

Unsigned Orthogonal Distance Fields: An Accurate Neural Implicit Representation for Diverse 3D Shapes

Yujie Lu[†], Long Wan[†], Nayu Ding[†], Yulong Wang[†], Shuhan Shen^{‡,¶}, Shen Cai^{†,*}, Lin Gao^{§,¶,*}

[†]Visual and Geometric Perception Lab, Donghua University

[‡]Institute of Automation, Chinese Academy of Sciences

[§]Institute of Computing Technology, Chinese Academy of Sciences

[¶]University of Chinese Academy of Sciences

*Corresponding author: hammer_cai@163.com, gaolin@ict.ac.cn

<https://github.com/cscvlab/UODFs>

Abstract

Neural implicit representation of geometric shapes has witnessed considerable advancements in recent years. However, common distance field based implicit representations, specifically signed distance field (SDF) for watertight shapes or unsigned distance field (UDF) for arbitrary shapes, routinely suffer from degradation of reconstruction accuracy when converting to explicit surface points and meshes. In this paper, we introduce a novel neural implicit representation based on unsigned orthogonal distance fields (UODFs). In UODFs, the minimal unsigned distance from any spatial point to the shape surface is defined solely in one orthogonal direction, contrasting with the multi-directional determination made by SDF and UDF. Consequently, every point in the 3D UODFs can directly access its closest surface points along three orthogonal directions. This distinctive feature leverages the accurate reconstruction of surface points without interpolation errors. We verify the effectiveness of UODFs through a range of reconstruction examples, extending from simple watertight or non-watertight shapes to complex shapes that include hollows, internal or assembling structures.

1. Introduction

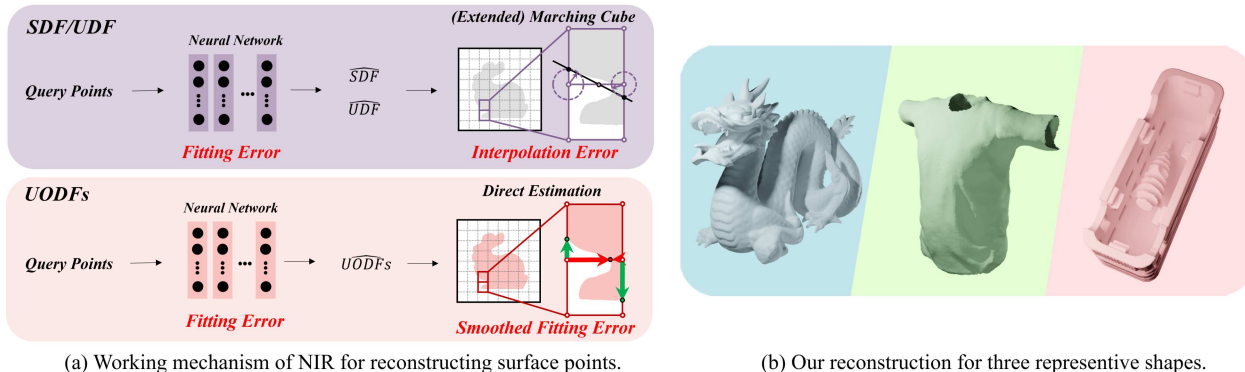
Neural implicit representation (NIR) of three-dimensional (3D) shapes has emerged as a noteworthy area of research in computer vision and graphics. Currently, the most prevalent NIR for 3D shapes is based on signed distance field (SDF), utilized in various applications such as shape reconstruction [6, 12, 24, 26, 29–32], or new view synthesis [22, 36, 37, 39, 41]. For any 3D shape, the SDF value of each spatial point reveals two properties: (1) whether the point is inside the shape (indicated by the SDF

sign); (2) the minimal distance among all directions from this point to the shape surface (depicted by its absolute value). Despite the advantages offered by SDF based NIR in multiple contexts, there still exist several scenarios where SDF is not applicable.

The first scenario arises when 3D shapes are not watertight or contain internal structures. In such instances, assigning an SDF sign to 3D points might be challenging or inaccurate. To address this limitation and cater to different types of 3D shapes, researchers have proposed the unsigned distance field (UDF) based NIR [8, 10, 14, 20, 35, 40]. Losing the sign, a point’s UDF value represents the minimal distance traversing all directions to the shape surface, irrespective of the point being inside or outside the model.

The second scenario to consider is the accurate reconstruction of surface points from the NIR of a 3D shape. A common approach for reconstructing surface points from SDF or UDF is the marching cubes (MC) algorithm [23]. This process is illustrated in the upper part in Fig. 1 (a). A neural network makes a prediction of the SDF or UDF values at grid (cube) corners, affected by fitting errors. The MC algorithm applicable to SDF then estimates grid edge points (GEP) through linear interpolation (see the zoom-in grids for detail). However, this approximation can be inaccurate for non-linear local shapes due to interpolation errors. As a result, while employing SDF based NIR, the reconstruction of surface points suffers from these two types of errors. Similar fitting and interpolation issues are also encountered in deep UDF methods, such as those detailed in [8, 14, 35], where the MC algorithm is extended in various ways to incorporate the sign in the UDF at grid corners. Furthermore, in downstream tasks like 3D object classification and part segmentation (see [15, 27, 28, 38]), shapes are often represented by 1024 surface points. The implicit representation and reconstruction of surface points at this scale requires large grids in the MC algorithm, further amplifying the interpolation error.

This work was supported in part by the National Natural Science Foundation of China (Grants 62322210, 62206046 and U2033218), Natural Science Foundation of Shanghai (Grant 22ZR1400200), and Fundamental Research Funds for the Central Universities (No. 2232023Y-01).



(a) Working mechanism of NIR for reconstructing surface points.

(b) Our reconstruction for three representative shapes.

Figure 1. Overview of unsigned orthogonal distance fields (UODFs) based NIR. Unlike SDF or UDF based NIR, which suffers from interpolation errors in surface point reconstruction, UODFs directly estimate surface points and mitigate fitting errors by averaging predictions for each GEP. The upper zoom-in grids illustrate the inaccuracies in traditional methods, where the middle grid edge point (colored purple) is approximately estimated with the distance values (denoted by two dotted circles) of two grid corners. This estimated GEP is far from its true position (colored red) which is jointly predicted by values of our UODFs (two red arrows), as shown in the lower zoom-in grids.

In this paper, we propose the unsigned orthogonal distance fields (UODFs) based NIR. Compared with conventional SDF and UDF, the most distinguishable feature of UODFs is that the minimal distances to the shape surface are defined along three orthogonal directions in 3D space. This means every point in UODFs can directly access its closest surface points along three orthogonal directions, when they exist. Fig. 1 (a) depicts how SDF (or UDF) and UODFs based NIR work in surface points reconstruction. All predicted distance values denoted by $\dots\widehat{DF}(s)$ at grid corners are affected by the fitting error of neural network. However, in contrast to SDF or UDF, the estimation of GEP for UODFs does not introduce interpolation errors. Instead, the fitting error can be mitigated through averaging multiple predictions for each GEP. Therefore, UODFs based NIR naturally leverages accurate reconstruction for diverse 3D shapes, three of which are shown in Fig. 1 (b).

The key contributions of this paper are as follows:

- UODFs based NIR is proposed, which allows for the representation and reconstruction of diverse 3D shapes (such as watertight, non-watertight, multi-layer, or assembling models) in a unified manner.
- UODFs based NIR enables the direct estimation of surface points from multiple distant sample points along three orthogonal directions, facilitating the interpolation-free reconstruction of grid edge points.
- UODFs based NIR significantly outperforms traditional SDF or UDF based NIR in terms of reconstruction accuracy, especially for open shapes or when reconstructing small point clouds.

2. Related Works

The focus in this paper lies in the neural distance field representation of 3D shapes, which could be broadly grouped into three categories for review: SDF based NIR, UDF based NIR, and other types of neural distance field.

2.1. SDF based Neural Implicit Representation

Signed distance field (SDF), which implicitly represents a 3D shape as the zero level-set of spatial positions, is popular in the domain of neural implicit representation. A number of methods, including but not limited to DeepSDF [26], IGR [13], SAL [1], FFN [31], SIREN [29], NI [12], and GC [3], generally use a multi-layer perceptron (MLP) network to globally fit SDF of 3D shapes. Another kind of methods adopts a local fitting strategy aimed at learning finer shape details, such as OctField [32], DeepLS [6], NGLoD [30] and Instant-NGP [24]. The marching cubes (MC) algorithm is conventionally employed to extract the zero isosurface of the SDF, enabling simultaneous computation of grid edge points and meshing. In addition, SDF is also adopted to implicitly represent geometric shapes in the task of new view synthesis, such as NeuS [36], MonoSDF [41], and VOXURF [39]. These SDF based methods come with the limitation that they can only represent models that are watertight.

2.2. UDF based Neural Implicit Representation

Unsigned distance field (UDF) based neural implicit methods are proposed to represent arbitrary shapes, such as NDF [10], CSP [34], GIFS [40], DeepCurrents [25], 3PSDF [8], NDC [9], MeshUDF [14], and HSDF [35]. Within the context of new view synthesis, there also exist works to learn UDF from multi-view images, such as NeUDF [20] and Neural UDF [21]. Although all these approaches represent open shapes, the prediction and meshing processes possess unique distinctions. For instance, NDF firstly obtains discrete points from the predicted UDF, and then employs the ball-pivoting algorithm [4] (BPA) to accomplish meshing; NeUDF follows a similar process but uses SPSR [18] for meshing surface points. MeshUDF classifies the signs of grid corners utilizing the UDF gradient, unlike HSDF that predicts the signs and UDF values of grid corners simultaneously. Both methods extend MC [23] for

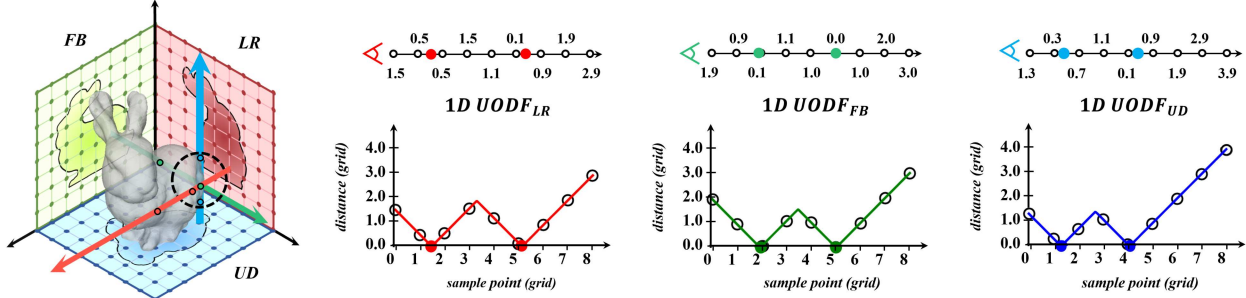


Figure 2. Sketch of UODFs. For understanding the three orthogonal components, 1D derivative, and possible discontinuity between adjacent rays of UODFs, refer to the three characteristics concluded in Sec. 3.2.

reconstructing surface points and meshing from UDF, and thus suffer from the fitting error of neural network and the interpolation error of adjacent corners, as shown in Fig. 1.

2.3. Other Types of Distance Field

A traditional distance field representation, named directed distance fields (DDFs) [19], was proposed in 2001 to better extract surface points on sharp edges. DDFs are calculated explicitly and used for subsequent contour reconstruction by extending the MC algorithm. Although our proposed UODFs based NIR relates to DDFs in terms of distance field definition, the focal points of our study involve fitting UODFs using neural networks and reconstructing GEP, which distinguishes our work from [19].

There exist some works [2, 17] studying deep distance fields along arbitrary directions for 3D shape representation. While these approaches exhibit flexibility in reconstructing surface points, the accuracy does not match the levels achieved by SDF or UDF based NIR.

3. Method

3.1. Definition of UODFs

The unsigned orthogonal distance of a point is defined as the closest distance from this point to the shape surface, along one orthogonal direction. The term ‘unsigned’ means the distance value can never be negative, and there is no distinction between the inside and outside of the shape. In 3D space, there are three orthogonal directions, denoted by ‘left-right (LR)’, ‘front-back (FB)’ and ‘up-down (UD)’ respectively. Therefore, UODFs portray each shape as a combination of three distance fields along these orthogonal directions, which are denoted by $UODF_{LR}$, $UODF_{FB}$ and $UODF_{UD}$, respectively. The ground truth values of UODFs can be computed by ray stabbing along the three orthogonal directions. For a normalized model \mathcal{A} , its surface is denoted by \mathcal{S} . UODFs are limited in the normalized cube space \mathcal{R} ($\mathcal{R} \in [-1, 1]^3$). Denote the ray crossing a point p along one orthogonal direction X by $ray(p, X)$ ($X \in \{LR, FB, UD\}$). The subset $D(p, X)$ denotes all intersections of the surface \mathcal{S} with $ray(p, X)$. UODFs of p is expressed by

$$UODFs(p) = \{UODF_{LR}(p), UODF_{FB}(p), UODF_{UD}(p)\} \quad (1)$$

where each UODF of the point is defined by




Figure 3. One slice from the $UODF_{UD}$ of the 3D model ‘Dragon’. Discontinuity between adjacent ‘UD’ rays occurs when there is a change in intersection with the surface. The grayscale areas denote undefined $UODF_{UD}$ in ground truth calculation and our prediction.

$$UODF_X(p) = \arg \min_{q \in D(p, X)} (|p - q|), \quad \text{if } D(p, X) \neq \emptyset \quad (2)$$

If $D(p, X)$ is empty, indicating $ray(p, X)$ does not intersect with \mathcal{S} , $UODF_X(p)$ is not defined by a specific value.

3.2. Characteristics of UODFs

UODFs possess several characteristics that significantly distinguish them from SDF or UDF.

Characteristic 1. UODFs consist of three unsigned distance fields along orthogonal directions. Any ray in an orthogonal direction corresponds to a 1D unsigned distance field. Thus, $UODF_X$ can be seen as a collection of 1D unsigned distance fields for all parallel rays in the orthogonal direction X . Fig. 2 depicts three orthogonal rays (red, green, and blue) of a point, their intersections with the shape (two surface points for each ray), and their 1D unsigned distance fields. Here rays and sample points (hollow circles) locate at grid corners, which will be explained in detail in Sec. 3.4 and Sec. 3.5.

Characteristic 2. The absolute value of 1D derivative of each UODF for each point is equal to 1. Fig. 2 also demonstrates that the UODF values of equally-spaced points on one side of the surface point form an arithmetic sequence. Each line segment in the 1D UODF diagrams is at an angle of ± 45 degrees. This characteristic can be utilized to mitigate estimation errors of surface points from UODF predictions of multiple sample points.

Characteristic 3. $UODF_X$ may display discontinuity between adjacent parallel rays, which distinguishes it from the SDF or UDF that is continuous everywhere. Fig. 3 depicts

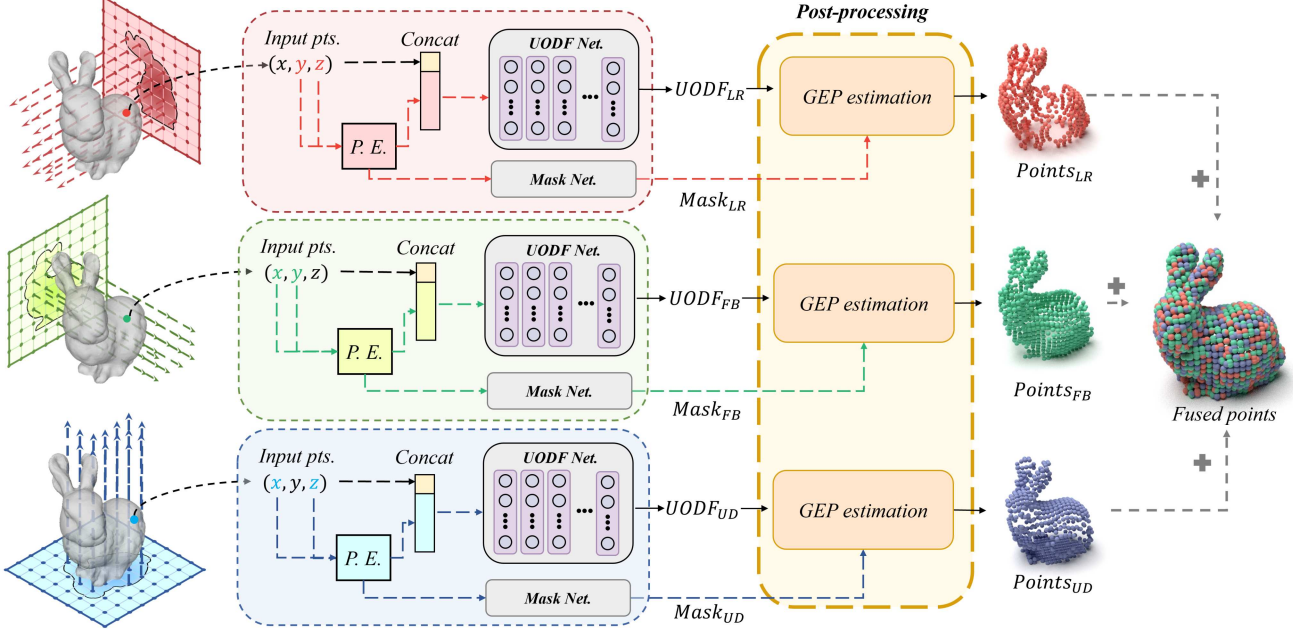


Figure 4. Network architecture and processing pipeline. Each UODF is individually regressed using a UODF and mask network. The grid edge points in one orthogonal direction can be estimated from its UODF, followed by points fusion to form final reconstruction results.

a slice of $UODF_{UD}$ from the 3D model ‘Dragon’. It is evident that the $UODF_{UD}$ values of two neighborly points on adjacent rays can differ greatly. This unique characteristic mirrors the behavior observed in laser sensors and facilitates the direct and accurate prediction of fine details.

In summary, these three characteristics of UODFs are beneficial to the accurate estimation of surface points, especially for diverse complex shapes or small point clouds. Since each UODF is defined on parallel rays, the UODFs in the non-occluded outside regions are basically equivalent to the depth maps captured from six facets of a unit cube. However, UODFs additionally manage the occluded outside and inside regions of 3D models in a unified manner.

3.3. Network Architecture and Processing Pipeline

Fig. 4 depicts the proposed network architecture and processing pipeline. Parallel rays are defined on each orthogonal plane. For a spatial point, its 2D coordinates on one orthogonal plane are employed to extract a 42-dimensional feature via position encoding (‘P. E.’ module). This ray feature, combined with the 3D coordinates, forms the 45-dimensional input for a multi-layer perceptron (MLP) network to regress the UODF in this orthogonal direction. Meanwhile, to avoid regressing the UODF of non-intersected rays which are not defined in Eq. 2, the ray feature is utilized to regress a 2D mask on this orthogonal plane (see the area in each silhouette). Subsequently, the surface points in this orthogonal direction can be estimated from rays that intersect the shape, which is illustrated in Sec. 3.5.

The networks to regress the UODF and the mask in each orthogonal direction are a 10×256 MLP and 3×256 MLP,

respectively. The supervised mask values are set to be ‘0’ and ‘1’, denoting the outside and inside of the silhouette respectively. The loss function of $UODF_X$ is composed of three component loss terms:

$$\mathcal{L}_{all} = \lambda_1 * \mathcal{L}_{value} + \lambda_2 * \mathcal{L}_{der} + \lambda_3 * \mathcal{L}_{pred}. \quad (3)$$

where the weights are experimentally set to: $\lambda_1 = 3000$, $\lambda_2 = 50$, and $\lambda_3 = 1000$.

For any point p , the term \mathcal{L}_{value} means that its predicted UODF value $\widehat{UODF}_X(p)$ should be equal to the ground truth $UODF_X(p)$, which is formulated as

$$\mathcal{L}_{value} = | \widehat{UODF}_X(p) - UODF_X(p) |. \quad (4)$$

According to ‘Characteristic 2’ analyzed in Sec. 3.2, the term \mathcal{L}_{der} denotes the absolute value of UODF derivative for any point should be equal to 1, which is expressed by

$$\mathcal{L}_{der} = | |UODF'_X(p)| - 1 |. \quad (5)$$

Inspired by the recent works [3, 35], the term \mathcal{L}_{pred} is extra added to constrain the UODF value of the predicted surface point of p , which is expressed by

$$\mathcal{L}_{pred} = | UODF_X(p - \text{sign}(UODF'_X(p)) * UODF_X(p)) |. \quad (6)$$

where $\text{sign}(UODF'_X(p))$ denotes the direction of derivative. Consequently, $p - \text{sign}(UODF'_X(p)) * UODF_X(p)$ represents the predicted surface point based on p and $UODF_X$.

3.4. Sampling During Inference and Training

The sampling process entails two steps: rays sampling upon orthogonal planes and points sampling along rays. During inference, both rays and points sampling are discretized. For the sake of comparison with MC [23] and MeshUDF [14], rays and points are sampled at grid corners with equal intervals in our experiments. However, UODFs

also facilitate an arbitrary 2D and 1D pattern for rays and points sampling, respectively.

During training, due to the possible discontinuity between adjacent rays analyzed as ‘Characteristic 3’ in Sec. 3.2, rays sampling is retained as discrete. We densely sample each ray on three orthogonal planes with a resolution 257*257. For points sampling along rays, we uniformly sample 256 points per 10 epochs. This accommodates the continuous nature of 1D UODF along each ray, allowing for any given point on the ray to accurately predict its closest surface point.

3.5. Surface Points Estimation for Each Ray

Although points sampling along rays in inference could be random for UODFs, we sample points with equal intervals in alignment with MC’s resolution. Subsequently, surface points along each ray are estimated from these equally spaced sample points. Fig. 5 illustrates how the surface points along a ray are estimated. Three intersection points, labelled ‘A’, ‘B’, and ‘C’, are represented by solid circles along this ray. The lower part of the figure depicts the 1D UODF along this ray, with nine equally spaced points (hollow) sampled. Unlike SDF or UDF, each sample point directly predicts its nearest surface point. As a result, our UODFs based NIR can potentially estimate a surface point from several sample points exhibiting consistent gradients. For instance, the nine sample points on this ray are divided into four segments, each of which estimates one surface point by averaging the predictions of all sample points within this segment. Both segments ‘S1’ and ‘S2’ estimate a surface point ‘A’. If the distance between the two estimated points is smaller than a threshold τ (with $\tau=1/512$ in our experiments), these two estimated points will be averagely merged into a new one. For segments ‘S3’ and ‘S4’, if the distance between their estimation points is greater than τ , both the surface points ‘B’ and ‘C’ are reconstructed, although they are on the same grid edge. This phenomenon occurs when two closed models are close together or a thin plate model is represented by two planes. However, the MC algorithm is limited to estimating at most one intersection point on each grid edge, accompanying the introduction of the interpolation error. In contrast, our UODF based estimation can reconstruct two intersection points on one grid edge and smooth out the predictions of multiple sample points, influenced by the fitting error of the network.

Although the surface point estimation in UODFs makes use of distance values from sample points on a ray, it is quite different from another point estimation technique suitable to SDF and UDF known as sphere tracing [16]. Sphere tracing adopts a conservative strategy, only permitting ray stepping in accordance with the unsigned distance from the current position. On the contrary, one point in UODFs directly predicts the closest point in an orthogonal direction, irrelevant to the ray tracing scheme.




Figure 5. Surface points estimation along a ray.

After estimating surface points in three orthogonal edges, a fusion operation is engaged to derive the final GEP, as shown in Fig. 4. Specifically, if there are estimated points on at least 3 adjacent edges, those points are retained. Otherwise, they will be deleted. This strategy is similar to MC and allows for the effective filtering of isolated points.

3.6. Mesh Extraction

In the aforementioned processing pipeline, UODFs based NIR reconstructs surface points for arbitrary shapes in a unified manner. If required, an additional meshing algorithm can be applied to extract mesh from surface points. We apply a masked Poisson method for mesh extraction from UODFs, as recommended by NeUDF [33]. Specifically, the screened Poisson surface reconstruction (SPSR) technique [18] is used to extract a watertight mesh, followed by masking out the spurious triangles that have a considerable distance from the reconstructed points.

Despite the SPSR method, without the fine-tuning of the default hyper-parameters, generates superior mesh results in our experiments, we explore another way for simultaneous GEP reconstruction and mesh extraction, based on the MeshUDF [14] algorithm. These findings are preliminary and will be refined in subsequent research.

4. Experiments

4.1. Dataset and Metrics

Dataset. Over 50 models from a variety of datasets are used to verify the reconstruction accuracy of diverse shapes. Specifically, 36 shapes come from Thingi10K [42], including 32 shapes in the Thingi32 subset, which is used in NI [12] and NGLoD [30]. Other shapes come from the Stanford 3D Scanning Repository [11], ShapeNet [7], the MGN dataset [5], and our own self-generation.

Metrics. The reconstruction accuracy of grid edge points (GEP) is evaluated by a L_2 Chamfer distance (CD), abbreviated as ‘GEP-CD’. For mesh, a L_2 CD and normal consistency (NC) of 100,000 evenly distributed points are measured, abbreviated as ‘Mesh-CD’ and ‘Mesh-NC’ respectively. For the sake of fair comparison, all reconstructed meshes from tested NIR methods employ predicted distance

values of 257^3 grid corners, except the experiments with multiple grid resolutions as shown in Sec. 4.6.

4.2. Reconstruction of Watertight Shapes

The first experiment is to verify the reconstruction accuracy of watertight shapes, of which the SDF sign is typically computed with high accuracy.

Table 1 exhibits the reconstruction metrics on the Thingi32 dataset, with comparative baselines established by SIREN [29] and NGLOD [30]. For SIREN, the normal information is not utilized to make the same supervision signal across all methods. For NGLOD, we test two levels of detail: LOD3 and LOD5, each with prior octree knowledge at maximum resolutions of 16^3 and 64^3 , respectively. In both CD metrics, UODFs outperforms the SOTA NGLOD5, even though it has the advantage of locally fitting SDF. From the table, it is also evident that the meshing technique significantly amplifies the CD, resulting in a decrease in reconstruction accuracy, whether using the MC algorithm used by NGLOD or the masking Poisson method employed in our work.

Fig. 6 visualizes the reconstructed meshes for one Thingi32 shape ‘Nandi the Bull’ and one Stanford Scanned shape ‘Dragon’. Except the three SDF based methods, the NIR method GIFS [40] dealing with arbitrary shapes also participates in the comparison. The zoom-in figures indicate that our UODFs based NIR can accurately reconstruct the fine details of these watertight shapes, but the UDF based GIFS fails.

4.3. Reconstruction of Non-watertight Shapes

In this subsection, we evaluate the reconstruction performance on non-watertight shapes. The comparison involves three UDF based methods, including NDF [10], HSDF [35], GIFS [40]. For the training of the UDF of individual shapes, $2M$ spatial points are sampled. Table 2 shows the metrics of 10 representative garments from the MGN dataset. The performance of the proposed UODFs based method substantially surpasses the other UDF based methods. It is worth noting that the accuracy metrics of the non-watertight models reconstructed by us are on the same level as those of the watertight models shown in Table 1. However, compared with the SOTA SDF based NGLOD, UDF based methods generally perform unsatisfactorily, due to different fitting strategies and the introduction of additional errors. For example, HSDF utilizes the shape feature extracted by 3D CNN to globally fit UDF and additionally requires to classify the inside and outside of grid corners.

As displayed in Fig. 7, we show the reconstructed meshes for the model ‘Bunny’ from the Stanford dataset, which is unclosed at its bottom. Although the three UDF based methods reconstruct the complete model without significant holes that should not exist, only our method is capable of reconstructing finer details. The NGLOD5 method

Table 1. Average reconstruction accuracy on Thingi32 dataset.

Metric \ Method	SDF			UODFs
	SIREN	NGLOD3	NGLOD5	Ours
CD-GEP ($*10^5$) ↓	149	0.684	0.432	0.378
CD-Mesh ($*10^5$) ↓	147	2.99	2.83	2.69
NC-Mesh ↑	92.3	98.0	98.5	98.4

Table 2. Average reconstruction accuracy on MGN10 dataset.

Metric \ Method	UDF			UODFs
	NDF	HSDF	GIFS	Ours
CD-GEP ($*10^5$) ↓	128	13.0	4.95	0.227
CD-Mesh ($*10^5$) ↓	128	14.4	6.1	1.93
NC-Mesh ↑	90.8	88.0	92.0	99.6

dealing with watertight models also participates in this comparison. However, as the SDF sign is incorrectly calculated near the bottom of the model, there exist obvious artifacts.

4.4. Reconstruction of Complex Shapes

The shapes in the above two subsections are structurally simple even though they have rich surface details. In this subsection, we test various shapes with complex structure. Reconstruction results of two complex shapes as examples are shown in Fig. 8. The first one is ‘Hilbert Cube’ from the Thingi10K dataset. Although this shape is watertight, the topology is challenging since it contains cavities on a large scale. Moreover, the SDF values may be calculated incorrectly in local areas. The second shape is an assembling model we generate, consisting of the internal multi-layer non-watertight fish and the external hollow yet watertight box. It can be seen from the zoom-in views that the three UDF based methods reconstruct the rough surface of these two complex shapes. NGLOD5 reconstructs the details of the watertight shape, but fails for the internal multi-layer fish. Our UODFs significantly outperforms the others, displaying its superiority.

4.5. Ablation of Loss Function

An ablation study is conducted for the proposed loss function. Specifically, we use three distinct shapes, which are the watertight ‘Nandi the Bull’ (in Fig. 6), the non-watertight ‘Bunny’ (in Fig. 7), and the complex ‘Hilbert Cube’ (in Fig. 8), to evaluate the reconstruction accuracy under varying loss configurations. Table 3 reveals that most results differ slightly, mainly because the loss term \mathcal{L}_{value} plays a most important role in our supervised training.

4.6. Reconstruction at Multi-resolution Grids

In this subsection, we conduct reconstructions at varying grid resolutions, ranging from 32^3 to 256^3 . Fig. 9 separately displays the reconstructed meshes and grid edge points (GEP) of a T-shirt. To demonstrate the superior reconstruction accuracy of our UODF-based NIR, we compare it with MeshUDF [14], which is an extended marching cube algorithm designed for UDF. MeshUDF is employed to use ground truth UDF values and gradients at grid cor-

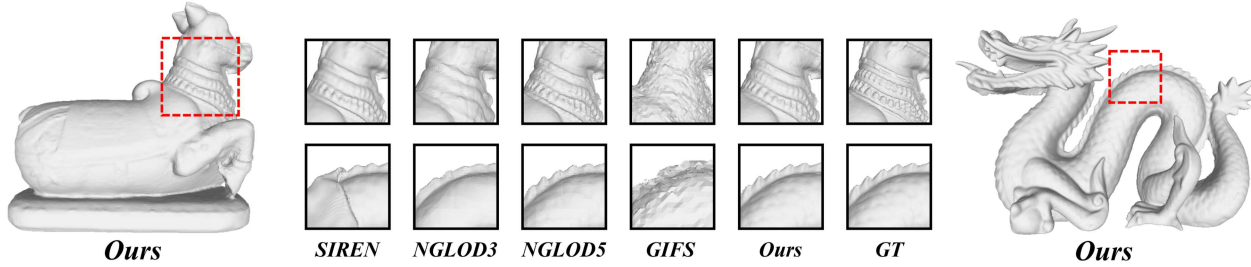


Figure 6. Reconstructed meshes for two watertight shapes.




Figure 7. Reconstructed meshes for the non-watertight ‘Bunny’.

Table 3. Ablation study of loss function on three shapes. The metrics in each cell are CD-GEP($\times 10^5$) \downarrow , CD-Mesh($\times 10^5$) \downarrow , and NC-Mesh \uparrow , respectively.

Loss \ Shape	‘Nandi the Bull’	‘Bunny’	‘Hilbert Cube’
All	0.228 / 2.77 / 98.5	0.170 / 2.73 / 99.2	0.00532 / 12.0 / 92.1
w/o \mathcal{L}_{der}	0.238 / 2.79 / 98.5	0.161 / 2.93 / 99.1	0.00663 / 12.0 / 92.0
w/o \mathcal{L}_{pred}	0.297 / 2.79 / 98.4	0.152 / 2.93 / 99.2	0.256 / 12.3 / 91.9

ners as input. It is worth noting that in this scenario, the reconstruction results of MeshUDF do not suffer from fitting errors of neural network, but are only affected by interpolation errors. For this non-watertight shape, our method consistently outperforms MeshUDF at all tested grid resolutions, even though MeshUDF uses ground truth UDF as input. Additionally, MeshUDF tends to create spurious meshes around the edge of the clothing, particularly noticeable at lower resolutions. We present the GEP reconstruction results in the lower part of the figure. The metric values shown, along with the zoomed-in views (located in the middle and right areas), clearly illustrate that our reconstructed GEP are in closer proximity to the model surface.

4.7. Discussion and Limitation

Additional experimental results are available in the supplementary material, where UODFs based NIR continues to exhibit exceptional reconstruction performance. This achievement can be attributed to several key factors discussed below:

1) Each UODF enjoys a similar working mechanism to a planar laser, capable of accurately measuring distances to the closest surface points along a direction. Notably, UODF further works within occluded or internal regions.

2) It is verified by us that a discontinuous UODF can be

fitted well by an MLP network. This discovery challenges the traditional belief that only continuous fields (e.g., SDF) are suitable for neural network fitting.

3) Interpolation-free estimation of GEP in one orthogonal direction, as well as the fusion of points across three directions, is robust to any model tested. If a mesh is required, the adopted masked SPSR also performs robustly.

A notable limitation of UODFs is the necessity of three neural networks, each needed to fit one UODF. This requirement leads to a relatively large number of parameters in neural networks at present. While certain correlation has been observed between distance fields in different directions, the prospect of fitting three fields within just one neural network is an area of interest for future exploration.

5. Conclusion

In this paper, we propose unsigned orthogonal distance fields (UODFs) based NIR for accurate reconstruction of diverse 3D shapes. UODFs diverge from conventional SDF and UDF in their unique characteristics, which include combining each UODF across three orthogonal directions, estimating surface points directly from distant sample points, and exhibiting discontinuities between rays. Consequently, specific neural networks are required for UODFs regression, along with post-processing methods for surface points reconstruction. Thorough experiments on more than 50 diverse shapes validate that UODFs consistently outrank all competitors in a unified reconstruction manner. UODFs could even outperform traditional meshing methods that use ground truth SDF or UDF input data. In the future, we plan to explore more applications of UODFs, such as in real-time rendering of shapes. The feasibility of fitting UODFs within a single neural network is also worth pursuing.

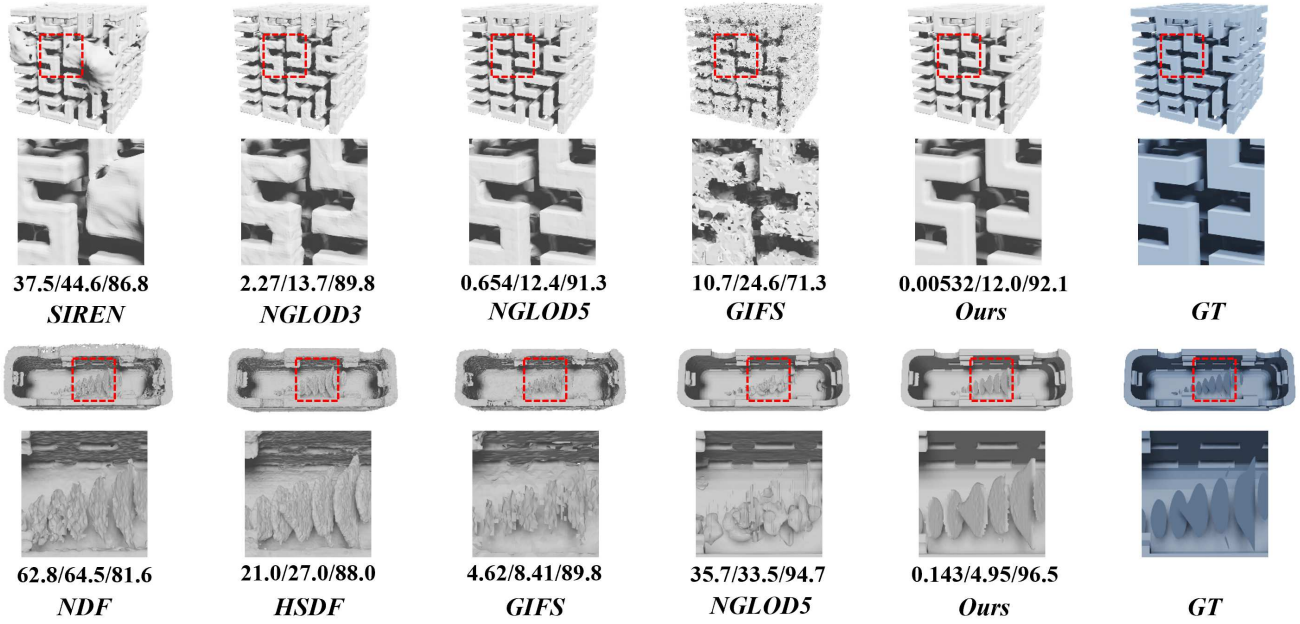


Figure 8. Reconstructed meshes for two representative shapes with complex structure. The metrics from left to right below each shape are CD-GEP(*10⁵) ↓, CD-Mesh(*10⁵) ↓, and NC-Mesh ↑, respectively.




Figure 9. Reconstruction results at various resolutions of grids. The metrics from left to right below each mesh are CD-GEP(*10⁵) ↓ and NC-Mesh ↑ respectively. The metric below each GEP result (along with red ground truth model) is CD-GEP(*10⁵) ↓. The zoom-in views in the middle highlight the spurious GEP results of MeshUDF. The zoom-in 3D views on the right draw local grids and ground truth GEP (red spheres), demonstrating the deviation of GEP reconstruction caused by the interpolation error (highlighted by black dotted boxes).

References

- [1] M. Atzmon and Y. Lipman. Sal: Sign agnostic learning of shapes from raw data. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [2] T. Aumentado-Armstrong, S. Tsogkas, S. Dickinson, and A. Jepson. Representing 3d shapes with probabilistic directed distance fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19321–19332, 2022. 3
- [3] M. Baorui, J. Zhou, Y. Liu, and Z. Han. Towards better gradient consistency for neural signed distance functions via level set alignment. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2, 4
- [4] F. Bernardini, J. Mittleman, H. Rushmeier, C. Silva, and G. Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE transactions on visualization and computer graphics (TVCG)*, 5(4):349–359, 1999. 2
- [5] B. Bhatnagar, G. Tiwari, C. Theobalt, and G. Pons-Moll. Multi-garment net: Learning to dress 3d people from images. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 5
- [6] R. Chabra, J. E. Lenssen, E. Ilg, T. Schmidt, J. Straub, S. Lovegrove, and R. Newcombe. Deep local shapes: Learning local sdf priors for detailed 3d reconstruction. In *European Conference on Computer Vision (ECCV)*, 2020. 1, 2
- [7] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 5
- [8] W. Chen, C. Lin, W. Li, and B. Yang. 3psdf: Three-pole signed distance function for learning surfaces with arbitrary topologies. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2
- [9] Z. Chen, A. Tagliasacchi, T. Funkhouser, and H. Zhang. Neural dual contouring. *ACM Transactions on Graphics (TOG)*, 41(4):1–13, 2022. 2
- [10] J. Chibane, A. Mir, and G. Pons-Moll. Neural unsigned distance fields for implicit function learning. In *Neural Information Processing Systems (NeurIPS)*, 2020. 1, 2, 6
- [11] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *Computer graphics and interactive techniques*, 1996. 5
- [12] T. Davies, D. Nowrouzezahrai, and A. Jacobson. On the effectiveness of weight-encoded neural implicit 3d shapes. *arXiv preprint arXiv:2009.09808*, 2020. 1, 2, 5
- [13] A. Gropp, L. Yariv, N. Haim, M. Atzmon, and Y. Lipman. Implicit geometric regularization for learning shapes. In *Proceedings of Machine Learning and Systems*, pages 3569–3579. 2020. 2
- [14] B. Guillard, F. Stella, and P. Fua. Meshudf: Fast and differentiable meshing of unsigned distance field networks. In *European Conference on Computer Vision (ECCV)*, 2022. 1, 2, 4, 5, 6
- [15] M. Guo, J. Cai, Z. Liu, T. Mu, R. R. Martin, and S. Hu. Pct: Point cloud transformer. *Computational Visual Media*, 7(2):187–199, Apr. 2021. 1
- [16] J. C. Hart. Sphere tracing: A geometric method for the antialiased ray tracing of implicit surfaces. *Visual Computer*, 12(10):527–545, 1996. 5
- [17] T. Houchens, C. Lu, S. Duggal, R. Fu, and S. Sridhar. Neuralodf: Learning omnidirectional distance fields for 3d shape representation. 06 2022. 3
- [18] M. Kazhdan and H. Hoppe. Screened poisson surface reconstruction. *ACM Transactions on Graphics (TOG)*, 32(3):1–13, Jul. 2013. 2, 5
- [19] L. Kobbelt, M. Botsch, U. Schwanecke, and H. Seidel. Feature sensitive surface extraction from volume data. In *SIGGRAPH*, 2001. 3
- [20] Y. Liu, L. Wang, J. Yang, W. Chen, X. Meng, B. Yang, and L. Gao. Neudf: Leaning neural unsigned distance fields with volume rendering. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 2
- [21] X. Long, C. Lin, L. Liu, Y. Liu, P. Wang, C. Theobalt, T. Komura, and W. Wang. Neuraludf: Learning unsigned distance fields for multi-view reconstruction of surfaces with arbitrary topologies. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [22] X. Long, C. Lin, P. Wang, T. Komura, and W. Wang. Sparseneus: Fast generalizable neural surface reconstruction from sparse views. In *European Conference on Computer Vision (ECCV)*, 2022. 1
- [23] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. *Computer Graphics*, 21(4):163–169, 1987. 1, 2, 4
- [24] T. Müller, A. Evans, C. Schied, and A. Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (TOG)*, 41(4):102:1–102:15, Jul. 2022. 1, 2
- [25] D. Palmer, D. Smirnov, S. Wang, A. Chern, and J. Solomon. DeepCurrents: Learning implicit representations of shapes with boundaries. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [26] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2
- [27] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1
- [28] C. R. Qi, L. Yi, H. Su, and L. J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Neural Information Processing Systems (NeurIPS)*, 2017. 1
- [29] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, and G. Wetzstein. Implicit neural representations with periodic activation functions. In *Neural Information Processing Systems (NeurIPS)*, 2020. 1, 2, 6
- [30] T. Takikawa, J. Litalien, K. Yin, K. Kreis, C. Loop, D. Nowrouzezahrai, A. Jacobson, M. McGuire, and S. Fidler. Neural geometric level of detail: Real-time rendering with implicit 3d shapes. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2, 5, 6

- [31] M. Tancik, P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. Barron, and R. Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In *Neural Information Processing Systems (NeurIPS)*, 2020. 1, 2
- [32] J. Tang, W. Chen, J. Yang, B. Wang, S. Liu, B. Yang, and L. Gao. Octfield: Hierarchical implicit functions for 3d modeling. In *Neural Information Processing Systems (NeurIPS)*, 2021. 1, 2
- [33] I. Ueda, Y. Fukuhara, H. Kataoka, H. Aizawa, H. Shishido, and I. Kitahara. Neural density-distance fields. *arXiv preprint arXiv:2207.14455*, 2022. 5
- [34] R. Venkatesh, T. Karmali, S. Sharma, A. Ghosh, R. V. Babu, L. A. Jeni, and M. Singh. Deep implicit surface point prediction networks. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12633–12642, 2021. 2
- [35] L. Wang, J. Yang, W. Chen, X. Meng, B. Yang, J. Li, and L. Gao. Hsdf: Hybrid sign and distance field for modeling surfaces with arbitrary topologies. In *Neural Information Processing Systems (NeurIPS)*, 2022. 1, 2, 4, 6
- [36] P. Wang, L. Liu, Y. Liu, C. Theobalt, T. Komura, and W. Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In *Neural Information Processing Systems (NeurIPS)*, 2021. 1, 2
- [37] Y. Wang, I. Skorokhodov, and P. Wonka. HF-NeuS: Improved Surface Reconstruction Using High-Frequency Details. In *Neural Information Processing Systems (NeurIPS)*, 2022. 1
- [38] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon. Dynamic graph cnn for learning on point clouds. *ACM Transactions On Graphics (TOG)*, 38(5):1–12, Oct. 2019. 1
- [39] T. Wu, J. Wang, X. Pan, X. Xu, C. Theobalt, Z. Liu, and D. Lin. Voxurf: Voxel-based efficient and accurate neural surface reconstruction. In *International Conference on Learning Representations (ICLR)*, 2023. 1, 2
- [40] J. Ye, Y. Chen, N. Wang, and X. Wang. Gifs: Neural implicit function for general shape representation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2, 6
- [41] Z. Yu, S. Peng, M. Niemeyer, T. Sattler, and A. Geiger. MonoSDF: Exploring Monocular Geometric Cues for Neural Implicit Surface Reconstruction. In *Neural Information Processing Systems (NeurIPS)*, 2022. 1, 2
- [42] Q. Zhou and A. Jacobson. Thingi10k: A dataset of 10,000 3d-printing models. *arXiv preprint arXiv:1605.04797*, 2016. 5