting “global” and “local” perspectives, are inherently complementary, and  
 056 their integration produces a more compact and discriminative region feature, as illustrated in Fig. 1.  
 057 Furthermore, since those points are often distributed across discrete surface patches, the neighborhood  
 058 point cloud can be decomposed into finer primitives **that are deliberately constructed to preserve**  
 059 **the underlying surface structure. Theoretically, complex structures can be constructed through the**  
 060 **aggregation of simple primitives. It** enables a more detailed characterization of local geometry, while  
 061 aggregating the primitives preserves information about the neighborhood as a whole.

062 In this work, we propose a *discriminative region representation* that models a query neighborhood by  
 063 fusing broad contextual point-wise features with narrow region features extracted from its surrounding  
 064 neighbors. Naturally, narrow region features can be obtained by treating the query neighbors as an  
 065 independent point cloud, enabling the use of standard point cloud analysis techniques. This offers a  
 066 principled alternative to weighted interpolation for encoding the local region relative to the query. We  
 067 further propose a *primitive-based region representation*. **The neighborhood is partitioned into multiple**  
 068 **triplets, each of which uniquely defines a planar patch. Through spherical projection and angular**  
 069 **sorting, primitives are constructed to preserve surface structure, thereby enabling clear discrimination**  
 070 **of adjacent layers when the query point lies within a multi-layer structure.** We first extract primitive  
 071 features using the discriminative region representation with triplet points and point-wise features.  
 072 Each primitive is then treated as a point, with its primitive feature serving as the corresponding  
 073 point-wise feature. Finally, these primitives are aggregated to construct the query region feature.

074 Building on the above ideas, we develop a **Region-Aware Unsigned Distance Fields** framework, **Re-**  
 075 **gionUDF**, and evaluate it on both synthetic and real-world datasets. We further assess its cross-domain  
 076 generalization by varying datasets and scene scales. In addition, existing learning-based UDF studies  
 077 have seldom conducted quantitative evaluations in room-level scenarios, typically providing only  
 078 qualitative visual results. Yet, room-level scenarios are inherently more complex and variable, making  
 079 them a more rigorous and informative benchmark for assessing learning capability. Therefore, we  
 080 place particular emphasis on evaluation in room-level scenarios. Overall, our results demonstrate that  
 081 RegionUDF consistently outperforms existing methods. Our main contributions are as follows:  
 082

- We design a discriminative region representation that incorporates narrow, fine-grained region features for more accurate modeling of the query neighborhood.
- We propose a primitive-based region representation that decomposes the region into primitives, **enabling detailed encoding of local geometry and clear distinction of multi-layer structures.**
- We validate our model on both synthetic and real-world datasets, demonstrating superior performance in both intra-domain and cross-domain evaluations. Comprehensive ablation studies further confirm the effectiveness of our proposed approaches.

## 090 2 RELATED WORK

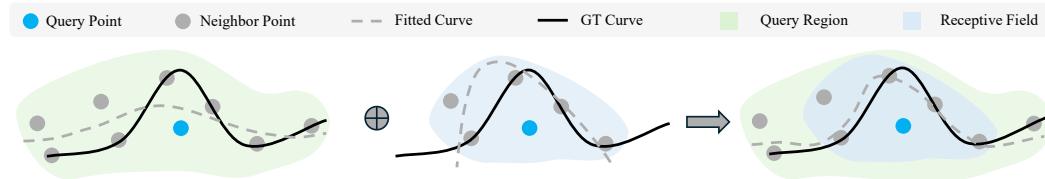
092 Surface reconstruction has long been studied, with discrete representations (Schönberger et al., 2016;  
 093 Schönberger & Frahm, 2016) limited by spatial resolution and memory. Implicit representations have  
 094 thus become mainstream, spanning objects to large-scale environments. Classic methods such as  
 095 Poisson reconstruction (Kazhdan & Hoppe, 2013; Kazhdan et al., 2006), radial basis functions (Carr  
 096 et al., 2001), and moving least-squares surfaces (Guennebaud & Gross, 2007; Levin, 1998) rely on  
 097 smoothness priors, whereas neural implicit approaches offer greater expressiveness and flexibility.

099 **Closed Surface.** In distance field modeling, SDFs parallel occupancy fields by mapping distances  
 100 to occupancy probabilities. Early neural implicit models such as DeepSDF (Park et al., 2019) and  
 101 ONet (Mescheder et al., 2019) predict signed distances or occupancy scores from global shape  
 102 codes. To better capture local geometry, voxel-based methods (Peng et al., 2020; Tang et al., 2021;  
 103 Chibane et al., 2020a) adopt voxel-based encodings, though interpolating latent features limits fine  
 104 detail. Spline-based methods (Williams et al., 2022; Huang et al., 2022) address this via local  
 105 patches. **NKSR (Huang et al., 2023) further refines the SDF by incorporating hierarchical information**  
 106 **to solve the global gradient optimization problem.** Optimization-based approaches (Gropp et al.,  
 107 2020; Sitzmann et al., 2020; Ma et al., 2021) instead solve the Eikonal equation with neural PDE  
 108 solvers (Sirignano & Spiliopoulos, 2018). A shared limitation is the assumption of closed surfaces,  
 109 restricting applicability to open surfaces in real scans.

108 **Open Surface.** To represent open surfaces, NDF (Chibane et al., 2020b) first introduced neural  
 109 unsigned distance fields, regressing unsigned distances from spatial queries to the surface. Subsequent  
 110 per-scene optimization methods refined this representation: CAP-UDF (Zhou et al., 2022) employs a  
 111 Chamfer-based pull objective, DUDF (Fainstein et al., 2024) enforces differentiability via hyperbolic  
 112 scaling, and LevelSetUDF (Zhou et al., 2023) stabilizes gradients through point-projection. Other  
 113 variants like PDDF (Aumentado-Armstrong et al., 2022), NeuralODF (Houchens et al., 2022), and  
 114 LineSeg (Ren & Hou, 2025), incorporate directional cues for direct mesh extraction. Specifically,  
 115 **UODF** (Lu et al., 2024) defines the minimal unsigned distance along three orthogonal directions, en-  
 116 abling each spatial point to directly access its closest surface point and thereby achieve high-precision  
 117 reconstruction without interpolation errors. While effective, these remain per-scene methods without  
 118 cross-domain generalization. Learning-based UDFs are less explored: GIFS (Ye et al., 2022) predicts  
 119 binary extension flags, NVF (Yang et al., 2023) regresses vectors to nearest surfaces, GeoUDF (Ren  
 120 et al., 2023) upsamples input points and normals for interpolation, and SALS (Ren & Hou, 2025)  
 121 learns line segment–surface relations beyond the UDF paradigm. Despite these advances, existing  
 122 approaches overlook that constructing query features with an emphasis on narrow neighborhood  
 123 features can yield more accurate representations.

### 3 METHOD

124 We propose a region-aware unsigned distance field framework that integrates point-wise and narrow  
 125 region features of each query neighborhood. Point-Wise features capture broad contextual information  
 126 but often oversmooth fine details, whereas narrow region features preserve local structures yet  
 127 are sensitive to noise. Their integration yields complementary and more discriminative region  
 128 representations. To further enhance local representation, we propose a primitive-based region  
 129 formulation, decomposing the neighborhood into finer primitives for more detailed characterization.



130 **Figure 1: Intuition of Discriminative Region Representation.** As an illustrative 2D example. The  
 131 query identifies neighbors via KNN, defining the query region (blue). Through the point-wise features  
 132 of these neighbors, the query also inherits a broader receptive field (green). Features restricted to the  
 133 query region capture finer local detail but are sensitive to noise, whereas point-wise features alone  
 134 oversmooth details. Their integration yields a more faithful/discriminative region representation.

#### 3.1 DISCRIMINATIVE REGION REPRESENTATION

135 We consider narrow region features to form complementary, more *discriminative region representation*  
 136 by fusing a narrow region feature  $r_p$  with the point-wise feature  $f_p$ . Specifically,  $f_p$  encodes multi-  
 137 scale surface context from the input point cloud at each neighboring point. In parallel,  $r_p$  encodes the  
 138 spatial configuration of the region defined by the set  $N_q$  of  $K$  neighboring points surrounding  $q$ .

139 Treating the neighborhood as an independent point cloud with the query point  $q$  as the origin of  
 140 its coordinate system allows the direct application of point cloud analysis techniques. We follow  
 141 the *pooling and propagation paradigm of PointNet++* (point-wise MLP → symmetric pooling →  
 142 skip(concat)). We formulate the whole distance learning process as:

$$143 \begin{aligned} r_p &= \phi(p - q) \oplus \left( \frac{1}{K} \sum_p \phi(p - q) \right) | p \in N_q, \\ 144 f_q &= \delta(r_p \oplus f_p) | p \in N_q, \\ 145 d_q &= \varphi(f_q \oplus q) \end{aligned} \tag{1}$$

146 where  $\oplus$  denotes element-wise concatenation,  $\phi$  denotes MLPs,  $\delta$  means attention aggregation,  $f_q$   
 147 represents the feature of the region relative to a query point  $q$ , and  $\varphi$  is the regression MLP.

162 As shown in Eq. 1, we treat the set of neighboring points as a local region centered at the query  
 163 point  $q$ . We apply a single layer of point-cloud abstraction followed by feature propagation to extract  
 164 the narrow region feature  $r_p$  as a narrow geometry representation relative to  $q$  for each neighbor  $p$ .  
 165 Concurrently, each neighbor is endowed with a point-wise feature  $f_p$ , which encodes a large receptive  
 166 field and provides contextual information independent of  $q$ . By fusing the narrow region features  
 167  $\{r_p\}$  with the query-independent features  $\{f_p\}$ , we construct a discriminative region representation  
 168 that effectively characterizes the region relative to the query point.

### 169 3.2 PRIMITIVE-BASED REGION REPRESENTATION

170 However, because the region defined by the reference points often comprises multiple discrete surface  
 171 patches, its holistic treatment in Eq. 1 may obscure fine-grained geometric details. **To address this,**  
 172 **we decompose the local region into a set of triplet primitives  $\{S_j\}$ , since three non-collinear points**  
 173 **uniquely define a plane and complex shapes can then be represented as compositions of planar**  
 174 **patches.** For each  $S_j$ , we apply the paradigm of Eq. 1 to extract a primitive feature  $f_{S_j}$ . We then  
 175 aggregate primitive features  $\{f_{S_j}\}$  to form a richer representation of the region surrounding  $q$ .

176 First, we decompose the query-centered region into a collection of triplet primitives. Let  $T = \{S_j\}_{j=1}^K$   
 177 denote the set of  $K$  triplet primitives that together define the region surrounding the query point  
 178  $q$ . Accurately recovering surface connectivity from discrete reference points is challenging due to  
 179 their sparse, unstructured nature. To mitigate this, we adopt an intuitive connectivity strategy based  
 180 on angular proximity: we select the  $K$  nearest neighbors of  $q$ , project them onto a virtual sphere  
 181 centered at  $q$ , sort by spherical angles, and connect each three consecutive points to form a primitive.  
 182 **Through these operations, primitives are constructed to preserve surface structure, thereby enabling**  
 183 **clear discrimination of adjacent layers when the query point lies within a multi-layer configuration.**

184 For each triplet primitive  $S = \{p_i \mid i = 0, 1, 2\}$  with the corresponding point-wise feature  $f_{p_i}$ , we  
 185 could extract a primitive feature  $f_S$  with Eq. 1 as following:

$$186 \begin{aligned} r_p &= \phi(p - q) \oplus \left( \frac{1}{3} \sum_p \phi(p - q) \right) | p \in S, \\ 187 f_S &= \frac{1}{3} \sum_{p \in S} \phi(r_p \oplus f_p), \end{aligned} \tag{2}$$

188 where  $\phi$  is an MLP,  $r_p$  is the hierarchical region feature at  $p$ , and  $f_S$  is obtained via MLP and mean  
 189 pooling. After computing each  $f_S$ , we aggregate these primitive features into a detailed region feature  
 190  $f_q$  relative to the query  $q$ , also with paradigm defined by Eq. 1:

$$191 \begin{aligned} c_S &= \frac{1}{3} \sum_{p \in S} (p - q), \\ 192 r_S &= \phi(c_S) \oplus \left( \frac{1}{K} \sum_S \phi(c_S - q) \right) | S \in T, \\ 193 f_q &= \delta(\{r_S \oplus f_S \mid S \in T\}), \\ 194 d_q &= \varphi(f_q \oplus q), \end{aligned} \tag{3}$$

195 where  $c_S$  denotes the center of the primitive  $S$ ,  $r_S$  means the hierarchical region feature at the  
 196 primitive  $S$ . In practice,  $\delta$  is the AttSet aggregation (Hu et al., 2020).

197 In practice, projecting reference points onto a sphere and sorting by angular coordinates via radix  
 198 sort ensures minimal angular separation and coherent primitive formation. Given the simplicity of  
 199 each triplet, a PointNet++-like structure suffices for feature extraction. Although each  $f_S$  is relative  
 200 to  $q$ , incorporating these region features yields a more discriminative representation of the region.  
 201 For implement details, please refer to Appendix section A.2.

## 211 4 EXPERIMENTS

### 212 4.1 OVERVIEW

213 In this Section, we present a thorough evaluation of our method on both synthetic and real-world  
 214 datasets. Specifically, we test on the watertight ShapeNet (Chang et al., 2015), the multi-layer open

surface “Car” subset of ShapeNet (Chang et al., 2015), on indoor scans from ScanNet (Dai et al., 2017) and Matterport3D (Chang et al., 2017). We also compare with the latest SALS on the ABC and non-manifold ABC datasets (Ren & Hou, 2025) it introduced. We further evaluate cross-domain performance across synthetic and real data and scene scales, supported by ablation studies on region-aware features and robustness tests under varying density and noise. For baseline reproduction details, see Appendix Section A.3.3. For cross-domain evaluation, see Appendix Section A.4.

## 4.2 WATERTIGHT SURFACES

Following established evaluation protocols, we assessed performance on watertight shapes from the ShapeNet 13-class dataset. For dataset details, please refer to Appendix Section A.3.1. For fairness, all experiments used 3K points as input. Given the large scale of the dataset, we directly adopted the comparative results reported in the GeoUDF (Ren et al., 2023) manuscript rather than reproducing them; consequently, the most recent SALS (Ren & Hou, 2025) is not included in this comparison. For evaluation metrics, please refer to Appendix Section A.2.3.

Table 1: **Watertight ShapeNet Comparison (Xu et al., 2019).**  $CD_{L_1} \times 10^{-2}$ , F-S. (%) with threshold 0.005 and 0.01. **Best** are in bold, and second-best are underlined. Note that Marching Cubes (MC) with a resolution of 128 was applied to all methods following the GeoUDF setting, while “+” indicates the higher resolution reported in their original papers.

	Clean				Noise (0.005)			
	$CD_{L_1} \downarrow$		F-S. $\uparrow$		$CD_{L_1} \downarrow$		F-S. $\uparrow$	
	Mean	Median	$F1^{0.5\%}$	$F1^{1\%}$	Mean	Median	$F1^{0.5\%}$	$F1^{1\%}$
NDF (Chibane et al., 2020b)	0.341	0.320	84.0	97.6	0.431	0.419	68.5	96.1
GIFS(Ye et al., 2022)	0.328	0.276	86.0	97.4	0.418	0.358	73.1	95.8
GIFS <sup>+</sup> (Ye et al., 2022)	0.281	0.243	91.4	98.5	0.376	0.348	78.0	96.8
GeoUDF (Ren et al., 2023)	0.234	0.226	93.8	99.2	0.289	0.278	89.3	98.7
<b>Ours</b>	<b>0.229</b>	<b>0.222</b>	<b>95.0</b>	<b>99.4</b>	<b>0.273</b>	<b>0.261</b>	<b>91.7</b>	<b>98.9</b>

We compare our model with mainstream UDF methods for watertight shapes following the GeoUDF protocol. As shown in Table 1, our model consistently outperforms existing learning-based UDF approaches across all metrics. Our method achieves consistent improvements over GeoUDF, particularly in F-score, with gains of 1.2% on clean data and 2.4% on noisy data at the 0.5% threshold, indicating closer surface approximation and fewer reconstruction artifacts. As shown in Fig. 2, GIFS produces scaly textures, while GeoUDF introduces boundary gaps and fragmented details. In contrast, our method attains higher fidelity, accurately recovering thin structures and separating adjacent objects.

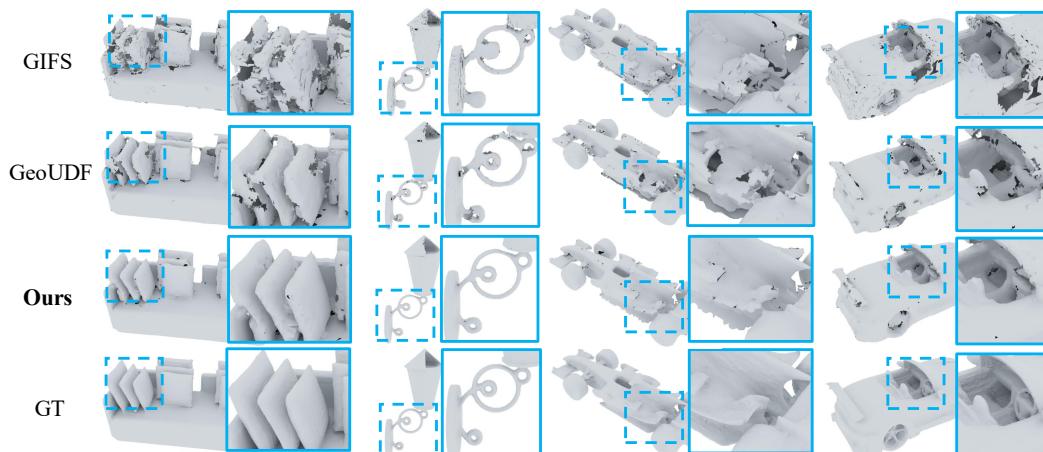


Figure 2: **ShapeNet (Chang et al., 2015) Visualization.** All methods are evaluated at a resolution of 128. Zoomed-in views highlight the regions with the most significant differences for comparison.

270 4.3 OPEN SURFACES  
271272 4.3.1 SURFACE RECONSTRUCTION OF SHAPES  
273

274 To assess our method’s capability on shapes with arbitrary topology, we follow prior work (Chibane  
275 et al., 2020b; Ye et al., 2022) and evaluate on the original “Car” category of ShapeNet, which features  
276 multi-layered and open-surface models. For evaluation metrics, please refer to Appendix Section  
277 A.2.3. Since GeoUDF (Ren et al., 2023) and SALS (Ren & Hou, 2025) do not provide training  
278 results on open surfaces, we retrain them on the non-watertight Car models using the default settings  
279 of their open source code and evaluate them under the same protocol. For SALS (Ren & Hou,  
280 2025), resolution 256 exceeds a single RTX 3090’s capacity, so we report results at 128; the default  
281 resolution is 256. We further compare against the recent learning-based open-surface reconstruction  
282 method SALS (Ren & Hou, 2025) on its proposed benchmark.

283 Table 2: **ShapeNet Cars (Chibane et al., 2020b) Comparison.** F-S. (%) with a distance threshold  
284 of 0.005 and 0.01,  $d_C = CD_{L_2} \times 10^4$ , and NC (%). Best are in bold, and second are underlined.

		$d_C$		F-S.		NC ↑
		Mean↓	Median↓	F-S. 0.005 ↑	F-S. 0.01 ↑	
287	GeoUDF <sup>–</sup> (Ren et al., 2023)	0.138	0.134	86.07	98.91	85.9
288	NVF <sup>–</sup> (Yang et al., 2023)	0.166	0.161	83.23	97.93	83.5
289	SALS (Ren & Hou, 2025)	0.179	0.175	83.05	97.49	77.4
290	<b>Ours<sup>–</sup></b>	<b>0.135</b>	<b>0.131</b>	<b>86.42</b>	<b>98.93</b>	<b>86.4</b>
291	NDF (Chibane et al., 2020b)	0.126	0.120	88.09	<b>99.54</b>	-
292	NDF (Mesh) (Chibane et al., 2020b)	0.202	0.193	77.40	97.97	79.1
293	GIFS (Ye et al., 2022)	0.128	0.123	88.05	99.31	-
294	GeoUDF (Ren et al., 2023)	0.120	0.114	89.23	99.29	86.6
295	NVF (Yang et al., 2023)	0.134	0.126	87.39	98.93	84.0
296	<b>Ours</b>	<b>0.110</b>	<b>0.105</b>	<b>90.69</b>	99.50	<b>86.9</b>

297 Table 3: **ABC and Non-Manifold ABC (Ren & Hou, 2025) Comparison.** F-S. (%) with a distance threshold  
298 of 0.005 and 0.01,  $CD_{L_1} \times 10^{-2}$ ,  $CD_{L_2} \times 10^{-6}$ , and NC (%). Best are in bold, and second are underlined.  
299 Default Marching Cubes resolution is 128.

	Method	CD		F-S.		NC ↑
		$CD_{L_1}$	$CD_{L_2}$	F-S. 0.005 ↑	F-S. 0.01 ↑	
300 ABC	NDF (Chibane et al., 2020b)	0.324	15.4	86.31	98.66	80.6
	GIFS (Ye et al., 2022)	0.345	17.3	84.56	97.84	91.7
	GeoUDF (Ren et al., 2023)	0.257	9.02	<u>92.88</u>	<u>99.78</u>	97.2
	SALS (Ren & Hou, 2025)	<b>0.251</b>	8.87	92.66	99.73	97.3
	<b>Ours</b>	<b>0.251</b>	<b>8.62</b>	<b>93.46</b>	<b>99.85</b>	<b>97.8</b>
307 Non-Manifold 308 ABC	NDF (Chibane et al., 2020b)	0.395	21.1	75.62	97.52	77.1
	GIFS (Ye et al., 2022)	0.412	22.6	73.78	97.71	89.9
	GeoUDF (Ren et al., 2023)	0.333	14.4	<u>83.69</u>	<u>99.55</u>	95.0
	SALS (Ren & Hou, 2025)	0.330	<u>14.3</u>	83.63	99.52	93.9
	<b>Ours</b>	<b>0.320</b>	<b>13.2</b>	<b>85.65</b>	<b>99.74</b>	<b>96.8</b>

312 For ShapeNet “Cars” benchmark, we compare RegionUDF against recent learning-based surface  
313 reconstruction methods, NDF (Chibane et al., 2020b), GIFS (Ye et al., 2022), GeoUDF (Ren et al.,  
314 2023) and SALS (Ren & Hou, 2025), on the ShapeNet Cars benchmark. Table 2 reports quantitative  
315 metrics, ours consistently achieves the best scores across most metrics and is comparable to the best  
316 normal consistency, indicating that ours produce fewer outlier artifacts.

317 For ABC and Non-Manifold ABC benchmark, in line with SALS default settings, we use 40k points  
318 without normals as input for each shape. *Please refer to Appendix Section A.3.3 for more details.*  
319 As shown in Table 3, our model achieves comparable CD performance while surpassing others on  
320 F-score and NC by about 1% on F-S<sup>0.005</sup> and NC, particularly on the non-manifold ABC benchmark,  
321 indicating fewer outlier artifacts and higher reconstruction quality.

322 As shown in Fig. 3, most methods struggle with complex non-manifold edges, an inherent limitation  
323 of UDFs. Our approach performs better on manifold regions, preserving sharp details along crisp



Figure 3: **ABC and Non-Manifold ABC (Ren & Hou, 2025) Visualization.** All methods are evaluated at a resolution of 128 with non-manifold edges preserved and holes filled. The first row shows non-manifold objects, while the subsequent three rows depict the original ABC objects.

edges. In contrast, GIFS produces scaly surfaces, GeoUDF shows missing parts and artifacts, and SALS introduces excessive edge artifacts. This demonstrates that discriminative features help recover sharp details, counteracting the smoothing effects of point-wise features.

#### 4.3.2 SURFACE RECONSTRUCTION OF ROOMS

To further assess our method on real-world scene scans, we evaluate it on two large-scale indoor datasets: ScanNet (Dai et al., 2017) and Matterport3D (Chang et al., 2017). For evaluation metrics, please refer to Appendix Section A.2.3. As no prior work reports room-level results on these benchmarks, we retrain them under default experimental settings reported in their papers and apply a consistent evaluation protocol for fair comparison. Due to SALS meshing limitations, experiments were conducted at two resolutions, with SALS results reported only at 128.

Table 4: **ScanNet and Matterport3D Comparison (Dai et al., 2017; Chang et al., 2017).**  $CD_{L_1} \times 10^{-3}$ ,  $CD_{L_2} \times 10^{-6}$ , F-S. (%) with threshold 0.005, and NC (%). Best results are in bold, and second-best results are underlined. Default MC resolution is 256; “–” denotes 128.

	ScanNet				Matterport3D			
	$CD_{L_1} \downarrow$	$CD_{L_2} \downarrow$	F-S. $\uparrow$	NC $\uparrow$	$CD_{L_1} \downarrow$	$CD_{L_2} \downarrow$	F-S. $\uparrow$	NC $\uparrow$
GeoUDF <sup>–</sup> (Ren et al., 2023)	2.36	11.9	91.0	88.7	2.35	10.1	91.7	93.9
NVF <sup>–</sup> (Yang et al., 2023)	2.20	11.0	92.9	<u>88.8</u>	<u>2.31</u>	10.0	92.2	93.7
SALS <sup>–</sup> (Ren & Hou, 2025)	<b>2.02</b>	<u>7.49</u>	<b>94.7</b>	87.1	2.81	15.8	87.8	89.5
<b>Ours<sup>–</sup></b>	<b>2.02</b>	<u>7.13</u>	<u>94.3</u>	<b>89.4</b>	<b>2.20</b>	<b>7.95</b>	<b>93.2</b>	<b>94.3</b>
NDF(Chibane et al., 2020b)	2.31	8.60	93.3	–	2.54	9.81	91.8	–
NDF(Mesh)(Chibane et al., 2020b)	2.83	48.7	87.0	88.1	2.84	12.6	87.9	83.5
GIFS (Ye et al., 2022)	2.20	7.92	94.6	87.7	2.56	10.3	91.3	92.3
GeoUDF (Ren et al., 2023)	2.05	7.84	93.8	<u>89.3</u>	<u>2.21</u>	<u>8.33</u>	<u>92.9</u>	<u>94.2</u>
NVF (Yang et al., 2023)	2.03	11.7	94.6	88.9	2.83	11.9	89.5	93.6
<b>Ours</b>	<b>1.86</b>	<b>6.03</b>	<b>95.7</b>	<b>90.1</b>	<b>2.09</b>	<b>7.29</b>	<b>94.1</b>	<b>95.5</b>

As shown in Table 4, our method substantially outperforms existing approaches on real-world room-level surface reconstruction. NDF (Mesh) fails to produce coherent meshes on ScanNet. We achieve a **10%** improvement in  $CD_{L_1}$  and nearly **20%** in  $CD_{L_2}$  on ScanNet, and 5% and 12% improvements on Matterport3D, reflecting better recovery of fine details and closer alignment with ground-truth

surfaces. Additionally, NC increases by 1% and F-Score by nearly 2%, indicating improved geometric structure preservation. At resolution 128, our model performs a comparable performance as SALS on ScanNet while clearly outperforming it on Matterport3D, demonstrating the effectiveness of our region-aware features in capturing expressive details.

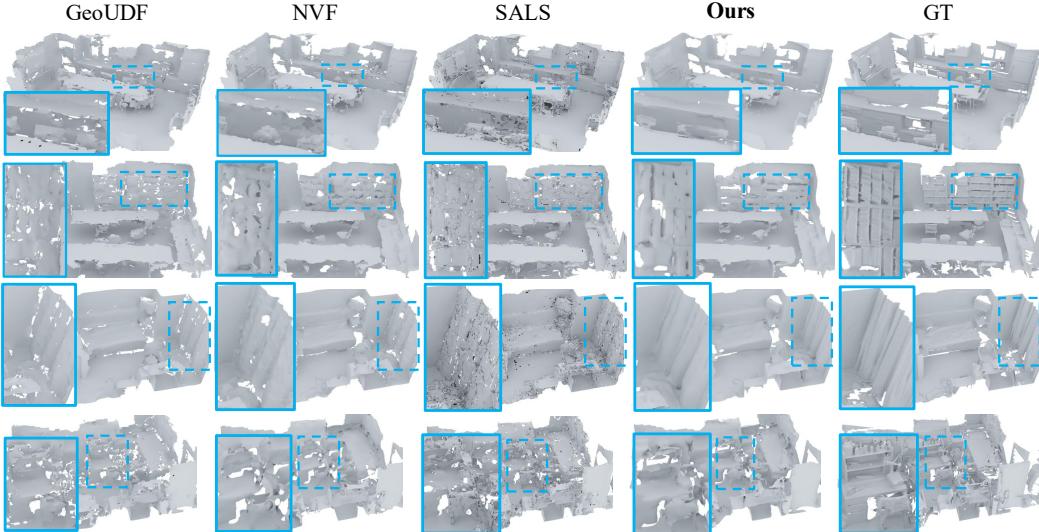


Figure 4: **ScanNet (Dai et al., 2017) Visualization.** All methods use a resolution of 128. Portions of the walls are removed to highlight interior details and more distinguishable structures.

Fig. 4 and Fig. 5 demonstrate the effectiveness of our model in room-level reconstruction. SALS often generates spurious structures and incomplete geometry, mainly due to the reduced number of input points and the limited capacity of its simple network. GeoUDF recovers the overall layout but produces noticeable distortions and fragmented artifacts near object boundaries, which cannot be corrected through standard post-processing because of incorrect local topology. NVF alleviates some of these issues by reducing missing regions, though it still struggles with complex multi-layered structures such as bookshelves, curtains, and clusters of pillows. In contrast, our method yields smooth and contiguous surfaces, reliably separates adjacent objects, and faithfully reconstructs fine details—including thin structures like curtains and pillowcases, while consistently preserving large planar surfaces such as beds and desks in densely cluttered environments.

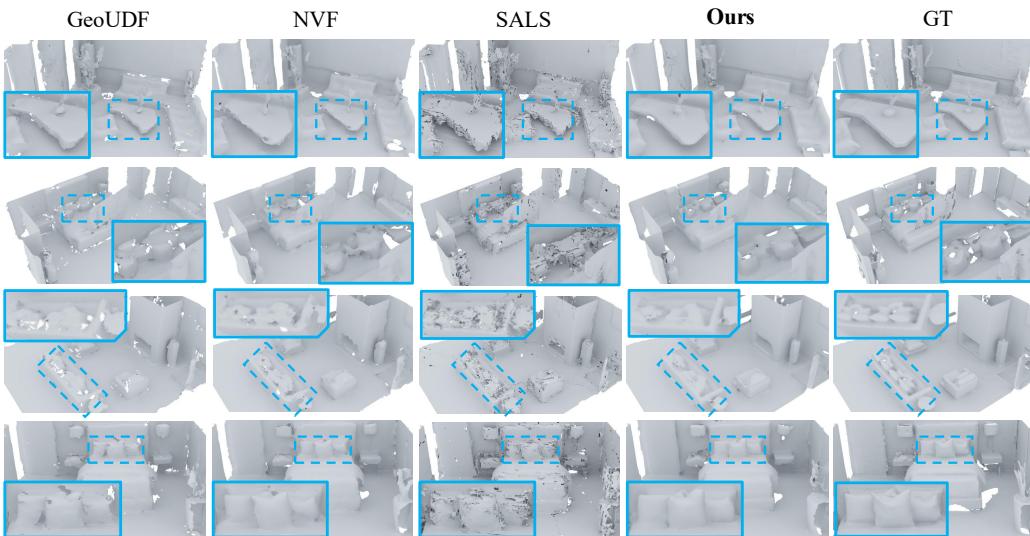


Figure 5: **Matterport3D (Chang et al., 2017) Visualization.** All methods use a resolution of 128. Portions of the walls are removed to highlight interior details and more distinguishable structures.

432 4.4 ABLATION STUDY  
433434 We conducted all ablation experiments on ScanNet, with default meshing resolution 256.  
435436 **Table 5: Method Ablation.** Report  $CD_{L_1}$  and  $CD_{L_2}$ , F-S. with thresholds 0.005, 0.01, and NC.

Methods	$CD_{L_1} \times 10^{-3}$		$CD_{L_2} \times 10^{-6}$		F-S. (%)		NC (%)	Params. (M)
	Mean↓	Median↓	Mean↓	Median↓	F-S. <sup>0.005</sup> ↑	F-S. <sup>0.01</sup> ↑		
$B_1$	1.92	1.91	6.33	6.02	95.3	99.5	89.0	11.32
$B_2$	1.92	1.91	6.30	6.00	95.3	99.5	89.2	11.36
$C_0$	2.10	2.07	8.20	7.73	93.7	99.1	88.6	11.37
$C_1$	1.90	1.90	6.33	6.06	95.3	99.5	89.8	11.36
$C_2$	<b>1.86</b>	<b>1.87</b>	<b>6.03</b>	<b>5.82</b>	<b>95.7</b>	<b>99.6</b>	<b>90.1</b>	<b>11.37</b>

444  
445 **Method Design** All variants in Table 5 are controlled experiments sharing the exact same backbone  
446 (**PointTransformer V2**) and similar network capacity, with the only differences being the small  
447 modules designed to produce query features. We denote POCO-style baselines (represent positional  
448 encoding interpolation) include: (1) Direct UDF regression, which failed to converge after 300  
449 epochs; (2) Incorporating query point locations as skip connections, similar to NDF and our method;  
450 (3) Replacing POCO’s attention module with our AttSet module, based on (2), denoted as  $B_1$  and  $B_2$ .  
451 We denote  $C_0$  as the interpolation baseline (Eq. 4),  $C_1$  as the architecture in Eq. 1 (Sec. 3.1), and  $C_2$   
452 as the primitive-based architecture (Eqs. 2–3). Notably, without the paradigm of Eq. 1,  $C_2$  reduces to  
453 a weighted sum of point-wise features.454 Specifically,  $C_0 \rightarrow C_1$  isolates the effect of narrow region aggregation versus pure interpolation,  
455  $C_1 \rightarrow C_2$  isolates the contribution of primitive features, and  $C_1$  versus  $B_1/B_2$  isolates the impact  
456 of region-aware design versus positional-encoding enhancements. As shown in Table 5, the metrics  
457 for  $C_0$ ,  $C_1$ , and  $C_2$  exhibit a clear upward trend.  $C_0$  serves as a baseline, whereas  $C_1$  significantly  
458 surpasses it, demonstrating the benefit of narrow region features.  $C_2$  further improves upon  $C_1$ ,  
459 indicating that primitive-based aggregation effectively captures finer details. Comparisons between  
460  $B_1$  and  $B_2$  reveal minimal impact from the AttSet module, while  $C_1$  outperforms both, highlighting  
461 that narrow region features complement point-wise features to achieve superior representations.462 **Table 6: Robustness Ablation.** Report  $CD_{L_1}$  and  $CD_{L_2}$ , F-S. with thresholds 0.005, 0.01, and NC.

Methods	Conditions	$CD_{L_1} \times 10^{-3}$		$CD_{L_2} \times 10^{-4}$		F-S. (%)		NC (%)
		Mean↓	Median↓	Mean↓	Median↓	F-S. <sup>0.005</sup> ↑	F-S. <sup>0.01</sup> ↑	
SALS	Clean	2.02	2.02	0.075	0.069	94.7	99.3	87.1
	Noise	4.80	4.74	0.370	0.349	61.4	91.9	84.0
	Noiser	8.49	8.43	1.06	1.04	23.2	62.3	74.4
	Sparse	4.98	4.92	0.467	0.426	61.5	88.9	83.2
GeoUDF	Clean	2.05	2.04	0.078	0.075	93.8	99.0	89.3
	Noise	4.03	4.00	0.240	0.236	69.8	95.7	75.5
	Noiser	10.4	10.4	1.84	1.85	24.6	51.6	56.2
	Sparse	2.45	2.42	0.130	0.122	90.2	97.5	87.6
NVF	Clean	2.03	1.98	0.117	0.069	94.6	99.3	88.9
	Noise	2.88	2.84	0.171	0.142	87.2	97.6	85.6
	Noiser	5.20	5.23	0.503	0.469	62.0	86.9	78.6
	Sparse	2.77	2.73	0.199	0.175	87.2	96.1	85.1
Ours	Clean	1.86	1.87	0.060	0.058	95.7	99.6	90.1
	Noise	2.51	2.46	0.112	0.104	90.9	98.6	88.4
	Noiser	4.91	4.89	0.418	0.412	64.8	89.0	80.0
	Sparse	2.28	2.25	0.125	0.105	91.6	98.1	89.4

481 **Robustness** We evaluated robustness on ScanNet under three challenging conditions: (1) Gaussian  
482 noise with  $\sigma = 0.005$  (“Noise”), (2) Gaussian noise with  $\sigma = 0.025$  (“Noiser”), and (3) sparse input  
483 with 3,500 points, i.e., one-third of the original input (“Sparse”).484 As shown in Table 6, under the *Noise* setting, NVF exhibits a 41.9% degradation on  $CD_{L_1}$ , whereas  
485 our method achieves a noticeably smaller decline, outperforming NVF by 7%. Under the *Sparse*

486 setting, NVF’s degradation increases to 36.5%, while our method again demonstrates enhanced  
 487 robustness with a **13.9%** smaller performance drop. In the *Noiser* setting, the degradation of our  
 488 method becomes *comparable* to that of NVF. GeoUDF, however, displays substantial instability in  
 489 noisy environments: its performance deteriorates by 96.5% under Noise and 407% under Noiser,  
 490 whereas our method reduces these declines by **61.6%** and **243%**, respectively. Under sparse sampling,  
 491 our degradation remains *close* to that of GeoUDF.

492 These results highlight that our formulation effectively exploits *independent neighborhood geometric*  
 493 *cues* to suppress both noise and sparsity as long as local geometry remains partially preserved.  
 494 When the local geometry is heavily destroyed by extreme noise (Noiser), our architecture naturally  
 495 transitions to relying more on *broader-receptive-field point-wise features via attention*, yielding a  
 496 controlled and acceptable degradation. Overall, the method demonstrates **balanced robustness across**  
 497 **multiple distribution drifts**, maintaining strong performance in both noisy and sparse environments.

498  
 499 Table 7: **Primitive Construction.** Report  $CD_{L_1}$  and  $CD_{L_2}$ , F-S. with 0.005, 0.01, and NC.

	$CD_{L_1} \times 10^{-3}$		$CD_{L_2} \times 10^{-6}$		F-S. (%)		NC (%)
	Mean↓	Median↓	Mean↓	Median↓	F-S. <sup>0.005</sup> ↑	F-S. <sup>0.01</sup> ↑	
Euclidean Distance-Based	2.11	2.11	8.43	8.21	93.1	99.1	88.6
Spherical Projection	<b>1.86</b>	<b>1.87</b>	<b>6.03</b>	<b>5.82</b>	<b>95.7</b>	<b>99.6</b>	<b>90.1</b>

500  
 501  
 502  
 503  
 504  
 505 **Primitive Construction** To demonstrate the effectiveness of our proposed primitive construction  
 506 method, we compare it with a basic Euclidean distance-based partitioning method on the ScanNet.  
 507

508 As shown in the Table 7, Euclidean distance-based partitioning markedly degrades structural fidelity  
 509 and thus harms reconstruction quality. Our spherical-projection grouping can produce ambiguity when  
 510 query points lie outside a multilayer structure, but this drawback is mitigated because query points  
 511 sampled between layers enable unambiguous separation of the two surfaces. During aggregation,  
 512 attention further downweights ambiguous primitives and upweights those originating from between-  
 513 layer queries, so the net effect of such ambiguity on the final reconstruction is small.

514 Table 8: **Primitive Construction.** Report  $CD_{L_1}$  and  $CD_{L_2}$ , F-S. with 0.005, 0.01, and NC.

	$CD_{L_1} \times 10^{-3}$		$CD_{L_2} \times 10^{-6}$		F-S. (%)		NC (%)
	Mean↓	Median↓	Mean↓	Median↓	F-S. <sup>0.005</sup> ↑	F-S. <sup>0.01</sup> ↑	
Segment	1.96	1.95	6.92	6.47	95.2	99.4	89.7
Triplet Plane	<b>1.86</b>	<b>1.87</b>	<b>6.03</b>	<b>5.82</b>	<b>95.7</b>	<b>99.6</b>	<b>90.1</b>
Four-Point Patch	2.06	2.04	7.54	7.05	94.3	99.3	88.3

521  
 522 **Primitive Type** To demonstrate the effectiveness of our proposed triplet-plane primitive, we  
 523 compare it with segment and 4-point patch on the ScanNet.

524 The choice of three points is motivated by geometric principles: in 3D space, the simplest non-  
 525 degenerate local structure is defined by three non-collinear points, which uniquely determines a  
 526 plane. Theoretically, simple 3D primitives provide one additional geometric degree of freedom while  
 527 avoiding the excessive expressiveness of large patches. This intermediate complexity allows them  
 528 to be learned reliably with low sample complexity and then composed to approximate rich surface  
 529 structures. Results shown in Table 8 could support above claims.

## 530 5 CONCLUSION

531  
 532 In this paper, we present RegionUDF, a region-aware UDF framework that explicitly incorporates  
 533 neighborhood information for each query region. We first design a discriminative region repre-  
 534 sentation that fuses broad contextual point-wise features with narrow region features, providing  
 535 complementary information. Building on this, we propose a primitive-based region representation  
 536 that decomposes neighborhoods into triplet-defined primitives, enabling finer characterization of local  
 537 geometry. Extensive experiments on synthetic and real-world benchmarks show that RegionUDF  
 538 achieves superior reconstruction accuracy at both object and scene levels, with strong cross-domain  
 539 generalization across diverse datasets and scales.

540     **Reproducibility Statement** The formulas in the main text Section 3 directly correspond to our  
 541 code. In Appendix Section A.2, we detail the backbone architecture, MLP configuration, and feature  
 542 dimensions. We also provide training parameters and evaluation settings, including the meshing  
 543 strategy, all based on official implementations. Furthermore, we describe the official baseline  
 544 implementations used and how we retrained them in Appendix Section A.3.3.

545  
 546     **REFERENCES**  
 547

548     Tristan Aumentado-Armstrong, Stavros Tsogkas, Sven Dickinson, and Allan D Jepson. Representing  
 549     3D Shapes with Probabilistic Directed Distance Fields. Proc. IEEE Int. Conf. Comput. Vis. Pattern  
 550     Recognit., pp. 19343–19354, 2022.

551     Abhishek Badki, Orazio Gallo, Jan Kautz, and Pradeep Sen. Meshlet Priors for 3D Mesh Reconstruc-  
 552     tion. Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit., 2020.

553  
 554     Alexandre Boulch and Renaud Marlet. POCO: Point Convolution for Surface Reconstruction. In  
 555     Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit., 2022.

556     Jonathan C Carr, Rick K Beatson, J B Cherrie, T J Mitchell, W Richard Fright, Bruce C McCallum,  
 557     and T R Evans. Reconstruction and Representation of 3D Objects with Radial Basis Functions. In  
 558     ACM SIGGRAPH, 2001.

559  
 560     Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Nießner, Manolis Savva,  
 561     Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3D: Learning from RGB-D Data in Indoor  
 562     Environments. arXiv:1709.06158, 2017.

563  
 564     Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li,  
 565     Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu.  
 566     ShapeNet: An Information-Rich 3D Model Repository. arXiv:1512.03012, 2015.

567  
 568     Julian Chibane, Thimo Alldieck, and Gerard Pons-Moll. Implicit functions in feature space for 3d  
 569     shape reconstruction and completion. In Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.,  
 2020a.

570  
 571     Julian Chibane, Aymen Mir, and Gerard Pons-Moll. Neural Unsigned Distance Fields for Implicit  
 572     Function Learning. In Proc. Adv. Neural Inf. Process. Syst., 2020b.

573  
 574     Christopher B. Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3D-R2N2:  
 575     A Unified Approach for Single and Multi-view 3D Object Reconstruction. In Proc. Eur. Conf.  
 576     Comput. Vis., 2016.

577  
 578     Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias  
 579     Nießner. ScanNet: Richly-annotated 3D Reconstructions of Indoor Scenes. In Proc. IEEE Int.  
 580     Conf. Comput. Vis. Pattern Recognit., 2017.

581  
 582     Miguel Fainstein, Viviana Siless, and Emmanuel Iarussi. DUDF: Differentiable Unsigned Distance  
 583     Fields with Hyperbolic Scaling. Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit., pp. 4484–  
 4493, 2024.

584  
 585     Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit Geometric  
 586     Regularization for Learning Shapes. arXiv, 2020.

587  
 588     Gael Guennebaud and Markus Gross. Algebraic Point Set Surfaces. In ACM SIGGRAPH, 2007.

589  
 590     Benoit Guillard, Federico Stella, and Pascal Fua. MeshUDF: Fast and Differentiable Meshing of  
 591     Unsigned Distance Field Networks. In Proc. Eur. Conf. Comput. Vis., 2022.

592  
 593     Rana Hanocka, Gal Metzer, Raja Giryes, and Daniel Cohen-Or. Point2Mesh : A Self-Prior for  
 594     Deformable Meshes. ACM SIGGRAPH, 39(4), 2020.

595  
 596     Trevor Houchens, Cheng-You Lu, Shivam Duggal, Rao Fu, and Srinath Sridhar. Neuralodf: Learning  
 597     Omnidirectional Distance Fields for 3D Shape Representation. arXiv, 2022.

594 Qingyong Hu, Bo Yang, Linhai Xie, Stefano Rosa, Yulan Guo, Zhihua Wang, Niki Trigoni, and  
 595 Andrew Markham. RandLA-Net: Efficient Semantic Segmentation of Large-Scale Point Clouds.  
 596 Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit., 2020.

597

598 Jiahui Huang, Hao-Xiang Chen, and Shi-Min Hu. A Neural Galerkin Solver for Accurate Surface  
 599 Reconstruction. ACM Trans. Graph., 41(6):1–16, 2022.

600 Jiahui Huang, Zan Gojcic, Matan Atzmon, Or Litany, Sanja Fidler, and Francis Williams. Neural  
 601 Kernel Surface Reconstruction. In Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit., 2023.

602

603 Michael Kazhdan and Hugues Hoppe. Screened Poisson Surface Reconstruction. ACM Trans.  
 604 Graph., 32(3):1–13, 2013.

605 Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson Surface Reconstruction. In Proc.  
 606 4th SGP, 2006.

607

608 Sebastian Koch, Albert Matveev, Zhongshi Jiang, Francis Williams, Alexey Artemov, Evgeny  
 609 Burnaev, Marc Alexa, Denis Zorin, and Daniele Panozzo. ABC: A Big CAD Model Dataset for  
 610 Geometric Deep Learning. In Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit., 2019.

611 David Levin. The Approximation Power of Moving Least-Squares. Math. Comput., 67(224):  
 612 1517–1531, 1998.

613

614 Jiehong Lin, Xian Shi, Yuan Gao, Ke Chen, and Kui Jia. CAD-PU: A Curvature-Adaptive Deep  
 615 Learning Solution for Point Set Upsampling. arXiv:2009.04660, 2020.

616 Yujie Lu, Long Wan, Nayu Ding, Yulong Wang, Shuhan Shen, Shen Cai, and Lin Gao. Unsigned  
 617 Orthogonal Distance Fields: An Accurate Neural Implicit Representation for Diverse 3D Shapes.  
 618 Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit., pp. 20551–20560, 2024.

619

620 Baorui Ma, Zhizhong Han, Yu-shen Liu, and Matthias Zwicker. Neural-Pull : Learning Signed  
 621 Distance Functions from Point Clouds by Learning to Pull Space onto Surfaces. In Proc. Int. Conf.  
 622 Mach. Learn., 2021.

623 Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger.  
 624 Occupancy Networks: Learning 3D Reconstruction in Function Space. In Proc. IEEE Int. Conf.  
 625 Comput. Vis. Pattern Recognit., 2019.

626 Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF:  
 627 Learning Continuous Signed Distance Functions for Shape Representation. In Proc. IEEE Int.  
 628 Conf. Comput. Vis. Pattern Recognit., 2019.

629

630 Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolu-  
 631 tional Occupancy Networks. In Proc. Eur. Conf. Comput. Vis., 2020.

632 Siddhant Ranade, Gonçalo Dias Pais, Ross Tyler Whitaker, Jacinto C Nascimento, Pedro Miraldo,  
 633 and Srikumar Ramalingam. SurfR: Surface Reconstruction with Multi-scale Attention. arXiv,  
 634 2025.

635

636 Siyu Ren and Junhui Hou. Shape as line segments: Accurate and flexible implicit surface representa-  
 637 tion. "Proc. Int. Conf. Learn. Representations", 2025.

638

639 Siyu Ren, Junhui Hou, Xiaodong Chen, Ying He, and Wenping Wang. GeoUDF: Surface Recon-  
 640 struction from 3D Point Clouds via Geometry-guided Distance Representation. In Proc. IEEE Int.  
 641 Conf. Comput. Vis., 2023.

642

643 Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-Motion Revisited. In Proc.  
 644 IEEE Int. Conf. Comput. Vis. Pattern Recognit., 2016.

645

646 Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise  
 647 View Selection for Unstructured Multi-View Stereo. In Proc. Eur. Conf. Comput. Vis., 2016.

648

649 Justin Sirignano and Konstantinos Spiliopoulos. DGM: A deep learning algorithm for solving partial  
 650 differential equations. J Comput Phys, 375:1339–1364, 2018.

648 Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit  
 649 Neural Representations with Periodic Activation Functions. *Proc. Adv. Neural Inf. Process. Syst.*,  
 650 33:7462–7473, 2020.

651

652 Jiapeng Tang, Jiabao Lei, Dan Xu, Feiying Ma, Kui Jia, and Lei Zhang. SA-ConvOnet: Sign-Agnostic  
 653 Optimization of Convolutional Occupancy Networks. In *Proc. IEEE Int. Conf. Comput. Vis.*, 2021.

654

655 Zhen Wang, Shijie Zhou, Jeong Joon Park, Despoina Paschalidou, Suya You, Gordon Wetzstein,  
 656 Leonidas Guibas, and Achuta Kadambi. Alto: Alternating Latent Topologies for Implicit 3D  
 657 Reconstruction. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.*, pp. 259–270, 2023.

658

659 Francis Williams, Zan Gojcic, Sameh Khamis, Denis Zorin, Joan Bruna, Sanja Fidler, and Or Litany.  
 660 Neural Fields as Learnable Kernels for 3D Reconstruction. In *Proc. IEEE Int. Conf. Comput. Vis.  
 661 Pattern Recognit.*, 2022.

662

663 Xiaoyang Wu, Yixing Lao, Li Jiang, Xihui Liu, and Hengshuang Zhao. Point Transformer V2:  
 664 Grouped Vector Attention and Partition-based Pooling. In *Proc. Adv. Neural Inf. Process. Syst.*,  
 665 2022.

666

667 Qiangeng Xu, Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. DISN: Deep  
 668 Implicit Surface Network for High-quality Single-view 3D Reconstruction. *Proc. Adv. Neural Inf.  
 669 Process. Syst.*, 2019.

670

671 Xianghui Yang, Guosheng Lin, Zhenghao Chen, and Luping Zhou. Neural Vector Fields: Implicit  
 672 Representation by Explicit Learning. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.*,  
 673 2023.

674

675 Jianglong Ye, Yuntao Chen, Naiyan Wang, and Xiaolong Wang. GIFS: Neural Implicit Function for  
 676 General Shape Representation. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.*, 2022.

677

678 Junsheng Zhou, Baorui Ma, Yu-Shen Liu, Yi Fang, and Zhizhong Han. Learning Consistency-Aware  
 679 Unsigned Distance Functions Progressively from Raw Point Clouds. In *Proc. Adv. Neural Inf.  
 680 Process. Syst.*, 2022.

681

682 Junsheng Zhou, Baorui Ma, Shujuan Li, Yu-Shen Liu, and Zhizhong Han. Learning a More  
 683 Continuous Zero Level Set in Unsigned Distance Fields through Level Set Projection. In *Proc.  
 684 IEEE Int. Conf. Comput. Vis.*, 2023.

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 **A APPENDIX**  
703704 **A.1 PRELIMINARY**  
705706 Naturally many works (Ye et al., 2022; Chibane et al., 2020b) follow the ONet (Mescheder et al.,  
707 2019) paradigm and the distance  $d_q$  between query point  $q$  and nearest surface is formulated as:

708 
$$d_q = \varphi \left( \sum_{g \in N_q} w_{gq} f_g \right), \quad (4)$$
  
709  
710

711 where  $f_g$  denotes the grid feature of grid  $g$  derived from point clouds  $\mathcal{P}$  through the 3D Conv  
712 backbone,  $w_{pq}$  is a scalar obtained from positions of the query point  $q$  and the grid  $g$ , as trilinear  
713 interpolation,  $N_q$  means the neighbors of the query point  $q$ , and  $\varphi$  means a Regression Network,  
714 e.g. MLPs. It is well recognized that grid division imposes a fixed discretization, which inevitably  
715 introduces quantization errors and leads to the loss of fine-grained geometric details.716 Recent point-based POCO (Boulch & Marlet, 2022) introduce attention-based interpolation to  
717 compute weights for feature aggregation:

718 
$$d_q = \varphi(A(\phi(f_p|p))), \quad (5)$$
  
719

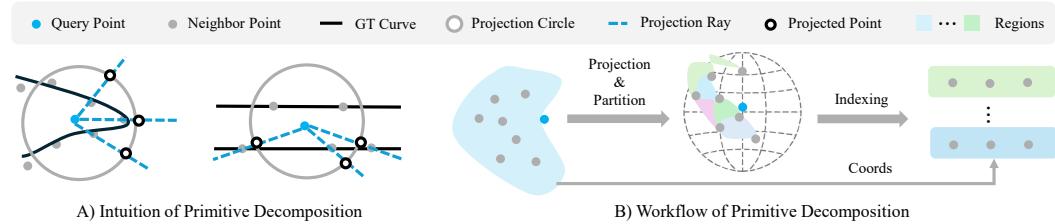
720 where  $p$  is a neighbor of query  $q$ ,  $A$  means attentive pooling, and  $\phi$  is a MLP to transform positions  
721 and point-wise features. However, these approaches largely ignore the inherent structural relationships  
722 within query neighborhoods, resulting a limited characterization of local regions. Point-Wise  
723 features capture broad contextual information over large receptive fields but often oversmooth local  
724 neighborhood details. In contrast, narrow region features faithfully describe neighborhood structures  
725 but may be prone to local extremes. These two levels of representation, global and local, are inherently  
726 complementary, and their integration yields a more compact and informative regional descriptor.  
727 Building on this, we propose a framework that explicitly combines point-wise features with narrow  
728 region representations to construct a discriminative representation of the query neighborhood.729 **A.2 IMPLEMENTATION**  
730731 We present the implementation of our method. For clarity, the extraction of point-wise features via  
732 PointTransformer V2 is omitted to highlight the core implementation (shown in Fig. 7).733 **Implementation of Discriminative Region Representation.** In practical implementation, we  
734 uniformly sample 8 reference points per query and encode the fused geometric and contextual  
735 features  $\{f_p, r_p\}$  through a lightweight two-layer MLP as shown in Eq.1. These per-point feature  
736 pairs are then aggregated using the attention-based AttSet mechanism (Hu et al., 2020) to form  
737 the final query feature  $f_q$ , which serves as input to the unsigned distance regressor. Since  $r_p$  is  
738 computed from relative positions and  $f_p$  is extracted by a translation-invariant backbone, the resulting  
739  $f_q$  inherits translation invariance, thereby ensuring robustness and stability in region representation.  
740 Additionally, the query point  $q$  is incorporated via a skip connection to retain explicit positional  
741 information throughout the unsigned distance regression process.742 **Implementation of Primitive-Based Region Representation.** We first reinforce the intuition  
743 behind primitive-based region description. As discussed in Section 3.1, narrow region features con-  
744 fined to the query region complement point-wise features, yielding more faithful and discriminative  
745 representations. A natural solution is to adopt finer-grained partitioning and feature learning, where  
746 the key challenge lies in constructing these fine-grained regions, termed primitives. While Euclidean  
747 distance-based neighbor selection is the simplest strategy, Fig. 6 illustrates that spherical projection  
748 more effectively aggregates patches on the same surface. The key lies in forming triplet primitives by  
749 angle-sorted adjacent points. While it introduces ambiguity in multi-layer structures, where patches  
750 may span different surfaces, such limitations are inherent to Euclidean distance-based method:751 

- 752 • When the query point is outside two layers, the primitives may treat them as a single surface,  
753 yet this limitation is not unique to our method and also appears in Euclidean distance-based  
754 partitioning.
- 755 • When the query point lies between layers, our formulation can effectively separate them,  
which Euclidean distance-based method cannot achieve.

756  
 757  
 758  
 759  
 760

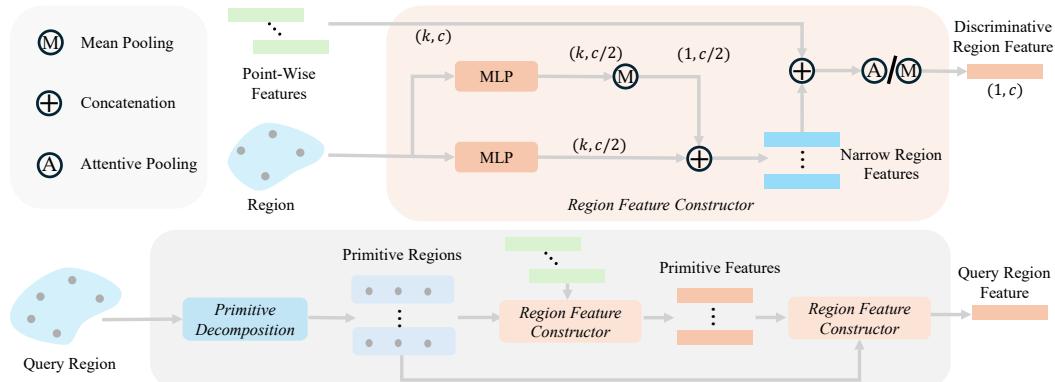
- Moreover, during primitive aggregation, the attention mechanism operates over primitive centers, enabling the model to reinterpret ambiguous primitives as representing the “in-between” region of two layers. These primitives are then compensated by those constructed from query points located between the layers, which provide clear and consistent cues.

761 Thus, ambiguity may introduce minor artifacts, such as closer surfaces or a few unintended connec-  
 762 tions, but overall the method retains clear advantages over Euclidean distance-based partitioning.



770  
 771  
 772 **Figure 6: Primitive Decomposition.** Subfigure A illustrates primitive decomposition in 2D. The  
 773 projection method preserves sharp details and discrete patches by prioritizing surface proximity over  
 774 Euclidean distance. Subfigure B shows that points are angle-sorted for primitive construction, while  
 775 projection serves only for partitioning—the final regions remain defined in Euclidean space.

776 Next, we describe the network framework for primitive-based region representation. As illustrated  
 777 in Fig. 7, the query region is first decomposed into individual primitives, each corresponding to  
 778 a distinct subregion. These primitives, together with point-wise features, are input into the region  
 779 feature constructor to learn primitive features. The primitives are then treated as points within the  
 780 query region, with their features corresponding to the associated point-wise features, and re-input  
 781 into the region feature constructor to produce the final query region feature. Note that, given the  
 782 simplicity of triplet primitives, mean pooling is sufficient in the region feature constructor, whereas  
 783 learning the final query region features requires attentive pooling.



800 **Figure 7: Primitive-Based Region Description.** A query region is decomposed into primitives, each  
 801 defined by its points and associated point-wise features. Primitive features are built via mean pooling  
 802 in Region Feature Constructor, while final query region features need attention pooling. Primitive  
 803 features serve as point-wise features, while primitives serve as points within the region.

804 **Loss Function** Both the UDF value and UDF gradient are essential for optimization. The distances  
 805 are optimized using  $\ell_1$  loss, while the gradient of the field is used to compute the cosine similarity:

$$l_d = \|d_{pd} - d_{gt}\|_1, \quad (6)$$

$$l_g = 1 - |\langle \vec{g}_{pd} \cdot \vec{g}_{gt} \rangle|,$$

809 where  $l_d$  denotes the loss of query distances,  $l_g$  means the similarity loss of gradient. Since the UDF  
 is not differential at the iso-surface, the normal vector constraint on the surface is omitted. The loss

810 function is the combination of the two, as  $l = \lambda_1 l_d^{off} + \frac{\lambda_1}{10} l_d^{on} + \lambda_2 l_g^{off}$ . In our experiments,  $\lambda_1, \lambda_2$   
 811 are hyper-parameters for weighting, *on* means on-surface queries, *off* means off-surface queries.  
 812

813 **Evaluation Setting** During testing, we input sparse point clouds and adopt the meshing strategy  
 814 introduced in MeshUDF (Guillard et al., 2022). During evaluation, we follow the protocol established  
 815 in GIFS (Ye et al., 2022), because the MeshUDF meshing strategy can produce artifacts at low  
 816 resolutions, especially when query points lie between multiple layers, we adopt the more robust,  
 817 though less efficient, GeoUDF (Ren et al., 2023) meshing strategy at a resolution of 128. As evaluation  
 818 metrics, including Chamfer Distance, F-Score, and Normal Consistency, are computed by uniformly  
 819 sampling points from both the predicted meshes (or dense point clouds) and the ground truth meshes.  
 820

821 **Params. Setting** Our framework employs PointTransformer V2 (Wu et al., 2022) as the backbone  
 822 with default settings except initial grid size and downsample ratio. Two configurations are used: for  
 823 simple watertight shapes with 3k points, the initial grid size is 0.015 with a downsample ratio of 2  
 824 per layer; for complex open-surface reconstruction, the initial grid size is 0.01 with a downsample  
 825 ratio of 2.5. Each layer uses 16 neighbors to ensure a large receptive field, and the number of nearest  
 826 neighbors  $K$  per query point is fixed at 8 across all experiments.  
 827

828 The point-wise features are 48-dimensional vectors derived from the input point cloud through  
 829 backbone. As MLP in Region Feature Constructor, we set it with one linear layer and one ReLU layer,  
 830 transforming a 4-dimensional position vector (with the length of the vector) to a 24-dimensional  
 831 feature vector. Note that the second optional mean pooling in the region feature constructor comprises  
 832 two Linear layers, two ReLU layers, followed by a mean pooling layer. The attentive pooling module  
 833 is formulated as:  $f = softmax(linear(f)) * f$ . The final regression network is 4-layers MLP  
 834 with channels [51,256,32,32,1]. Note that, aside from the point-based backbone which uses batch  
 835 normalization for feature extraction, our core implementation does not rely on BN.  
 836

### A.2.1 TRAINING SETTING

837 We adopt the ADAMW optimizer with default parameters, initializing the learning rate at  $10^{-3}$ . A  
 838 warm-up phase of 2000 steps is applied at the beginning of training. Furthermore, the learning rate is  
 839 decayed to 30% of its current value at epochs 30, 70, 110, and 200 for all datasets. For dataset-related  
 840 settings refer to the Experiment section 4. For loss,  $\lambda_1 = 1, \lambda_2 = 1e^{-3}$ . For Result 2, we trained on  
 841 2 RTX3090 GPUs with batch 4 for 56 hours, with other two datasets results 4 trained on 2 RTX3090  
 842 GPUs with batch 4 for 24 hours. The CPU core is Intel Xeon Platinum 8383C CPU @ 2.70 GHz.  
 843

### A.2.2 MESHING STRATEGY

844 At the zero level set of an unsigned distance field (UDF), the derivative is undefined, leading to  
 845 unstable gradients as the predicted values approach this boundary. This instability complicates  
 846 the simultaneous optimization of distance and gradient losses (Eq. 6) in the vicinity of the zero  
 847 level set. Mesh extraction methods for UDFs, such as the Marching Cubes adaptation used in  
 848 MeshUDF (Guillard et al., 2022), depend on reliable gradient computations to assign relative signs  
 849 via gradient dot products. When gradients fluctuate excessively near surface, mesh extraction suffers  
 850 from missing or erroneous edge intersections, producing holes and spurious artifacts. As a result, *the  
 851 meshing strategies can amplify gradient extremes, yielding visible gaps or distorted geometry*.  
 852

853 In our implementation, we adopt the MeshUDF pipeline to extract iso-surfaces from the predicted  
 854 UDF. For object-level reconstruction, we use a voxel grid of resolution 128 and discard vertices  
 855 whose distance estimates exceed  $\frac{1}{3}$  or  $\frac{1}{1.5}$  of the voxel size from the zero level set. For room-level  
 856 reconstruction, we employ resolutions of 128 and 256. Since MeshUDF often introduces artifacts at  
 857 low resolutions—particularly when query segments are inclined to the plane without intersection—we  
 858 instead use GeoUDF at resolution 128 despite its higher cost.  
 859

### A.2.3 EVALUATION METRICS

860 Evaluation metrics included Chamfer Distance (CD) and F-Score (F-S.) and Normal Consistency (NC).  
 861 For object-level evaluation, they are computed by sampling 100k points from both the reconstructed  
 862 surfaces and the ground truth. For room-level evaluation, they are computed by sampling 500k  
 863

864 points from both the reconstructed surfaces and 100k points from the ground truth. **All objects are**  
 865 **normalized to the unit cube for comparison.**  
 866

867 **A.3 EXPERIMENT SETTINGS**  
 868

869 **A.3.1 DATASET INTRODUCTION**  
 870

871 We include 1 watertight and 4 open-surface reconstruction benchmarks as follows:

872 **ShapeNet** (Chang et al., 2015) is an object-level 13 classes watertight dataset processed by the DISN  
 873 (Xu et al., 2019), with train/val/test split according to 3D-R2N2 (Choy et al., 2016).

874 **ShapeNet Cars** (Chang et al., 2015) is an object-level dataset of non-watertight car models featuring  
 875 complex internal architectures of ShapeNet core dataset. Following the original NDF split (Chibane  
 876 et al., 2020b), we use 5,249 scans for training, 749 for validation, and 1,499 for testing, while  
 877 employing 10k points as input rather than occupancy data.

878 **ABC & Non-Manifold ABC** (Ren & Hou, 2025) is an object-level open-surface reconstruction  
 879 benchmark introduced by SALS (Ren & Hou, 2025). Following its protocol, the model is trained on  
 880 100 shapes from Thingi10K (Zhou & Jacobson, 2016) and evaluated on 50 shapes from ABC (Koch  
 881 et al., 2019), referred to as manifold ABC, as well as on 50 randomly selected shape pairs from ABC  
 882 that are intersected to construct non-manifold ABC.

883 **ScanNet** (Dai et al., 2017) comprises 1,513 RGB-D scans of indoor scenes with complex topology  
 884 and noisy, open surfaces. We adopt the standard split of 1,201 training and 312 validation scans;  
 885 since test-set GT meshes are unavailable, all quantitative results are performed on the validation set.

886 **Matterport3D** (Chang et al., 2017) contains 10,800 panoramic RGB-D views derived from 194,400  
 887 images across 90 building-scale scenes. We follow the published partitioning of 61 scenes for training,  
 888 11 for validation, and 18 for testing. Each building-scale scene is further divided into room-level  
 889 segments by Officials, and our experiments are conducted on the segmented dataset.

890  
 891 **A.3.2 TRAINING DATA GENERATION**  
 892

893 For both object-level and indoor datasets, we adopt the preprocessing pipeline like NDF (Chibane  
 894 et al., 2020b), normalizing each ground-truth mesh to fit within **the unit cube**. During training, we  
 895 sample query points both on and off the surface: specifically, we generate 10K on-surface points  
 896 and 100K off-surface points per scene following the NDF sampling strategy. For each query point,  
 897 we compute the nearest surface point to derive the corresponding distance and directional vector.  
 898 Notably, for watertight shapes, we follow the convention of using 3k input points. For the ABC  
 899 benchmark introduced by SALS, 40k points are sampled to maintain consistency.

900  
 901 **A.3.3 REPRODUCE BASELINES**  
 902

903 Since many learning-based UDF approaches do not provide published results on room-scale scans,  
 904 we reproduce their scene-level performance by retraining each method using the default settings  
 905 reported in their papers and evaluating under a consistent protocol. All methods are compared using  
 906 the same evaluation pipeline; specific implementation details are provided below.

907 NDF (Chibane et al., 2020b): Official implementation. We retrain NDF using its default settings and  
 908 employ MeshUDF (Guillard et al., 2022) as the meshing strategy to extract meshes.

909  
 910 GIFS (Ye et al., 2022): Official implementation. We retrain GIFS with default settings. For mesh  
 911 extraction, we set the resolution to 128 and 256 instead of the original 160.

912 NVF (Yang et al., 2023): Official implementation. We retrain NVF using its default settings. As the  
 913 journal version code is not yet available, we evaluate only the conference version.

914 GeoUDF (Ren et al., 2023): Official implementation. We retrain GeoUDF with default settings and  
 915 align the number of input points in its upsampling module to match our experiments.

916  
 917 SALs (Ren & Hou, 2025): Official implementation. We retrain SALs under default settings. *This*  
 918 *method introduces the ABC and non-manifold ABC benchmarks. Since the authors did not release*

918 *their pretrained model or detailed dataset, we reproduce the experiments with their offical codes*  
 919 *under a consistent setting and report our reproduced results.*

920  
 921 POCO (Boulch & Marlet, 2022): Official implementation. We directly adopt the query feature  
 922 construction code to implement  $B_1$  and  $B_2$  in Section 4.4.

923 **A.3.4 VISUALIZATION ANALYSIS**  
 924

925 For all figures, we use Blender 4.5.3 for rendering. Meshes are post-processed in MeshLab with three  
 926 steps: Remove Non-Manifold Edges, Fill Holes, and HC Laplacian Smooth. For GeoUDF and SALS,  
 927 which produce many non-manifold structures, these steps have limited effect, so we additionally  
 928 apply the Smooth with Angle filter ( $30^\circ$  threshold). *Overlapping and non-manifold structures can*  
 929 *yield erroneous normals in Blender, producing black shadows that are artifacts rather than missing*  
 930 *geometry, with negligible impact on distance-based metrics such as CD and F-Score.* The scaly  
 931 texture observed in GIFS is further accentuated by the rendering method.

932 **A.4 MORE RESULTS**  
 933

934 **A.4.1 CROSS-DOMAIN RECONSTRUCTION**  
 935

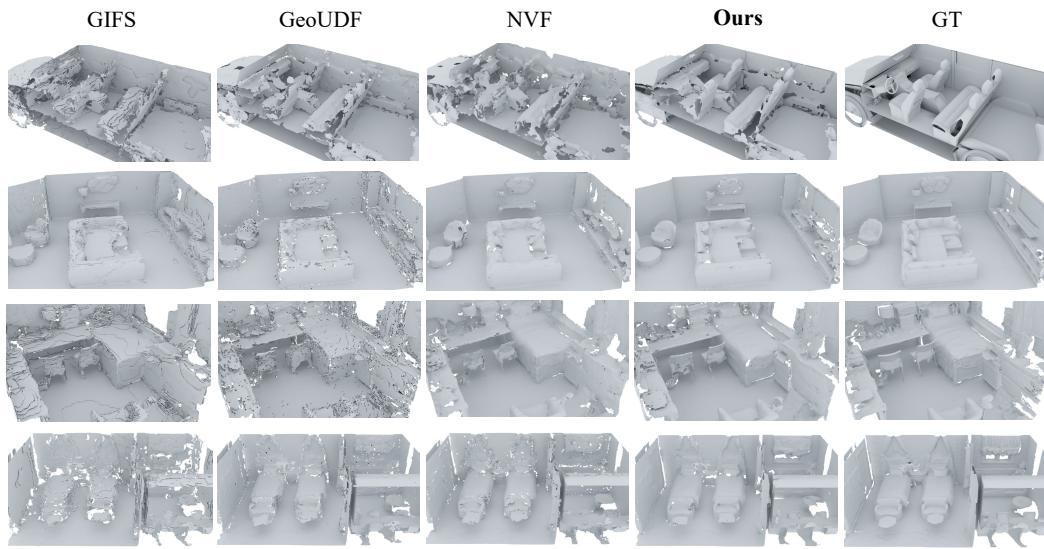
936 Beyond intra-domain reconstruction, we evaluate cross-domain generalization by training on one  
 937 dataset and testing on two others. We consider shape-to-scene, scene-to-scene, and scene-to-shape  
 938 transfers to quantify performance. These experiments demonstrate the robustness and adaptability of  
 939 our method across diverse datasets and varying scales of reconstruction scenarios.

940  
 941 **Table 9: Cross-Domain Evaluation Results.**  $CD_{L_1} \times 10^{-3}$ ,  $CD_{L_2} \times 10^{-6}$ , F-S. (%) with a  
 942 threshold of 0.01, and NC (%). Best results are in bold, and second-best results are underlined.

Trained on	$CD_{L_1} \downarrow$	$CD_{L_2} \downarrow$	F-S. $\uparrow$	NC $\uparrow$	$CD_{L_1} \downarrow$	$CD_{L_2} \downarrow$	F-S. $\uparrow$	NC $\uparrow$
ShapeNet Cars	Tested on ScanNet				Tested on Matterport3D			
NDF (Chibane et al., 2020b)	2.64	12.6	98.0	-	2.80	11.8	<b>98.9</b>	-
NDF (Mesh) (Chibane et al., 2020b)	2.76	15.0	98.3	78.0	3.32	26.2	<u>97.7</u>	81.6
GIFS (Ye et al., 2022)	2.50	11.7	98.2	87.3	2.79	12.6	<b>98.7</b>	<u>92.4</u>
NVF (Yang et al., 2023)	<b>2.11</b>	<u>9.19</u>	<b>98.8</b>	<u>87.6</u>	<u>2.33</u>	<u>10.0</u>	98.8	<u>91.9</u>
GeoUDF (Ren et al., 2023)	2.44	13.0	97.3	86.0	2.48	11.6	98.3	91.5
<b>Ours</b>	<b>2.11</b>	<b>8.92</b>	<b>98.8</b>	<b>88.0</b>	<b>2.23</b>	<b>8.59</b>	<b>99.2</b>	<b>93.3</b>
Matterport3D	Tested on ScanNet				Tested on ShapeNet Cars			
NDF (Chibane et al., 2020b)	2.40	9.60	98.9	-	3.20	13.4	<b>99.4</b>	-
NDF (Mesh) (Chibane et al., 2020b)	2.80	15.6	97.1	88.2	-	-	-	-
GIFS (Ye et al., 2022)	2.31	9.09	98.9	87.8	3.44	16.4	98.2	83.0
NVF (Yang et al., 2023)	1.98	7.90	99.2	89.0	3.31	16.2	98.0	81.3
GeoUDF (Ren et al., 2023)	2.06	<u>7.91</u>	<u>99.0</u>	<u>89.0</u>	<u>2.99</u>	<u>12.4</u>	99.2	<b>85.2</b>
<b>Ours</b>	<b>1.88</b>	<b>6.65</b>	<b>99.4</b>	<b>91.1</b>	<b>2.89</b>	<b>11.4</b>	<b>99.5</b>	<u>85.1</u>
ScanNet	Tested on ShapeNet Cars				Tested on Matterport3D			
NDF (Chibane et al., 2020b)	3.25	13.9	<u>99.2</u>	-	2.61	10.2	<u>99.2</u>	-
NDF (Mesh) (Chibane et al., 2020b)	3.76	130	98.4	74.4	-	-	-	-
GIFS (Ye et al., 2022)	3.49	17.4	97.7	81.6	2.68	11.5	98.7	91.1
NVF (Yang et al., 2023)	3.45	17.9	97.5	79.2	3.17	38.5	98.1	92.0
GeoUDF (Ren et al., 2023)	<u>3.03</u>	<u>12.7</u>	<u>99.2</u>	<u>84.6</u>	<u>2.26</u>	<u>8.69</u>	<u>99.2</u>	<u>93.7</u>
<b>Ours</b>	<b>2.90</b>	<b>11.5</b>	<b>99.4</b>	<b>85.2</b>	<b>2.14</b>	<b>7.61</b>	<b>99.5</b>	<b>94.6</b>

958 As shown in Table 9, NDF (Mesh) exhibits almost complete failure when generalizing from scene-  
 959 level to shape-level reconstruction on Matterport3D, due to large errors in NDF predicted UDF  
 960 gradients. By contrast, our method consistently outperforms existing approaches across almost all  
 961 cross-domain scenarios, and matches their Normal Consistency when transferring from scene-level  
 962 to shape-level. In the shape-to-scene setting, our model achieves performance comparable to NVF,  
 963 while lower  $CD_{L_2}$  and higher NC. Notably, our model outperforms all baselines across all metrics in  
 964 the scene-to-scene transfer setting. They demonstrate that our reconstruction results more closely  
 965 align with GT surfaces, and the larger  $CD_{L_2}$  improvement indicates a reduction in outlier artifacts.

972 The visual results Fig. 8 further demonstrate that our method excels at recovering sharp and fine-  
 973 grained details. At the scene-to-object level, it accurately reconstructs two closely positioned seats  
 974 within a vehicle. At the object-to-scene level, it yields noticeably sharper and more coherent object  
 975 boundaries. At the scene-to-scene level, it recovers challenging structures such as wall-mounted bed  
 976 curtains and crisp window-sill edges—features that competing methods fail to preserve.  
 977



996 **Figure 8: Cross-domain evaluation results.** The first row shows scene-to-object transfer from  
 997 Matterport3D to ShapeNet-Cars. The second and third rows illustrate object-to-scene transfer from  
 998 ShapeNet-Cars to Matterport3D and ScanNet, respectively. The fourth row presents scene-to-scene  
 999 transfer from ScanNet to Matterport3D. All views highlight the most distinctive regions.

#### A.4.2 ABLATION STUDY

1003 **Gradient Loss** To further illustrate the sensitivity of our method to gradient supervision, we  
 1004 conducted experiments on ScanNet, adjusting the coefficients  $\lambda_2$ . We observe that our method  
 1005 performs best with a gradient weight of  $1e^{-3}$ ; however, further reduction or even the absence of  
 1006 gradient supervision results in only minor performance degradation. This indicates that incorporating  
 1007 **region-aware** features improves robustness to the gradient supervision term.

1010 **Table 10: Gradient Weight Ablation.**  $d_C = CD_{L_2} \times 10^6$ , F-S. (%) with threshold 0.005, NC (%).

	$\lambda_2 = 0$			$\lambda_2 = 1e^{-2}$			$\lambda_2 = 1e^{-3}$			$\lambda_2 = 1e^{-4}$		
Methods	$d_C \downarrow$	F-S.↑	NC↑	$d_C \downarrow$	F-S.↑	NC↑	$d_C \downarrow$	F-S.↑	NC↑	$d_C \downarrow$	F-S.↑	NC↑
$C_1$	7.56	94.9	89.4	9.52	92.7	88.0	<b>6.33</b>	<b>95.3</b>	<b>89.8</b>	7.71	94.0	88.7
$C_2$	6.16	95.5	89.7	6.87	94.9	89.3	<b>6.03</b>	<b>95.7</b>	<b>90.1</b>	6.24	95.4	89.8

1016 Our region-aware features offer stronger representational capacity than alternative designs, enabling  
 1017 the model to learn a more accurate distance field and thereby reducing reliance on gradient supervision.  
 1018 Since the gradient field is directly derived from the distance field, once the latter is well captured,  
 1019 gradients can be reliably obtained without heavy explicit supervision. In this way, a more accurate  
 1020 distance field naturally **reduces the marginal utility** of the gradient loss as shown in Table 10.  $C_1$   
 1021 exhibits relatively *weaker feature representation capabilities*, so it **relies more heavily** on gradient  
 1022 supervision to converge to a better distance field.

1023 **KNN Ablation** Since we emphasize the complementarity of narrow region features with point-wise  
 1024 features that provide larger receptive fields, the effect of KNN size—which directly determines the  
 1025 query region—on narrow region features warrants ablation.

1026

1027 Table 11: **KNN Size.** Report  $CD_{L_1}$  and  $CD_{L_2}$ , F-S. with thresholds 0.005, 0.01, and NC.

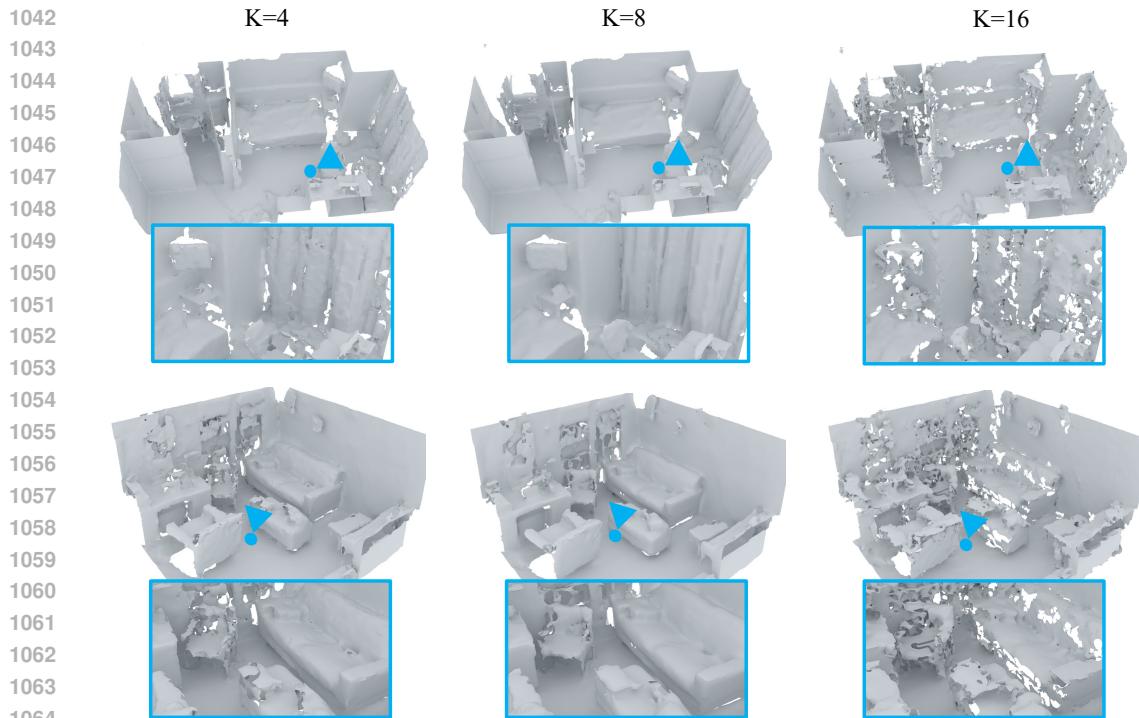
K	$CD_{L_1} \times 10^{-3}$		$CD_{L_2} \times 10^{-6}$		F-S. (%)		NC (%)
	Mean↓	Median↓	Mean↓	Median↓	$F-S.^{0.005} \uparrow$	$F-S.^{0.01} \uparrow$	
4	2.00	2.00	7.40	7.06	94.2	99.4	88.6
8	<b>1.86</b>	<b>1.87</b>	<b>6.03</b>	<b>5.82</b>	<b>95.7</b>	<b>99.6</b>	<b>90.1</b>
16	2.50	2.50	11.3	11.1	88.0	98.3	82.4

1034

1035

1036 As shown in Table 11, setting  $K = 4$  results in significantly worse performance across all metrics,  
 1037 indicating that a narrow region with too few points cannot effectively enhance query region feature  
 1038 representation. Conversely,  $K = 16$  causes the query region to cover a broader area, reducing its  
 1039 ability to capture fine-grained details and produce more misleading guidance. Setting  $K = 8$  strikes  
 1040 a balance between expressiveness and scope, achieving the best performance.

1041



1065 Figure 9: **KNN Size Ablation.** We selected scenes containing both simple structures, such as planes,  
 1066 and complex, varied geometries, highlighting the complex regions for more meaningful comparisons.  
 1067

1068 Fig. 9 illustrates the effect of  $K$  on reconstruction quality. With  $K = 4$ , performance is adequate in  
 1069 planar regions but fails in complex areas (e.g., curtain undulations), causing abnormal protrusions  
 1070 or loss of detail. In contrast,  $K = 16$  captures overly broad neighborhoods, introducing redundant  
 1071 information and misleading feature guidance. Instead of complementing point-wise features, it  
 1072 becomes detrimental.  $K = 8$  provides sufficient context to capture moderately complex structures  
 1073 without excessively enlarging the query scope or introducing redundant information.

1074

1075

1076

1077 **Meshing Strategy** The relative advantage of our method remains stable across both extractors.  
 1078 GIFS and SALS are not UDF-based and cannot be meaningfully re-meshed with UDF-specific  
 1079 pipelines. For UDF work ( NVF, GeoUDF, and ours) we ran controlled comparisons on ScanNet  
 using both MeshUDF and GeoUDF (E-MC).

1080

1081 Table 12: **Meshing Strategy.** Report  $CD_{L_1}$  and  $CD_{L_2}$ , F-S. with thresholds 0.005, and NC.

Methods	$CD_{L_1} \times 10^{-3}$	$CD_{L_2} \times 10^{-6}$	F-S. (%)	NC (%)
GeoUDF (MeshUDF)	8.96	89.0	9.6	79.4
GeoUDF (E-MC)	2.05	7.84	93.8	89.3
NVF (MeshUDF)	2.03	11.7	94.6	88.9
NVF (E-MC)	2.05	12.0	94.4	86.8
Ours (MeshUDF)	1.86	6.03	95.7	90.1
Ours (E-MC)	<b>1.81</b>	<b>5.54</b>	<b>96.5</b>	<b>91.2</b>

1089

1090

1091

The GeoUDF paper indicates that its **E-MC method is more accurate than MeshUDF**, and our practical experience confirms this holds true in most cases. The remaining differences arise from field definitions: NVF’s directions are not strictly equivalent to the UDF gradient and can exhibit directional noise near surfaces; MeshUDF’s voting step partially suppresses that noise while E-MC is more sensitive to it, producing different modes. Additionally, MeshUDF is completely unsuitable for GeoUDF and will generate a double layer near the surface. We omitted E-MC for other UDF methods at 256 resolution due to its high computational cost and instead used each method’s default reconstruction procedure from the original papers.

1099

1100

1101

1102

**Efficiency and Effectiveness** We report performance comparisons on ScanNet clean and noiser (added Gaussian noise with  $\sigma=0.025$ ). Performance metrics are based on results at 256 resolution. Here we summarize the efficiency metrics and measurement protocol concisely: inference time= wall-clock for the meshing pipeline (including distance inference + surface extraction); memory = peak GPU memory during inference with 100k query points.

1107

1108

1109

Table 13: Comparison of methods under clean/noiser settings and efficiency metrics.

Methods	$CD_{L_1} \times 10^{-3}$ (clean)	$CD_{L_1} \times 10^{-3}$ (noiser)	Inference time (128)	Inference time (256)	Params. (M)	Memory (G/100k pts)
GeoUDF	2.05	10.4	15.3	86.6	0.775	7.4
NVF	2.03	5.20	2.9	17.0	10.30	8.9
Ours	<b>1.86</b>	<b>4.91</b>	3.1	18.6	11.37	20.9

1113

1114

It can be observed that although GeoUDF has the smallest number of parameters and memory footprint, it takes too long to reconstruct a room and completely fails in high-noise environments. Our approach achieves a **balance between efficiency and effectiveness**.

1116

1117

1118

1119

1120

#### A.4.3 FAILURE CASE ANALYSIS

1121

1122

We consider a challenging, real-world scenario: multi-tier cabinets commonly found in indoor scenes. Sparse point sampling often causes closely stacked layers to be perceived as a single surface, which blurs inter-tier boundaries and degrades reconstruction quality. In the example shown in Fig.10, all methods struggle to recover a complete, clean cabinet geometry; reconstructions exhibit substantial noise and structural collapse in the inter-tier regions.

1123

1124

1125

1126

Despite producing some local artifacts and irregular layer geometry, our method preserves clear separations between tiers: reconstructed layers show distinct gaps where the point evidence indicates separation, rather than collapsing into a single thick shell. This contrast highlights the role of our primitive construction, by encoding local structure and aggregating query-centric primitives, the method is better able to maintain layer distinctions when query points fall between surfaces. By comparison, competing methods produce largely undifferentiated or chaotic interior geometry (GIFS in this case even yields a fully enclosed volume), which demonstrates that differences in primitive/patch construction materially affect the ability to resolve multi-layer structure.

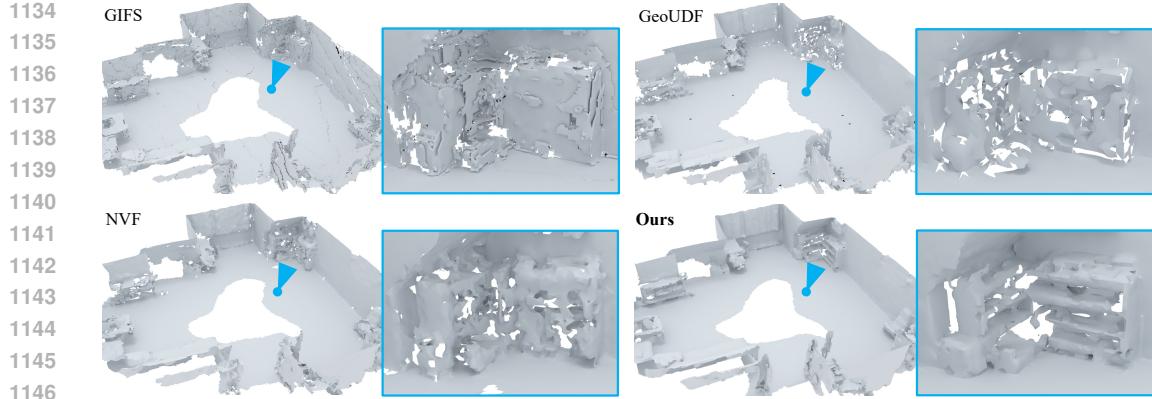


Figure 10: **Failure Case.** A typical failure case: the multi-layer cabinet. Our method preserves clear inter-layer gaps, while other approaches collapse multi-layer structures into one volume.

1150

## 1151 A.5 OTHERS

### 1153 A.5.1 LIMITATION AND FUTURE WORK

1154 Our RegionUDF demonstrates robust performance in reconstructing 3D scene surfaces from sparse  
 1155 point clouds, however, its current mesh extraction relies on adapting the MeshUDF Marching  
 1156 Cubes pipeline applied to a pseudo-signed distance field. A more promising approach would involve  
 1157 developing a dedicated algorithm for directly extracting manifolds from unsigned distance predictions,  
 1158 thereby avoiding the limitations of pseudo-SDF conversion. Furthermore, given that semantic  
 1159 labels naturally extend over continuous surfaces, integrating our implicit surface reconstruction with  
 1160 concomitant semantic segmentation via a unified network that predicts both geometry and semantics  
 1161 constitutes a compelling avenue for future research.

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187