

# GIFS: Neural Implicit Function for General Shape Representation

Jianglong Ye<sup>1</sup>

Yuntao Chen<sup>2</sup>  
<sup>1</sup>UC San Diego

Naiyan Wang<sup>2</sup>  
<sup>2</sup>TuSimple

Xiaolong Wang<sup>1</sup>

## Abstract

Recent development of neural implicit function has shown tremendous success on high-quality 3D shape reconstruction. However, most works divide the space into inside and outside of the shape, which limits their representing power to single-layer and watertight shapes. This limitation leads to tedious data processing (converting non-watertight raw data to watertight) as well as the incapability of representing general object shapes in the real world. In this work, we propose a novel method to represent general shapes including non-watertight shapes and shapes with multi-layer surfaces. We introduce *General Implicit Function for 3D Shape (GIFS)*, which models the relationships between every two points instead of the relationships between points and surfaces. Instead of dividing 3D space into predefined inside-outside regions, GIFS encodes whether two points are separated by any surface. Experiments on ShapeNet show that GIFS outperforms previous state-of-the-art methods in terms of reconstruction quality, rendering efficiency, and visual fidelity. Project page is available at <https://jianglongye.com/gifs>.

## 1. Introduction

3D shape reconstruction is a fundamental problem in computer vision, robotics, and graphics. The choice of shape representation heavily affects the reconstruction quality. In recent years, Neural Implicit Function (NIFs) have achieved high quality and high resolution reconstruction with a compact representation [9, 11, 31, 47, 48]. Most NIF approaches employ a neural network to classify whether a 3D point is inside or outside the surface [47] as shown in Figure 1(a) or compute the signed distance from a 3D point to one surface [48]. However, most shapes in the real world cannot be simply represented by one watertight surface.

Considering multi-layer shapes (e.g., cars with seats inside), binary classification of the 3D space is insufficient. To tackle this problem, one possible approach is to extend the original definition from binary classification to multi-class classification. We take the double-layer circle as an example and illustrate this idea in Figure 1(b). The lim-

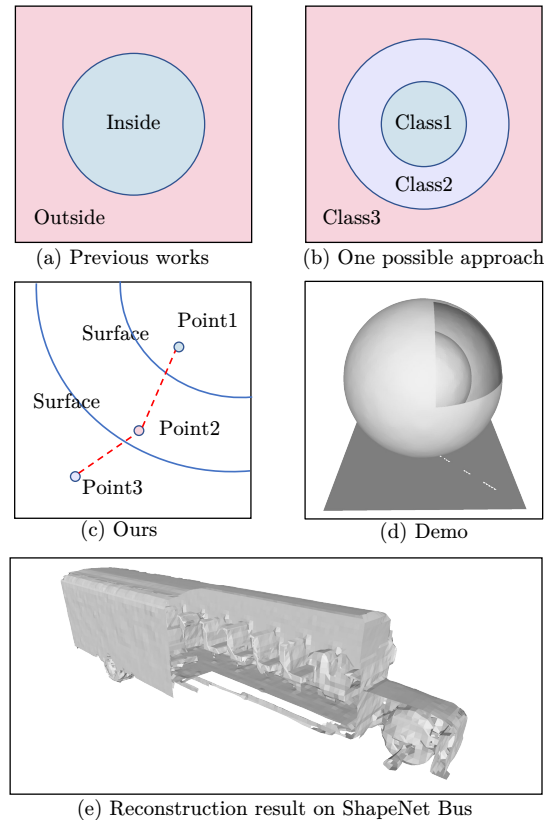


Figure 1. **Comparison of different methods.** Previous works divide the space into inside and outside like (a). One possible approach for multi-layer shape is to perform multi-class classification like (b). Our method no longer assigns predefined labels to points but classifies whether whether two points are separated by any surface, as shown in (c). In (d) and (e), we present a demo shape and a reconstructed bus from ShapeNet respectively.

itation of the extension is that shapes with different layers require different models. Therefore it can not represent general shapes in a unified approach. This method will also fail to represent non-watertight shapes (e.g., garments) since they cannot be partitioned into any number of regions.

In this paper, we introduce a *General Implicit Function for 3D Shape (GIFS)*. Along with this representation, we further design an algorithm to extract explicit surfaces from GIFS. Instead of dividing a 3D space into predefined cate-

gories like Figure 1(a) and (b), our representation focuses on the relationships between every two 3D points. Specifically, we define a binary flag to indicate whether two points are on the same side of object surfaces (see Sec. 3.1 for formal definition). As shown in Figure 1(c), three points are separated by different object surfaces, and our method no longer assigns predefined labels to points but compares between them. We embed 3D points into latent space and employ neural networks to approximate binary flags based on their embeddings. In this way, our method is still able to separate different regions, and the shape is implicitly represented by the decision boundary of the binary flags.

Compared to all previous NIFs that model the relationship between points and surfaces (occupancy [47], signed distance [48] and unsigned distance [13]), modeling the relationships between different points with GIFS allows representing general shapes including multi-layer and non-watertight shapes. To recover the shape from GIFS, we modify the Marching Cubes algorithm to build triangle faces from sampled binary flags. We show a demo shape represented by our method in Figure 1(d), it is a double-layered ball on top of a non-watertight plane. The larger ball (outside surface) is cut for better visualization.

Experiments on ShapeNet [10] demonstrate that GIFS achieves state-of-the-art performance, and can be trained directly on the raw data. Directly trainable on raw data not only avoids tedious data processing but also prevents the loss of accuracy during pre-processing. Compared with NDF [13], another generalized shape representation outputs point cloud, our GIFS is far more efficient and visually appealing due to the straightforward mesh generation. We show an example of the reconstruction results on ShapeNet in Figure 1(e), which is a bus with multiple seats inside.

We highlight our main contributions as follows:

- We introduce a novel generalized shape representation based on spatial relationship between points.
- We design a learning-based implementation GIFS and a surface extraction algorithm for it.
- Our method achieves state-of-the-art performance in 3D reconstruction for watertight and more general shapes and shows advantages in rendering efficiency and visual fidelity over existing approaches.

## 2. Related Work

### 2.1. Learning Explicit Shape Representations

Explicit shape representations for 3D shape learning can be roughly classified into three categories: point-based, mesh-based, and voxel-based methods.

**Point-based methods.** As direct outputs of many sensors (e.g., LiDAR, depth cameras), point clouds are a popular representation for 3D learning. PointNet based architectures [50, 51] utilize maxpool operator to keep permu-

tation invariance and are widely used to extract geometry features from point clouds. In addition, recent research for learning with point clouds includes kernel point convolution [20, 58, 64], graph-based architecture [42, 61], transformer [27, 29] and so on. Point clouds are also served as a output representation in 3D reconstruction [13, 19]. However, unlike other representations, point clouds do not describe topology and require non-trivial post-processing steps [5, 8, 36, 37] to generate renderable surfaces.

**Mesh-based methods.** Mesh consists of a set of vertices and edges, which can be defined as a graph. Hence, graph convolution can be directly applied on mesh for geometry learning [7, 26]. Mesh can also be considered as an output representation in shape reconstruction. Most methods [34, 44, 52, 60] deform a template and thereby are limited to a fixed topology. More recent methods predict vertices and faces directly [16, 23, 25], but do not guarantee surface continuity and are prone to generating self-intersecting faces. Mesh is also popular in human representations and is utilized to estimate human shape, pose [33, 39, 40], and garments [1, 6]. The underlying mesh model [45, 49] still restricts the topology and details.

**Voxel-based methods.** Voxels are an intuitive 3D extension of pixels in 2D images. The common learning paradigm (i.e., convolution) in the 2D domain could be extended to voxels naturally. Consequently, voxels are widely used for shape learning [30, 32, 35]. Occupancy grid is the simplest use case of voxel representation. However, due to the cubically growing memory footprint, the resolution of the grid is practically limited to a certain level [14, 43, 65]. Therefore, the voxel-based methods struggle to reconstruct high-fidelity shapes with fine details. Higher resolutions can be achieved at the cost of limited training batches and slow training [62, 63], or complex multi-resolution implementations [28, 57]. Replacing occupancy with Truncated Signed Distance functions [15] for learning [17, 41] allows for sub-voxel precision, although the resolution is still limited by the fixed grid.

### 2.2. Learning Implicit Shape Representations

In the past few years, a lot of advances have been made in 3D shape learning with neural implicit function [11, 21, 24, 47, 48, 53, 54, 66]. Despite the differences in output (occupancy [47, 53], signed distance [22, 48], or unsigned distance [13]), they all learn a continuous function to predict the relationship between the querying point (x-y-z) and surface. Pioneering methods [11, 47, 48] represent shape with a global latent code. An encoder [11, 47] or an optimization method [48] is adopted to embed the shape. And then a decoder is used to reconstruct shape from the latent code. For constructing a smooth latent space, some techniques such as curriculum learning [18], and adversarial training [38] are exploited. Following works [55, 56] also use periodic acti-

vation functions to map the input position to high dimensional space to preserve high-frequency details. Unlike the global methods, the local methods divide the 3D space into uniform grids [9, 12, 31] or local parts [21, 22], and the latent code varies by local grids/parts. Since each latent code only needs to represent the shape of a local area, better details can be preserved and generalizability is improved.

The methods mentioned above rely on inside and outside partition and thus always require tedious data processing to artificially close the shapes, which results in non-trivial loss of details and inner structures. Recent works [3, 4, 68] introduce sign agnostic learning to learn from the raw data. However, the output is again signed distance and hence can not reconstruct general shapes. To model general shapes, NDF [13] proposes to predict unsigned distance fields (UDFs) to represent general shapes. However, the output of NDF is point clouds, which does not only rely on an expensive post-processing step but also struggles to obtain high-quality final meshes. Following works [59, 67] incorporate normals and gradients to improve reconstruction accuracy and efficiency.

In contrast, we propose to predict whether two different points are separated by any surface, and it allows general shape representation. To the best of our knowledge, our method is the first to utilize a neural network to model the relationship between different points in shape representation. Since the different regions are separated in our representation, a Marching Cubes-like algorithm can be employed to directly convert our representation into a mesh rather than a point cloud. Due to this, we outperform NDF in terms of efficiency and visual effect.

### 3. Method

#### 3.1. Generalized Shape Representation

Instead of modeling the relationship between points and surfaces such as [13, 22, 47, 48, 53], we propose a novel shape representation focusing on the relationship between different points. We utilize a binary flag to indicate whether two points are separated by any surface. Specifically, the pair of points are considered to be separated if their line segment intersects any surface.

Let the surface of object be a subset  $S \subset \mathbb{R}^3$ . Given two points  $p_1 \in \mathbb{R}^3$  and  $p_2 \in \mathbb{R}^3$ , we can get a segment  $e$  formed by these two points:  $e(p_1, p_2) = \{p_1 + k * (p_2 - p_1) \mid k \in [0, 1]\}$ . If the segment intersects the surface, there is at least one point that belongs to both the segment and the surface. The binary flag  $b$  used to indicate whether the segment intersects the surface is defined as:

$$b(p_1, p_2, S) = \begin{cases} 1, & \exists p \in e(p_1, p_2), p \in S \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

Figure 2 (a) shows the example of the binary flags: the flag




Figure 2. **Generalized shape representation.** In (a), we show examples of the binary flag. The flag of the points separated by the surface is positive, otherwise, it is negative. In (b) and (c), we show how a double-layer circle and a non-watertight curve are represented. The positive flags are bolded red.

of Point 1 and Point 2 which are separated by any surface is positive (1), and the flag of Point 1 and Point 3 is negative (0).

The definition differs from the connected components in computer vision. Our binary flag is not transitive in 3D space, i.e.,  $b(p_1, p_2, S) = b(p_2, p_3, S)$  does not imply that  $b(p_1, p_2, S) = b(p_2, p_3, S) = b(p_1, p_3, S)$ . This is the reason that our definition allows to represent non-watertight shapes.

Similar to the previous implicit function, the shape is still represented by the decision boundary. Different regions are separated, and the binary flag changes when crossing the surface. Figure 2 (b) and (c) show how a double-layer circle and a non-watertight curve are represented. The flags of the points crossing the circle and the curve are positive (bolded red in the figure), other flags are negative (gray in the figure). Note that the input points of the binary flag are not limited to horizontal or vertical points, and the resolution will be much higher when extracting surfaces.

#### 3.2. Learning General Implicit Function for 3D Shape

We propose to learn a *General Implicit Function for 3D Shape* (GIFS) to estimate the binary flag  $b$ . It takes two 3D points  $p_1$  and  $p_2$  plus an observation  $X$  as inputs. In this paper, we study the task of shape reconstruction from sparse point cloud  $X \in \mathbb{R}^{n \times 3}$ . For completeness, we first explain the encoding module, which follows IF-Net [12]. Then we describe our decoding module and learning procedure. The overview of our method is shown in Figure 3.

**Point embedding.** The input point cloud is first converted to a discrete voxel grid. Then a 3D CNN is applied to get the multi-scale grid features  $F_1, \dots, F_n, F_k \in \mathcal{F}_k^{K \times K \times K}$ , where  $K$  is the grid size varies with scale and  $\mathcal{F}_k \in \mathbb{R}^C$  is a deep feature with channels  $C$ . Refer to IF-Net [12] for further details. Given a query point  $p$ , a set of deep features is sampled from the multi-scale encoding and is concatenated as the embedding  $z \in \mathcal{F}_1 \times \dots \times \mathcal{F}_n$  of the

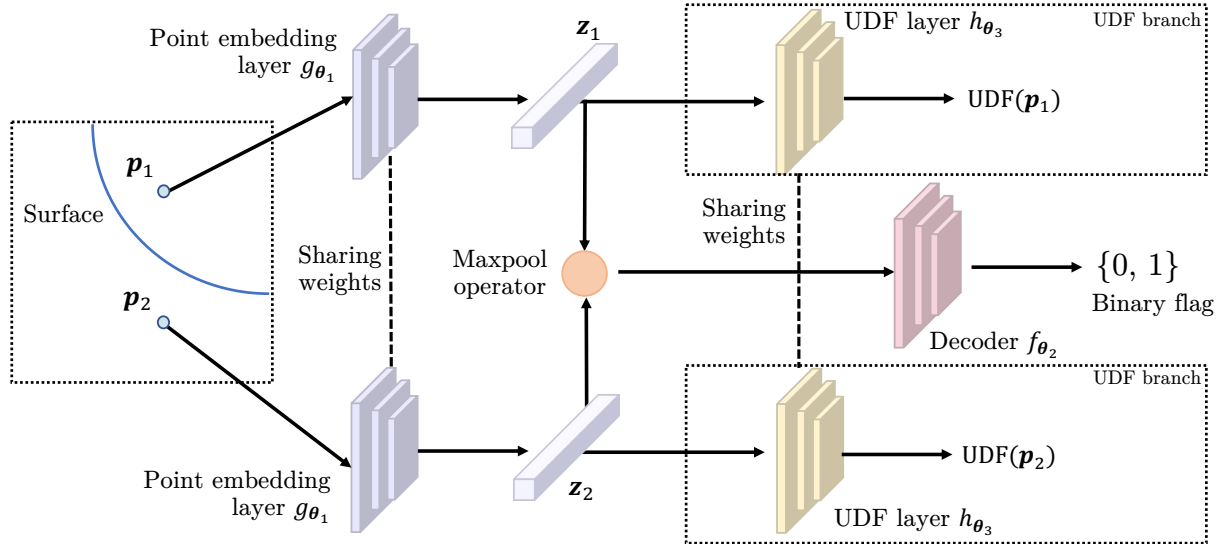


Figure 3. **Overview of GIFS.** We learn a neural implicit function to classify whether two points are on the same side of object surfaces. First, two 3D points are passed to the embedding layer to get corresponding features. Then a maxpool operator is utilized to keep the permutation invariance. The decoder takes the fused feature as input and approximates the binary flag between points. An extra UDF branch can be used to enhance the spatial perception of the feature.

point. We denote the point embedding layer as:

$$z = g_{\theta_1}(\mathbf{X}, \mathbf{p}): \mathbb{R}^n \times \mathbb{R}^3 \mapsto \mathcal{F}_1 \times \dots \times \mathcal{F}_n, \quad (2)$$

where  $\theta_1$  is the learnable parameters for embedding layer. The purple block in Figure 3 indicates the embedding layer. **Binary flag prediction.** A decoder is implemented to predict the binary flag of points. To keep the permutation invariance, a maxpool operator is applied to the points embeddings  $z_1$  and  $z_2$ :  $z_f = \text{MAX}(z_1, z_2)$ , where  $z_f$  is the fused feature of the two points. Then the fused feature  $z_f$  is passed to the decoder to predict the binary flag  $\mathbf{b}$ . The decoder  $f$  is a multilayer perceptron parameterized by  $\theta_2$ :

$$f_{\theta_2}(z_f) \approx \mathbf{b}: \mathcal{F}_1 \times \dots \times \mathcal{F}_n \mapsto [0, 1], \quad (3)$$

The pink block in Figure 3 indicates the decoder.

**UDF branch.** Although the above decoder is sufficient for shape reconstruction, we additionally incorporate a UDF branch to enhance the spatial perception of the point feature and accelerate the surface extraction (see Sec. 3.3). Unsigned distance function (UDF) is defined as the distance between the querying point and its closet point on the object surface. Therefore it does not introduce the traditional inside/outside partition. NDF [13] is the first to use neural implicit function to approximate UDF. Given a embedding  $z$  of point  $\mathbf{p}$ , the UDF layer  $h$  is denoted as:

$$h_{\theta_3}(z) \approx \text{UDF}(\mathbf{p}): \mathcal{F}_1 \times \dots \times \mathcal{F}_n \mapsto \mathbb{R}_0^+, \quad (4)$$

where  $\theta_3$  is the learnable parameters for the UDF layer and  $\text{UDF}(\mathbf{p})$  is the ground truth UDF for point  $\mathbf{p}$ . The yellow

block in Figure 3 indicates the UDF layer. Ablation studies on the UDF branch are shown in Sec. 4.4.

**Learning procedure.** To train the point embedding layer  $g_{\theta_1}(\cdot)$ , the decoder  $f_{\theta_2}(\cdot)$  and the UDF layer  $h_{\theta_3}(\cdot)$ , pairs  $\{\mathbf{X}_i, \mathcal{S}_i | i \in 1, \dots, N\}$  of observation  $\mathbf{X}_i$  and corresponding ground truth shape  $\mathcal{S}_i$  are required, where  $i$  is the index of observation and  $N$  is the number of training examples. For each training example, we sample a pair  $\{\mathbf{p}_1, \mathbf{p}_2\}$  of two points and compute their ground truth  $\text{UDF}(\mathbf{p}, \mathcal{S})$  as well as binary flags  $\mathbf{b}(\mathbf{p}_1, \mathbf{p}_2, \mathcal{S})$ . During training, two loss terms are exploited:  $\mathcal{L}_b$  for binary flag prediction and  $\mathcal{L}_u$  for UDF regression. The  $\mathcal{L}_b$  is defined as

$$\mathcal{L}_b(\mathbf{X}, \mathbf{p}_1, \mathbf{p}_2) = |f_{\theta_2}(\text{MAX}(g_{\theta_1}(\mathbf{X}, \mathbf{p}_1), g_{\theta_1}(\mathbf{X}, \mathbf{p}_2))) - \mathbf{b}(\mathbf{p}_1, \mathbf{p}_2, \mathcal{S}_\mathbf{X})|, \quad (5)$$

where  $\mathcal{S}_\mathbf{X}$  is the corresponding ground truth shape for observation  $\mathbf{X}$ . And the  $\mathcal{L}_u$  is defined as:

$$\mathcal{L}_u(\mathbf{X}, \mathbf{p}) = |\min(h_{\theta_3}(g_{\theta_1}(\mathbf{X}, \mathbf{p})), \delta) - \min(\text{UDF}(\mathbf{p}, \mathcal{S}_\mathbf{X}), \delta)|, \quad (6)$$

where  $\delta > 0$  is a threshold which concentrates the model capacity to represent the near space of the surface.

The parameters are optimized by minimizing following mini-batch loss:

$$\mathcal{L}_B(\theta_1, \theta_2, \theta_3) = \sum_{\mathbf{X} \in \mathcal{B}} \sum_{\mathbf{p}_1, \mathbf{p}_2 \in \mathcal{P}} \mathcal{L}_b(\mathbf{X}, \mathbf{p}_1, \mathbf{p}_2) + \lambda(\mathcal{L}_u(\mathbf{X}, \mathbf{p}_1) + \mathcal{L}_u(\mathbf{X}, \mathbf{p}_2)), \quad (7)$$

where  $\mathcal{B}$  is a mini-batch of input point cloud observations,  $\mathcal{P}$  is a subset of sampled pairs for training and  $\lambda$  is the weight for  $\mathcal{L}_u$ .

### 3.3. Surface Extraction from GIFS

Since the proposed method is different from previous NIFs, it’s non-trivial to design an algorithm for surface extraction. We adapt Marching Cubes [46] algorithm to extract explicit surface (mesh) from our implicit representation. Marching Cubes divides the space into the 3D grid and locates cubes that intersect the object surface. For the intersected cube, triangles are created based on the occupancies of 8 vertices. A lookup table consisting of  $2^8 = 256$  different assignments is used to increase the speed.

Although there is no occupancy in our method, we make an assumption that on a micro-scale, the space near the surface can always be divided into two classes. Based on that, we divide the space into the 3D grid and binary classify the 8 vertices in each cell according to the binary flags between them. Then we utilize the lookup table in Marching Cubes to create triangles. Let  $c_i \in \{0, 1\}$  be the assignment for vertex  $i \in \{0, 1 \dots 7\}$  and  $b_{ij} \in [0, 1]$  be the flag between vertex  $i$  and vertex  $j$ . The cost function of the assignment problem is:

$$\mathcal{L} = \sum_i \sum_{j \neq i} |c_i - c_j| |1 - b_{ij}|. \quad (8)$$

The optimal assignment can be found by minimizing the cost function. We simply apply a brute-force algorithm in implementation since the speed is still acceptable.

**Coarse-to-fine surface extraction.** Similar to OccNet [47], a coarse-to-fine paradigm is adopted to speed up the extraction process. We start from a low-resolution grid and evaluate the UDF at the center of each cube. If the UDF value is smaller than the cube size multiples some threshold  $\tau$ , we assume the cube intersects the surface and subdivide it into 8 subcubes. The evaluation is repeated on the new cubes until the desired final resolution is reached. We perform our adapted Marching Cubes on final cubes.

**Mesh refinement.** The initial mesh extracted by adapted Marching Cubes is only an approximation of the decision boundary and can be further refined with the UDF value. We randomly sample  $N$  points  $\mathbf{p}_i$  from each face of the output mesh and refine it by minimizing the loss:

$$\mathcal{L} = \sum_i^N |h_{\theta_3}(g_{\theta_1}(\mathbf{X}, \mathbf{p}_i))|. \quad (9)$$

Basically, we further push the mesh surface to the decision boundary by minimizing the UDF value of the surface. We report an ablation study for the mesh refinement in Sec. 4.4.

## 4. Experiments

We focus on the task of 3D shape reconstruction from sparse point clouds to validate GIFS. We first show that GIFS can reconstruct watertight shapes on par with the state-of-the-art methods in Sec. 4.2, and then show GIFS can reconstruct general shapes including multi-layer, non-watertight shapes in Sec. 4.3. Ablation studies on architecture, grid size and sampling strategy are shown in Sec. 4.4.

### 4.1. Experimental Settings

**Datasets.** In our experiments, we follow the setting of NDF and choose the “Car” category of the ShapeNet [10] dataset, which consists of 7498 synthetic objects. The “Car” category has the largest amount of multi-layer shapes as well as non-watertight shapes. For watertight shapes, we use the common training/test split from 3D-R2N2 [14] and the processed watertight meshes from DISN [66]. For general shapes, we use the training/test split from NDF [13]. During training and evaluation, all meshes are normalized and centered to a unit cube. Besides, we employ the MGN [6] dataset, which consists of 307 garments, to show the representation power of our method.

**Evaluation Metrics.** To measure the reconstruction quality, we adopt established metrics defined in LDIF [21]: the Chamfer distance and F-Score. The Chamfer distance measures the average errors of all points and is sensitive to outliers. F-Score measures the percentage of good predictions. For evaluation, we randomly sample 100k points from the surface to calculate all metrics. The thresholds for F-Score are 0.01 (same as LDIF) and 0.005 (for better comparison). Lower is better for Chamfer distance; higher is better for F-Score.

**Baselines.** For the experiment of watertight shapes, we compare our method to Point Set Generation Networks [19] (PSGN), Deep Marching Cubes [43] (DMC), Occupancy Network [47] (OccNet), and IF-Net [12], the state-of-the-art in shape reconstruction. For the experiment of general shapes, we compare our method to NDF [13]. For methods that the authors do not provide a pretrained model, we retrain it until the minimum validation error is reached.

**Implementation details.** The resolution of the voxel grid in the point embedding layer is  $128^3$  for 3000 points input and  $256^3$  for 10000 points input. The decoder for binary flag prediction is a 5-layer multilayer perceptron. All internal layers are 256-dimensional and have ReLU non-linearities. The UDF layer is exactly the same as the decoder in NDF. The  $\delta$  in Eq. 6 is set to 0.1, the  $\lambda$  in Eq. 7 is set to 10. During training, we employ the Adam optimizer with a learning rate of  $1 \times 10^{-4}$ . During inference, the initial resolution of the grid is  $20^3$  and is subdivided 3 times. The final resolution is  $160^3$ . The  $\tau$  in Sec. 3.3 is set to 2.

During data generation, we mainly follow the same strategy as NDF [13]. Specifically, we first sample points

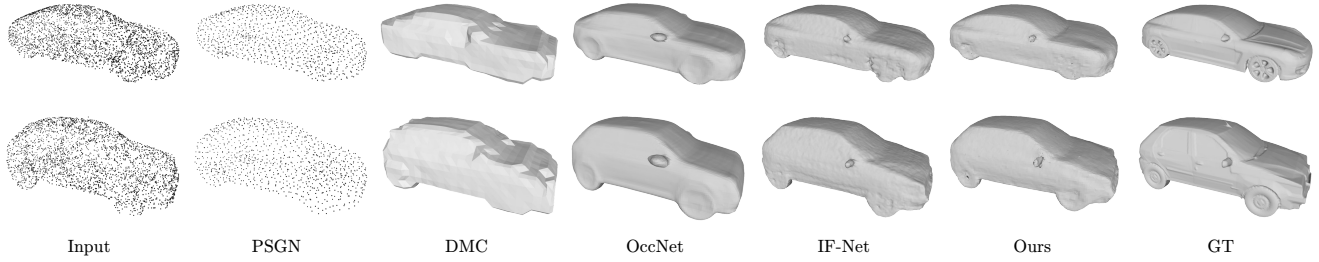


Figure 4. **Comparison of methods trained on watertight shapes.** Our method can reconstruct watertight shapes with the same accuracy as the state-of-the-art method.




Figure 5. **Reconstruction results of different methods on general shapes.** Our method can reconstruct smooth, continuous surfaces, and achieve a better visual effect than previous methods.

$p_i^S \in \mathbb{R}^3$  on the ground truth surfaces and add two random displacements  $\{n_i^0, n_i^1\} \sim \mathcal{N}(0, \Sigma)$  to produce pairs of points  $\{p_i^0, p_i^1\}$ , i.e.,  $p_i^0 = p_i^S + n_i^0, p_i^1 = p_i^S + n_i^1$ , where  $\Sigma \in \mathbb{R}^{3 \times 3}$  is a diagonal covariance matrix with entries  $\Sigma_{i,i} = \sigma$ . We adopt 3 different  $\sigma$  of 0.005, 0.01, 0.03. We also randomly sample 10% of the data in the 3D grid. We report performances with different combinations of  $\sigma$  in Sec. 4.4. We employ the intersection algorithm from CGAL [2] to calculate the ground truth binary flags.

## 4.2. Shape Reconstruction of Watertight Shapes

We first show the representation power of our method on watertight shapes. We train and evaluate with the processed data provided by DISN. Their data processing step closes shape and removes all interior structures. In this experiment, the input for all methods is 3000 points sampled from the watertight data of the ShapeNet “Car” category. For a fair comparison, the resolution of the encoding voxel grid in both IF-Net and our method is  $128^3$ .

Figure 4 shows that our method can reconstruct a watertight shape with the same accuracy as the state-of-the-art method. Note that OccNet uses a global code to represent

Method	Chamfer distance ↓		F-Score ↑	
	Mean	Median	F1 <sup>0.005</sup>	F1 <sup>0.01</sup>
Input	0.782	0.754	26.15	65.00
PSGN [19]	2.251	2.018	4.33	24.34
DMC [43]	5.963	3.654	32.22	58.06
OccNet [47]	3.251	2.386	28.75	65.41
IF-Net [12]	0.260	0.125	86.70	96.50
NDF [13]	0.152	0.125	85.70	<b>98.32</b>
Ours	<b>0.146</b>	<b>0.114</b>	<b>88.75</b>	98.09

Table 1. **Quantitative evaluation on watertight shapes.** We train and evaluate our method on the watertight data of the ShapeNet “Car” category. We report the mean and median value of the Chamfer distance, as well as F-Score at different thresholds. The Chamfer distance results  $\times 10^{-4}$ . Our method outperforms baselines in all but one metric.

the whole shape and outputs a smoother result which can be wrong. As shown in the figure, the rear-view car mirrors of the OccNet results differ significantly from the GT shape. We also quantitatively compare our method to baselines in Table 1 and show that GIFS outperforms baselines in all but one metric.




Figure 6. **Zoom-in comparison of different methods on general shapes.** The surface in the reconstruction result of NDF is far from smooth and contains holes.




Figure 7. **Reconstruction results of multi-layer shapes on other classes.** Our method can reconstruct internal structures of various shapes.

### 4.3. Shape Reconstruction of General Shapes

We show the reconstruction results of our method on general shapes including multi-layer, and non-watertight shapes. Since our method does not require a pre-processing step (e.g., deep fusion), it can directly train and evaluate with the raw ShapeNet data. We also do not lose accuracy due to this step. In this experiment, we mainly compare our method to the NDF, the previous method aimed at modeling general shapes. For a fair comparison, the resolution of the encoding voxel grid in both NDF and our method is  $256^3$ . The input is 10000 points sampled from the raw data of the ShapeNet “Car” category.

**Qualitative analysis.** Figure 5 shows reconstruction results of two multi-layer cars from the test set. Traditional implicit neural representations are limited to single-layer, watertight shapes, so all the internal structure of the car is lost in the IF-Net results. NDF can represent general shapes, but its network output is a point cloud and relies on the ball-pivoting algorithm to obtain mesh. Even with carefully selected thresholds, the surface in the reconstruction result of NDF is still far from smooth and has a number of holes. In contrast, our method is able to reconstruct a flat, continuous surface, and thereby achieve a better visual effect. This can be observed more clearly in the zoom-in comparison in Figure 6. Note that the NDF visualization code also includes a filter to close holes, while our method does not apply any filter.

Figure 7 shows the reconstruction results of multi-layer shapes on other classes. Our method can reconstruct the chairs in the bus and the airplane or the bowl in the mi-




Figure 8. **Reconstruction results of non-watertight shapes on other classes.** The non-watertight shapes are difficult for traditional implicit neural representations to reconstruct.

Method	Chamfer distance ↓		F-Score ↑	
	Mean	Median	F1 <sup>0.005</sup>	F1 <sup>0.01</sup>
Input	0.363	0.355	48.50	88.34
Watertight GT	2.628	2.293	68.82	81.60
NDF [13]	0.126	0.120	88.09	99.54
Ours	0.128	0.123	88.05	99.31

Table 2. **Quantitative evaluation on general shapes.** We train and evaluate our method on the raw data of the ShapeNet “Car” category. Our method achieves comparable performance with the state-of-the-art method.

crowave. Figure 8 shows the reconstruction results of non-watertight shapes. Our method can represent cans with a hole and bag, which are difficult for the traditional implicit function to represent.

We also apply our model to garment reconstruction. Without retraining nor fine-tuning the model, we directly inference on MGN [6] dataset. The reconstruction results are shown in Figure 9. From the top view, we can see the garments are non-watertight as well.

**Quantitative evaluation.** As shown in Table 2, we also quantitatively evaluate our method and find that GIFS achieves comparable performance with the state-of-the-art method NDF. Here, the evaluation of NDF is performed on point clouds without topologies rather than meshes. We also sample points from the watertight ground truth shapes and report its metric (watertight GT in table). We consider this as the upper bound of the traditional implicit function.

**Speed comparison.** The limitation of NDF is not only the lack of continuity on the surface but also the extremely slow running speed. This is due to that the ball-pivoting [5] algorithm used by NDF requires a large number of nearest-neighbor searches, normal calculation, and surface fitting. In the experiment, we use the default  $1 \times 10^6$  points and the parameters provided by NDF.<sup>1</sup> The platform for NDF

<sup>1</sup>Using downsampled points or different parameters will reduce the runtime, but it is also more likely to create non-smooth surfaces and holes.




Figure 9. **Reconstruction results on the MGN [6] dataset.** Our method allows the reconstruction of the non-watertight garments.

Method	Runtime ↓	Space ↓
NDF [13]	2593 s	40.57 MB
Ours	<b>53 s</b>	<b>2.43 MB</b>

Table 3. **Algorithm efficiency analysis.** Our method has a tremendous improvement in temporal and spatial efficiency over the previous method without losing accuracy.

mesh conversion is a PC with a 6-Core 2.6 GHz CPU and the platform for the network inference is the RTX 2080 Ti. We report the runtime of NDF and our method in Table 3. Note that the runtime here consists of both network inference and mesh generation. With the unoptimized code, our method has a tremendous advantage over the NDF.

**Space comparison.** Similarly, a large number of points or the complex mesh generated by the ball-pivoting algorithm is also spatially inefficient. We statistics the average size of the files generated by both methods and report the result in Table 3.

#### 4.4. Ablation Study

**Architecture.** We mainly ablate the UDF branch and the mesh refinement and report the results in Table 4. The experiment is conducted on watertight ShapeNet data, and the resolution of the encoding voxel grid is  $128^3$ . We observe that both modules play an important role in our framework. The unary UDF loss is complementary to our pairwise GIFS loss. The mesh refinement enables the generated mesh to be closer to the decision boundary.

**Surface extraction.** We evaluate the performance of different 3D grid sizes in surface extraction. Table 5 shows how the accuracy and speed change with the different resolution of the 3D grid. We achieve better reconstruction with a higher resolution of the 3D grid. Thanks to the coarse-to-fine paradigm, the memory and computation cost do not have an exponential growth but instead only increase for a reasonable amount.

**Sampling strategy.** We compare the performance of different  $\sigma$  during data generation (see Sec 4.1) and report the results in Table 6. Our method is robust to different sampling strategies.

Method	Chamfer distance ↓		F-Score ↑	
	Mean	Median	F1 <sup>0.005</sup>	F1 <sup>0.01</sup>
w/o UDF	0.227	0.206	73.38	97.42
w/o Refine	0.152	0.121	87.98	97.93
Full model	<b>0.146</b>	<b>0.114</b>	<b>88.75</b>	<b>98.09</b>

Table 4. **Ablation study.** Both UDF branch and mesh refinement boost the performance.

Res	CD ↓		F-score ↑		Efficiency ↓	
	Mean	Median	F1 <sup>0.005</sup>	F1 <sup>0.01</sup>	Time	Space
$80^3$	0.445	0.231	70.30	97.66	12 s	0.43 MB
$160^3$	0.128	0.123	88.05	99.31	53 s	2.43 MB
$240^3$	0.125	0.121	88.83	99.36	121 s	7.00 MB
$320^3$	0.123	0.119	89.10	99.36	182 s	11.51 MB

Table 5. **Surface extraction.** We compare the performance of different 3D grid sizes in surface extraction. Our method achieves better reconstruction with a higher resolution of the 3D grid.

$\sigma$	Chamfer distance ↓		F-score ↑	
	Mean	Median	F1 <sup>0.005</sup>	F1 <sup>0.01</sup>
0.005, 0.01, 0.03	0.128	0.123	88.05	99.31
0.003, 0.01, 0.03	0.131	0.126	87.22	99.37
0.01, 0.03	0.137	0.131	87.36	98.90

Table 6. **Sampling strategy.** We compare the performance of different  $\sigma$  during data generation (see Sec 4.1). Our method is robust to different sampling strategies.

## 5. Discussion

**Conclusion.** In this paper, we introduce a generalized shape representation and its corresponding neural network GIFS. GIFS allows the reconstruction of high-quality general object shapes including watertight, non-watertight, and multi-layer shapes. We also devise an algorithm to directly extract mesh from GIFS. Experiments demonstrate that GIFS not only achieves state-of-the-art performance in both watertight and general shapes but also shows advantages in terms of efficiency and visual effect. We believe that GIFS takes a step towards general shape representation.

**Limitations and future work.** We describe two limitations: (i) In our surface extraction algorithm, an assumption exists that at sufficiently small scales the surface can be divided into two classes. For general shapes, however, the assumption does not guarantee to hold when the resolution of the grid is low. (ii) Our running speed, limited by surface extraction based on brute-force strategy and simple loops, is still far from real-time. We also point out two possible research directions. (i) Improving the smoothness of the generated surface using the first and second-order gradients of the implicit field. (ii) Accelerating network inference and surface extraction using parallel computing.

**Acknowledgements.** This work was supported, in part, by gifts from TuSimple.



## References

- [1] Thiemo Alldieck, Gerard Pons-Moll, Christian Theobalt, and Marcus Magnor. Tex2shape: Detailed full human body geometry from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2293–2303, 2019. 2
- [2] Pierre Alliez, Stéphane Tayeb, and Camille Wormser. 3D fast intersection and distance computation. In *CGAL User and Reference Manual*. CGAL Editorial Board, 5.3 edition, 2021. 6
- [3] 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. 3
- [4] Matan Atzmon and Yaron Lipman. Sald: Sign agnostic learning with derivatives. *arXiv preprint arXiv:2006.05400*, 2020. 3
- [5] Fausto Bernardini, Joshua Mittleman, Holly Rushmeier, Claudio Silva, and Gabriel Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE transactions on visualization and computer graphics*, 5(4):349–359, 1999. 2, 7
- [6] Bharat Lal Bhatnagar, Garvita Tiwari, Christian Theobalt, and Gerard Pons-Moll. Multi-garment net: Learning to dress 3d people from images. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5420–5430, 2019. 2, 5, 7, 8
- [7] Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017. 2
- [8] Fatih Calakli and Gabriel Taubin. Ssd: Smooth signed distance surface reconstruction. In *Computer Graphics Forum*, volume 30, pages 1993–2002. Wiley Online Library, 2011. 2
- [9] 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, 2020. 1, 3
- [10] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 2, 5
- [11] 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. 1, 2
- [12] Julian Chibane, Thiemo Alldieck, and Gerard Pons-Moll. Implicit functions in feature space for 3d shape reconstruction and completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6970–6981, 2020. 3, 5, 6
- [13] Julian Chibane, Aymen Mir, and Gerard Pons-Moll. Neural unsigned distance fields for implicit function learning. In *Advances in Neural Information Processing Systems*, 2020. 2, 3, 4, 5, 6, 7, 8
- [14] Christopher B Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *European conference on computer vision*, pages 628–644, 2016. 2, 5
- [15] Brian Curless and Marc Levoy. A volumetric method for building complex models from range images. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 303–312, 1996. 2
- [16] Angela Dai and Matthias Nießner. Scan2mesh: From unstructured range scans to 3d meshes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5574–5583, 2019. 2
- [17] Angela Dai, Charles Ruizhongtai Qi, and Matthias Nießner. Shape completion using 3d-encoder-predictor cnns and shape synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5868–5877, 2017. 2
- [18] Yueqi Duan, Haidong Zhu, He Wang, Li Yi, Ram Nevatia, and Leonidas J Guibas. Curriculum deepsdf. In *European Conference on Computer Vision*, pages 51–67, 2020. 2
- [19] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3d object reconstruction from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 605–613, 2017. 2, 5, 6
- [20] Hehe Fan, Xin Yu, Yuhang Ding, Yi Yang, and Mohan Kankanhalli. Pstnet: Point spatio-temporal convolution on point cloud sequences. In *International conference on learning representations*, 2020. 2
- [21] Kyle Genova, Forrester Cole, Avneesh Sud, Aaron Sarna, and Thomas A. Funkhouser. Local deep implicit functions for 3d shape. In *CVPR*, pages 4856–4865, 2020. 2, 3, 5
- [22] Kyle Genova, Forrester Cole, Daniel Vlasic, Aaron Sarna, William T Freeman, and Thomas Funkhouser. Learning shape templates with structured implicit functions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7154–7164, 2019. 2, 3
- [23] Georgia Gkioxari, Jitendra Malik, and Justin Johnson. Mesh r-cnn. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9785–9795, 2019. 2
- [24] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. *arXiv preprint arXiv:2002.10099*, 2020. 2
- [25] Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. A papier-mâché approach to learning 3d surface generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 216–224, 2018. 2
- [26] Kan Guo, Dongqing Zou, and Xiaowu Chen. 3d mesh labeling via deep convolutional neural networks. *ACM Transactions on Graphics (TOG)*, 35(1):1–12, 2015. 2
- [27] Meng-Hao Guo, Jun-Xiong Cai, Zheng-Ning Liu, Tai-Jiang Mu, Ralph R Martin, and Shi-Min Hu. Pct: Point cloud transformer. *Computational Visual Media*, 7(2):187–199, 2021. 2

- [28] Christian Häne, Shubham Tulsiani, and Jitendra Malik. Hierarchical surface prediction for 3d object reconstruction. In *2017 International Conference on 3D Vision (3DV)*, pages 412–420. IEEE, 2017. [2](#)
- [29] Le Hui, Hang Yang, Mingmei Cheng, Jin Xie, and Jian Yang. Pyramid point cloud transformer for large-scale place recognition. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6078–6087, 2021. [2](#)
- [30] Mengqi Ji, Juergen Gall, Haitian Zheng, Yebin Liu, and Lu Fang. SurfacerNet: An end-to-end 3d neural network for multi-view stereopsis. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2307–2315, 2017. [2](#)
- [31] 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. [1](#), [3](#)
- [32] Danilo Jimenez Rezende, SM Eslami, Shakir Mohamed, Peter Battaglia, Max Jaderberg, and Nicolas Heess. Unsupervised learning of 3d structure from images. *Advances in Neural Information Processing Systems*, 29:4996–5004, 2016. [2](#)
- [33] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7122–7131, 2018. [2](#)
- [34] Angjoo Kanazawa, Shubham Tulsiani, Alexei A Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 371–386, 2018. [2](#)
- [35] Abhishek Kar, Christian Häne, and Jitendra Malik. Learning a multi-view stereo machine. In *Advances in Neural Information Processing Systems*, pages 365–376, 2017. [2](#)
- [36] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, volume 7, 2006. [2](#)
- [37] Michael Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. *ACM Transactions on Graphics (ToG)*, 32(3):1–13, 2013. [2](#)
- [38] Marian Kleineberg, Matthias Fey, and Frank Weichert. Adversarial generation of continuous implicit shape representations. *arXiv preprint arXiv:2002.00349*, 2020. [2](#)
- [39] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2252–2261, 2019. [2](#)
- [40] Nikos Kolotouros, Georgios Pavlakos, and Kostas Daniilidis. Convolutional mesh regression for single-image human shape reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4501–4510, 2019. [2](#)
- [41] Lubor Ladicky, Olivier Saurer, SoHyeon Jeong, Fabio Maninchedda, and Marc Pollefeys. From point clouds to mesh using regression. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3893–3902, 2017. [2](#)
- [42] Loic Landrieu and Martin Simonovsky. Large-scale point cloud semantic segmentation with superpoint graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4558–4567, 2018. [2](#)
- [43] Yiyi Liao, Simon Donne, and Andreas Geiger. Deep marching cubes: Learning explicit surface representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2916–2925, 2018. [2](#), [5](#), [6](#)
- [44] Chen-Hsuan Lin, Oliver Wang, Bryan C Russell, Eli Shechtman, Vladimir G Kim, Matthew Fisher, and Simon Lucey. Photometric mesh optimization for video-aligned 3d object reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 969–978, 2019. [2](#)
- [45] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6):1–16, 2015. [2](#)
- [46] 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. [5](#)
- [47] 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](#), [3](#), [5](#), [6](#)
- [48] 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](#), [3](#)
- [49] Gerard Pons-Moll, Sergi Pujades, Sonny Hu, and Michael J Black. Clothcap: Seamless 4d clothing capture and retargeting. *ACM Transactions on Graphics (ToG)*, 36(4):1–15, 2017. [2](#)
- [50] 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](#)
- [51] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J. Guibas. PointNet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in Neural Information Processing Systems*, pages 5099–5108, 2017. [2](#)
- [52] Anurag Ranjan, Timo Bolkart, Soubhik Sanyal, and Michael J Black. Generating 3d faces using convolutional mesh autoencoders. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 704–720, 2018. [2](#)
- [53] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2304–2314, 2019. [2](#), [3](#)
- [54] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for

- high-resolution 3d human digitization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 84–93, 2020. 2
- [55] 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, 2020. 2
- [56] Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In *Advances in Neural Information Processing Systems*, 2020. 2
- [57] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Octree generating networks: Efficient convolutional architectures for high-resolution 3d outputs. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2088–2096, 2017. 2
- [58] Hugues Thomas, Charles R Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J Guibas. Kpconv: Flexible and deformable convolution for point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6411–6420, 2019. 2
- [59] 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. 3
- [60] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 52–67, 2018. 2
- [61] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *Acm Transactions On Graphics (tog)*, 38(5):1–12, 2019. 2
- [62] Jiajun Wu, Yifan Wang, Tianfan Xue, Xingyuan Sun, Bill Freeman, and Josh Tenenbaum. Marrnet: 3d shape reconstruction via 2.5d sketches. In *Advances in Neural Information Processing Systems*, pages 540–550, 2017. 2
- [63] Jiajun Wu, Chengkai Zhang, Tianfan Xue, William T Freeman, and Joshua B Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In *Advances in Neural Information Processing Systems*, pages 82–90, 2016. 2
- [64] Wenxuan Wu, Zhongang Qi, and Li Fuxin. Pointconv: Deep convolutional networks on 3d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9621–9630, 2019. 2
- [65] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015. 2
- [66] Qiangeng Xu, Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. *arXiv preprint arXiv:1905.10711*, 2019. 2, 5
- [67] 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. 3
- [68] Wenbin Zhao, Jiabao Lei, Yuxin Wen, Jianguo Zhang, and Kui Jia. Sign-agnostic implicit learning of surface self-similarities for shape modeling and reconstruction from raw point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10256–10265, 2021. 3