

Details Enhancement in Unsigned Distance Field Learning for High-fidelity 3D Surface Reconstruction

Cheng Xu^{1, 2, 3}, Fei Hou^{1, 2*}, Wencheng Wang^{1, 2}, Hong Qin⁴, Zhebin Zhang⁵, Ying He⁶

¹Key Laboratory of System Software (CAS), State Key Laboratory of Computer Science, Institute of Software, Chinese Academy of Sciences, China

²University of Chinese Academy of Sciences, China

³School of Advanced Interdisciplinary Sciences, University of Chinese Academy of Sciences, China

⁴Department of Computer Science, Stony Brook University, USA.

⁵InnoPeak Technology, USA

⁶College of Computing and Data Science, Nanyang Technological University, Singapore
{xucheng, houfei, whn}@ios.ac.cn, qin@cs.stonybrook.edu, zhebin.zhang@oppo.com, yhe@ntu.edu.sg

Abstract

While Signed Distance Fields (SDF) are well-established for modeling watertight surfaces, Unsigned Distance Fields (UDF) broaden the scope to include open surfaces and models with complex inner structures. Despite their flexibility, UDFs encounter significant challenges in high-fidelity 3D reconstruction, such as non-differentiability at the zero level set, difficulty in achieving the exact zero value, numerous local minima, vanishing gradients, and oscillating gradient directions near the zero level set. To address these challenges, we propose Details Enhanced UDF (DEUDF) learning that integrates normal alignment and the SIREN network for capturing fine geometric details, adaptively weighted Eikonal constraints to address vanishing gradients near the target surface, unconditioned MLP-based UDF representation to relax non-negativity constraints, and DCUDF for extracting the local minimal average distance surface. These strategies collectively stabilize the learning process from unoriented point clouds and enhance the accuracy of UDFs. Our computational results demonstrate that DEUDF outperforms existing UDF learning methods in both accuracy and the quality of reconstructed surfaces.

Introduction

While signed distance fields (SDF) are favored for their capability to represent watertight surfaces, unsigned distance fields (UDFs) provide a means to model both open surfaces and objects with complex inner structures. However, achieving high-quality UDFs that accurately reconstruct 3D surfaces with fine geometric details is challenging for several reasons. Firstly, UDFs struggle to precisely achieve a zero value, making it difficult to generate open boundaries. Secondly, UDFs are theoretically non-differentiable at the zero level set, resulting in vanishing gradients near the target surface. This issue leads to numerous undesired local minima, complicating the extraction of the zero level set. Thirdly, the gradient directions of UDFs tend to oscillate

near the surface, causing the reconstructed surfaces to be fragmented (Guillard, Stella, and Fua 2022).

Due to the inherently low accuracy of learned UDFs, the extracted zero level sets are typically over-smoothed and lack crucial geometric details. Several studies have aimed to enhance the precision of UDF learning. For instance, NDF (Chibane, Mir, and Pons-Moll 2020) trains a shape encoder and a decoder from 3D surfaces of various types, including point clouds, meshes and mathematical functions. As a supervised method, its performance heavily relies on the quality and diversity of the training dataset. Unsupervised approaches, such as CAP-UDF (Zhou et al. 2022) and LevelSetUDF (Zhou et al. 2023), offer greater flexibility in handling a wider range of 3D models. Despite advancements in UDF learning techniques, all existing methods still suffer from relatively low accuracy in the learned distance fields compared to SDFs. This limitation significantly diminishes their practical usage in real-world applications.

This paper introduces a new method, called Details Enhanced UDF (DEUDF) learning, aimed at enhancing the accuracy of UDF learning from unoriented point clouds to ensure that learned UDFs can capture the fine geometric details of target surfaces. A key observation is the significant role normal directions play in learning fine details. Although obtaining globally consistent orientations (Xu et al. 2023) is challenging due to its combinatorial and global optimization nature, acquiring normal directions locally, for instance, through principal component analysis (Hoppe et al. 1992), is feasible. Consequently, we constrain the UDF gradients to align with normal directions to enhance detail capture, while disregarding normal orientations.

Existing methods strictly constrain UDFs to non-negative values, so the UDFs are hard to achieve exact zero values. Meanwhile, such strict constraint has brought some other problems as illustrated in Figure 1. To overcome this limitation, we relax the strict requirements that UDFs must be non-negative and that the surface must precisely correspond to the zero iso-surfaces. Though the distances are not strictly non-negative, we still call it UDF. This adaptation enables the use of an unconditioned multilayer perceptron (MLP),

*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

meaning an MLP that outputs its value directly without additional operations to make the output positive. Unlike traditional methods that generate UDFs by taking the absolute value of a learned SDF (Zhou et al. 2023) – prone to inducing oscillating gradients – or by using the softplus activation function in MLPs to eliminate negative values (Liu et al. 2023) – leading to vanishing gradients – our relaxation not only addresses the vanishing gradients but also stabilizes the oscillation of gradient directions near the surface.

While SDFs maintain well-behaved gradients with consistent unit length throughout 3D space, UDFs often experience vanishing gradients at the zero level set, diminishing the effectiveness of uniformly applied Eikonal constraints for UDF learning. To address this issue, we propose an adaptively weighted Eikonal constraint, specifically tailored to align with the unique properties of UDFs. Moreover, we incorporate the SIREN network (Sitzmann et al. 2020) to represent high-frequency details in UDFs, thereby enhancing the encoding capabilities of our model. We consider the intended surface should be around zero and the distance values—both positive and negative—should be as small as possible. To achieve this, we adopt DCUDF (Hou et al. 2023), an optimization-based surface extraction algorithm.

By integrating normal alignment, unconditioned MLPs with SIREN activation functions, adaptively weighted Eikonal constraints, and UDF-tailored surface extraction techniques, DEUDF significantly improves the accuracy of UDF learning. Evaluations on benchmark datasets demonstrate our method outperforms baseline methods in terms of UDF accuracy and quality of reconstructed surfaces.

Related Work

Surface reconstruction from point clouds has been studied extensively for the last three decades. The field has seen significant evolution, from computational geometry methods (Amenta and Bern 1998; Dey and Goswami 2003) to implicit function techniques (Hoppe et al. 1992; Ohtake et al. 2003; Kazhdan, Bolitho, and Hoppe 2006; Kazhdan and Hoppe 2013; Hou et al. 2022; Liu et al. 2024), and more recently to deep learning approaches (Park et al. 2019; Chibane, Mir, and Pons-Moll 2020; Zhou et al. 2023; Ren et al. 2023; Wang et al. 2023a; Fainstein, Siless, and Iarussi 2024). Due to space constraints, this section primarily focuses on deep learning-based 3D reconstruction techniques.

Both signed distance fields and occupancy fields effectively represent closed surfaces. An occupancy field defines whether each point in space is inside or outside a given shape. ONet (Mescheder et al. 2019) employs a deep neural network classifier to implicitly represent 3D surfaces as a continuous decision boundary, while IF-Net (Chibane, Alldieck, and Pons-Moll 2020) and CONet (Peng et al. 2020) use encoders to capture shape. Compared to occupancy fields, SDFs provide additional information about the distance of a point from the surface of the object, making them favored for applications that require accurate shape representation, such as reconstruction, shape interpolation and completion. DeepSDF (Park et al. 2019) introduces an innovative implicit encoder that defines the boundary of a

3D shape as the zero level set of a learned implicit function. Following this, numerous neural SDF-based works have been developed. For example, DeepLS (Chabra et al. 2020) utilizes a grid structure to store latent codes for local shape features, SIREN (Sitzmann et al. 2020) introduces a novel activation function for increasing the network’s capability to capture high-frequency signals, and IDF (Wang, Rahmann, and Sorkine-Hornung 2022) employs displacement maps to enhance the representation of fine details. Additionally, SDFs have been utilized to represent geometric shapes for neural rendering tasks, such as NeuS (Wang et al. 2021) and VolSDF (Yariv et al. 2021), which leverage SDFs for 3D reconstruction from multi-view images.

To model general non-watertight surface, Chibane et al. (Chibane, Mir, and Pons-Moll 2020) introduced neural UDFs, which predict the unsigned distance from a query point to the nearest surface point. GIFS (Ye et al. 2022) models the relationship between points rather than between points and surfaces, while NVF (Yang et al. 2023) learns a vector as an alternative to calculating gradient from UDFs, representing the direction from query points to the target surface. Unlike these methods, which utilize separate neural networks to extract supplementary information that aids UDF learning, CAP-UDF (Zhou et al. 2022) and GeoUDF (Ren et al. 2023) focus on enhancing the density of the input point clouds via adopting upsampling techniques. Despite these advancements, the challenge of ambiguous gradients near the zero level set remains. To address the non-differentiability issue of UDFs at zero, LevelSetUDF (Zhou et al. 2023) projects the properties of the non-zero level set to the zero level set to learn a continuous and smooth UDF, while DUDF (Fainstein, Siless, and Iarussi 2024) adopts a new representation to maintain differentiability at points close to the target surface. Although LevelSetUDF and DUDF tackle the non-differentiable issues, they still struggle to match the quality of reconstruction—particularly for surfaces with fine details—achieved SDF learning methods. Additionally, similar to SDFs, UDFs are also used to implicitly represent 3D shapes in neural rendering tasks, such as 3D reconstruction from multi-view images (e.g., NeuralUDF (Long et al. 2023), NeUDF (Liu et al. 2023), 2S-UDF (Deng et al. 2024)) and 3D generation (Yu et al. 2025; Zhou et al. 2024). See Table 1 for a qualitative comparison of existing UDF learning methods.

Extracting the zero level set from UDFs is technically non-trivial, as it is rare for the learned UDFs to precisely reach zero values. There are several research efforts aiming at addressing this issue. Gradient-based methods such as CAP-UDF (Zhou et al. 2022), MeshUDF (Guillard, Stella, and Fua 2022) and GeoUDF (Ren et al. 2023) use both gradient directions and UDF values to detect zero crossings, while optimization-based techniques, such as DCUDF (Hou et al. 2023), optimize the surface locally to minimize the average distance value on the surface.

Method

Let $\mathcal{P} = \{\mathbf{p}_i \in \mathbb{R}^3\}_{i=1}^n$ represent the input raw point cloud, which has been uniformly scaled to fit within the cube domain $\Omega = [-1, 1]^3$. We employ an MLP to parameterize

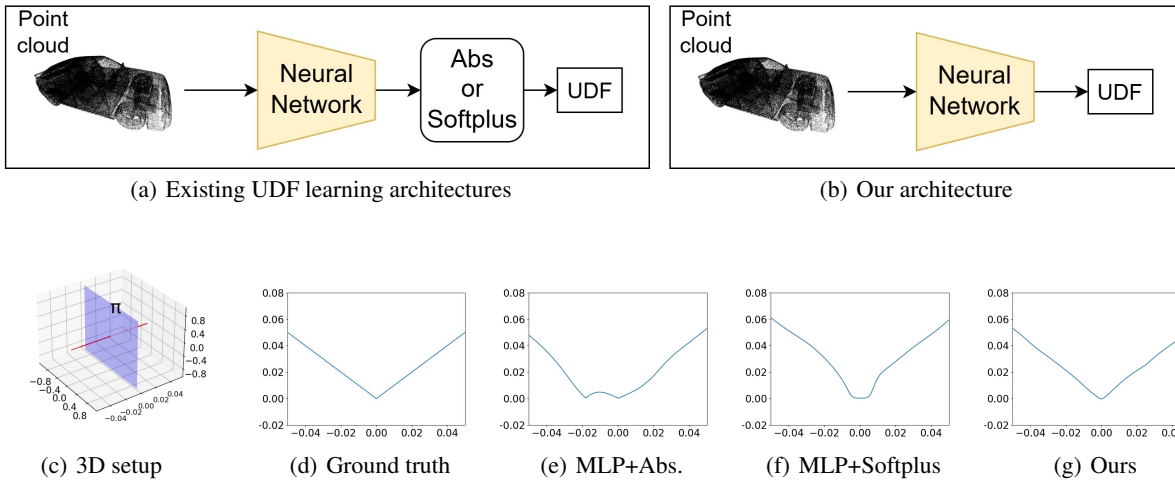


Figure 1: Illustration of UDF learning with various neural representations. (a) Existing neural network architectures often use an absolute value or softplus function to prevent negative distances. (b) In contrast, our method relaxes the non-negative condition and employs an unconditioned MLP with the SIREN activation function for predicting the distances. (c) To show the differences between existing representations and ours, we consider a plane π and a line perpendicular to π . We plot the unsigned distance to π on this line learned by our method, but with different MLP output layers. (d) The horizontal axis represents a signed distance range from -0.05 to 0.05 , while the vertical axis measures the learned unsigned distance. Ideally, the unsigned distance should exhibit a perfect “V” shape relative to the signed distance. (e) However, UDFs parameterized by conditioned MLPs can present defects. For example, learning a negative value followed by taking the absolute value results in a “W”-shaped distance field around the zero level set. (f) Employing the softplus activation function to eliminate negative values yields learned UDFs with vanishing gradients across a relatively large distance range near the zero value. (g) In contrast, our method, which employs unconditioned MLPs, significantly narrows this range of vanishing gradients.

Method	Input	MLP	Eikonal	Non-negativity	Learning
NeUDF	multi-view images	softplus+PE	uniform	softplus	unsupervised
NeuralUDF	multi-view images	softplus+PE	uniform	ABS	unsupervised
2S-UDF	multi-view images	softplus+PE	uniform	softplus	unsupervised
NDF	sparse point clouds	ReLU	-	ABS	supervised
GIFS	sparse point clouds	ReLU	-	ABS	supervised
GeoUDF	sparse point clouds	LeakyReLU	-	ABS	supervised
DUDF	dense point clouds	SIREN	uniform	ABS+HS	supervised
CAP-UDF	sparse point clouds	ReLU+PE	-	ABS	unsupervised
LevelSetUDF	dense point clouds	ReLU+PE	-	ABS	unsupervised
Ours	dense point clouds	SIREN	adaptive	no	unsupervised

Table 1: Qualitative comparison of existing UDF learning methods. HS: hyperbolic scaling; PE: positional encoding; ABS: absolute value.

the UDF for \mathcal{P} , denoted by f . Our objective is to accurately learn f in order to extract a high-fidelity mesh that represents the geometric structure of \mathcal{P} .

Relaxation of non-negative constraints

Traditional methods for learning UDFs generally ensure non-negative distance values by adopting specific strategies, such as taking the absolute value or using the softplus in the last layer. However, as illustrated in Figure 1, these approaches have significant drawbacks in accurately representing distances near zero. For example, using the absolute value results in UDFs exhibiting a “W” shape, leading to changes in gradient directions and the presence of multiple minimum values. Moreover, when the absolute value is applied to an SDF, the resulting UDF exhibits characteristics similar to

those of an SDF. This leads to the unintended consequence of gap filling even in point clouds that represent open surfaces. See Figure 4 for an example. On the other hand, employing the softplus activation function helps avoid the W-shaped artifacts associated with the absolute value approach. Nonetheless, this method tends to generate a U-shaped distance field, characterized by a relatively wide bandwidth around the zero value, approximately between 0 and 0.04. This occurs because for $x \in (-\infty, 0)$, $\text{softplus}(x)$ yields a small positive value with almost zero derivatives. Consequently, this results in vanishing gradients for query points near the target surface, which can significantly hinder the effectiveness of network training that relies on gradient-based optimization techniques.

Observing both vanishing gradients and oscillating gradient directions stem from the strict non-negative constraint on distance values, we propose relaxing the conditions that require UDFs to be non-negative and the surface to coincide precisely with zero iso-surface. We use unconditioned MLPs to represent UDFs and consider the local minimal distance surface of the UDF value around zero, which may be either positive or negative, as the intended surface. As illustrated in Figure 1 (e), this relaxation results in a distance function with a significantly narrower bandwidth compared to using the softplus activation function, thereby providing a high-quality approximation to the ground truth distance field, which exhibits a V-shaped profile.

With UDFs parameterized by unconditioned MLPs, we

define the following loss functions for learning UDFs without ground-truth supervision:

$$\mathcal{L}_{\text{dist}} = \sum_{\mathbf{p}_i \in \mathcal{P}} |f(\mathbf{p}_i)|, \quad (1)$$

and

$$\mathcal{L}_{\text{positive}} = \sum_{\mathbf{x} \in \Omega} \exp(-100f(\mathbf{x})). \quad (2)$$

The distance term $\mathcal{L}_{\text{dist}}$ encourages the zero level set of the learned UDFs to pass through the input points \mathbf{p}_i . The positivity enforcement term $\mathcal{L}_{\text{positive}}$ is designed to ensure that values of $f(\mathbf{x})$ for off-surface points \mathbf{x} are large and positive. This term encourages the majority of sample points are assigned positive values, effectively preventing the generation of negative distance values and ensuring the function behaves like a true UDF. Additionally, it helps to maintain a clear distinction between surface and non-surface regions, crucial for accurate surface reconstruction.

Remark. In NeuralUDF (Long et al. 2023), a similar loss term in the form $\exp(-100|f|)$ was used. It is important to note that our loss term does not include the absolute value. This subtle difference significantly impacts the behavior of the learned distance field f . With the absolute value, their loss encourages $|f|$ being a large positive value, which consequently reduces the occurrence of points with zero distance values. This reduction minimizes the presence of small disconnected components in the reconstructed surfaces (Wang, Rahmann, and Sorkine-Hornung 2022; Long et al. 2023). Therefore, the \exp in their loss functions acts as a regularizer to smooth the learned distance fields. As mentioned above, the use of the absolute value $|f|$ in the loss function can lead to undesired side effects, such as a W-shaped profile in the learned UDFs, which may consequently result in watertight models. In sharp contrast, our loss term, which omits the absolute value, serves as a soft non-negative constraint. This encourages f to remain positive as much as possible, thus differentiating it from an SDF, and enabling f to mimic a true UDF. Even though DCUDF (Hou et al. 2023) used an unconditional MLP to represent UDF, it needs the ground truth UDF to supervise and keep the distances almost positive. While, DEUDF is unsupervised learning.

Normal alignment

Normal directions are critical for enhancing surface details in the reconstruction process. Let $\mathcal{N} = \{\mathbf{n}_i\}_{i=1}^n$ represent the set of unit normals for the point set \mathcal{P} . Following (Hoppe et al. 1992), we apply principal component analysis to each point \mathbf{p}_i to determine its normal direction \mathbf{n}_i . Since UDF gradients typically vanish on the surface, it is impractical to directly constrain the gradients of \mathcal{P} .

To address this issue, we generate a set of sample point pairs $\mathcal{Q} = \{(\mathbf{q}_i^1, \mathbf{q}_i^2)\}_{i=1}^n$ in each training epoch, where each point \mathbf{q}_i is strategically displaced from the surface. Specifically, $\mathbf{q}_i^1 = \mathbf{p}_i + \lambda_i \mathbf{n}_i$ and $\mathbf{q}_i^2 = \mathbf{p}_i - \lambda_i \mathbf{n}_i$, with the displacement λ_i being randomly chosen in the range $(0, 0.003]$. This ensures that \mathcal{Q} contains samples on both sides of the surface, enabling a balanced evaluation on both sides of the geometric structure of interest. We then impose constraints on the

UDF gradient directions at points in \mathcal{Q} , using the following normal alignment loss term:

$$\mathcal{L}_{\text{normal}} = \sum_{\substack{\mathbf{n}_i \in \mathcal{N} \\ k \in \{1,2\}}} \left(1 + (-1)^k \frac{\nabla f(\mathbf{q}_i^k) \cdot \mathbf{n}_i}{\|\nabla f(\mathbf{q}_i^k)\|_2 \cdot \|\mathbf{n}_i\|_2} \right). \quad (3)$$

Adaptively weighted Eikonal constraints

The Eikonal constraint, expressed as $\|\nabla f\| = 1$, is extensively utilized in the learning processes for SDFs. However, when applied to UDFs, this approach faces challenges due to the diminished gradient magnitudes near the zero level set. Direct application of Eikonal constraints to regularize UDFs may cause the actual surface to deviate from the input point cloud \mathcal{P} and may also increase the minima of the learned UDF, as illustrated in Table 4 and Figure 3. To address this issue, we propose a formulation for an adaptively weighted Eikonal loss term:

$$\mathcal{L}_{\text{eikonal}} = \sum_{\mathbf{x} \in \mathcal{Q} \cup \Omega} \delta(f(\mathbf{x})) \left| \|\nabla f(\mathbf{x})\|_2 - 1 \right|, \quad (4)$$

where the weight function, $\delta(\cdot)$, is designed to reduce the contribution from points close to the target surface. A U-shaped function with controllable bandwidth serves this purpose effectively. In our implementation, we employ the attenuation function used in IDF (Wang, Rahmann, and Sorkine-Hornung 2022) as our weight as $\delta(d) = \left(1 + \left(\frac{\xi}{d}\right)^4\right)^{-1}$, where ξ represents the threshold beyond which the influence of the Eikonal constraint begins to diminish significantly. In our experiments, we initially set ξ to 0.01 and gradually decrease ξ to 0.002 over the course of the learning process, following the learning rate. This adjustment is made to enhance the attenuation effect. We evaluate the Eikonal loss $\mathcal{L}_{\text{eikonal}}$ for points in the set \mathcal{Q} , which serves as a proxy of the target geometry, as well as for randomly sampled points throughout the entire domain Ω .

UDF learning

For the network architecture, we employ a 5-layer SIREN network (Sitzmann et al. 2020), which consists of 256 units per layer. The network utilizes a sinusoidal activation function $\sin(\omega x)$, with a default frequency¹ $\omega = 60$ for clean point clouds in our implementation, effectively encoding fine geometric details. Our network takes spatial coordinates (x, y, z) as inputs and outputs the predicted unsigned distance. The training process aims to minimize the following loss function:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{\text{dist}} + \lambda_2 \mathcal{L}_{\text{positive}} + \lambda_3 \mathcal{L}_{\text{normal}} + \lambda_4 \mathcal{L}_{\text{eikonal}} \quad (5)$$

where λ_i s are weights assigned to balance the contributions from the four loss terms. We empirically set $\lambda_1 = 400$, $\lambda_2 = 50$, $\lambda_3 = 40$ and $\lambda_4 = 10$ in our implementation. We train the neural network using the Adam (Kingma and Ba 2015) optimizer, starting with a learning rate of 5×10^{-5} . The learning rate decays to zero following a cosine annealing schedule (Loshchilov and Hutter 2017).

¹For noisy point clouds, using a lower frequency enhances resilience against noise.

		Stanford 3D Scene				Stanford 3D Scan				ShapeNet-Cars			
		Chamfer-L1 (\downarrow)		F-score (\uparrow)		Chamfer-L1 (\downarrow)		F-score (\uparrow)		Chamfer-L1 (\downarrow)		F-score (\uparrow)	
Method	Distance	Mean	Median	$F1^{0.005}$	$F1^{0.0025}$	Mean	Median	$F1^{0.005}$	$F1^{0.0025}$	Mean	Median	$F1^{0.005}$	$F1^{0.0025}$
CAP-UDF	Unsigned	3.37	3.33	98.96	84.51	4.12	3.87	99.12	69.02	4.97	4.63	95.37	56.42
DUDF	Unsigned	3.79	3.26	97.33	79.43	4.20	3.95	99.07	68.10	6.05	5.51	89.02	44.18
LeverSetUDF	Unsigned	3.16	2.90	99.17	85.92	4.12	3.87	99.04	68.83	5.03	4.63	95.01	55.57
NSH	Signed	-	-	-	-	4.21	3.96	99.12	68.18	-	-	-	-
IDF	Signed	-	-	-	-	4.07	3.83	99.14	69.66	-	-	-	-
Ours	Unsigned	3.09	2.85	99.41	86.38	4.08	3.83	99.14	69.59	4.91	4.58	95.53	56.98

Table 2: Quantitative results on Stanford 3D Scene, Stanford 3D Scans (watertight) and ShapeNet-Cars. Chamfer distances are measured in the unit of $\times 10^{-3}$. For CAP-UDF and LevelSetUDF, their results often include small isolated components. To compute the Chamfer distance, we cleaned their meshes by removing components that were distant from the ground truth point cloud after mesh extraction using DCUDF.

Surface extraction

After obtaining the UDFs, we proceed to extract the surface from the learned UDFs. Due to the relaxation of the non-negative constraints, the target geometry does not align precisely with the zero level set. Instead, we identify the target surface as the local minimal distance surface (which can be either positive or negative) near the zero values.

One possible method for extracting the target geometry from UDFs involves explicitly using the UDF gradient, such as MeshUDF (Guillard, Stella, and Fua 2022) and GeoUDF (Ren et al. 2023), both of which are variants of the standard Marching Cubes algorithm. On each cube edge, if the gradient directions of the UDF at the two endpoints are opposite and their UDF values are below a specified threshold, a zero crossing is marked on that edge. However, this approach is not suitable for our purpose because our learned UDFs do not ensure the positivity of the distance value, and the distance values on the surface may exceed the specified UDF threshold. If this occurs, the iso-surfacing method would exclude these cubes, leading to the creation of extracted surfaces with undesired holes.

To tackle this challenge, we adopt the optimization-based method DCUDF (Hou et al. 2023), which initiates by extracting a double cover using the Marching Cubes algorithm at a small positive iso-value on the UDF. Subsequently, it shrinks the double cover to the local minimal distance surface. This method does not require the UDF to be strictly positive nor does it depend on a threshold to select candidate cubes. As a result, it effectively identifies the local minimal distance surface, yielding a high-quality triangle mesh that accurately represents the target surface.

Experiments

Datasets

We evaluate our method using three datasets: ShapeNet-Cars (Chang et al. 2015) with 108 models², the Stanford 3D Scene Dataset (Zhou and Koltun 2013) with 5 models and the Stanford 3D Scan Repository³ with 8 models. For

²We select all models whose names start with “1”.

³<https://graphics.stanford.edu/data/3Dscanrep/>

each shape, we randomly sample 300K points as input. After learning the UDFs, we employ DCUDF with a resolution of 512^3 to extract the target surface. To evaluate the accuracy, We use Chamfer distance (CD) and F-score as quantitative measures. For F-score, we set the thresholds to 0.5% and 0.25%. Following previous methods (Zhou et al. 2022, 2023), we randomly sample 100K points from both the reconstructed surfaces and the ground truth meshes for computing CD and F-score. All points and meshes are normalized to bounding boxes whose longest edges are 2. We test on an NVIDIA Tesla V100 GPU with 32GB memory (about 5GB used for a UDF learning). It takes about 30 minutes to learn a UDF.

Results & comparisons

We compare our method with three state-of-the-art UDF learning methods: LevelSetUDF (Zhou et al. 2023), CAP-UDF (Zhou et al. 2022) and DUDF (Fainstein, Silless, and Iarussi 2024). Since we adopt DCUDF (Hou et al. 2023) for surface extraction, we also test DCUDF for the three baselines to ensure fairness of comparisons. For CAP-UDF and LevelSetUDF, we observe that DCUDF could produce better results than their original implementations in terms of Chamfer distances and visual effects. But for DUDF, the results of DCUDF are not as good as the original results. Therefore, to report the best results of the baseline methods, we choose to adopt DCUDF for extracting the zero level set from the UDF outputs from both CAP-UDF and LevelSetUDF. While DUDF uses their original results for comparisons. For the ShapeNet-Cars and Stanford 3D Scene datasets, the outputs of DCUDF remain as double-layered meshes, bypassing the min-cut based double-layer segmentation post-processing. In addition, we assess our approach against IDF (Wang, Rahmann, and Sorkine-Hornung 2022) and NSH (Wang et al. 2023b), two state-of-the-art SDF learning methods, on the watertight surfaces with fine geometric details from the Stanford 3D Scan Repository. The results are illustrated in Table 2 and Figure 2.

3D objects with fine geometric details To explore the ability of our method for representing 3D objects with fine geometric details, we evaluate our method on Stanford 3D

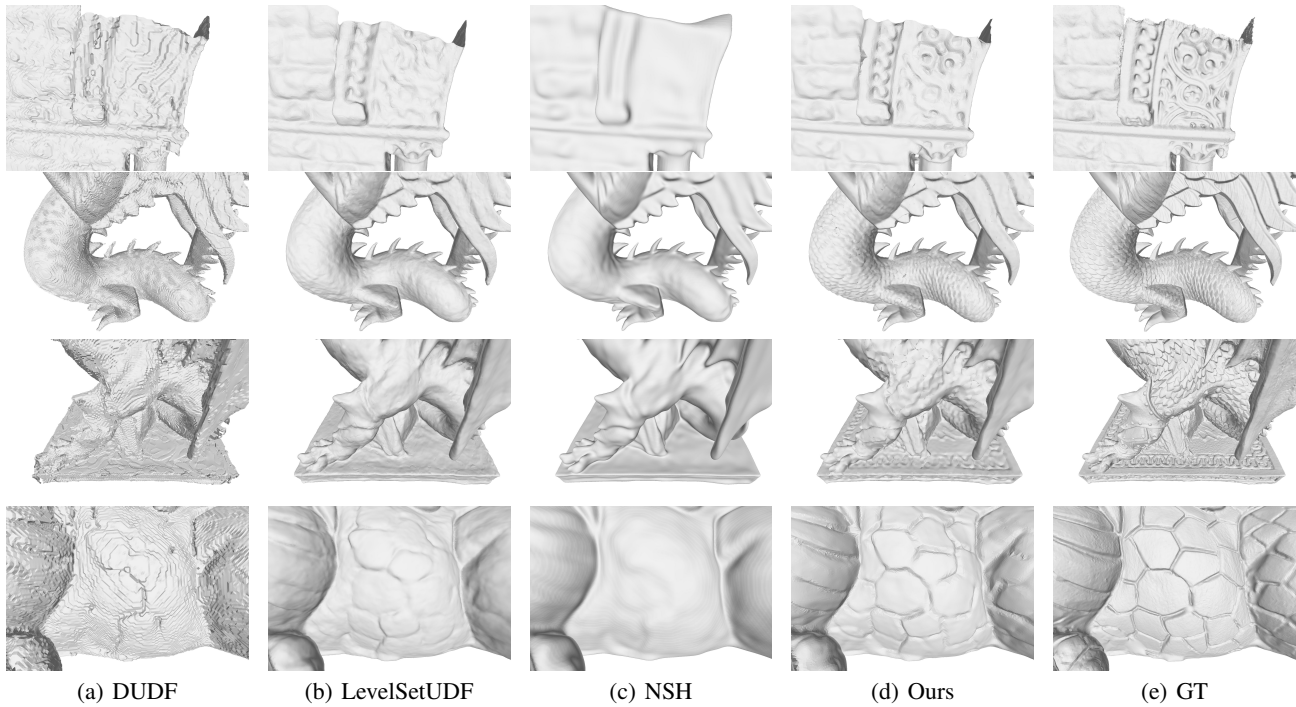


Figure 2: Visual comparison with two recent UDF learning approaches, DUDF (Fainstein, Siless, and Iarussi 2024) and LevelSetUDF (Zhou et al. 2023), and one recent SDF learning method, NSH (Wang et al. 2023b), on surfaces with fine geometric details. Our method yields visually pleasing results, reconstructing significantly more details than the other methods.

Scene dataset and Stanford 3D Scan dataset. As shown in Table 2 and Figure 2, our method achieves the best performance in UDF-based methods, and performs close to SDF-based methods.

3D objects with complex inner structures We further explore our method for representing 3D objects and scenes with complex inner structures. We evaluate our method on Stanford 3D Scene and ShapeNet-Cars datasets. As shown in Table 2 and Figure 4, our method is more stable on complex structures and performs optimally for keeping open boundaries.

Method	Gaussian Noise- $N(0,0.0025)$				Gaussian Noise- $N(0,0.005)$			
	Chamfer-L1 (\downarrow)		F-score (\uparrow)		Chamfer-L1 (\downarrow)		F-score (\uparrow)	
	Mean	Median	$F1^{0.005}$	$F1^{0.0025}$	Mean	Median	$F1^{0.005}$	$F1^{0.0025}$
CAP-UDF	5.76	5.49	92.68	43.39	6.26	5.99	87.82	37.54
LeverSetUDF	5.63	5.37	93.39	45.60	6.42	6.18	86.45	36.05
DUDF	4.27	3.83	97.91	70.57	4.61	4.11	96.25	65.04
Ours	4.07	3.79	98.74	71.89	5.64	5.35	92.74	46.18

Table 3: Quantitative results for reconstruction on noisy point clouds from Stanford 3D Scenes and 3D Scans. We add Gaussian noise $N(0, 0.0025)$ and $N(0, 0.005)$ to the input point clouds and use the SIREN with a frequency of 30 to resist noises.

Noisy point clouds As illustrated in Table 3, we add $N(0, 0.0025)$ and $N(0, 0.005)$ Gaussian noise to the nor-

malized point clouds. To improve the noise resistance, we reduce the SIREN frequency to 30. Our method is best at resisting small noises. DUDF is the best for large noise because it always oversmooths the surface. However, for the same reason, DUDF is the worst at reconstructing details. Apart from DUDF, our DEUDF is better than others for large noisy data. Meanwhile, DEUDF is the best at reconstructing details.

Deviation from zero It is important to note that the distances of the learned UDF of a surface are unlikely to be exactly zero due to inherent learning errors. Interestingly, even though we relax the non-negativity constraint in our DEUDF, the deviation from zero in our results is actually smaller than that of UDFs learned by other methods enforcing strict non-negativity. We evaluate the average absolute distance values of the reconstructed mesh points in the learned UDF to evaluate the deviation. We test on the Stanford 3D Scene and 3D Scan Repository datasets. The average deviations are as follows: DUDF: 2.60×10^{-3} , CAP-UDF: 2.63×10^{-4} , LevelSetUDF: 1.39×10^{-4} , and our DEUDF: 1.34×10^{-4} .

Ablation studies

We conduct ablation studies to demonstrate the effectiveness of each component within our method.

Unconditioned MLPs We assess the impact of using an unconditioned SIREN network on the performance of our

method by comparing it to other versions of the SIREN network that utilize absolute value and softplus function in the output layer, respectively. As shown in Figure 3 and Table 4, the absolute output layer results in large fitting error. On the other hand, due to the vanishing gradient effect of softplus, the reconstructed mesh using the SIREN network with a softplus output is typically over-smoothed.




Figure 3: Visual results of the ablation studies: (a) Applying the absolute value to the output of the SIREN network. (b) Applying the softplus function to the output of the SIREN network. (c) Using uniform Eikonal constraints. (d) Removing normal alignment. (e) Replacing estimated normals with random vector. (f) Applying all components.




Figure 4: Visual comparisons with DUDF, CAP-UDF and LevelSetUDF on an indoor scene of the Stanford 3D Scene dataset featuring noise, an imperfect scan, and two car models of the ShapeNet-Cars dataset showcasing complex structures. Our method remains the open boundaries, such as, the bookshelves, inner structures of the vehicle, car window and car exhaust vent.

Normal alignment To evaluate the contribution of normal alignment to reconstruct geometric details, we compare our

method with and without the normal alignment loss. In addition, we test the robustness of our method when the estimated normals are unreliable. Figure 3 and Table 4 show the effects of the normal alignment loss. Interestingly, we observe that even fully random normals could also improve the results significantly. While the normals are unreliable, the \mathcal{L}_{normal} also locates the local minimal distance surface between the two sample points \mathbf{q}_i^1 and \mathbf{q}_i^2 . This is also beneficial to minimize the error.

	CD-mean	CD-median	$F1^{0.005}$	$F1^{0.0025}$
SIREN+Abs.	5.27	3.07	87.65	77.87
SIREN+softplus	2.91	2.71	99.79	91.08
w/o weighted Eikonal	2.87	2.54	99.26	90.98
w/o normal alignment	3.20	2.97	99.36	88.68
random normal alignment	2.72	2.53	99.87	92.12
DEUDF	2.69	2.51	99.95	93.52

Table 4: Ablation studies on the model ‘‘Stonewall’’ (Figure 3) of the Stanford 3D Scene Dataset.

Weighted Eikonal We explore the effects of different configurations of the Eikonal loss. Specifically, we compare our method using a standard Eikonal loss that is applied uniformly to all sample points with our adaptively weighted Eikonal loss. We observe that the standard Eikonal loss results in learned UDFs with lower accuracy near the zero level sets, leading to numerous small holes in the extracted meshes. In contrast, our adaptively weighted Eikonal loss more effectively addresses the vanishing gradient problem, and stabilizes the learning process, thereby yielding meshes with higher quality. In Figure 3 and Table 4, due to the uniform Eikonal constraint that impedes the UDF learning process at the zero level set, slightly higher distance values are observed around the target surface, which can result in holes in the extracted mesh, even though a larger DCUDF (Hou et al. 2023) threshold has been applied in this example.

Conclusions and Limitations

This paper presents an improved UDF learning method for high fidelity 3D surface reconstruction. The method integrates novel UDF representation, normal alignment, adaptively weighted Eikonal constraint and SIREN network to learn more accurate UDFs. Our DEUDF can not only learn geometry details but also keep boundaries, thereby maintaining better topology. Extensive experiments illustrate that our method produces low Chamfer distances and better topology outperforming state-of-the-art methods.

Our method is primarily designed for dense, evenly sampled point clouds, as sparse point clouds often fail to capture fine geometric details. If the input point cloud is highly uneven, in some sparse locations, the learned unsigned distance values could be high resulting in small holes in the reconstructed mesh. In the future, we aim to extend DEUDF to handle point clouds with non-uniform densities, where points are sparse in smooth areas and dense in regions with fine details.

Acknowledgments

This work was supported in part by the National Key R&D Program of China under Grant 2023YFB3002901, the Basic Research Project of ISCAS under Grant ISCAS-JCMS-202303, the Major Research Project of ISCAS under Grant ISCAS-ZD-202401, the Ministry of Education, Singapore, under its Academic Research Fund Grants (MOE-T2EP20220-0005 & RT19/22), and an OPPO gift fund.

References

- Amenta, N.; and Bern, M. 1998. Surface Reconstruction by Voronoi Filtering. In *Proceedings of SoCG*, 39–48.
- Chabra, R.; Lenssen, J. E.; Ilg, E.; Schmidt, T.; Straub, J.; Lovegrove, S.; and Newcombe, R. 2020. Deep Local Shapes: Learning Local SDF Priors for Detailed 3D Reconstruction. In *Proc. of ECCV*, 608–625.
- Chang, A. X.; Funkhouser, T.; Guibas, L.; Hanrahan, P.; Huang, Q.; Li, Z.; Savarese, S.; Savva, M.; Song, S.; Su, H.; Xiao, J.; Yi, L.; and Yu, F. 2015. ShapeNet: An Information-Rich 3D Model Repository. [arXiv:1512.03012](https://arxiv.org/abs/1512.03012).
- Chibane, J.; Alldieck, T.; and Pons-Moll, G. 2020. Implicit Functions in Feature Space for 3D Shape Reconstruction and Completion. In *Proc. of CVPR*, 6968–6979.
- Chibane, J.; Mir, A.; and Pons-Moll, G. 2020. Neural Unsigned Distance Fields for Implicit Function Learning. In *Proc. of NeurIPS*, 21638–21652.
- Deng, J.; Hou, F.; Chen, X.; Wang, W.; and He, Y. 2024. 2S-UDF: A Novel Two-stage UDF Learning Method for Robust Non-watertight Model Reconstruction from Multi-view Images. In *Proc. of CVPR*.
- Dey, T. K.; and Goswami, S. 2003. Tight Cocone: A Water-Tight Surface Reconstructor. In *Proc. of ACM SMA*, 127–134.
- Fainstein, M.; Siless, V.; and Iarussi, E. 2024. DUDF: Differentiable Unsigned Distance Fields with Hyperbolic Scaling. In *Proc. of CVPR*.
- Guillard, B.; Stella, F.; and Fua, P. 2022. MeshUDF: Fast and Differentiable Meshing of Unsigned Distance Field Networks. In *Proc. of ECCV*, 576–592.
- Hoppe, H.; DeRose, T.; Duchamp, T.; McDonald, J.; and Stuetzle, W. 1992. Surface reconstruction from unorganized points. *SIGGRAPH Comput. Graph.*, 26(2): 71–78.
- Hou, F.; Chen, X.; Wang, W.; Qin, H.; and He, Y. 2023. Robust Zero Level-Set Extraction from Unsigned Distance Fields Based on Double Covering. *ACM Trans. Graph.*, 42(6).
- Hou, F.; Wang, C.; Wang, W.; Qin, H.; Qian, C.; and He, Y. 2022. Iterative Poisson Surface Reconstruction (iPSR) for Unoriented Points. *ACM Trans. Graph.*, 41(4).
- Kazhdan, M.; Bolitho, M.; and Hoppe, H. 2006. Poisson Surface Reconstruction. In *Proc. of SGP*, 61–70.
- Kazhdan, M.; and Hoppe, H. 2013. Screened Poisson Surface Reconstruction. *ACM Trans. Graph.*, 32(3).
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *Proc. of ICLR*.
- Liu, W.; Li, J.; Chen, X.; Hou, F.; Xin, S.; Wang, X.; Wu, Z.; Qian, C.; and He, Y. 2024. Diffusing Winding Gradients (DWG): A Parallel and Scalable Method for 3D Reconstruction from Unoriented Point Clouds. [arXiv:2405.13839](https://arxiv.org/abs/2405.13839).
- Liu, Y.-T.; Wang, L.; Yang, J.; Chen, W.; Meng, X.; Yang, B.; and Gao, L. 2023. NeUDF: Learning Neural Unsigned Distance Fields with Volume Rendering. In *Proc. of CVPR*, 237–247.
- Long, X.; Lin, C.; Liu, L.; Liu, Y.; Wang, P.; Theobalt, C.; Komura, T.; and Wang, W. 2023. NeuralUDF: Learning Unsigned Distance Fields for Multi-view Reconstruction of Surfaces with Arbitrary Topologies. In *Proc. of CVPR*, 20834–20843.
- Loshchilov, I.; and Hutter, F. 2017. SGDR: Stochastic Gradient Descent with Warm Restarts. In *Proc. of ICLR*.
- Mescheder, L.; Oechsle, M.; Niemeyer, M.; Nowozin, S.; and Geiger, A. 2019. Occupancy Networks: Learning 3D Reconstruction in Function Space. In *CVPR*, 4455–4465.
- Ohtake, Y.; Belyaev, A.; Alexa, M.; Turk, G.; and Seidel, H.-P. 2003. Multi-Level Partition of Unity Implicits. *ACM Trans. Graph.*, 22(3): 463–470.
- Park, J. J.; Florence, P.; Straub, J.; Newcombe, R.; and Lovegrove, S. 2019. DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. In *Proc. of CVPR*, 165–174.
- Peng, S.; Niemeyer, M.; Mescheder, L.; Pollefeys, M.; and Geiger, A. 2020. Convolutional Occupancy Networks. In *Proc. of ECCV*, 523–540.
- Ren, S.; Hou, J.; Chen, X.; He, Y.; and Wang, W. 2023. GeoUDF: Surface Reconstruction from 3D Point Clouds via Geometry-guided Distance Representation. In *Proc. of ICCV*, 14214–14224.
- Sitzmann, V.; Martel, J. N. P.; Bergman, A. W.; Lindell, D. B.; and Wetzstein, G. 2020. Implicit neural representations with periodic activation functions. In *Proc. of NeurIPS*.
- Wang, P.; Liu, L.; Liu, Y.; Theobalt, C.; Komura, T.; and Wang, W. 2021. NeuS: Learning Neural Implicit Surfaces by Volume Rendering for Multi-view Reconstruction. *NeurIPS*.
- Wang, R.; Wang, Z.; Zhang, Y.; Chen, S.; Xin, S.; Tu, C.; and Wang, W. 2023a. Aligning Gradient and Hessian for Neural Signed Distance Function. In *Proc. of NeurIPS*.
- Wang, Y.; Rahmann, L.; and Sorkine-Hornung, O. 2022. Geometry-Consistent Neural Shape Representation with Implicit Displacement Fields. In *Proc. of ICLR*.
- Wang, Z.; Zhang, Y.; Xu, R.; Zhang, F.; Wang, P.; Chen, S.; Xin, S.; Wang, W.; and Tu, C. 2023b. Neural-Singular-Hessian: Implicit Neural Representation of Unoriented Point Clouds by Enforcing Singular Hessian. *ACM Trans. Graph.*, 42(6): 274:1–274:14.
- Xu, R.; Dou, Z.; Wang, N.; Xin, S.; Chen, S.; Jiang, M.; Guo, X.; Wang, W.; and Tu, C. 2023. Globally Consistent Normal Orientation for Point Clouds by Regularizing the Winding-Number Field. *ACM Trans. Graph.*, 42(4).
- Yang, X.; Lin, G.; Chen, Z.; and Zhou, L. 2023. Neural Vector Fields: Implicit Representation by Explicit Learning. In *Proc. of CVPR*, 16727–16738.

- Yariv, L.; Gu, J.; Kasten, Y.; and Lipman, Y. 2021. Volume Rendering of Neural Implicit Surfaces. In *Proc. of NeurIPS*, 4805–4815.
- Ye, J.; Chen, Y.; Wang, N.; and Wang, X. 2022. GIFS: Neural Implicit Function for General Shape Representation. In *Proc. of CVPR*, 12819–12829.
- Yu, Z.; Dou, Z.; Long, X.; Lin, C.; Li, Z.; Liu, Y.; Müller, N.; Komura, T.; Habermann, M.; Theobalt, C.; et al. 2025. Surf-D: Generating High-Quality Surfaces of Arbitrary Topologies Using Diffusion Models. In *European Conference on Computer Vision (ECCV2024)*, 419–438. Springer.
- Zhou, J.; Ma, B.; Li, S.; Liu, Y.-S.; and Han, Z. 2023. Learning a More Continuous Zero Level Set in Unsigned Distance Fields through Level Set Projection. In *Proc. of ICCV*, 3158–3169.
- Zhou, J.; Ma, B.; Liu, Y.-S.; Fang, Y.; and Han, Z. 2022. Learning Consistency-Aware Unsigned Distance Functions Progressively from Raw Point Clouds. In *Proc. of NeurIPS*, 16481–16494.
- Zhou, J.; Zhang, W.; Ma, B.; Shi, K.; Liu, Y.-S.; and Han, Z. 2024. UDiFF: Generating Conditional Unsigned Distance Fields with Optimal Wavelet Diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Zhou, Q.-Y.; and Koltun, V. 2013. Dense Scene Reconstruction with Points of Interest. *ACM Trans. Graph.*, 32.