

# Learning a More Continuous Zero Level Set in Unsigned Distance Fields through Level Set Projection

Junsheng Zhou<sup>1\*</sup>, Baorui Ma<sup>1,2,\*</sup>, Shujuan Li<sup>1</sup>, Yu-Shen Liu<sup>1†</sup>, Zhizhong Han<sup>3</sup>

School of Software, Tsinghua University, Beijing, China<sup>1</sup>

Beijing Academy of Artificial Intelligence, Beijing, China<sup>2</sup>

Department of Computer Science, Wayne State University, Detroit, USA<sup>3</sup>

zhoujs21@mails.tsinghua.edu.cn, brma@baai.ac.cn, lisj22@mails.tsinghua.edu.cn

liuyushen@tsinghua.edu.cn, h312h@wayne.edu

## Abstract

*Latest methods represent shapes with open surfaces using unsigned distance functions (UDFs). They train neural networks to learn UDFs and reconstruct surfaces with the gradients around the zero level set of the UDF. However, the differential networks struggle from learning the zero level set where the UDF is not differentiable, which leads to large errors on unsigned distances and gradients around the zero level set, resulting in highly fragmented and discontinuous surfaces. To resolve this problem, we propose to learn a more continuous zero level set in UDFs with level set projections. Our insight is to guide the learning of zero level set using the rest non-zero level sets via a projection procedure. Our idea is inspired from the observations that the non-zero level sets are much smoother and more continuous than the zero level set. We pull the non-zero level sets onto the zero level set with gradient constraints which align gradients over different level sets and correct unsigned distance errors on the zero level set, leading to a smoother and more continuous unsigned distance field. We conduct comprehensive experiments in surface reconstruction for point clouds, real scans or depth maps, and further explore the performance in unsupervised point cloud upsampling and unsupervised point normal estimation with the learned UDF, which demonstrate our non-trivial improvements over the state-of-the-art methods. Code is available at <https://github.com/junshengzhou/LevelSetUDF>.*

## 1. Introduction

Surface reconstruction from point clouds plays an important role in vision, robotics and graphics. Recently, neural implicit functions [61, 58, 57] have achieved remarkable

results by representing 3D geometries with signed distance functions (SDFs) [51, 13, 73], occupancy [57, 64, 60] or unsigned distance functions (UDFs) [89, 20, 27]. SDFs and occupancy explore the internal-to-external relations, which are limited to reconstruct closed surfaces. Researchers therefore turn to UDFs, which is a more general function, to represent shapes with arbitrary topologies.

Learning an accurate zero level set of UDF is the key factor for representing complex geometries and supporting 3D applications [52, 42, 26, 17]. However, the differential neural networks struggle from learning the zero level set where the UDF is not differentiable. This leads to unstable optimizations, which results in large errors on unsigned distances and gradients around the zero level set, causing failures in downstream 3D applications. Taking surface reconstruction for example, SOTA methods [89, 27] extract surfaces from UDFs by judging gradient directions around the zero level set, where the unreliable gradients lead to highly fragmented and discontinuous surfaces.

To resolve this issue, we introduce two novel constraints on the zero level set in forms of losses for learning accurate UDF without ground truth unsigned distances. The key insight is from the observation that despite the unreliable zero level set, the non-zero level sets learned by UDF are quite convincing, which can serve as an accurate guidance for the optimization on zero level set. We therefore propose the level set projection constraint to achieve a smoother and more continuous zero level set of UDF by enforcing the parallelism between the gradient of a query in a non-zero level set and the gradient of its projection to the zero level set. Moreover, we introduce an adaptive weighting strategy to force the neural network to focus more on the optimization nearer to the zero level set.

The level set projection constraint aims to guide the gradients at zero level set, where we propose another unsigned distance loss as a supplement for correcting distance errors

\*Equal contribution.

†The corresponding author is Yu-Shen Liu.

at the zero level set. To bring a more robust guidance to the optimization of zero level set, we introduce an addition gradient constraint on queries to learn more accurate and regular non-zero level sets of UDF by encouraging the gradient orthogonal to the tangent plane of the surface.

Based on our learned stable UDF with an accurate zero level set, we explore three different 3D applications, including surface reconstruction, unsupervised point normal estimation and unsupervised point cloud upsampling. Our improvements over unsupervised or even supervised baselines justify both our effectiveness and the importance of an accurate and continuous zero level set to the representation ability of neural unsigned fields.

Our main contributions can be summarized as follows.

- We propose two novel constraints on UDF to achieve a smoother and more continuous zero level set for learning to represent shapes or scenes without ground truth unsigned distances.
- We justify that an accurate and continuous zero level set is the key factor to represent complex 3D geometries and supporting 3D downstream applications.
- We explore three different 3D applications with our learned UDF, including surface reconstruction, unsupervised point normal estimation and unsupervised point cloud upsampling, where we show our superiority over the state-of-the-art supervised / unsupervised methods.

## 2. Related Work

With the rapid development of deep learning, the neural networks have shown great potential in 3D applications [66, 30, 79, 75, 78, 74, 80, 76, 14, 39]. We mainly focus on learning Neural Implicit Functions with networks for representing 3D shapes or scenes.

**Neural Implicit Functions.** Recently, Neural Implicit Functions (NIFs) have shown promising results in surface reconstruction [61, 51, 57, 16, 29, 15, 38, 55, 53], novel view synthesis [58, 4, 65, 87, 32, 24, 59], image super-resolution [68, 59, 11], point normal estimation [44], etc. In the task of surface reconstruction, the NIFs approaches train a neural network to represent shapes and scenes with binary occupancy [57, 64, 19] or signed distance functions (SDFs) [61, 36, 22, 21, 63, 10, 84, 6], and then extract surfaces from the learned implicit functions with the marching cubes algorithm [50]. Previous methods [57, 61] learn a global latent code for each 3D shape with a shared MLP to decode geometries. Some advanced approaches [64, 36] propose to use more latent codes to represent detailed local geometries. Latest studies bring more priors learned from large scale datasets to enhance the representation ability of NIFs. PCP [54] introduces the predictive context priors to

represent large-scale point clouds. OnSurf [52] explores the on-surface prior to reconstruct smooth surfaces from sparse point clouds.

NeRF [58] and the following studies [59, 62, 69, 56, 37, 67, 33, 45] extends NIFs to the task of novel view synthesis. NeuS [73], VolSDF [81] and Unisurf [60] advance the rendering procedure of NeRF for surface reconstruction from multi-view images. ManhattanSDF [28] leverages the Manhattan assumption for representing indoor scenes. NeuRIS [71] and MonoSDF [86] brings the depth prior and normal prior for improving the surface quality.

Occupancy and SDFs are merely suitable to represent closed shapes. Recently, researchers explore the neural unsigned distances (UDFs) [20, 89, 88, 70, 18, 72, 49, 48] to represent shapes with arbitrary topology. NDF [20] trains a neural network to learn UDFs with the ground truth distance values. GIFS [82] jointly learns UDFs and the relationship of queries. MeshUDF [27] propose a differentiable meshing algorithm for UDF. However, all the previous UDF-based methods fail to reconstruct smooth and continuous surfaces with high-fidelity details due to the disability of learning an accurate and smooth zero level set of UDF.

**Implicit Learning from Raw Point Clouds.** With the ground truth signed/unsigned distances or occupancy labels as supervision, previous methods learn NIFs as a regression [20, 61] problem or classification [57] problem. However, these expensive supervisions come from the ground truth meshes which are hard to collect. Moreover, such methods also fail to generalize to unseen cases with large geometry variations to the training data, which limits the application scenarios of such methods. To learn implicit representations directly from raw point clouds without ground truth supervisions, current methods explore the supervision from sign agnostic [2, 3] learning with gradient constraints [25]. Neural-Pull [51] employs a new learning schema to pull space onto surfaces for optimizing SDFs.

Latest work CAP-UDF [89] learns UDF progressively with a field consistency loss. However, the UDF inferred by CAP-UDF are not accurate and the gradients are unreliable, especially in the zero level set which indicates the location of 3D geometries. The key problem is that the differential neural networks struggle from learning non-differentiable zero level set of UDF, leading to highly fragmented and discontinuous surfaces. Since the non-zero level sets of UDF are all differentiable, we propose to pull the non-zero level sets onto the zero level set with gradient constraints which align gradients and correct unsigned distance errors on the zero level set, leading to a smoother and more continuous zero level set of unsigned distance field. Our learning schema involves the constraints for both zero level set and non-zero level sets of UDF, which differentiates our method from the previous ones [89, 2, 3, 51, 20, 61].

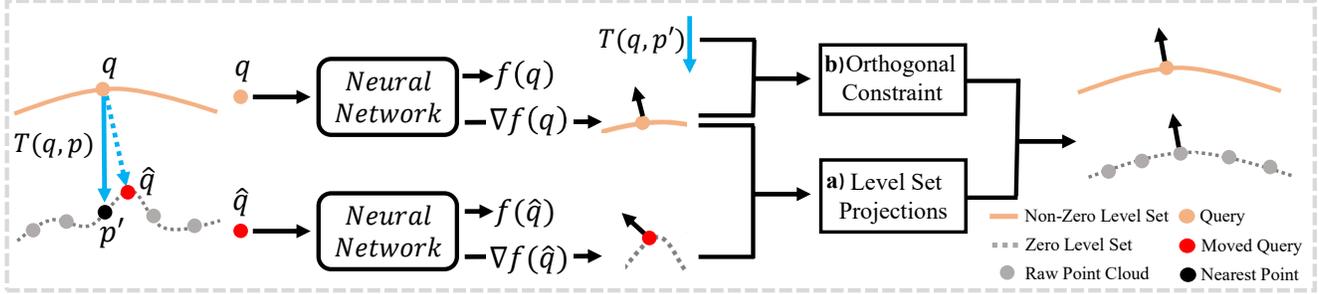


Figure 1. Overview of our method. (a) The level set projection constraint leverages the gradient  $\nabla f(q)$  at convincing non-zero level set as the guidance to gradient  $\nabla f(\hat{q})$  at unreliable zero level set. (b) The gradient-surface orthogonal constraint encourages the gradients  $\nabla f(q)$  at query  $q$  orthogonal to the tangent plane of the surface at  $p'$  which is the nearest point of  $q$  in the raw point cloud.

### 3. Method

**Neural UDFs and Level Sets.** We focus on learning unsigned distance functions for representing 3D shapes or scenes with arbitrary typology. Given a raw point cloud  $P = \{p_i \in R^3\}_{i=1}^N$  with  $N$  points, we sample a set of queries  $Q = \{q_j \in R^3\}_{j=1}^M$  around  $P$ . The UDF  $f$  predicts the unsigned distances  $s^P$  between queries  $\{q_j\}_{j=1}^M$  and the shape described by  $P$ , formulated as:

$$s_j^P = f(q_j). \quad (1)$$

A level set  $S_d$  of UDF is defined as a region of the 3D space containing a continuous set of 3D points with the same unsigned distance  $d$ , formulates as:

$$S_d = \{q | f(q) = d\}, \quad (2)$$

and the level sets of UDF are defined as  $\mathcal{S} = \{S_d\}_{d>0}$ .

**Learning UDFs from Point Clouds.** Inspired by CAP-UDF [89], we leverage the pulling optimization for learning UDFs from raw point clouds without ground truth unsigned distances as supervision. Specifically, given a query  $q_i$  as input, we move it against the direction of the gradient  $g_i = \nabla f(q_i)$  at  $q_i$ , the moving operation is formulated as:

$$\hat{q}_i = q_i - f(q_i) \cdot \nabla f(q_i) / \|\nabla f(q_i)\|_2, \quad (3)$$

where the chamfer distance between moved queries  $\hat{Q} = \{\hat{q}_i\}_{i=1}^M$  and the raw point cloud  $P$  is calculated as the loss function for optimization, formulated as:

$$\begin{aligned} \mathcal{L}_{CD} = & \frac{1}{M} \sum_{i \in [1, M]} \min_{j \in [1, N]} \|\hat{q}_i - p_j\|_2 \\ & + \frac{1}{N} \sum_{j \in [1, N]} \min_{i \in [1, M]} \|p_j - z_i\|_2, \end{aligned} \quad (4)$$

**Zero Level Set of UDF.** An accurate zero level set is the key factor for UDF to represent 3D shapes or scenes, since the zero level set indicates the location of 3D geometries. The key problem here is that the differential networks struggle

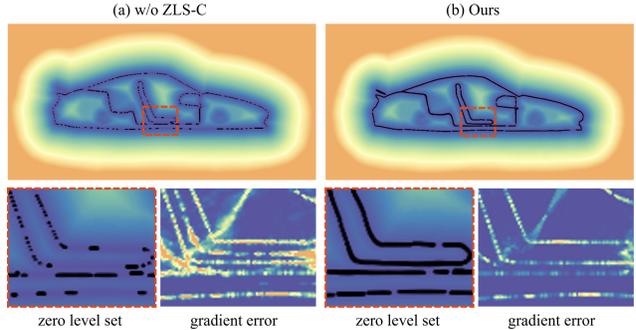


Figure 2. We visualize the fields and the zero level set of UDF learned with or without our designed constraints on the zero level set (ZLS-C). The bluer the color, the smaller the unsigned distances in the field and the errors on gradients.

from learning the zero level set where the UDF is not differentiable, which leads to large errors on unsigned distances and gradients around the zero level set.

We provide a visualization of fields and zero level set of UDF learned with neural networks in Fig. 2. As shown in Fig. 2 (a), the zero level set of UDF learned without proper constraints is highly fragmentary and discontinuous with large errors on the gradients at the space near to the zero level set. On the contrary, we achieve the accurate and continuous zero level set as shown in Fig. 2 (b) by introducing level set projection constraint and gradient-surface orthogonal constraint to the zero level set.

**Level Set Projections.** For learning accurate zero level set of UDF, we design level set projection constraint to guide the learning of zero level set using the rest non-zero level sets. Our idea comes from the observation that despite the unreliable zero level set, the other level sets of the learned UDF are quite convincing, which can serve as the guidance to the field optimization on the zero level set. As shown in Fig. 2, the gradient errors on the non-zero level set are much less than the errors on the zero level set.

We illustrate our method in Fig. 1, where we pull the non-zero level set onto the zero level set with gradient constraints which align gradients and correct unsigned distance errors on the zero level set. We achieve this by encouraging the parallelism between the gradient at queries  $Q = \{q_i\}_{i=1}^M$

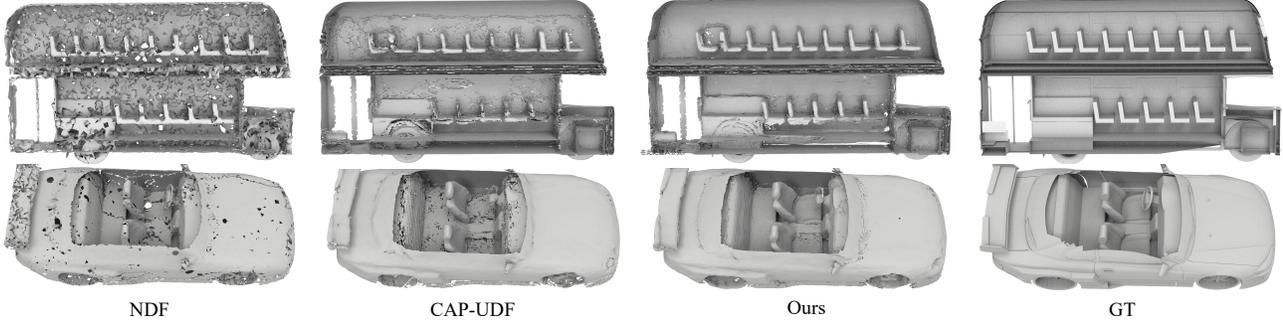


Figure 3. Visual comparisons of surface reconstruction on ShapeNet dataset.

Method	Chamfer-L2		F-Score		NC
	Mean	Median	$F1^{0.005}$	$F1^{0.01}$	
Input	0.363	0.355	48.50	88.34	-
Watertight GT	2.628	2.293	68.82	81.60	-
GT	0.076	0.074	95.70	99.99	-
NDF [20]	0.202	0.193	77.40	97.97	79.1
GIFS [82]	0.128	0.123	88.05	99.31	-
CAP-UDF <sub>BPA</sub>	0.141	0.138	84.84	99.33	81.8
CAP-UDF [89]	0.119	0.114	88.55	99.82	82.5
Ours	<b>0.098</b>	<b>0.097</b>	<b>92.18</b>	<b>99.90</b>	<b>85.0</b>

Table 1. Surface reconstruction for point cloud on ShapeNet cars dataset (Chamfer-L2 $\times 10^4$ ).

and the gradient at their projections  $\hat{Q} = \{\hat{q}_i\}_{i=1}^M$  to the zero level set, formulated as:

$$\mathcal{L}_{\text{proj}} = \frac{1}{M} \sum_{i=1}^M \gamma_i \left( 1 - \left| \frac{\nabla f(q_i) \cdot \nabla f(\hat{q}_i)}{\|\nabla f(q_i)\|_2 \cdot \|\nabla f(\hat{q}_i)\|_2} \right| \right), \quad (5)$$

where  $\gamma_i$  is the factor for weighting the gradient constraint at the query  $q_i$ , indicating the importance of optimizing query  $q_i$ , formulated as:

$$\gamma_i = \exp(-\lambda |f(q_i)|). \quad (6)$$

We set  $\gamma$  to increase when queries are closer to the zero level set where  $\lambda$  is a hyper-parameter to control the increasing speed. The nonuniform weighting encouraging the network to focus more on the optimizations near the zero level set of UDF.

While  $\mathcal{L}_{\text{proj}}$  can provide an effective guidance to the gradients at zero level set, we further explore a more accurate unsigned distance constraint for correcting the distance values at the zero level set of UDF. The implemented loss is defined as:

$$\mathcal{L}_{\text{dist}} = \frac{1}{N} \sum_{i=1}^N |f(p_i)|, \quad (7)$$

where  $p_i \in P$  is the raw point cloud indicating the exact surface to be represented.

**Gradient-Surface Orthogonal Constraint.** The key idea to guide the learning of zero level set using the rest level sets

is based on the observation that the non-zero level sets are much more reliable. To bring a more robust guidance to the optimization of zero level set, we introduce an addition gradient constraint on queries to learn a more stable and regular level sets at the non-zero level sets of UDF by enforcing the gradient orthogonal to the tangent plane of the surface. We design a novel loss to encourage the consistency between the gradient direction at a query  $q_i$  and the direction from  $q_i$  to its closest point  $p_i$  on  $P$ , formulated as:

$$\mathcal{L}_{\text{orth}} = \frac{1}{M} \sum_{i=1}^M \left( 1 - \left| \frac{\nabla f(q_i) \cdot T(q_i, p_i)}{\|\nabla f(q_i)\|_2 \cdot \|T(q_i, p_i)\|_2} \right| \right), \quad (8)$$

where  $T(q_i, p_i)$  is the direction from  $q_i$  to  $p_i$ . The constraint is illustrated in Fig. 1.

A more naive way of introducing the direction vector  $T(q_i, p_i)$  for learning is to explicitly substitute the gradient  $\nabla f(q_i)$  with the direction vector  $T(q_i, p_i)$ , as proposed in GenSDF [21]. However,  $T(q_i, p_i)$  is obviously not the exact direction for  $q_i$  to reach the approximated surface since the raw point cloud is the discrete representation of the surface, and  $p_i$  is not the exact nearest points of  $q_i$  in the approximated surface, which leads to wrong supervisions for GenSDF.

On the contrary, We introduce a more flexible way by treating direction vector  $T(q_i, p_i)$  as an optimization objective, which enables the networks to accommodate the inconsistency between the closest point  $p_i$  on  $P$  and the closest point  $\hat{q}_i$  predicted by the networks.

**Loss Function.** Finally, our loss function is formulated as:

$$\mathcal{L} = \mathcal{L}_{\text{CD}} + \alpha_1 \mathcal{L}_{\text{proj}} + \alpha_2 \mathcal{L}_{\text{dist}} + \alpha_3 \mathcal{L}_{\text{orth}}, \quad (9)$$

where  $\alpha_1$ ,  $\alpha_2$  and  $\alpha_3$  are the balance weights for our designed three losses to learn accurate zero level set of UDF.

**Applications.** The zero level set of UDF indicates the location of represented 3D geometries. Thus, a well-learned UDF with accurate and continuous zero level set can be adopt to different 3D applications.

**Surface Reconstruction.** Following previous methods [89, 27], the surfaces can be extracted from the zero level set of UDF by judging gradient directions, formulated as:

$$D(q_i, q_j) = \nabla f(q_i) \cdot \nabla f(q_j), \quad (10)$$

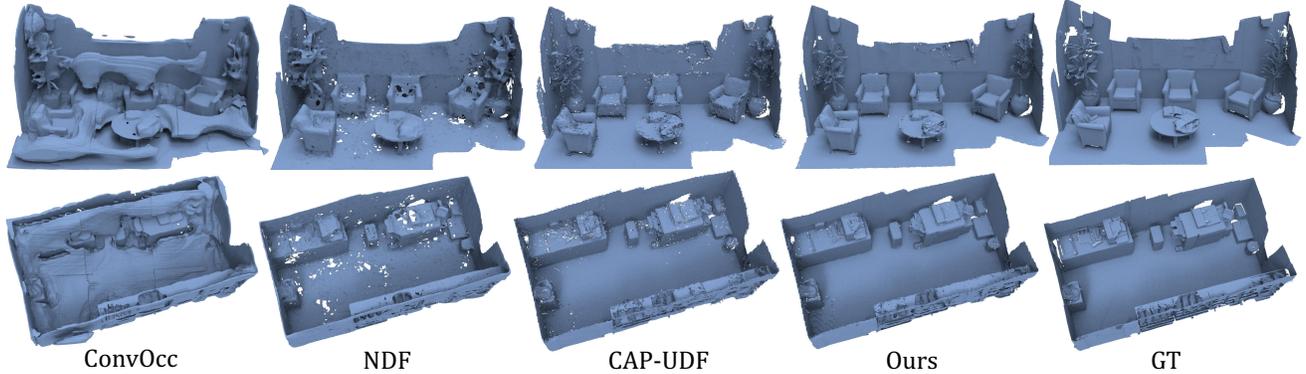


Figure 4. Visual comparisons of surface reconstruction on 3D Scene dataset. Input contains 1K points /  $m^2$

		Burghers			Lounge			Copyroom			Stonewall			Totempole		
		L2CD	LICD	NC												
500/ $m^2$	ConvONet [64]	26.97	0.081	0.905	9.044	0.046	0.894	10.08	0.046	0.885	17.70	0.063	0.909	2.165	0.024	0.937
	LIG [36]	3.080	0.046	0.840	6.729	0.052	0.831	4.058	0.038	0.810	4.919	0.043	0.878	9.38	0.062	0.887
	DeepLS [13]	0.714	0.020	0.923	10.88	0.077	0.814	0.552	0.015	0.907	0.673	0.018	0.951	21.15	0.122	0.927
	NDF [20]	0.546	0.018	0.917	0.314	0.012	0.921	0.242	0.012	0.907	0.226	0.012	0.949	1.049	0.025	0.939
	OnSurf [52]	0.609	0.018	<b>0.930</b>	0.529	0.013	0.926	0.483	0.014	0.908	0.666	0.013	0.955	2.025	0.041	0.954
	CAP-UDF [89]	0.192	0.011	0.911	0.099	0.009	0.911	0.120	0.009	0.902	0.069	0.008	0.958	0.131	0.010	0.954
	Ours	<b>0.161</b>	<b>0.009</b>	0.916	<b>0.087</b>	<b>0.006</b>	<b>0.933</b>	<b>0.102</b>	<b>0.007</b>	<b>0.920</b>	<b>0.061</b>	<b>0.007</b>	<b>0.964</b>	<b>0.114</b>	<b>0.007</b>	<b>0.960</b>
1000/ $m^2$	ConvONet [64]	27.46	0.079	0.907	9.54	0.046	0.894	10.97	0.045	0.892	20.46	0.069	0.905	2.054	0.021	0.943
	LIG [36]	3.055	0.045	0.835	9.672	0.056	0.833	3.61	0.036	0.810	5.032	0.042	0.879	9.58	0.062	0.887
	DeepLS [13]	0.401	0.017	0.920	6.103	0.053	0.848	0.609	0.021	0.901	0.320	0.015	0.954	0.601	0.017	0.950
	NDF [20]	1.168	0.027	0.901	0.393	0.014	0.910	0.269	0.013	0.908	0.509	0.019	0.936	2.020	0.036	0.922
	OnSurf [52]	1.339	0.031	<b>0.929</b>	0.432	0.014	0.934	0.405	0.014	0.914	0.266	0.014	0.957	1.089	0.029	0.954
	CAP-UDF [89]	0.191	0.010	0.910	0.092	0.008	0.927	0.113	0.009	0.911	0.066	0.007	0.962	0.139	0.009	0.955
	Ours	<b>0.146</b>	<b>0.009</b>	0.921	<b>0.068</b>	<b>0.006</b>	<b>0.941</b>	<b>0.093</b>	<b>0.007</b>	<b>0.925</b>	<b>0.050</b>	<b>0.005</b>	<b>0.970</b>	<b>0.107</b>	<b>0.007</b>	<b>0.960</b>

Table 2. Surface reconstruction for point clouds under 3D Scene. L2CD $\times$ 1000.

to classify whether two queries  $q_1$  and  $q_2$  are in the same  $\mathcal{D}(q_i, q_j) > 0$  or the opposite side  $\mathcal{D}(q_i, q_j) < 0$  of the approximated surface. The marching cubes-like algorithm is then used to extract surfaces.

**Unsupervised Point Normal Estimation.** With a well learned zero level set, we compute the gradients at the raw point clouds  $P$  as their estimated normals, formulated as:

$$\mathcal{N}(P) = \{\nabla f(p)\}, p \in P. \quad (11)$$

**Unsupervised Point Cloud Upsampling.** We randomly sample queries around  $P$ , and then move them to the zero level set with Eq. 3 as the upsampled points of  $P$ . We further filter the queries with unsigned distance values larger than a threshold  $\beta$  for avoiding outliers. The procedure for upsampling  $P$  with a factor of  $G$  is formulated as:

$$\mathcal{U}(P) = \{\hat{q}_i, i \in [1, N \times G] | f(q_i) < \beta\} \quad (12)$$

## 4. Experiments

For evaluating the performance of our method, we conduct experiments on surface reconstruction from raw point clouds of shapes (Sec. 4.1), room-level scenes (Sec. 4.2), and real scans (Sec. 4.4), where we also show our results on large-scale driving scenes obtained by fusing LiDAR

depths in Sec. 4.3. We further explore the ability of our learned UDF with accurate zero level set for unsupervised point normal estimation (Sec. 4.5) and unsupervised point upsampling (Sec. 4.6). Finally, we conduct ablation studies in Sec. 4.7.

### 4.1. Surface Reconstruction for Shapes

**Dataset and Metrics.** To explore the ability of our method in representing shapes with arbitrary topology, we follow previous methods [89, 82] to conduct experiments on the ‘‘Car’’ category of ShapeNet dataset which contains the largest number of shapes with multi-layer structures or open surfaces. We sample 10K points per shape as the input for surface reconstruction. For evaluating the performances, we follow the common setting [82] to sample 100K points from the reconstructed surfaces and leverage Chamfer Distance, Normal Consistency [57] and F-score with a threshold of 0.005 / 0.01 as the evaluation metrics.

**Comparisons.** We compare our method with the state-of-the-art shape reconstruction methods NDF [20], GIFS [82] and CAP-UDF [89]. We report the quantitative comparison results in Tab. 1, where ‘‘Watertight GT’’ is the result of sampling points from the watertight ground truth, indicating the best performance of the methods base on SDF or occupancy. We also show the superior limit of the dataset

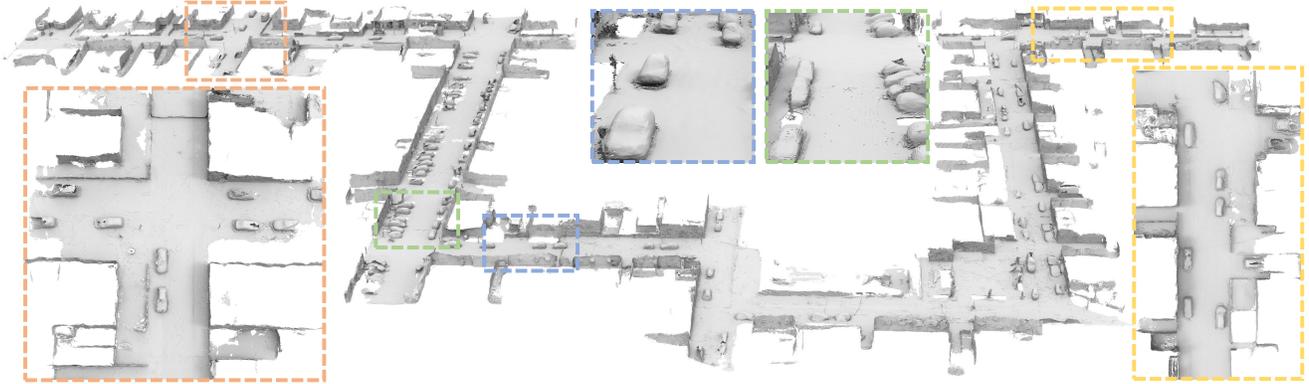


Figure 5. Surface reconstruction results on KITTI odometry dataset (Sequence 00, frame 3000 to 4000) .

by sampling points on the ground truth mesh and report the result as “GT”. Tab. 1 shows that our method achieves the best performance over all the baselines in terms of all metrics. The visual comparison in Fig. 3 further shows that our methods produce more continuous and accurate geometries over all baselines.

## 4.2. Surface Reconstruction for Room-Level Scenes

**Dataset and Metrics.** We further explore our performance on real scene scans to demonstrate the ability of our methods to scale to large-scale point clouds. Following On-Surf [52], we conduct experiments on the 3D Scene dataset consisting of complex scenes with open surfaces and scan noises. The input point cloud is achieved by sampling 500 or 1000 points per  $m^2$  uniformly for each scene. We follow the common setting [52, 89] to sample  $1M$  points on both the reconstructed and the ground truth meshes. We leverage L1 / L2 Chamfer Distance and Normal Consistency as the evaluation metrics.

**Comparisons.** With different point densities, we compare our method with the state-of-the-art methods ConvONet [64], LIG [36], DeepLS [13], OnSurf [52], NDF [20] and CAP-UDF [89]. The numerical comparison in Tab. 2 demonstrates our superior performance in all scenes. We further present a visual comparison in Fig. 4, where the input point clouds contains  $1K$  points/ $m^2$ . The visualization further shows that our method can reconstruct more detailed and smooth surfaces in complex scenes, where the other UDF-base methods NDF and CAP-UDF produce fragmented surfaces.

## 4.3. Large-Scale Driving Scene Reconstruction

**Dataset.** To further demonstrate our scalability to extremely large scenes, we follow NGS [35] to apply our method to KITTI dataset [23]. We use the LiDAR scans of frame 3000 to 4000 in Squeece00 subset of KITTI dataset as the full input, which are transformed into world coordinates using the provided camera trajectories. We leverage the slide-window strategy to divide the full scene into 15

Method	Chamfer-L1	F-Score	NC
IGR [25]	0.178	75.5	86.9
Point2Mesh [31]	0.116	64.8	-
SPSR [40]	0.232	73.5	-
SAP [63]	0.076	83.0	88.6
Neural-Pull [51]	0.106	79.7	87.2
NDF [20]	0.238	68.6	80.4
CAP-UDF [89]	0.073	84.5	88.6
Ours	<b>0.071</b>	<b>85.1</b>	<b>91.0</b>

Table 3. Surface reconstruction for point cloud on SRB dataset.

blocks of size  $51.2m^3$ , where we sample 300K points per block as the local input to train our network. The final scene is achieved by blending the reconstructed local scenes.

**Analysis.** The reconstruction results under KITTI scenes is shown in Fig. 5, where our method reproduces visually appealing performances with accurate geometries. Note that: 1) The LiDAR scans contain large amount of noises, especially in the difficult driving scenes. 2) Our method learns to reconstruct surfaces directly from the raw point cloud without learning priors from large-scale datasets.

## 4.4. Surface Reconstruction for Real Scans

**Dataset and Metrics.** We demonstrate our advantages for surface reconstruction from real scans by conducting experiments on the Surface Reconstruction Benchmarks [77] (SRB). We follow previous methods [63, 89] to leverage Chamfer Distance and F-Score with a threshold of 1% for evaluation. Note that the ground truth of SRB dataset is dense point cloud containing millions of points. For a comprehensive comparison on the reconstruction quality, we further evaluate the Normal Consistency by first converting the ground truth into meshes with BPA [8] algorithm.

**Comparisons.** We first report our numerical comparison under SRB dataset with the classic and learning-based methods in Tab. 3. The compared methods include IGR [25], Point2Mesh [31], Screened Poisson Surface Reconstruction (SPSR) [40], Shape As Points (SAP) [63], Neural-Pull [51], NDF [20] and CAP-UDF [89]. Our method

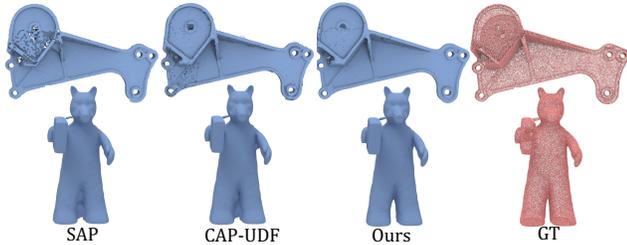


Figure 6. Visual comparisons of surface reconstruction on SRB dataset. SRB dataset provides only point clouds as GT.

	Method	Clean	Strip	Gradient	average
Supervised	PCPNet [26]	9.66	11.47	13.42	11.61
	Hough [9]	10.23	12.47	11.02	11.24
	Nesti-Net [7]	6.99	8.47	9.00	8.15
	IterNet [41]	6.72	7.73	7.51	7.32
	DeepFit [5]	6.51	7.92	7.31	7.25
Unsupervised	Jet [12]	12.23	13.39	13.13	12.92
	PCA [1]	12.29	13.66	12.81	12.92
	Ours (w/o ZLS-c)	7.30	7.74	7.78	7.61
	Ours	<b>6.51</b>	<b>7.45</b>	<b>7.64</b>	<b>7.20</b>

Table 4. Point normal estimation results on the PCPNet dataset.

achieves the best performance in all metrics as shown in Tab. 3. The visual comparison in Fig. 6 further shows that our methods produce more accurate and smooth geometries over all baselines.

#### 4.5. Unsupervised Point Normal Estimation

**Dataset and Metrics.** We adopt the widely-used PCPNet [26] dataset to evaluate our method on the task of point normal estimation. Each point cloud in the PCPNet dataset contains 100K points. We conduct experiments under three different settings. The inputs of “Clean” setting are sampled uniformly on shapes, and two additional settings with varying point densities (Stripes and Gradients) are added to explore our performance in handling irregular data. Following previous methods [5, 26], we sample 5000 points per shape and compute the angle root mean square error (RMSE) between the predicted normals and the ground truth normals as the evaluation metrics.

**Comparisons.** We conduct a comparison with traditional normal estimation methods including PCA [1], Jets [12] and the state-of-the-art learning-based methods PCPNet [26], HoughCNN [9], Nesti-Net [7], Iter-Net [41] and DeepFit [5]. Note that all the previous learning-based methods are designed in a supervised manner which requires a large scale dataset for training, while we extract normals directly from the unsigned distance fields learned from raw point cloud in an unsupervised way. The numerical results in Tab. 4 show that our method achieves a comparable or even better performance than the state-of-the-art supervised methods. The results also show the effectiveness of our method where the performance degrades from 7.20 to 7.61 without

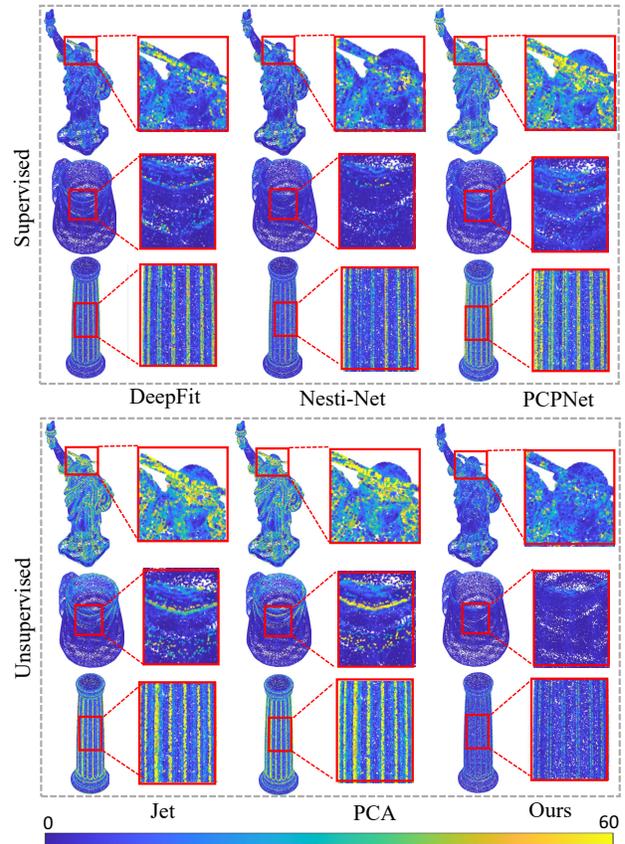


Figure 7. Visual comparisons on error maps of point normal estimation on PCPNet dataset.

Method	Supervised?	P2F	CD	HD
PU-Net [85]	Yes	6.84	0.72	8.94
MPU [83]	Yes	3.96	0.49	6.11
PU-GAN [42]	Yes	2.33	0.28	4.64
Dis-PU [43]	Yes	2.01	0.22	2.83
EAR [34]	No	5.82	0.52	7.37
L2G-AE [46]	No	39.37	6.31	63.23
SPU-Net (Tr2T) [47]	No	5.97	0.38	2.24
SPU-Net (All2T)	No	5.79	0.37	2.55
SPU-Net (Te2T)	No	6.85	0.41	2.18
Ours (w/o ZLS-C)	No	1.43	0.24	2.16
Ours	No	<b>1.02</b>	<b>0.19</b>	<b>1.55</b>

Table 5. Point cloud upsampling results on the PU-GAN dataset.

our constraints. The visual comparison in Fig. 7 shows that our method produce more accurate estimations compared to other methods, especially on the complex geometries.

#### 4.6. Unsupervised Point Cloud Upsampling

**Dataset and Metrics.** For the task of point cloud upsampling, we follow previous methods [47, 43] to conduct experiments on the dataset provided by PU-GAN [42]. We evaluate our method under the setting to generate 8,198

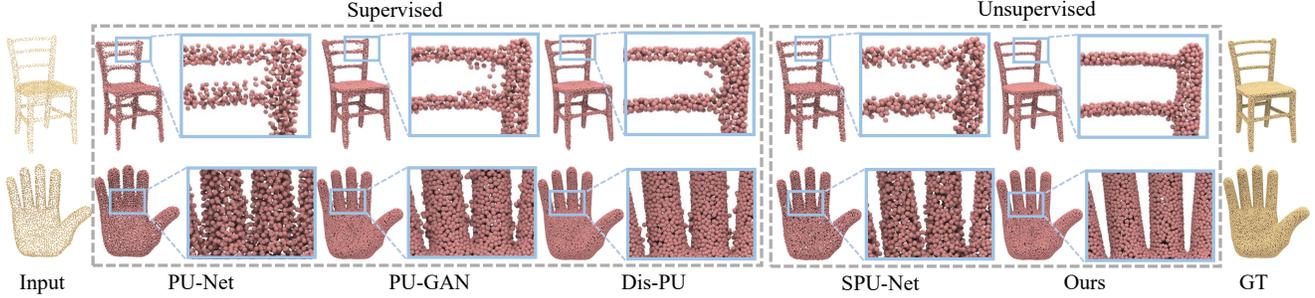


Figure 8. Visual comparisons of point cloud upsampling on PU-GAN dataset.

points for each sparse input with 2,048 points. We adopt Chamfer Distance (CD), Hausdorff Distance (HD) and Point-to-Surface distance (P2F) as the evaluation metrics.

**Comparisons.** We compare our method with the state-of-the-art supervised methods PU-Net [85], MPU [83], PU-GAN [42], Dis-PU [43] and unsupervised methods EAR [34], L2G-AE [46] and SPU-Net [47]. The quantitative comparison is shown in Tab. 5, where our method outperforms the previous supervised and unsupervised methods significantly. The results also show the effectiveness of our method where the performance degrades from 1.55 to 2.16 in terms of HD without our designed constraints. The visualization in Fig. 8 further shows that our upsampling results are more smooth and uniform even compared with the state-of-the-art supervised methods.

#### 4.7. Ablation Studies

We conduct ablation studies to demonstrate the effectiveness of each proposed constraint and explore the effect of some important hyper-parameters. We report the performance under a subset of ShapeNet Cars dataset in terms of L2 Chamfer Distance ( $\times 10^4$ ) and Normal Consistency.

**Effect of each Design.** We first justify the effectiveness of each proposed constraint in Tab. 6. We report the performance without all the three designed losses as “w/o ZLS-C”. We explore each loss by reporting the results without level set projection loss as “w/o  $\mathcal{L}_{proj}$ ”, surface unsigned distance loss as “w/o  $\mathcal{L}_{dist}$ ”, gradient-surface orthogonal loss as “w/o  $\mathcal{L}_{orth}$ ” and adaptive point weight as “w/o AW”. We replace our level set projection constraint with the GenSDF-style [21] optimization as “GenSDF-Style”. We further constrain the level set projection loss on a randomly selected non-zero level set other than the zero level set and report the performance as “Non-zero LSC”. The result demonstrates that the zero level set is the key factor that influence the representation ability of UDF.

**Weight of Level Set Projection Loss.** We report the performance with different weights of the level set projection loss  $\mathcal{L}_{proj}$  in Tab. 7. We found that a suitable weight of  $\mathcal{L}_{proj}$  will improve the reconstruction accuracy, and it may affect the optimization to converge if we weight too much.

**Weight of Surface Unsigned Distance Loss.** We further

	w/o ZLS-C	w/o $\mathcal{L}_{proj}$	w/o $\mathcal{L}_{dist}$	w/o $\mathcal{L}_{orth}$
L2CD	0.126	0.107	0.101	0.096
NC	82.7	82.9	83.9	84.1
	w/o AW	GenSDF-Style	Non-zero LSC	Ours
L2CD	0.096	0.101	0.098	<b>0.091</b>
NC	84.2	83.9	83.5	<b>84.4</b>

Table 6. Effect of framework design.

	0.001	0.002	0.01	0.1
L2CD	0.095	<b>0.091</b>	0.162	0.359
NC	84.2	<b>84.4</b>	83.4	79.6

Table 7. Weight of level set projection loss.

explore the effect of weights  $\alpha_2$  for the surface unsigned distance loss in Tab 8, where we found that a too large or too small weight will degenerate the performance.

	0.001	0.01	0.1	1
L2CD	0.101	0.100	<b>0.091</b>	0.449
NC	84.1	84.1	<b>84.4</b>	78.9

Table 8. Weight of surface unsigned distance loss.

**Weight of Gradient-Surface Orthogonal Loss.** We evaluate the effect of the weight  $\alpha_3$  for gradient-surface orthogonal loss in Tab. 9. A proper weight  $\alpha_3 = 0.01$  brings the best improvement while a too much weight will affect the convergence.

	0.001	0.01	0.1	1
L2CD	0.094	<b>0.091</b>	0.335	2.55
NC	84.2	<b>84.4</b>	77.7	72.9

Table 9. Weight of gradient-surface orthogonal loss.

**Effect of Adaptive per Point Weights.** We further report the performance under different  $\lambda$  which controls the increasing speed of adaptive weights for each point in Tab. 10. We found that  $\lambda$  has a relative small impact on the Chamfer Distance with large affect on the Normal Consistency, which indicates that the adaptive per point weight can improve the surface smoothness.

	1	10	100	1000
L2CD	0.093	<b>0.091</b>	0.092	0.094
NC	84.1	<b>84.4</b>	83.7	83.2

Table 10. Effect of adaptive per point weights.

## 5. Conclusion

In this paper, we justify that an accurate and continuous zero level set of unsigned distance functions is the key factor to represent complex 3D geometries. We therefore propose two novel constraints on UDF to achieve a more continuous zero level set by projecting non-zero level set to guide the optimization of zero level set. With the ability of learning consistent and accurate zero level set of UDF, we explore three different 3D applications including surface reconstruction, unsupervised point normal estimation and unsupervised point cloud upsampling, where the visual and numerical comparisons with the state-of-the-art supervised/unsupervised methods justify our effectiveness and show our superiority over the latest methods.

## 6. Acknowledgement

This work was supported by National Key R&D Program of China (2022YFC3800600), the National Natural Science Foundation of China (62272263, 62072268), and in part by Tsinghua-Kuaishou Institute of Future Media Data.

## References

- [1] Hervé Abdi and Lynne J Williams. Principal component analysis. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2(4):433–459, 2010. 7
- [2] Matan Atzmon and Yaron Lipman. SAL: Sign agnostic learning of shapes from raw data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2565–2574, 2020. 2
- [3] Matan Atzmon and Yaron Lipman. SALD: Sign agnostic learning with derivatives. In *International Conference on Learning Representations*, 2020. 2
- [4] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021. 2
- [5] Yizhak Ben-Shabat and Stephen Gould. DeepFit: 3d surface fitting via neural network weighted least squares. In *European Conference on Computer Vision*, pages 20–34. Springer, 2020. 7
- [6] Yizhak Ben-Shabat, Chamin Hwa Koneputugodage, and Stephen Gould. Digs: Divergence guided shape implicit neural representation for unoriented point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19323–19332, 2022. 2
- [7] Yizhak Ben-Shabat, Michael Lindenbaum, and Anath Fischer. Nesti-Net: Normal estimation for unstructured 3d point clouds using convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10112–10120, 2019. 7
- [8] Fausto Bernardini, Joshua Mittleman, Holly Rushmeier, Cláudio Silva, and Gabriel Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, 5(4):349–359, 1999. 6
- [9] Alexandre Boulch and Renaud Marlet. Deep learning for robust normal estimation in unstructured point clouds. In *Computer Graphics Forum*, volume 35, pages 281–290. Wiley Online Library, 2016. 7
- [10] Alexandre Boulch and Renaud Marlet. POCO: Point convolution for surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6302–6314, 2022. 2
- [11] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18392–18402, 2023. 2
- [12] Frédéric Cazals and Marc Pouget. Estimating differential quantities using polynomial fitting of osculating jets. *Computer Aided Geometric Design*, 22(2):121–146, 2005. 7
- [13] Rohan Chabra, Jan E Lenssen, Eddy Ilg, Tanner Schmidt, Julian Straub, Steven Lovegrove, and Richard Newcombe. Deep local shapes: Learning local sdf priors for detailed 3D reconstruction. In *European Conference on Computer Vision*, pages 608–625. Springer, 2020. 1, 5, 6
- [14] Chao Chen, Zhizhong Han, Yu-Shen Liu, and Matthias Zwicker. Unsupervised learning of fine structure generation for 3D point clouds by 2D projections matching. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 12466–12477, 2021. 2
- [15] Chao Chen, Yu-Shen Liu, and Zhizhong Han. Latent partition implicit with surface codes for 3d representation. In *European Conference on Computer Vision*, pages 322–343. Springer, 2022. 2
- [16] Chao Chen, Yu-Shen Liu, and Zhizhong Han. Gridpull: Towards scalability in learning implicit representations from 3d point clouds. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2023. 2
- [17] Chao Chen, Yu-Shen Liu, and Zhizhong Han. Unsupervised inference of signed distance functions from single sparse point clouds without learning priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17712–17723, 2023. 1
- [18] Weikai Chen, Cheng Lin, Weiyang Li, and Bo Yang. 3psdf: Three-pole signed distance function for learning surfaces with arbitrary topologies. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18522–18531, 2022. 2
- [19] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5939–5948, 2019. 2
- [20] 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. 1, 2, 4, 5, 6
- [21] Gene Chou, Ilya Chugunov, and Felix Heide. Gensdf: Two-stage learning of generalizable signed distance functions. In *Advances in Neural Information Processing Systems*. 2, 4, 8

- [22] Yueqi Duan, Haidong Zhu, He Wang, Li Yi, Ram Nevatia, and Leonidas J Guibas. Curriculum deepsf. In *European Conference on Computer Vision*, pages 51–67. Springer, 2020. 2
- [23] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3354–3361. IEEE, 2012. 6
- [24] Lily Goli, Daniel Rebain, Sara Sabour, Animesh Garg, and Andrea Tagliasacchi. nerf2nerf: Pairwise registration of neural radiance fields. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9354–9361. IEEE, 2023. 2
- [25] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In *International Conference on Machine Learning*, pages 3789–3799. PMLR, 2020. 2, 6
- [26] Paul Guerrero, Yanir Kleiman, Maks Ovsjanikov, and Niloy J Mitra. PCPNet: learning local shape properties from raw point clouds. In *Computer Graphics Forum*, volume 37, pages 75–85. Wiley Online Library, 2018. 1, 7
- [27] Benoit Guillard, Federico Stella, and Pascal Fua. MeshUDF: Fast and differentiable meshing of unsigned distance field networks. *European Conference on Computer Vision*, 2022. 1, 2, 4
- [28] Haoyu Guo, Sida Peng, Haotong Lin, Qianqian Wang, Guofeng Zhang, Hujun Bao, and Xiaowei Zhou. Neural 3d scene reconstruction with the manhattan-world assumption. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5511–5520, 2022. 2
- [29] Zhizhong Han, Chao Chen, Yu-Shen Liu, and Matthias Zwicker. Drwr: A differentiable renderer without rendering for unsupervised 3d structure learning from silhouette images. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, 2020. 2
- [30] Zhizhong Han, Chao Chen, Yu-Shen Liu, and Matthias Zwicker. Shapecaptioner: Generative caption network for 3d shapes by learning a mapping from parts detected in multiple views to sentences. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 1018–1027, 2020. 2
- [31] Rana Hanocka, Gal Metzer, Raja Giryes, and Daniel Cohen-Or. Point2mesh: a self-prior for deformable meshes. *ACM Transactions on Graphics (TOG)*, 39(4):126–1, 2020. 6
- [32] Ayaan Haque, Matthew Tancik, Alexei A Efros, Aleksander Holynski, and Angjoo Kanazawa. Instruct-nerf2nerf: Editing 3d scenes with instructions. *arXiv preprint arXiv:2303.12789*, 2023. 2
- [33] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022. 2
- [34] Hui Huang, Shihao Wu, Minglun Gong, Daniel Cohen-Or, Uri Ascher, and Hao Zhang. Edge-aware point set resampling. *ACM transactions on graphics (TOG)*, 32(1):1–12, 2013. 7, 8
- [35] Jiahui Huang, Hao-Xiang Chen, and Shi-Min Hu. A neural galerkin solver for accurate surface reconstruction. *ACM Transactions on Graphics (TOG)*, 41(6):1–16, 2022. 6
- [36] Chiyu Jiang, Avneesh Sud, Ameesh Makadia, Jingwei Huang, Matthias Nießner, Thomas Funkhouser, et al. Local implicit grid representations for 3D scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6001–6010, 2020. 2, 5, 6
- [37] Sijia Jiang, Jing Hua, and Zhizhong Han. Coordinate quantized neural implicit representations for multi-view reconstruction. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2023. 2
- [38] Chuan Jin, Tieru Wu, and Junsheng Zhou. Multi-grid representation with field regularization for self-supervised surface reconstruction from point clouds. *Computers & Graphics*, 2023. 2
- [39] Heewoo Jun and Alex Nichol. Shap-e: Generating conditional 3d implicit functions. *arXiv preprint arXiv:2305.02463*, 2023. 2
- [40] Michael Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. *ACM Transactions on Graphics (TOG)*, 32(3):1–13, 2013. 6
- [41] Jan Eric Lenssen, Christian Osendorfer, and Jonathan Masci. Deep iterative surface normal estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11247–11256, 2020. 7
- [42] Ruihui Li, Xianzhi Li, Chi-Wing Fu, Daniel Cohen-Or, and Pheng-Ann Heng. Pu-gan: a point cloud upsampling adversarial network. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7203–7212, 2019. 1, 7, 8
- [43] Ruihui Li, Xianzhi Li, Pheng-Ann Heng, and Chi-Wing Fu. Point cloud upsampling via disentangled refinement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 344–353, 2021. 7, 8
- [44] Shujuan Li, Junsheng Zhou, Baorui Ma, Yu-Shen Liu, and Zhizhong Han. NeAF: Learning neural angle fields for point normal estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023. 2
- [45] Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8456–8465, 2023. 2
- [46] Xinhai Liu, Zhizhong Han, Xin Wen, Yu-Shen Liu, and Matthias Zwicker. L2g auto-encoder: Understanding point clouds by local-to-global reconstruction with hierarchical self-attention. In *Proceedings of the 27th ACM International Conference on Multimedia*, pages 989–997, 2019. 7, 8
- [47] Xinhai Liu, Xinchun Liu, Yu-Shen Liu, and Zhizhong Han. SPU-Net: Self-supervised point cloud upsampling by coarse-to-fine reconstruction with self-projection optimization. *IEEE Transactions on Image Processing*, 31:4213–4226, 2022. 7, 8
- [48] Yu-Tao Liu, Li Wang, Jie Yang, Weikai Chen, Xiaoxu Meng, Bo Yang, and Lin Gao. Neudf: Learning neural unsigned distance fields with volume rendering. In *Proceedings of*

- the *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 237–247, 2023. 2
- [49] Xiaoxiao Long, Cheng Lin, Lingjie Liu, Yuan Liu, Peng Wang, Christian Theobalt, Taku Komura, and Wenping Wang. Neuraludf: Learning unsigned distance fields for multi-view reconstruction of surfaces with arbitrary topologies. *arXiv preprint arXiv:2211.14173*, 2022. 2
- [50] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3D surface construction algorithm. *ACM Siggraph Computer Graphics*, 21(4):163–169, 1987. 2
- [51] Baorui Ma, Zhizhong Han, Yu-Shen Liu, and Matthias Zwicker. Neural-Pull: Learning signed distance function from point clouds by learning to pull space onto surface. In *International Conference on Machine Learning*, pages 7246–7257. PMLR, 2021. 1, 2, 6
- [52] Baorui Ma, Yu-Shen Liu, and Zhizhong Han. Reconstructing surfaces for sparse point clouds with on-surface priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 1, 2, 5, 6
- [53] Baorui Ma, Yu-Shen Liu, and Zhizhong Han. Learning signed distance functions from noisy 3d point clouds via noise to noise mapping. In *International Conference on Machine Learning (ICML)*, 2023. 2
- [54] Baorui Ma, Yu-Shen Liu, Matthias Zwicker, and Zhizhong Han. Surface reconstruction from point clouds by learning predictive context priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [55] Baorui Ma, Junsheng Zhou, Yu-Shen Liu, and Zhizhong Han. Towards better gradient consistency for neural signed distance functions via level set alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17724–17734, 2023. 2
- [56] Xiaoxu Meng, Weikai Chen, and Bo Yang. Neat: Learning neural implicit surfaces with arbitrary topologies from multi-view images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 248–258, 2023. 2
- [57] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3D reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4460–4470, 2019. 1, 2, 5
- [58] B Mildenhall, PP Srinivasan, M Tancik, JT Barron, R Ramamoorthi, and R Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision*, 2020. 1, 2
- [59] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)*, 41(4):1–15, 2022. 2
- [60] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5589–5599, 2021. 1, 2
- [61] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019. 1, 2
- [62] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5865–5874, 2021. 2
- [63] Songyou Peng, Chiyu Jiang, Yiyi Liao, Michael Niemeyer, Marc Pollefeys, and Andreas Geiger. Shape as points: A differentiable poisson solver. *Advances in Neural Information Processing Systems*, 34, 2021. 2, 6
- [64] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *European Conference on Computer Vision*, pages 523–540. Springer, 2020. 1, 2, 5, 6
- [65] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10318–10327, 2021. 2
- [66] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017. 2
- [67] Radu Alexandru Rosu and Sven Behnke. Permutosdf: Fast multi-view reconstruction with implicit surfaces using permutohedral lattices. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8466–8475, 2023. 2
- [68] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems*, 33:7462–7473, 2020. 2
- [69] Matthew Tancik, Vincent Casser, Xinchun Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan, Jonathan T Barron, and Henrik Kretzschmar. Block-NeRF: Scalable large scene neural view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8248–8258, 2022. 2
- [70] Rahul Venkatesh, Tejan Karmali, Sarthak Sharma, Aurobrata Ghosh, R Venkatesh Babu, László A Jeni, and Maneesh Singh. Deep implicit surface point prediction networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12653–12662, 2021. 2
- [71] Jiepeng Wang, Peng Wang, Xiaoxiao Long, Christian Theobalt, Taku Komura, Lingjie Liu, and Wenping Wang. NeuRIS: Neural reconstruction of indoor scenes using normal priors. In *17th European Conference on Computer Vision*, pages 139–155. Springer, 2022. 2
- [72] Li Wang, Jie Yang, Weikai Chen, Xiaoxu Meng, Bo Yang, Jintao Li, and Lin Gao. Hsdf: Hybrid sign and distance field for modeling surfaces with arbitrary topologies. In *Advances in Neural Information Processing Systems*. 2
- [73] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. NeuS: Learning neural im-

- licit surfaces by volume rendering for multi-view reconstruction. *Advances in Neural Information Processing Systems*, 34:27171–27183, 2021. [1](#), [2](#)
- [74] Xin Wen, Peng Xiang, Zhizhong Han, Yan-Pei Cao, Pengfei Wan, Wen Zheng, and Yu-Shen Liu. PMP-Net: Point cloud completion by learning multi-step point moving paths. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7443–7452, 2021. [2](#)
- [75] Xin Wen, Peng Xiang, Zhizhong Han, Yan-Pei Cao, Pengfei Wan, Wen Zheng, and Yu-Shen Liu. PMP-Net++: Point cloud completion by transformer-enhanced multi-step point moving paths. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. [2](#)
- [76] Xin Wen, Junsheng Zhou, Yu-Shen Liu, Hua Su, Zhen Dong, and Zhizhong Han. 3D shape reconstruction from 2D images with disentangled attribute flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3803–3813, 2022. [2](#)
- [77] Francis Williams, Teseo Schneider, Claudio Silva, Denis Zorin, Joan Bruna, and Daniele Panozzo. Deep geometric prior for surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10130–10139, 2019. [6](#)
- [78] Peng Xiang, Xin Wen, Yu-Shen Liu, Yan-Pei Cao, Pengfei Wan, Wen Zheng, and Zhizhong Han. Snowflakenet: Point cloud completion by snowflake point deconvolution with skip-transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5499–5509, 2021. [2](#)
- [79] Peng Xiang, Xin Wen, Yu-Shen Liu, Yan-Pei Cao, Pengfei Wan, Wen Zheng, and Zhizhong Han. Snowflake point deconvolution for point cloud completion and generation with skip-transformer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(5):6320–6338, 2022. [2](#)
- [80] Peng Xiang, Xin Wen, Yu-Shen Liu, Hui Zhang, Yi Fang, and Zhizhong Han. Retro-fpn: Retrospective feature pyramid network for point cloud semantic segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2023. [2](#)
- [81] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34:4805–4815, 2021. [2](#)
- [82] Jianglong Ye, Yuntao Chen, Naiyan Wang, and Xiaolong Wang. GIFS: Neural implicit function for general shape representation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. [2](#), [4](#), [5](#)
- [83] Wang Yifan, Shihao Wu, Hui Huang, Daniel Cohen-Or, and Olga Sorkine-Hornung. Patch-based progressive 3d point set upsampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5958–5967, 2019. [7](#), [8](#)
- [84] Wang Yifan, Shihao Wu, Cengiz Oztireli, and Olga Sorkine-Hornung. Iso-points: Optimizing neural implicit surfaces with hybrid representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 374–383, 2021. [2](#)
- [85] Lequan Yu, Xianzhi Li, Chi-Wing Fu, Daniel Cohen-Or, and Pheng-Ann Heng. Pu-net: Point cloud upsampling network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2790–2799, 2018. [7](#), [8](#)
- [86] Zehao Yu, Songyou Peng, Michael Niemeyer, Torsten Sattler, and Andreas Geiger. Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. In *Advances in Neural Information Processing Systems*. [2](#)
- [87] Wenyuan Zhang, Ruofan Xing, Yunfan Zeng, Yu-Shen Liu, Kanle Shi, and Zhizhong Han. Fast learning radiance fields by shooting much fewer rays. *IEEE Transactions on Image Processing*, 2023. [2](#)
- [88] Fang Zhao, Wenhao Wang, Shengcai Liao, and Ling Shao. Learning anchored unsigned distance functions with gradient direction alignment for single-view garment reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12674–12683, 2021. [2](#)
- [89] Junsheng Zhou, Baorui Ma, Yu-Shen Liu, Yi Fang, and Zhizhong Han. Learning consistency-aware unsigned distance functions progressively from raw point clouds. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#)