

# Ladybird: Quasi-Monte Carlo Sampling for Deep Implicit Field Based 3D Reconstruction with Symmetry

Yifan Xu<sup>\*1</sup>, Tianqi Fan<sup>\*2,3</sup>, Yi Yuan<sup>†1</sup>, and Gurprit Singh<sup>3</sup>

<sup>1</sup> Netease Fuxi AI Lab, Hangzhou, China

{xuyifan, yuanyi}@corp.netease.com

<sup>2</sup> Saarland Informatics Campus, Saarbruecken, Germany

<sup>3</sup> MPI for Informatics, Saarbruecken, Germany

{tfan,gsingh}@mpi-inf.mpg.de

**Abstract.** Deep implicit field regression methods are effective for 3D reconstruction from single-view images. However, the impact of different sampling patterns on the reconstruction quality is not well-understood. In this work, we first study the effect of *point set discrepancy* on the network training. Based on Farthest Point Sampling algorithm, we propose a sampling scheme that theoretically encourages better generalization performance, and results in fast convergence for SGD-based optimization algorithms. Secondly, based on the reflective symmetry of an object, we propose a feature fusion method that alleviates issues due to self-occlusions which makes it difficult to utilize local image features. Our proposed system *Ladybird* is able to create high quality 3D object reconstructions from a single input image. We evaluate Ladybird on a large scale 3D dataset (ShapeNet) demonstrating highly competitive results in terms of Chamfer distance, Earth Mover’s distance and Intersection Over Union (IoU).

**Keywords:** 3D reconstruction, deep learning, sampling, symmetry

## 1 Introduction

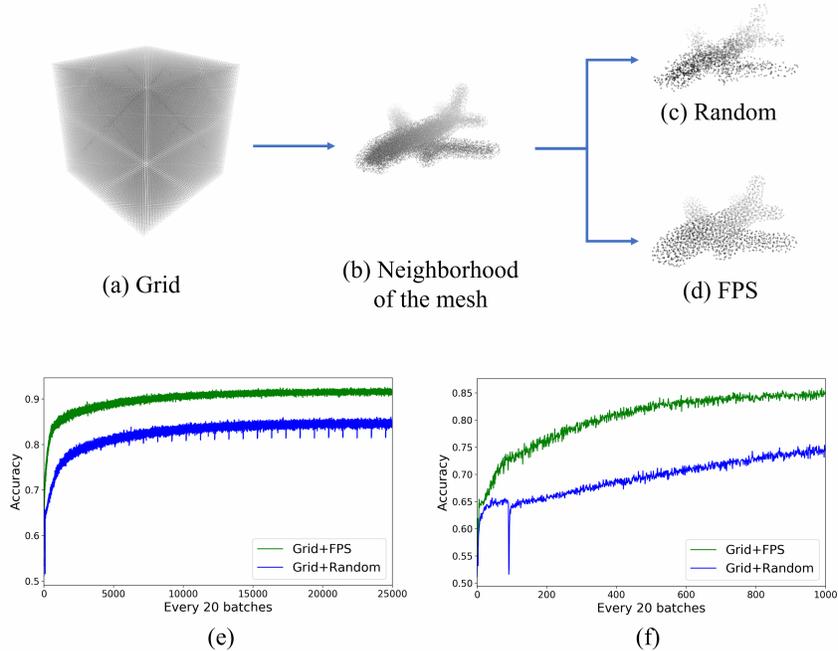
Due to the under-constrained nature of the problem, 3D object reconstruction from a single-view image has been a challenging task. Large shape and structure variations among objects make it difficult to define one dedicated parameterized model. Methods based on template deformation are often restricted by the initial topology of the template, and are not able to recover holes for instance. Recently, deep learning based implicit fields regression methods have shown great potential in monocular 3D reconstruction. Mescheder et al. [20] and DISN [32] create visually pleasing smooth shape reconstruction, with consistent normal and complex topology using implicit fields.

---

\*These two authors contribute equally.

†Corresponding author.

ORCID: Yi Yuan<sup>[0000-0003-2507-8181]</sup> Gurprit Singh<sup>[0000-0003-0970-5835]</sup>

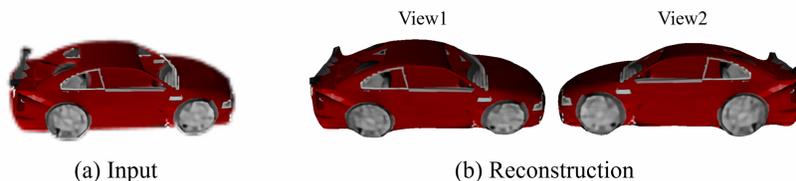


**Fig. 1.** Top: Demonstration of our sampling strategy for implicit field regression network training. A neighborhood of the mesh (b) is sampled from a set of dense grid points (a). A sparse set of points is sampled from (b) uniformly at random (c) or through FPS (d). Bottom: comparison of the training accuracy between Grid+FPS and Grid+Random sampling for the same network architecture during training. (e) is the plot of the training accuracy for the first 25 epoch. (f) is the plot of the training accuracy of the first epoch. Sampling with lower discrepancy results in faster convergence and better accuracy during training.

An implicit field is a real-valued function defined on  $\mathbb{R}^3$  whose iso-surface recovers the mesh of interest. Common choices of implicit field are signed distance field, truncated signed distance field, or occupancy probability field. A network  $g_w(I, p)$  is trained to predict the implicit field of point  $p \in \mathbb{R}^3$ , based on the input image  $I$ , where  $w$  are the parameters which are optimized with stochastic gradient descent (SGD) type algorithms. This is followed by post-processing methods like marching cube and sphere tracing to reconstruct the mesh.

The loss function for the implicit field regression problem is the  $L_2$  distance between the ground truth implicit field and the network  $g_w$  predicted output. During training, a sparse set of 3D points need to be sampled in a compact region containing the mesh to approximate the optimization objective. We formulate this empirical loss as a Monte Carlo estimator.

While most prior discussion on sampling [20] focuses on designing a probability measure for the integral that puts different weights for regions of different



**Fig. 2.** Ladybird is able to produce high quality 3D reconstruction from a single input image. The consideration of symmetry allows recovering of occluded geometry and texture completion.

distance to the mesh surface, we look at the problem from a point view of discrepancy of the sample sets. When approximating an integral, different samplers have different error convergence rates with respect to the sample size [24] [22]. Low discrepancy sequences/points or blue noise (in 2D) samples give better estimation, for instance, compared to random samples (white noise).

Given a set of locally uniform samples whose distance to the target mesh is bounded by a threshold, we show that farthest point sampling algorithm (FPS) can be used to select a sparse subset with low discrepancy for training  $g_w$ . An overview of our method is shown in Figure 1. Our proposed sampling scheme results in better generalization performance as it provides better approximation to the expected loss, thanks to the Koksma-Hlawka inequality [15]. Empirically our sampling scheme also results in faster convergence for SGD-based optimization algorithms, which speeds up the training process significantly as shown in Figure 1(e,f).

Many deep 3D implicit field reconstruction works [20] [5] explore the use of global shape encoding. While being good at capturing the general shape and obtaining interesting interpolation in the latent space, sometimes it is difficult to recover fine geometric details with only global features. Local features found via aligning image to mesh by modeling the camera are used to address the issue. However, for occluded points, it is ambiguous what local features should be used. Usually all the sampled points are projected to the images [32], and hence points in the back use features of the points that occlude them.

As most man-made objects are symmetric about a plane, we observe that this problem can be alleviated via the consideration of reflective symmetry. For a symmetric pair of points  $p$  and  $q$ , the implicit fields at  $p$  and  $q$  are the same, and often at least one of them is visible in the image. Hence we can use the local features of  $q$  to improve the implicit field predication of  $p$ , which can also be understood as utilizing two-view information. Our feature fusion method imposes a symmetry prior on the network  $g_w$ , which gives significant improvement of the reconstruction quality as shown in Figure 2. Unlike previous works [29] [33] that focus on the design of loss function, detection or encoding of symmetry, our method naturally integrates into the pixel-to-mesh alignment framework.

The advantage of spatially aligning the image to mesh and utilizing the corresponding local features is that the fine shape details and textures can be better

recovered. However, when  $p$  is occluded, the feature obtained by such alignment no longer has an intuitive meaning. Recently Front2Back [33] addresses such issues by detecting reflective symmetries from the data and synthesizing the opposite orthographic view. Our approach is simpler and does not depend on symmetry detection.

**Mesh** AtlasNet [11] represents a mesh as a locally parameterized surface and predicts the local patches from a latent shape representation learned via reconstruction objectives. Mitchell et al. [21] proposes to represent 3D shapes using higher order functions.

Pixel2Mesh [28] uses graph CNN to progressively deform an ellipsoid template mesh to fit the target. Features from different layers in the CNN are used to generate different resolution of details. 3DN [29] infers vertex offsets from a template mesh according to the image object’s category, and proposes differentiable mesh sampling operator to compute the loss function. SDM-NET [9] uses VAE to generate a spatial arrangement of deformable parts of an object. Pan et al. [23] proposes a progressive method that alternates between deforming the mesh and modifying the topology. Mesh R-CNN [10] unifies object detection and shape reconstruction, with a mesh prediction branch that first produces coarse cubified meshes which are refined with a graph convolution network.

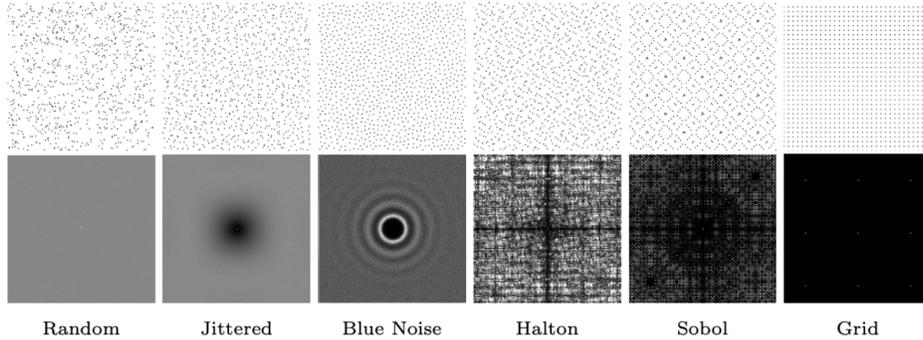
DIB-R [5], Soft Rasterizer [19] design differentiable rasterization layers that enable unsupervised training for reconstruction tasks. DIST [18] proposes an optimized differentiable sphere tracing layer for differentiable SDF rendering.

**Point Cloud and Voxel** Fan et al. [8] proposes a conditional shape sampler to predict multiple plausible point clouds from an input image. Lin et al. [17] uses an auto-encoder to synthesize partial point clouds from multiple views, which is combined as a dense point cloud. Then the loss is computed via rendering the depth images from multiple views. Li et al. [16] uses a CNN to predict multiple depth maps and corresponding deformation fields, which are fused to form the full 3D shape.

3D-R2N2 [6] uses recurrent neural networks to generate voxelized 3D reconstruction. Pixel2Vox [30] uses encoder-decoder structures to generate a coarse 3D voxel.

### 1.1 Sampling Methods in Monte Carlo Integration

Realistic image synthesis involves evaluating very high-dimensional light transport integrals. (Quasi-)Monte Carlo (MC) numerical methods are traditionally employed to approximate these integrals which is highly error prone. This error directly depends on the sampling pattern used to estimate the underlying integral [25]. These sampling patterns can be highly correlated. Fourier power spectra are commonly employed to characterize these correlations among samples (Figure 3).



**Fig. 3.** Top row shows different point patterns for different samplers for  $N = 1024$  samples. Bottom row shows the corresponding *expected* power spectra. Random samples are completely decorrelated which results in a flat power spectrum. 2D stratification (jittered) results in a power spectrum with small dark region around the center (DC frequency). For blue noise (Poisson Disk) sampler, this dark (no energy low-frequency) region is larger. However, for Halton and Sobol samplers, the corresponding power spectrum shows some spikes, but it preserves well the underlying stratification along dimensions which is characterized as a dark cross in the middle of the spectrum. Finally, a simple regular (Grid) pattern has a grid like power spectrum (zoom-in to the right-most bottom image to see the grid structure).

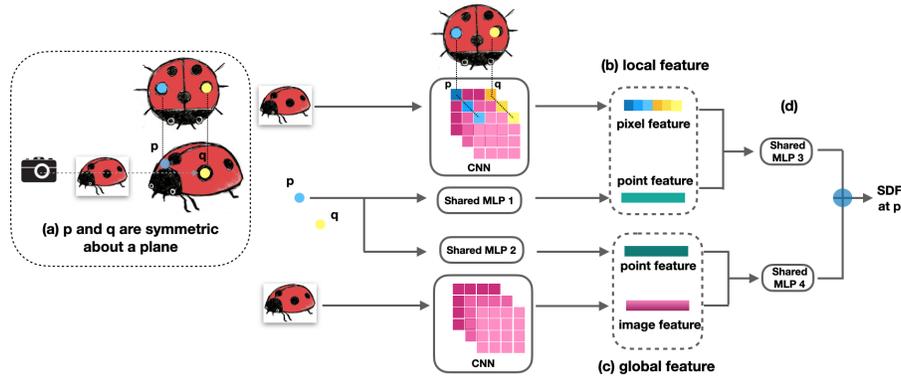
Blue noise samplers [25] are well-known to show good improvements for low-dimensional integration problems whereas low-discrepancy [22] samplers like Halton [12] and Sobol [13] are more effective for higher dimensional problems. In this work, we use farthest point selection strategy [7] from any given pointset to select our samples.

## 2 Our Approach

We first start with a theoretical motivation for our sampling methods. This is followed by the proposed symmetric feature fusion module and our 3D reconstruction pipeline (illustrated in Figure 4).

### 2.1 Preliminary

In Quasi-Monte Carlo integration literature, the equidistribution of a point set is tested by calculating the discrepancy of the set. This approach assigns a single quality number, the discrepancy, to every point set. The lower the discrepancy, the better (uniform) the underlying point set would be. We focus on the *star discrepancy* of a point set, which computes discrepancy with respect to rectangular axis-aligned sub-regions with one of their corners fixed to the origin. Mathematically, the star discrepancy can be defined as follows:



**Fig. 4.** Overview of Ladybird. a)  $p, q \in \mathbb{R}$  are symmetric about a plane. Their projections to the image are found via a camera model. b) Local feature consists of point feature of  $p$  and local image feature from pixels corresponding to  $p$  and  $q$ . c) Global feature consists of point feature of  $p$  and global image feature. d) Local feature and global feature are encoded through two MLPs whose parameters are shared among all  $p \in \mathbb{R}$ . In the end, marching cube is used to extract iso-surface.

**Definition 1** Let  $P = \{x_1, x_2, \dots, x_N\}$  be a set of points in  $\mathbb{R}^d$ , then the star discrepancy of  $P$  is

$$D_N^*(P) = \sup_{B \in J} \left\| \frac{A(B; P)}{N} - \lambda_d(B) \right\|, \quad (1)$$

where  $\lambda_d$  is the Lebesgue measure on  $\mathbb{R}^d$ ,  $A(B; P)$  is the number of points in  $P$  that are in  $B$ , and  $J = \left\{ \prod_{i=1}^d [0, u_i) \mid 0 < u_i \leq 1 \right\}$ .

For a given point set or a sequence  $P$  (stochastic or deterministic), the error due to sampling is directly related to the star discrepancy  $D_N^*(P)$  of the point set  $P$ . This relation is given by the Koksma-Hlawka inequality [15] as described below:

**Theorem 1** Let  $I = [0, 1]^d$  and  $f$  is a function on  $I$  with bounded variation  $V(f)$ . Then for any  $x_1, x_2, \dots, x_N \in I$ ,

$$\left\| \frac{1}{N} \sum_{i=1}^N f(x_i) - \int_I f(x) dx \right\| \leq V(f) D_N^*(\{x_1, \dots, x_N\}) \quad (2)$$

The above inequality states that for  $f$  with bounded variation, a point set with lower discrepancy gives less error when numerically integrating  $f$ .

The distance between two implicit fields is an integral, and a set of points in needs to be sampled to approximate such integral which appears in the expected loss for deep implicit fields regression. By triangle inequality, using lower discrepancy sampler indicates a better bound on the generalisation error.

## 2.2 Sampling

Given an input image  $I$ , we denote a neural network as  $g_w(I, p)$  that predicts an implicit field of point  $p \in \mathbb{R}^3$ . Let  $f_I(p)$  be the ground truth implicit field of the mesh  $M$  from which  $I$  is rendered, and let  $S$  be the training set. To estimate the expected loss, we need to estimate the following:

$$\sum_{I \in S} \int_{\mathbb{R}^3} (f_I(p) - g_w(I, p))^2 m(p) dp \quad (3)$$

where  $m(p)$  is a probability density function in  $\mathbb{R}^3$  supported in a compact region near the mesh  $M$ .

Instead of studying different choices for  $m(p)$  and their effects on training, we study the impact of different sampling patterns on the integral estimation.

The error convergence rate of an estimator is greatly influenced by the sampling pattern [22, 24]. Sparse sampling could result in aliasing following the Nyquist-Shannon theorem. A better sampling strategy would allow faster convergence to the true integral resulting in better generalisation performance. Following the Koksma-Hlawka inequality, in order to better approximate the  $L_2$  distance between  $g_w$  and  $f_M$ —which indicates better generalisation of the network on different input points  $p$ —sample sets of lower discrepancy should be preferred.

In consideration of the time efficiency, usually we pre-compute the implicit field of a dense set of points around the mesh surface, where a sparse subset is chosen uniformly during training. Hence we consider the following problem: given a set of points  $A$ , how to select a subset  $B \subseteq A$  consisting of  $N$  points with low discrepancy. It is natural to consider farthest point sampling algorithm(FPS): initially  $x_1 \in A$  is selected uniformly at random. Then iteratively,

$$x_i = \arg \max_{x \in A \setminus B} (\min_{y \in B} d(x, y)) \quad (4)$$

is added to  $B$ . In Section 3.3, we show that compared to randomly selecting a sample subset  $B$  from  $A$ , sampling using the FPS approach results in lower discrepancy.

## 2.3 Feature Fusion Based on Symmetry

For a fixed camera model, let  $\pi$  be the corresponding projection that maps 3D points to the image plane. Assume that the target mesh  $M$  is symmetric about  $xy$  plane<sup>4</sup>, and  $A$  is the rigid transformation such that the input image is formed via the composition  $\pi \circ A$ . In practice, either  $A$  is known or  $A$  is predicted via a camera network from input image.

For a point  $p$  not too far from  $M$ , let  $I_p$  be the pixel in the image that corresponds to  $\pi(p)$ . A convolution neural network (CNN) is used to extract

<sup>4</sup> ShapeNet data set is aligned, and most objects are symmetric about  $xy$  plane.

features from the input image  $I$ . Let  $F_p$  be the concatenation of feature vectors at  $I_p$  in different layers of the CNN.

We can use  $F_p$  to guide the regression of the implicit field at  $p$ . However when  $p$  is occluded, the pixel value of  $I_p$  is not determined by  $p$  but by  $r \in M$  with smallest z-buffer value whose projection  $\pi(r)$  also lies in the pixel  $I_p$ . There is no clear relation between the implicit field at  $p$  and that at  $r$ .

For a point  $v = (x, y, z)$ , such that  $p = Av$ , the symmetric point  $q$  of  $p$  is  $A\bar{v}$  where  $\bar{v} = (x, y, -z)$ . The implicit field at  $p$  should equal to that at  $q$ . Hence it is reasonable to include  $F_q$  as part of the local feature of  $p$ , which we call feature fusion. One straight-forward and effective way to implement feature fusion is to concatenate  $F_p$  and  $F_q$ .

### 3 Experiments

To show the effectiveness of our proposed system Ladybird, we provide quantitative as well qualitative comparisons to other methods. Our backbone network architecture is based on DISN [32]. Our implementation of Ladybird is in Tensorflow 1.9 [2], and the system is tested on Nvidia GTX 1080Ti with Cuda 9.0. In all our experiments, Adam optimizer [14] is used with  $\beta = 0.5$  and an initial learning rate of  $1e-4$ .

#### 3.1 Data Processing

For dataset, we use ShapeNet Core v1 [3], and use the official train/test split. There are 13 categories of objects. For each object, 24 views are rendered as in 3D-R2N2 [6]. We randomly select 6000 images from the training set as the validation set, and our training set contains 726,600 images. The data is aligned and most objects (about 80 percent) are symmetric about  $xy$  plane. We normalize the object mesh such that its center of mass is at the origin and the mesh lies in the unit sphere.

To efficiently and accurately compute the SDF values, we use polygon soup algorithm [31] to compute the SDF on  $256^3$  grid points. After that, non-grid point SDF values are obtained through tri-linear interpolation.

For each mesh object, first we sample  $256^3$  points  $P_1$  using Grid, Jitter, or Sobol sampler [1] and compute the corresponding SDF values. In Jitter, each grid point jitters with Gaussian noise of mean 0 and standard deviation 0.02. We then sample a subset  $P_2 \subset P_1$  consisting of 32,768 points from  $P_1$  in the following way: from each SDF range  $[-0.10, -0.03]$ ,  $[-0.03, 0.00]$ ,  $[0.00, 0.03]$ , and  $[0.03, 0.10]$ , 1/4-th of points are sampled uniformly at random. During training time, a subset  $P_3 \subset P_2$  consisting of 2048 points are sampled from  $P_2$  uniformly at random or through FPS at each epoch. Depending on the sampling pattern used to sample  $P_1$  (say  $A$ ) and  $P_3$  (say  $B$ ), the resulting sampling pattern is denoted by  $A + B$ .

At test time, the SDF of  $256^3$  grid points are predated and marching cube is used to extract the iso-surface.

**Table 1.** Mean ( $\times 0.01$ ) and standard deviation ( $\times 0.01$ ) of star discrepancy of different samplers. A+B means we first sample  $256^2$  points using sampling method A and then select a subset of size  $n = 1024, 2048, 4096$  with method B.

Sample size	Metric	Grid+Random	Grid+FPS	Jitter+FPS	Sobol+FPS
1024	Mean	4.51	3.86	6.49	4.35
	Std	0.66	0.19	0.49	1.43
2048	Mean	2.98	2.48	6.07	2.51
	Std	0.26	0.16	0.77	0.47
4096	Mean	2.34	1.96	6.10	1.65
	Std	0.33	0.08	0.30	0.37

### 3.2 Network Details

We use a pre-trained camera pose estimation network from DISN [32], to predict a rigid transformation matrix  $A$  described in Section 2.3. VGG-16 is used as a CNN module to extract features from the input image. For a given point  $p$ ,  $F_p$  is the concatenation of features (Section 2.3) at different layers of VGG-16 at pixel  $I_p$  that  $p$  projects under the known or predicted camera intrinsics. Assuming  $q$  and  $p$  being symmetric about a plane, the pixel feature of  $p$  is one of the following:

1. Base:  $F_p$  which is of dimension 1472.
2. Symm(Near):  $F_p$  or  $F_q$  depending on the one having smaller z-buffer value.
3. Symm(Avg): The average of  $F_p$  and  $F_q$ .
4. Symm(Concat): The concatenation of  $F_p$  and  $F_q$ .

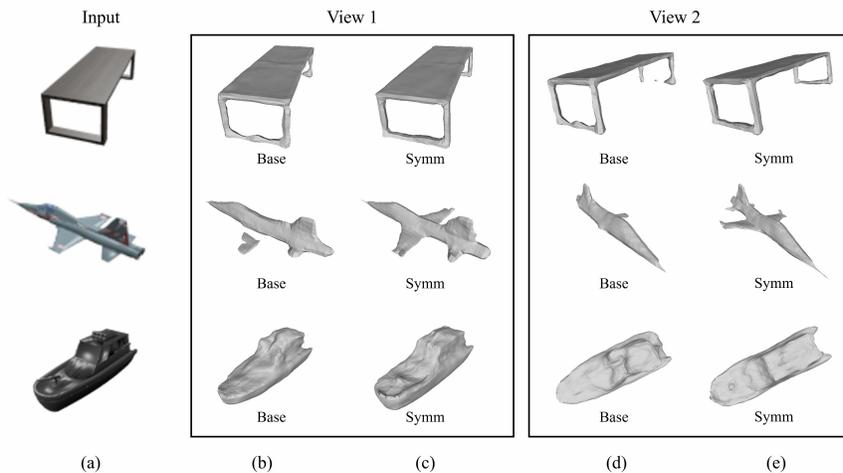
As shown in Figure 4, the image feature is the output of VGG-16 (of dimension 1024). Two stream of point features are processed with two MLPs, each of parameters (64, 256, 512). Each stream is concatenated with pixel feature and image feature respectively, to form a local and a global feature. These global and local features are encoded through two MLPs, each of parameters (512, 256, 1), and the encoded values are added as the predicted SDF at  $p$ .

### 3.3 Samplers Impact on Training

To assess the effect of different samplers on training, we set our pixel features to Base (see Section 3.2), use ground truth camera parameters and keep the batch size to 20. In Table 1, we report the star discrepancy of different samplers in 2D. We first sample  $256^2$  points using Grid, Jitter, Sobol sampler in  $[0, 1]^2$ , then selecting 1024, 2048, 4096 points uniformly at random or through FPS. In Jitter, each grid point jitters with Gaussian noise of mean 0 and standard deviation 0.01. We experimentally verify in 2D that Grid+FPS sampling has lower discrepancy and lower variance compared to Grid+Random.

**Table 2.** Effect of different samplers on the reconstruction results on ShapeNet test set. Method A+B means  $P_1$  is sampled with A and  $P_3$  is sampled with B. Metrics are class mean of CD ( $\times 0.001$ ), and class mean of EMD ( $\times 100$ ), computed on 2048 points. Grid+FPS outperforms other methods.

Metric	Grid+Random	Grid+FPS	Jitter+FPS	Sobol+FPS
CD	10.17	8.43	19.88	11.33
EMD	2.71	2.57	2.92	2.84



**Fig. 5.** Impact of feature fusion based on reflective symmetry. (a) indicates the input images. (b) and (d) are the reconstruction results using Base in two different views. (c) and (e) are the reconstruction results using Symm(Concat). We see that Symm(Concat) helps to improve the reconstruction quality.

*Grid vs. Sobol:* The SDF validation accuracy of Sobol+FPS (0.914) is similar to that of Grid+FPS (0.917), which is higher than Grid+Random (0.825). However, SDF prediction is an intermediate step for the reconstruction task. Marching Cube is used to recover the mesh from the SDF, which requires SDF values at grid points. Due to this grid restriction imposed by Marching Cube, Grid sampling ensures better training/test data consistency. In addition, Grid+FPS and Grid+Random leads to more stable training results (cf. Sobol+FPS) due to lower std. Our work advocates that Grid+FPS is suitable for 3D reconstruction based on deep implicit fields and marching cube.

In Table 2, we report the comparison of reconstruction using different samplers in terms of Chamfer distance (CD)<sup>5</sup> and Earth Mover’s distance (EMD) [27]. We see that Grid+FPS outperforms Grid+Random, Jitter+FPS, as well

<sup>5</sup> For two point set  $S_1$  and  $S_2$ , CD is defined to be  $\sum_{x \in S_1} \min_{y \in S_2} \|x-y\|_2^2 + \sum_{y \in S_2} \min_{x \in S_1} \|x-y\|_2^2$ .

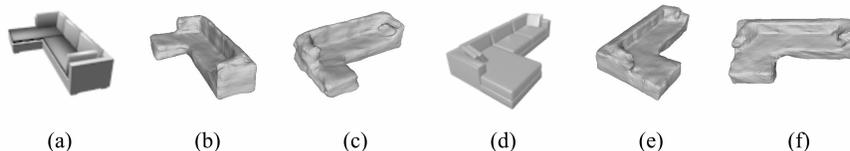
**Table 3.** Comparison between different feature fusion operation evaluated on ShapeNet test set. Metrics are CD ( $\times 0.001$ ), and EMD ( $\times 100$ ), computed on 2048 points. Ground truth camera parameters are used.

Metric	Local image feature	plane	bench	box	car	chair	display	lamp	speaker	rifle	sofa	table	phone	boat	Mean
CD	Base	5.33	5.37	9.33	4.42	7.73	7.07	24.36	13.65	<b>3.32</b>	5.78	9.37	8.13	5.79	8.43
	Symm(Avg)	7.27	17.00	12.29	4.97	14.83	15.83	58.77	23.76	6.72	11.15	12.06	61.73	5.96	19.41
	Symm(Near)	4.73	5.50	9.13	4.12	6.70	7.05	18.43	12.26	3.62	6.70	11.49	4.49	5.37	7.66
	Symm(Concat)	<b>3.86</b>	<b>4.30</b>	<b>8.04</b>	<b>4.11</b>	<b>5.43</b>	<b>6.09</b>	<b>14.10</b>	<b>10.53</b>	3.51	<b>5.05</b>	<b>8.13</b>	<b>4.16</b>	<b>4.92</b>	<b>6.33</b>
EMD	Base	2.35	2.30	2.91	2.47	2.66	2.44	4.21	3.19	1.69	2.29	2.78	1.95	2.14	2.57
	Symm(Avg)	2.14	2.36	2.98	2.42	2.56	2.54	4.69	3.41	1.71	2.45	2.77	3.25	2.12	2.72
	Symm(Near)	2.24	2.22	2.95	2.40	2.53	2.42	4.11	3.14	<b>1.65</b>	2.38	2.85	1.89	2.11	2.53
	Symm(Concat)	<b>2.07</b>	<b>2.06</b>	<b>2.80</b>	<b>2.38</b>	<b>2.32</b>	<b>2.28</b>	<b>3.59</b>	<b>2.98</b>	1.73	<b>2.18</b>	<b>2.57</b>	<b>1.85</b>	<b>2.07</b>	<b>2.38</b>
IoU	Base	63.4	56.3	52.0	77.8	58.1	60.2	41.7	58.4	70.4	71.3	53.8	75.7	66.0	61.9
	Symm(Near)	64.7	56.3	54.3	79.0	60.2	61.1	43.4	58.5	71.6	69.8	52.8	76.4	67.5	62.7
	Symm(Concat)	<b>66.6</b>	<b>60.3</b>	<b>56.4</b>	<b>80.2</b>	<b>64.7</b>	<b>63.7</b>	<b>48.5</b>	<b>61.5</b>	<b>71.9</b>	<b>73.5</b>	<b>58.1</b>	<b>78.1</b>	<b>68.8</b>	<b>65.6</b>

as Sobol+FPS. Jitter+FPS performs the worst and its 2D analogue also has the highest star discrepancy. We observe that Grid+FPS reduces noisy phantom blocks around the mesh, and hence reduces the need for post-processing and cleaning. This property is highly desired, because sometimes the cleaning algorithm cannot distinguish between small components and noise. In addition, Grid+FPS encourages faster training convergence as shown in Figure 1.

### 3.4 Effect of Feature Fusion Based on Symmetry

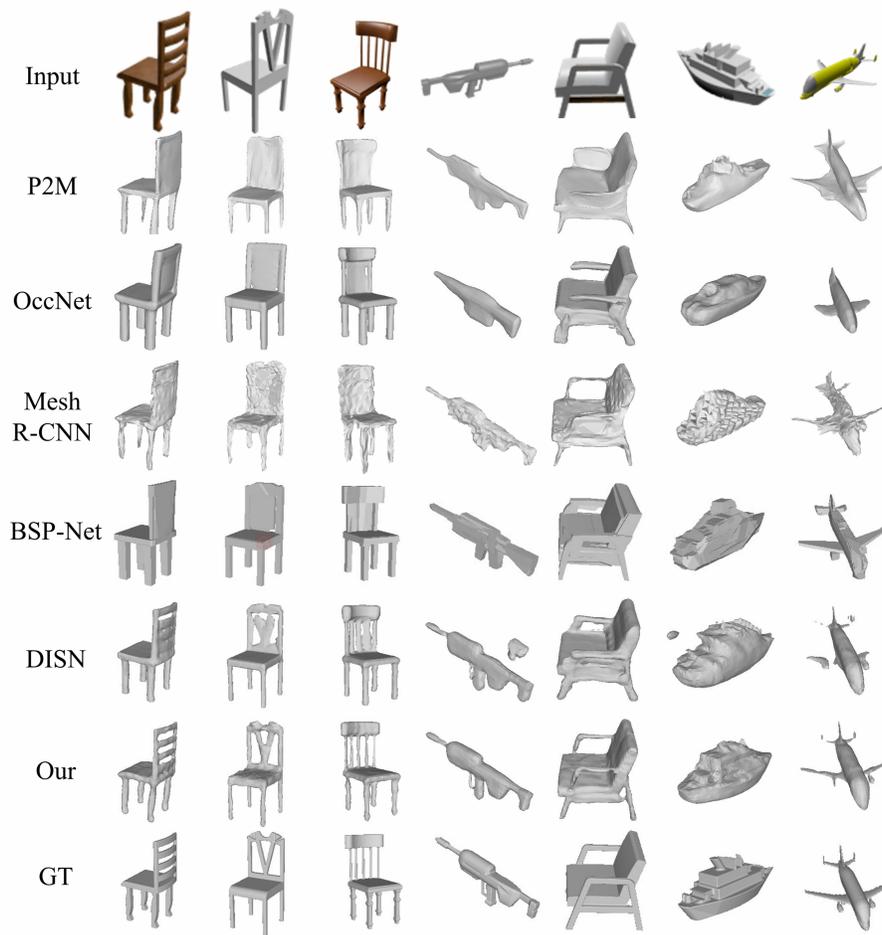
To analyze the effect of symmetry-based feature fusion, we choose Grid+FPS sampling method. The corresponding batch size for this experiment is kept 16.



**Fig. 6.** Symm(Concat) can produce good reconstruction result for non-symmetrical object, without ground truth camera parameters. (a) and (d) are input images. (b) and (c) are reconstruction result of (a) rendered from 2 different views. (e) and (f) are reconstruction result of (d) from two different views.

In Table 3 and Figure 5, we compare the effects of different feature fusion operations that are defined in Section 3.2 on the reconstruction result from ShapeNet. Ablation study shows that Symm(Near) and Symm(Concat) improve the reconstruction results. We see that concatenation of features from symmetrical pair performs the best. The reason is that Symm(Concat) better utilizes additional information comparing to Symm(Near) and Symm(Avg). When both

$p$  and its symmetry point  $q$  are visible in the image, the pixel features of  $p$  and  $q$  are both helpful for recovering the local shape at  $p$ . We observe that Symm(Concat) is able to produce reconstruction result for non-symmetrical object as shown in Figure 6. It has the interpretation of adding the most promising additional local feature based on a symmetry prior.



**Fig. 7.** Qualitative comparison with other methods. The first row contains the input image. The released model of P2M (Pixel2Mesh)[28], OccNet [20], Mesh-R-CNN [10], BSP-NET [4], DISN [32] are used to generate the results. The last row GT contains the ground truth meshes.

### 3.5 Comparison with Other Methods

In this subsection, the sampling method is Grid+FPS. The pixel feature is Symm(Concat). Camera parameters are estimated using the network mentioned in Section 4.2.

We report comparison with other state-of-the-art methods in terms of CD, EMD, and IoU. From Table 4, we see that Ladybird outperforms other methods. Figure 7 shows qualitative comparison of Ladybird with other methods. We see that Ladybird is able to reconstruct high quality mesh with fine geometric details from a single input image. Note that due to the difference between the train/test split of OccNet [20] and that of ours, we evaluate OccNet [20] on the intersection between two test sets.

**Table 4.** Evaluations on ShapeNet Core test set for various methods. Metrics are CD ( $\times 0.001$ ), EMD ( $\times 100$ ) and IoU (%), the larger the better), computed on 2048 points. *Ours<sub>cam</sub>* is Ladybird with estimated camera parameters, and *Ours* is Ladybird with ground truth camera parameters.

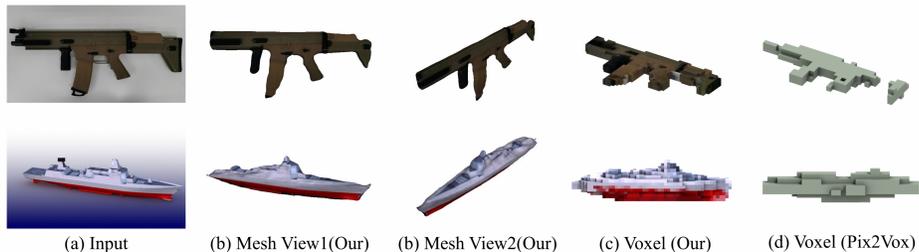
Metric	Method	plane	bench	box	car	chair	display	lamp	speaker	rifle	sofa	table	phone	boat	Mean
CD	AtlasNet [11]	5.98	6.98	13.76	17.04	13.21	7.18	38.21	15.96	4.59	8.29	18.08	6.35	15.85	13.19
	Pixel2Mesh [28]	6.10	6.20	12.11	13.45	11.13	<b>6.39</b>	31.41	<b>14.52</b>	4.51	<b>6.54</b>	15.61	6.04	12.66	11.28
	3DN [29]	6.75	7.96	<b>8.34</b>	7.09	17.53	8.35	<b>12.79</b>	17.28	<b>3.26</b>	8.27	14.05	5.18	10.20	9.77
	IMNET [5]	12.65	15.10	11.39	8.86	11.27	13.77	63.84	21.83	8.73	10.30	17.82	7.06	13.25	16.61
	3DCNN [32]	10.47	10.94	10.40	5.26	11.15	11.78	35.97	17.97	6.80	9.76	13.35	6.30	9.80	12.30
	OccNet [20]	7.70	6.43	9.36	5.26	7.67	7.54	26.46	17.30	4.86	6.72	10.57	7.17	9.09	9.70
	DISN [32]	9.96	8.98	10.19	5.39	7.71	10.23	25.76	17.90	5.58	9.16	13.59	6.40	11.91	10.98
	<i>Ours<sub>cam</sub></i>	<b>5.85</b>	<b>6.12</b>	9.10	<b>5.13</b>	<b>7.08</b>	8.23	21.46	14.75	5.53	6.78	<b>9.97</b>	<b>5.06</b>	<b>6.71</b>	<b>8.60</b>
	<i>Ours</i>	3.86	4.30	8.04	4.11	5.43	6.09	14.10	10.53	3.51	5.05	8.13	4.16	4.92	6.33
	EMD	AtlasNet [11]	3.39	3.22	3.36	3.72	3.86	3.12	5.29	3.75	3.35	3.14	3.98	3.19	4.39
Pixel2Mesh [28]		2.98	2.58	3.44	3.43	3.52	2.92	5.15	3.56	3.04	2.70	3.52	2.66	3.94	3.34
3DN [29]		3.30	2.98	3.21	3.28	4.45	3.91	3.99	4.47	2.78	3.31	3.94	2.70	3.92	3.56
IMNET [5]		2.90	2.80	3.14	2.73	3.01	2.81	5.85	3.80	2.65	2.71	3.39	2.14	2.75	3.13
3DCNN [32]		3.36	2.90	3.06	<b>2.52</b>	3.01	2.85	4.73	3.35	2.71	2.60	3.09	2.10	2.67	3.00
OccNet [20]		2.75	2.43	3.05	2.56	2.70	<b>2.58</b>	<b>3.96</b>	3.46	2.27	<b>2.35</b>	2.83	2.27	2.57	2.75
DISN [32]		2.67	2.48	3.04	2.67	2.67	2.73	4.38	3.47	2.30	2.62	3.11	2.06	2.77	2.84
<i>Ours<sub>cam</sub></i>		<b>2.48</b>	<b>2.29</b>	<b>3.03</b>	2.65	<b>2.60</b>	2.61	4.20	<b>3.32</b>	<b>2.22</b>	2.42	<b>2.82</b>	<b>2.06</b>	<b>2.46</b>	<b>2.71</b>
<i>Ours</i>		2.07	2.06	2.80	2.38	2.32	2.28	3.59	2.98	1.73	2.18	2.57	1.85	2.07	2.38
IoU		AtlasNet [11]	39.2	34.2	20.7	22.0	25.7	36.4	21.3	23.2	45.3	27.9	23.3	42.5	28.1
	Pixel2Mesh [28]	51.5	40.7	43.4	50.1	40.2	55.9	29.1	52.3	50.9	60.0	31.2	69.4	40.1	47.3
	3DN [29]	54.3	39.8	49.4	59.4	34.4	47.2	35.4	45.3	57.6	60.7	31.3	71.4	46.4	48.7
	IMNET [5]	55.4	49.5	51.5	74.5	52.2	56.2	29.6	52.6	52.3	64.1	45.0	70.9	56.6	54.6
	3DCNN [32]	50.6	44.3	52.3	<b>76.9</b>	52.6	51.5	36.2	58.0	50.5	67.2	50.3	70.9	57.4	55.3
	OccNet [20]	54.7	45.2	<b>73.2</b>	73.1	50.2	47.9	<b>37.0</b>	<b>65.3</b>	45.8	67.1	<b>50.6</b>	70.9	52.1	56.4
	DISN [32]	57.5	52.9	52.3	74.3	54.3	56.4	34.7	54.9	59.2	65.9	47.9	72.9	55.9	56.9
	<i>Ours<sub>cam</sub></i>	<b>60.0</b>	<b>53.4</b>	50.8	74.5	<b>55.3</b>	<b>57.8</b>	36.2	55.6	<b>61.0</b>	<b>68.5</b>	48.6	<b>73.6</b>	<b>61.3</b>	<b>58.2</b>
	<i>Ours</i>	66.6	60.3	56.4	80.2	64.7	63.7	48.5	61.5	71.9	73.5	58.1	78.1	68.8	65.6

Since ShapeNet is a synthesized dataset, we further provide quantitative evaluation on Pix3D [26] (Table 5), and some qualitative examples of in-the-wild images which are randomly selected from the internet (Figure 8). These results show that Ladybird generalizes well to natural images. For the experiment on Pix3D, we fine-tune Ladybird and DISN [32] (both pre-trained on ShapeNet) on

**Table 5.** Evaluations on Pix3D [26] test set. Metrics are CD ( $\times 0.001$ ), and EMD ( $\times 100$ ), computed on 2048 points. Ground truth camera parameters are used.

Metric	Method	bed	bookcase	chair	desk	misc	sofa	table	tool	wardrobe	Mean
CD	DISN [32]	12.74	35.29	23.82	18.70	31.18	3.85	18.46	46.00	<b>4.23</b>	18.51
	Ours	<b>5.73</b>	<b>15.89</b>	<b>13.03</b>	<b>10.38</b>	<b>30.34</b>	<b>3.28</b>	<b>8.38</b>	<b>28.39</b>	5.58	<b>10.02</b>
EMD	DISN [32]	2.84	4.65	3.97	4.04	<b>4.53</b>	1.99	3.85	5.66	2.11	3.53
	Ours	<b>2.35</b>	<b>3.07</b>	<b>3.23</b>	<b>2.77</b>	4.96	<b>1.84</b>	<b>2.42</b>	<b>3.68</b>	<b>1.99</b>	<b>2.75</b>
IoU	DISN [32]	71.2	43.0	59.0	53.7	48.8	89.4	57.8	37.3	85.6	64.4
	Ours	<b>78.2</b>	<b>67.8</b>	<b>66.5</b>	<b>67.5</b>	<b>49.5</b>	<b>91.8</b>	<b>74.2</b>	<b>58.4</b>	<b>86.8</b>	<b>73.3</b>

Pix3D train set, and use the ground truth camera poses and the segmentation masks.

**Fig. 8.** Reconstruction results for online images. (a) indicates input images. (b) and (c) are our reconstruction results in mesh and voxel representation respectively. (d) shows the reconstruction results of Pixel2Vox [30]. Ladybird naturally produces accurate uv-map for texturing.

## 4 Conclusion

We study the impact of sample set discrepancy on the training efficiency of implicit field regression networks, and proposes to use FPS instead of Random sampling to select training points. We also propose to explore local feature fusion based on reflective symmetry to improve the reconstruction quality. Qualitatively and quantitatively we verify the efficiency of our methods through extensive experiments on large-scale dataset ShapeNet.

## Acknowledgement

We would like to thank the anonymous reviewers for their helpful feedback and suggestions. We would like to thank Zilei Huang for his help in accelerating the data processing and debugging.

## References

1. Uni(corn—form) tool kit, <https://utk-team.github.io/utk/>
2. Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al.: Tensorflow: A system for large-scale machine learning. In: 12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16). pp. 265–283 (2016)
3. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al.: Shapenet: An information-rich 3d model repository. arXiv preprint arXiv:1512.03012 (2015)
4. Chen, Z., Tagliasacchi, A., Zhang, H.: Bsp-net: Generating compact meshes via binary space partitioning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 45–54 (2020)
5. Chen, Z., Zhang, H.: Learning implicit fields for generative shape modeling. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5939–5948 (2019)
6. Choy, C.B., Xu, D., Gwak, J., Chen, K., Savarese, S.: 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In: European conference on computer vision. pp. 628–644. Springer (2016)
7. Eldar, Y., Lindenbaum, M., Porat, M., Zeevi, Y.Y.: The farthest point strategy for progressive image sampling. *IEEE Transactions on Image Processing* **6**(9), 1305–1315 (1997)
8. Fan, H., Su, H., Guibas, L.J.: 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. pp. 605–613 (2017)
9. Gao, L., Yang, J., Wu, T., Yuan, Y.J., Fu, H., Lai, Y.K., Zhang, H.: Sdm-net: Deep generative network for structured deformable mesh. *ACM Transactions on Graphics (TOG)* **38**(6), 1–15 (2019)
10. Gkioxari, G., Malik, J., Johnson, J.: Mesh r-cnn. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 9785–9795 (2019)
11. Groueix, T., Fisher, M., Kim, V.G., Russell, B.C., Aubry, M.: Atlasnet: A papier-mâché approach to learning 3d surface generation. arXiv preprint arXiv:1802.05384 (2018)
12. Halton, J.H.: Algorithm 247: Radical-inverse quasi-random point sequence. *Communications of the ACM* **7**(12), 701–702 (1964)
13. Joe, S., Kuo, F.Y.: Constructing sobol sequences with better two-dimensional projections. *SIAM Journal on Scientific Computing* **30**(5), 2635–2654 (2008)
14. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arxiv:1412.6980 (2014)
15. Kuipers, L., Niederreiter, H.: Uniform distribution of sequences. Courier Corporation (2012)
16. Li, K., Pham, T., Zhan, H., Reid, I.: Efficient dense point cloud object reconstruction using deformation vector fields. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 497–513 (2018)
17. Lin, C.H., Kong, C., Lucey, S.: Learning efficient point cloud generation for dense 3d object reconstruction. In: Thirty-Second AAAI Conference on Artificial Intelligence (2018)
18. Liu, S., Zhang, Y., Peng, S., Shi, B., Pollefeys, M., Cui, Z.: Dist: Rendering deep implicit signed distance function with differentiable sphere tracing. arXiv preprint arXiv:1911.13225 (2019)

19. Liu, S., Li, T., Chen, W., Li, H.: Soft rasterizer: A differentiable renderer for image-based 3d reasoning. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 7708–7717 (2019)
20. Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4460–4470 (2019)
21. Mitchell, E., Engin, S., Isler, V., Lee, D.D.: Higher-order function networks for learning composable 3d object representations. arXiv preprint arXiv:1907.10388 (2019)
22. Niederreiter, H.: Low-discrepancy and low-dispersion sequences. *Journal of number theory* **30**(1), 51–70 (1988)
23. Pan, J., Han, X., Chen, W., Tang, J., Jia, K.: Deep mesh reconstruction from single rgb images via topology modification networks. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 9964–9973 (2019)
24. Pilleboue, A., Singh, G., Coeurjolly, D., Kazhdan, M., Ostromoukhov, V.: Variance analysis for monte carlo integration. *ACM Trans. Graph. (Proc. SIGGRAPH)* **34**(4), 124:1–124:14 (2015)
25. Singh, G., ztireli, C., Ahmed, A.G., Coeurjolly, D., Subr, K., Deussen, O., Ostromoukhov, V., Ramamoorthi, R., Jarosz, W.: Analysis of sample correlations for Monte Carlo rendering. *Computer Graphics Forum* **38**(2), 473–491 (2019)
26. Sun, X., Wu, J., Zhang, X., Zhang, Z., Zhang, C., Xue, T., Tenenbaum, J.B., Freeman, W.T.: Pix3d: Dataset and methods for single-image 3d shape modeling. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2974–2983 (2018)
27. Villani, C.: *Optimal transport: old and new*, vol. 338. Springer Science & Business Media (2008)
28. Wang, N., Zhang, Y., Li, Z., Fu, Y., Liu, W., Jiang, Y.G.: Pixel2mesh: Generating 3d mesh models from single rgb images. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 52–67 (2018)
29. Wang, W., Ceylan, D., Mech, R., Neumann, U.: 3dn: 3d deformation network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1038–1046 (2019)
30. Xie, H., Yao, H., Sun, X., Zhou, S., Zhang, S.: Pix2vox: Context-aware 3d reconstruction from single and multi-view images. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2690–2698 (2019)
31. Xu, H., Barbič, J.: Signed distance fields for polygon soup meshes. In: Proceedings of Graphics Interface 2014, pp. 35–41 (2014)
32. Xu, Q., Wang, W., Ceylan, D., Mech, R., Neumann, U.: Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. arXiv preprint arXiv:1905.10711 (2019)
33. Yao, Y., Schertler, N., Rosales, E., Rhodin, H., Sigal, L., Sheffer, A.: Front2back: Single view 3d shape reconstruction via front to back prediction. arXiv preprint arXiv:1912.10589 (2019)