

ZeroMesh: Zero-shot Single-view 3D Mesh Reconstruction

Xianghui Yang, Guosheng Lin, Luping Zhou

Abstract—Single-view 3D object reconstruction is a fundamental and challenging computer vision task that aims at recovering 3D shapes from single-view RGB images. Most existing deep learning based reconstruction methods are trained and evaluated on the same categories, and they cannot work well when handling objects from novel categories that are not seen during training. Focusing on this issue, this paper tackles Zero-shot Single-view 3D Mesh Reconstruction, to study the model generalization on unseen categories and encourage models to reconstruct objects literally. Specifically, we propose an end-to-end two-stage network, ZeroMesh, to break the category boundaries in reconstruction. Firstly, we factorize the complicated image-to-mesh mapping into two simpler mappings, *i.e.*, image-to-point mapping and point-to-mesh mapping, while the latter is mainly a geometric problem and less dependent on object categories. Secondly, we devise a local feature sampling strategy in 2D and 3D feature spaces to capture the local geometry shared across objects to enhance model generalization. Thirdly, apart from the traditional point-to-point supervision, we introduce a multi-view silhouette loss to supervise the surface generation process, which provides additional regularization and further relieves the overfitting problem. The experimental results show that our method significantly outperforms the existing works on the ShapeNet and Pix3D under different scenarios and various metrics, especially for novel objects.

Index Terms—Mesh, Single-view reconstruction, Generalization, Zero-shot.

I. INTRODUCTION

HUMANS are able to imagine a rough 3D shape from a given RGB image even though the particular object was not seen before. Modeling such a reconstruction process is an interesting research topic in computer vision, known as the single-view 3D reconstruction. This task is highly ill-posed because it is unable to know the object portion unseen from the camera perspective. Humans can accomplish the task because we store plenty of objects in mind, which enables us to establish the correlation between 2D images and 3D shapes. The correlation can further be applied to novel images carrying familiar objects. Based on this observation, computer vision researchers turn to deep learning techniques to mimic the 3D reconstruction process of humans. Recently, deep neural networks have achieved remarkable progress on a variety of 3D object representations, *e.g.*, voxels [1]–[5], point clouds [6]–[8], and implicit functions [9]–[11], which are friendly to current network architectures. However, compared with the above representations, the mesh based representation has not yet been fully explored even though it is more efficient to store geometry information and more widely used in industries.

Existing single-view 3D mesh reconstruction works mostly adopt the encoder-decoder architectures, where the encoder

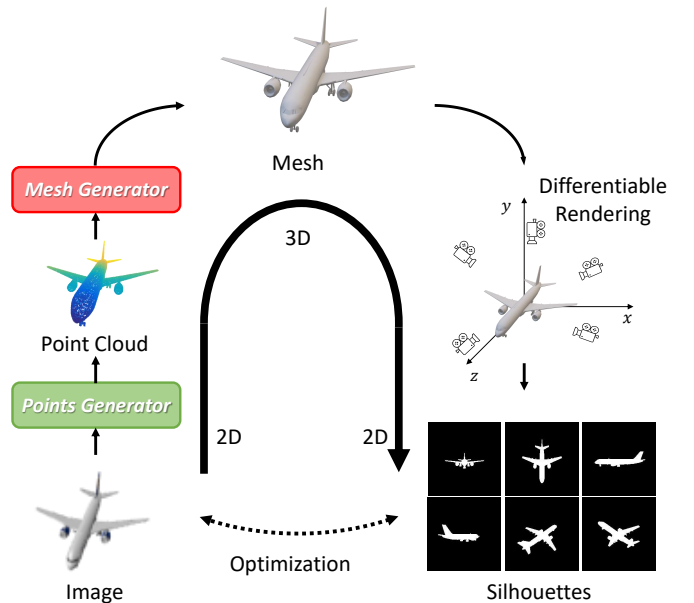


Fig. 1. Overview of the proposed ZeroMesh framework for zero-shot single-view 3D mesh reconstruction. It consists of two stages jointly trained: i) point cloud generation from the single input image, and ii) mesh generation from the intermediate point cloud.

extracts perceptual features from the input images while the decoder deforms a template (2D square [12] or sphere [13], [14]) to warp the target 3D shape. These mesh reconstruction networks are normally trained and evaluated on the same categories with encouraging performance. However, when generalizing them to novel categories, we observe a substantial performance decline as these training protocols suffer from the overfitting problem on the seen classes. The class-specific models are difficult to deploy in practice since it is inherently infeasible to collect a training set covering all types of object shapes around the world. Moreover, handling objects from various categories will also cause performance degradation, even when these categories are included in the training set. Thus, an important question was raised: what do single-view 3D reconstruction networks indeed learn? Tatarchenko *et al.* [15] argue that the existing models actually function by recognition rather than reconstruction, which is supported by an observation that a benchmark work [12] performs even worse than the simple nearest neighbor searching or classification baselines, both quantitatively and qualitatively.

To improve model generalization in single-view 3D mesh reconstruction, this paper aims to tackle the zero-shot single-

view 3D mesh reconstruction. We develop a new framework, **ZeroMesh**, to encourage the model to reconstruct 3D shapes literally instead of retrieving objects from memory. Under the zero-shot setting, the model is trained on seen classes, *i.e.*, **base classes**, and tested on unseen classes, *i.e.*, **novel classes**. ZeroMesh comprises a two-stage pipeline to generate the triangular mesh from a single-view RGB image, with three strategies to break the boundaries of object categories for better model generalization.

First, we disentangle the complicated reconstruction process, *i.e.*, image-to-mesh mapping, into two simpler mappings, *i.e.*, image-to-point mapping and point-to-mesh mapping. This operation decreases the complexity of the learning task and makes the training process easier. Moreover, the point-to-mesh mapping is more of a geometric problem than a recognition problem and has less dependency on object categories. Evidently, network models for 3D point meshing (point-to-mesh generation) showed good performance on novel classes [12]. Specifically, this strategy is implemented in a two-stage framework trained in an end-to-end manner. At the first stage, a point cloud generator estimates the point clouds of the target objects guided by image features. At the second stage, a mesh generator deforms triangular templates into target meshes by vertex shifts towards the intermediate point cloud, while preserving raw connections during deformation. **Second**, we explore the feature representations, *i.e.*, local features [13], [16] and global features [1], [12], [14], [17]–[21], for the generalization, proving the superiority of feature sampling strategy instead of the widely-used global features, and extend the 2D local feature sampling into 3D to prevent the model from overfitting. We argue that global features describing overall shapes easily overfit training object categories and thus cannot generalize well to novel categories. In contrast, local features focus on local geometry shared across categories (seen and unseen) and prevent recognition by the limited receptive field. Although local features have been sporadically seen in single-view 3D reconstruction [13], [16], they were mainly proposed to improve reconstruction quality, rather than being explicitly linked to model generalization. Evidently, most reconstruction methods studying model generalization [17]–[19] still adopt global features. This paper would like to emphasize that local features could not only improve the details of reconstruction but also matter to model generalization, especially for novel unseen object categories, as carefully evaluated in our work. It attempts to partially answer the question in [15]: how to encourage models to really learn reconstruction instead of retrieval. Specifically, the 2D Feature Sampler obtains local features from the extracted 2D feature maps by 3D-to-2D projection according to camera intrinsic matrix and pose, while the 3D Feature Sampler extracts local features from the encoded 3D feature groups of the generated point cloud by nearest neighbor searching. **Third**, we design a multi-view silhouette loss in 2D space to inject the face quality supervision, which serves as an additional regularization for reconstruction. It is noted that the widely used point-to-point distance, Chamfer Distance, is not sufficient to supervise the edges and faces generation due to the lack of a direct measure for evaluation. Thus, we introduce additional supervision

based on the consistency of the 2D projections of the predicted and the ground-truth meshes from different camera views. It is noteworthy that although 2D local features and silhouette loss have been sporadically seen in 3D reconstruction literature, their importance to zero-shot reconstruction has never been identified and evaluated as in our work. For example, silhouette loss is only adopted by unsupervised 3D reconstruction [22]–[25] as the second choice when 3D shapes are agnostic, while we point out its advantages for supervised learning in this paper.

Our main contributions are summarized as follows. **First**, we develop an end-to-end two-stage framework for zero-shot single-view 3D mesh reconstruction by disentangling the task into point cloud generation and surface recovery, which simplifies the problem and alleviates overfitting on seen object categories. **Second**, we argue that the coarse global feature cannot guide the accurate generation and tend to overfit the training set. To this end, we introduce the local feature sampling strategy in both 2D and 3D spaces and demonstrate the benefits for zero-shot reconstruction. **Third**, we regularize reconstruction by a multi-view silhouette loss to measure surface reconstruction quality within 2D space, which was ignored by point-to-point distances, and furthermore improves the performance with regard to the Chamfer Distance and Earth Mover Distance. **Fourth**, we compare the generalization capacity of our model with the existing single-view 3D mesh reconstruction methods on novel objects. The experimental results on ShapeNet.V1 [26] and Pix3D [27] demonstrate the significant improvement by our ZeroMesh on various metrics.

II. RELATED WORK

Single-view 3D Reconstruction. Single-view 3D reconstruction is a challenging task. Accurate reconstruction requires the integration of strong geometric priors about our 3D world, which are, however, limited in the wild scenario [28]–[31]. Learning based methods, therefore, become dominant in this field due to their robustness and accessibility. According to the employed 3D representations, the deep learning based methods can be classified into voxel based [1]–[5], point cloud based [6]–[8], mesh based [16], [23]–[25], [32]–[36], and implicit function [9]–[11] based frameworks. Among them, mesh reconstruction [12]–[14] is most related to our work. The majority of the existing single-view 3D mesh reconstruction methods adapts the encoder-decoder framework, where the encoder extracts perceptual features from the input image, and the decoder deforms a template (sphere or squire) to the target 3D shape. Wang *et al.* [13] firstly applied deep learning networks to this task, where VGG network [37] was used as the encoder and a graph convolutional network (GCN) was used as the decoder. Groueix *et al.* [12] represented a 3D shape as a collection of parametric surface elements to flexibly represent shapes with arbitrary topology. Pan *et al.* [14] focused on the topology changes and proposed a topology modification network by adaptively deleting faces. It is noteworthy that these methods are trained and evaluated on the same object categories.

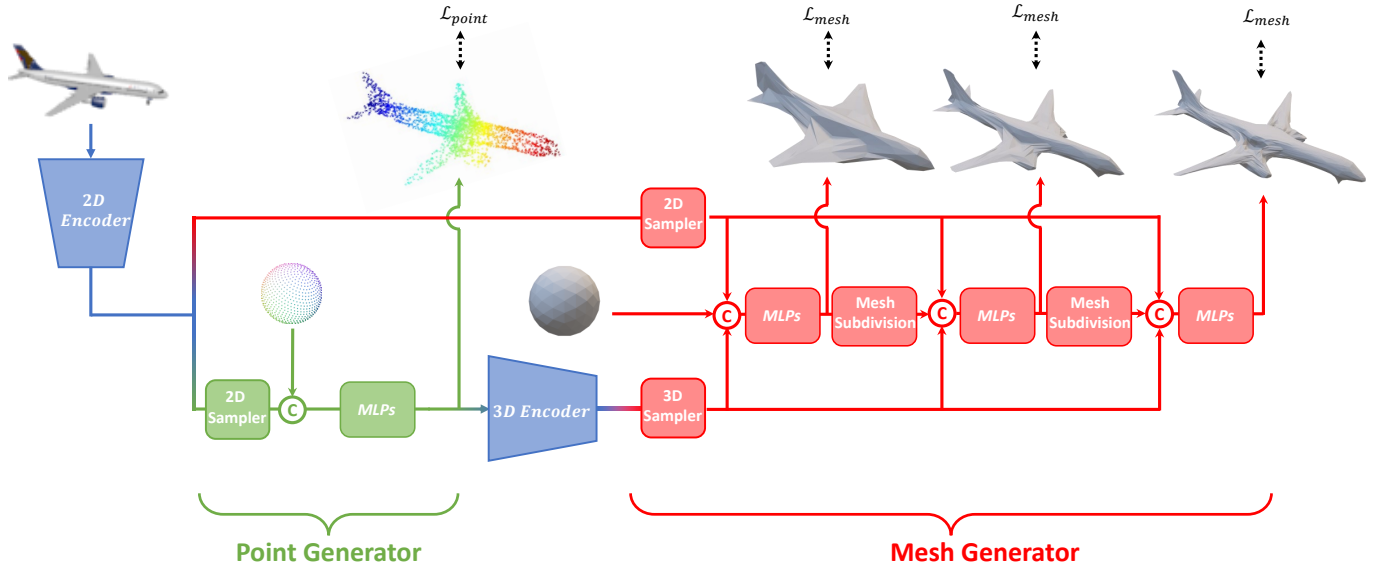


Fig. 2. Network architecture of the proposed ZeroMesh. The framework includes two main modules: Point Cloud Generator (green) and Mesh Generator (red). Both generators sample local features by point/vertex coordinates by 2D and 3D samplers, respectively. The mesh generator employs a coarse-to-fine process to improve mesh generation iteratively.

Generalized Single-view 3D Reconstruction. There are also a few research works catering for the generalization capacity of 3D reconstruction models. To the best of our knowledge, there are only three voxel based works [17], [18], [38] and one signed distance function based work [19] exploring novel class object reconstruction, elaborated as follows. Zhang *et al.* [17] pioneered the generalized single-view voxel reconstruction, where the 2D-3D mapping was decomposed into 2D-2.5D-3D mappings with the use of depth and normals as intermediate representations. Wang *et al.* [18] followed that work, and more importantly, introduced shape interpretation into reconstruction and jointly learned interpretation and reconstruction to capture more generic geometry. Bautista *et al.* [38] studied feature description bias for generalization and emphasized the reconstruction from multiple views. Thai *et al.* [19] shared the similar intermediate representation as [17], [18] but transferred the pipeline into the signed distance function by replacing the decoder with the conditional batch normalization network. Besides, there are two papers related to few-shot single-view 3D reconstruction [20], [21], which rely on additional 3D inputs, *i.e.*, support shapes.

III. METHODS

In this section, we present our ZeroMesh model that consists of a point cloud generator and a mesh generator in sequence as shown in Fig. 2. Specifically, both of the two generators move the points by predicting per-point offset under the guidance of local features. The local 2D and 3D features are generated by 2D and 3D samplers to enhance model understanding to object local geometry rather than object category, which will be illustrated in the Sec. III-B. To complement 3D point distance in surface supervision, our pipeline differs from previous methods in introducing additional supervision by our proposed multi-view silhouette loss, which will be discussed in Sec. III-C.

A. Two-stage reconstruction

ZeroMesh disentangles the complicated image-to-mesh generation into two simpler generation tasks, *i.e.*, point cloud generation and surface recovery, through a Point Cloud Generator and a Mesh Generator, elaborated as follows.

Point Cloud Generator. The point cloud generator takes the input 2D image as the guidance and transforms the point cloud template sampled from a unit sphere into the target point cloud $P \in \mathbb{R}^{p \times 3}$, where p is the number of points. Specifically, the input 2D image is passed to a 2D encoder, such as ResNet, to extract feature maps $F \in \mathbb{R}^{c \times h \times w}$, upon which, a 2D sampler is applied to sample local features. These local features are then sent into a multi-layer perceptron (MLP) to predict per-point offset. The whole process could be formulated as:

$$\hat{p}_i = p_i + MLP([p_i | f_i]), \quad (1)$$

where p_i and \hat{p}_i indicate the i -th point before and after the transformation, f_i is the local feature of p_i sampled by the 2D sampler, and $[\cdot]$ denotes a concatenation operation. The point cloud generator is supervised by Chamfer Distance between the generated and the ground-truth point clouds.

Mesh Generator. Taking the 2D image features and the 3D point cloud features as guidance, our mesh generator reconstructs the target object $M = (V, E)$, where $V \in \mathbb{R}^{v \times 3}$ and $E \in \mathbb{R}^{e \times 2}$ denote the vertices and the edges, respectively. The reconstruction is achieved by moving the vertices of a unit sphere mesh template towards the ground-truth vertices while maintaining the edge connections. Specifically, the intermediate point cloud output by the point cloud generator is sent into a 3D encoder, such as PointNet⁺⁺ [39], to extract feature groups $G \in \mathbb{R}^{n \times c}$ as the 3D information source of the mesh generator. We then sample the per-vertex 2D feature f_i from the 2D image feature maps $F \in \mathbb{R}^{c \times h \times w}$, and the per-vertex 3D feature g_i from the 3D feature groups $G \in \mathbb{R}^{n \times c}$




Fig. 3. Local feature Sampler. (a) 2D local features are sampled by 2D projection according to camera intrinsic matrix, camera pose and point/vertex coordinates. (b) 3D local features are sampled by nearest neighbor searching according to vertex coordinates and feature group centroids.

(n is the number of groups), and then concatenate them with the template coordinate v_i to recover the final mesh. Please see Sec. III-B for feature sampling. Our mesh generation is a coarse-to-fine process composed of a sequence of modules for refinement and subdivision to decrease the overfitting and intersection, similar to Pixel2Mesh [13]. Specifically, in the i -th module, the current predicted mesh M_i is firstly refined according to the output of the $(i-1)$ -th module through an MLP, and then subdivided by breaking each triangle face into four faces via adding three vertices at the mid-point of the triangle edges. The updated mesh M_i is then output by the i -th module and used as the input of the $(i+1)$ -th module for another round of refinement. Our mesh generator predicts the per-vertex offset through an MLP so that after deformation the i -th vertex v_i is,

$$\hat{v}_i = v_i + MLP([v_i | f_i | g_i]), \quad (2)$$

where v_i and \hat{v}_i denote the i -th vertex before and after the deformation, f_i and g_i denote the local features by 2D sampling and 3D sampling, respectively, and $[\cdot | \cdot | \cdot]$ is a concatenation operation. The mesh losses in Sec. III-D are applied on all module outputs $\{M_1, M_2, \dots, M_n\}$ for supervision.

B. Feature Sampling

Local features associated with each 3D point are extracted by sampling the RGB image features $F \in \mathbb{R}^{c \times h \times w}$ through the 2D sampler or sampling the intermediate point cloud features $G \in \mathbb{R}^{n \times c}$ through the 3D sampler.

2D sampler. For each point $x_i = (x_i, y_i, z_i)$ from the point cloud $P \in \mathbb{R}^{p \times 3}$ or the vertices $V \in \mathbb{R}^{v \times 3}$, we calculate its 2D projection on the image plane and sample local features from the image feature maps by bilinear interpolation. Specifically, for 2D projection, we first transfer the point coordinates from the world coordinate x_i into the camera coordinate x'_i using camera intrinsic matrix $K \in \mathbb{R}^{3 \times 4}$, and then calculate the point position p_i on the image plane. That is, $[x'_i | 1] = K [R | T] [x_i | 1]$ and $p_i = (\frac{x'_i}{z'_i}, \frac{y'_i}{z'_i})$, where $R \in SO(3)$ and $T \in \mathbb{R}^{3 \times 1}$ are the rotation and the translation matrices. After projection, the local features are interpolated from the four nearby pixels around the position p_i on the 2D feature

maps $F \in \mathbb{R}^{c \times h \times w}$. The 2D local features corresponding to the point x_i are defined as

$$f_i = \sum_{j \in \mathcal{N}(p_i)} \omega_j f_j, \quad (3)$$

where $\mathcal{N}(p_i)$ denotes the nearest four corner pixels of the position p_i ; $f_j \in \mathbb{R}^c$ is the feature at the position p_j extracted from $F \in \mathbb{R}^{c \times h \times w}$, and the bilinear interpolation weights ω_j are calculated according to the pixel position p_j and the projected point position p_i .

3D sampler. A given 3D point x_i is assigned 3D local features by nearest neighbor searching on the 3D feature groups $G \in \mathbb{R}^{n \times c}$ extracted by the 3D encoder from the intermediate point cloud. Each feature group encodes the local points within the sphere with radius r . During encoding, to guarantee translation invariant and capture relative correlation, the coordinates of points in a local region are firstly translated into a local frame relative to the centroid point: $x_i = x_i - \bar{x}$ for $i = 1, 2, \dots, p$, where \bar{x} is the coordinate of the centroid and p is the point number belonging to the group. For more details about the feature groups, please refer to [39]. Given the 3D feature groups with the group centroids $\{(g_1, \bar{x}_1), (g_2, \bar{x}_2), \dots, (g_n, \bar{x}_n)\}$ and the given point x_i , the 3D local feature sampling is defined as,

$$g_i = \min_{\|x_i - \bar{x}_j\|_2} \{g_j\}, \quad j = 1, 2, \dots, n \quad (4)$$

where x_i is query point and \bar{x}_j is the j -th group centroid. Note that the 3D sampler is only deployed at the second stage, *i.e.*, mesh generation.

C. Multi-view Silhouette Loss

To introduce face quality supervision and regularization into training, we render the predicted mesh and the ground-truth mesh into silhouettes by multiple virtual cameras and calculate the Intersection over Union (IoU) between the corresponding predicted and ground-truth binary masks. The multi-view silhouette loss is formulated as,

$$\mathcal{L}(P, Q) = \frac{1}{|V|} \sum_V \sum_i^{h \times w} 1 - \frac{p_i^v q_i^v}{p_i^v + q_i^v - p_i^v q_i^v}, \quad (5)$$




Fig. 4. Multi-view silhouette rendering. Multiple cameras are set at different positions to render silhouette maps.

where $P, Q \in \mathbb{R}^{v \times h \times w}$ are the predicted and the ground-truth silhouettes, respectively, from the v camera views by differentiable rendering [22].

D. Overall Loss

In addition to our proposed multi-view silhouette loss, we adopt the widely used losses in 3D space, *i.e.*, Chamfer Distance $\mathcal{L}_{chamfer}$ and the normal loss \mathcal{L}_{normal} . The Chamfer Distance is to measure the distance between two point sets, and the normal loss minimizes the normal directions between predicted and ground-truth mesh. The two losses are defined as follows:

$$\mathcal{L}_{chamfer} = \sum_p \min_q \|p - q\|_2^2 + \sum_q \min_p \|q - p\|_2^2$$

$$\mathcal{L}_{normal} = \sum_p \min_q \|n_p - n_q\|_2^2 + \sum_q \min_p \|n_q - n_p\|_2^2, \quad (6)$$

where p and q are the points from the predicted and the ground-truth point sets, n_p and n_q are the normals corresponding to the points p and q . The points and normals are randomly sampled from the surface of the generated and ground-truth mesh $M = (V, E)$. Moreover, to improve mesh quality, we also apply the edge length loss \mathcal{L}_{edge} and the vertex move loss \mathcal{L}_{move} to penalize too long edges and dramatic moving for vertices during deformation, as follows:

$$\mathcal{L}_{edge} = \sum_{e \in E} \|e_1 - e_2\|_2^2, \quad \mathcal{L}_{move} = \sum_{v \in V} \|\hat{v} - v\|_2^2, \quad (7)$$

where e_1 and e_2 are the two vertices of an edge, and v and \hat{v} indicate the vertices before and after the deformation.

Putting it together, the overall loss function is,

$$\mathcal{L} = \mathcal{L}_{point} + \mathcal{L}_{mesh}, \quad (8)$$

where \mathcal{L}_{point} is Chamfer Distance $\mathcal{L}_{chamfer}$, and the \mathcal{L}_{mesh} is defined as,

$$\mathcal{L}_{mesh} = \mathcal{L}_{chamfer} + \lambda_{normal} \mathcal{L}_{normal} + \lambda_{IoU} \mathcal{L}_{IoU} + \lambda_{edge} \mathcal{L}_{edge} + \lambda_{move} \mathcal{L}_{move}, \quad (9)$$

where the hyper-parameters λ_{normal} , λ_{IoU} , λ_{edge} , and λ_{move} are simply set to balance these loss terms.

E. Implementations

We adopt ResNet18 [40] and PointNet⁺⁺ [39] as our 2D and 3D encoders. To obtain the local features with different sizes of the receptive field, the 2D feature sampling is conducted on the feature maps from ResBlock 2, 3, and 4, and the 3D feature sampling is on the feature groups from the grouping layer 1, 2, and 3. Final pooling layers and fully connected layers in both are not used. The resolution of the input images is 224×224 . We set the batch size as 12, the total training epoch as 250, and the learning rate as 1×10^{-3} with the decay of 0.3 at the 100-th, 200-th, and 230-th epoch. The values of hyperparameters used in the overall loss are set as $\lambda_{normal} = 10^{-3}$, $\lambda_{edge} = 10^{-1}$, $\lambda_{move} = 10^{-3}$, $\lambda_{IoU} = 10^{-2}$ without any elaborated adjustments. The multi-view silhouette loss \mathcal{L}_{IoU} is only applied on the final mesh after 100 epochs due to time cost.

IV. RESULTS

Dataset. We demonstrate the effectiveness of our model on the ShapeNetCore v1.0 [26] dataset and the Pix3D [27] dataset. The **ShapeNet** dataset contains 55 shape classes and we only take 16 classes with relatively large number of objects, and split them into the base classes and the novel classes. The base classes include *car*, *chair*, *monitor*, *plane*, *rifle*, *speaker*, *table* and *telephone*, and the novel classes include *bench*, *bus*, *cabinet*, *lamp*, *pistol*, *sofa*, *train* and *watercraft*. To reduce the class imbalance, we only randomly sample 200 shapes from each class for testing. Among the 16 classes, the RGB images of 13 classes are provided by [1], and the images of the remaining 3 classes are rendered by ourselves using Blender with the consistent rendering setting. The **Pix3D** dataset only contains 9 classes, 10069 real-world images and 395 unique 3D models. Among the 9 classes, 2 classes, *i.e.*, *tools* and *misc*, only consists of 47 and 68 images. Splitting the remaining 7 classes into base classes and novel classes and training on them is not suitable for our zero-shot setting due to the diversity and size. Thus, we directly test the generalization of our ZeroMesh and other methods on the Pix3D [27] dataset by using the model trained on the ShapeNet [26] without any further training or refinement. Due to the large domain gaps, we randomly sample 100 images from the 7 classes *bed*, *bookcase*, *chair*, *desk*, *sofa*, *table*, *wardrobe* and take them all as novel classes for testing.

Evaluation criteria. The reconstruction performance is evaluated by five criteria, *i.e.*, Chamfer Distance (CD), Earth Mover Distance (EMD), Multi-view Silhouette Intersection over Union (IoU), Normal Consistency (NC) and F-score. Specifically, the CD and EMD are distance metrics between point sets [6]. The F-score is defined as the harmonic mean between the precision and the recall, based on if a prediction/ground-truth point can find any other ground-truth/prediction point within the threshold $\tau = 10^{-3}$. The NC measures the cosine error of the prediction normal from its ground-truth [13], [15], [16]. To calculate the point-based metrics, we first randomly sample 2500 points and normals from each generated surface and its ground-truth, respectively, and then measure the distance upon the two sampled point

TABLE I

PERFORMANCE COMPARISON ON SHAPENET UNDER CHAMFER DISTANCE (CD), EARTH MOVER DISTANCE (EMD), MULTI-VIEW SILHOUETTE IOU (IOU), NORMAL CONSISTENCY (NC), AND F-SCORE (F). BEST RESULTS ARE BOLD. THE SYMBOLS v AND o REPRESENT THE OBJECT-CENTERED AND THE VIEWER-CENTERED COORDINATES, RESPECTIVELY.

Methods	Base Classes					Novel Classes				
	CD↓ $\times 10^{-3}$	EMD↓ $\times 10^{-2}$	IoU↑ $\times 10^{-2}$	NC↓ $\times 10^{-1}$	F↑ $\times 10^{-2}$	CD↓ $\times 10^{-3}$	EMD↓ $\times 10^{-2}$	IoU↑ $\times 10^{-2}$	NC↓ $\times 10^{-1}$	F↑ $\times 10^{-2}$
AtlasNet [12]	5.35	7.56	78.63	4.72	52.84	23.91	11.60	67.17	7.14	39.65
Pixel2Mesh [13]	7.55	8.87	77.14	5.30	45.71	11.81	11.51	72.73	6.52	39.06
OccNet [2]	8.63	5.09	73.68	4.41	41.00	40.48	10.59	65.31	6.60	29.32
Mesh R-CNN [16]	6.04	4.68	77.25	6.50	47.46	8.84	5.60	74.10	7.50	40.50
TMNet [14]	6.07	5.85	78.81	4.68	53.12	32.79	11.40	66.23	7.20	37.51
SDFNet [19]	13.22	7.41	73.73	4.79	34.37	23.20	9.34	68.51	6.20	25.09
Ours (v)	4.18	5.54	83.91	4.67	57.23	6.71	6.16	79.85	5.96	47.61
Ours (o)	3.96	5.36	85.85	4.51	59.57	6.69	5.96	81.50	5.80	50.09

TABLE II

QUANTITATIVE COMPARISON ON PIX3D, WITH THE MODEL TRAINED ON SHAPENET, UNDER CHAMFER DISTANCE (CD), EARTH MOVER DISTANCE (EMD), MULTI-VIEW SILHOUETTE IOU (IOU), NORMAL CONSISTENCY (NC), AND F-SCORE (F). THE SYMBOLS v AND o REPRESENT THE VIEWER-CENTERED AND THE OBJECT-CENTERED COORDINATES.

Methods	CD↓ $\times 10^{-3}$	EMD↓ $\times 10$	IoU↑ $\times 10^{-2}$	NC↓ $\times 10$	F↑ $\times 10^{-2}$
AtlasNet [12]	59.75	21.41	52.14	8.72	15.79
Pixel2Mesh [13]	87.44	20.79	46.49	9.16	8.96
OccNet [2]	57.37	14.05	55.32	7.51	14.32
Mesh R-CNN [16]	58.74	16.84	46.95	9.27	12.97
TMNet [14]	74.64	20.38	46.08	8.49	12.34
SDFNet [19]	99.82	21.28	49.58	8.98	6.49
Ours (v)	30.00	16.22	52.38	7.43	21.24
Ours (o)	34.87	13.64	52.03	7.62	19.33

sets. Since neither CD nor EMD takes into account the surface/mesh connectivity, our proposed IoU loss is used as another evaluation criterion to further account for mesh quality. Specifically, we render the output meshes into silhouettes (binary masks) from different camera views and calculate their mean Intersection over Union (IoU) broadly-utilized in segmentation.

Viewer-centered vs Object-centered. There are two choices of coordinate system: **viewer-centered** [13], [16], [19] and **object-centered** [12], [14] coordinates according to pose of reconstruction. The viewer-centered reconstruction improves the generalization on unseen objects but leads to scale-depth ambiguity. The object-centered reconstruction is more friendly to downstream tasks, *e.g.*, analysis, editing, rendering, and generalize to new domains (different camera setting) but relies on camera pose to obtain 2D local features. Accordingly, we provide two variants of our frameworks under both coordinates, respectively. Our viewer-centered variant takes *azimuth* = 0 and *elevation* = 0 as camera pose and reconstructs objects aligned with input images. We evaluate all methods with canonical pose by rotation, translation and normalization using ground-truth pose to eliminate depth-scale error.

A. Quantitative Results.

We evaluate our model and compare it to five state-of-the-art methods for single-view 3D reconstruction, *i.e.*, Pixel2Mesh [13], Mesh R-CNN [16], AtlasNet [12], Topology Modification Network (TMNet) [14] and SDFNet [19], while the first four methods are mesh-based reconstruction and

SDFNet is a signed distance function method addressing novel classes. Under both the normal and the zero-shot settings, we use the samples from the base classes for training, whilst the tests are conducted on the seen base classes under the normal setting and on the unseen novel classes under the zero-shot setting. We re-run the shared codes for the four mesh-based methods for a fair comparison, and employ the released SDFNet-Img code [19] for SDFNet test. Note that Pixel2Mesh [13], Mesh R-CNN [16] and SDFNet [19] use the viewer-centered coordinates. To align their output with other methods, we transform all their results into object-centered coordinates by using camera pose.

As illustrated by Tab. I, our approach using either of the two coordinates consistently outperforms the state-of-the-art methods on all metrics under both settings, and such advantage is especially pronounced on novel classes. Although Mesh R-CNN outperforms Pixel2Mesh quantitatively, we noticed that the reconstructed surface from Mesh R-CNN is not as smooth as Pixel2Mesh as proved by NC. Specifically, on novel classes, compared with the second best performer Mesh R-CNN, our ZeroMesh significantly decreases CD from 8.84×10^{-3} to 6.69×10^{-3} , NC from 7.50×10^{-1} to 5.80×10^{-1} , and increases IoU from 74.10×10^{-2} to 81.50×10^{-2} , F-score from 40.50×10^{-2} to 50.09×10^{-2} . It is noted that on seen base classes, AtlasNet and TMNet perform better than the early work Pixel2Mesh, possibly due to their elaborated mesh deformation mechanism. However, these two methods need extensive training on a large amount of data, which means they heavily rely on priors, and are only able to process objects with similar structures, leading to their poor

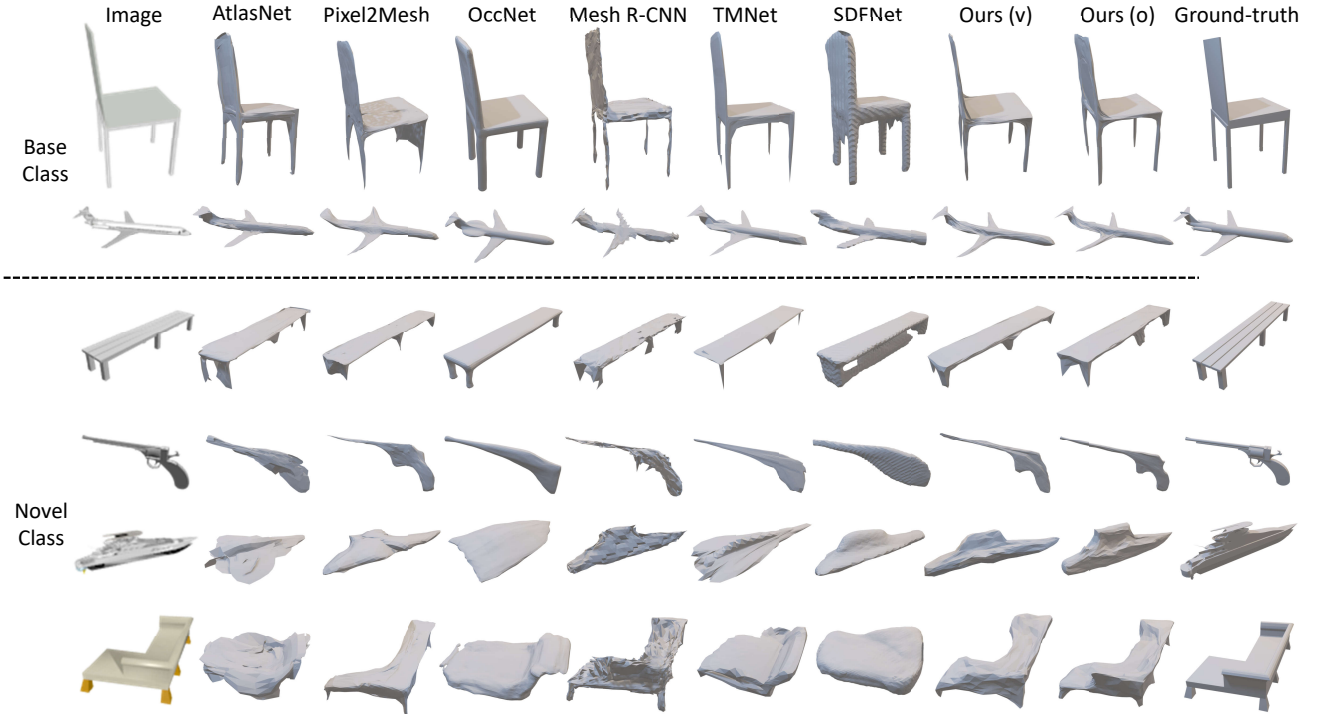


Fig. 5. Six visual examples (in rows) from ShapeNet [26]. The 1st and 2nd rows are results on base classes, while the 3rd-6th rows are results on novel classes. The 1st and 9th columns show the input images and ground-truth shapes, and the 2nd-8th columns show the shapes reconstructed by Pixel2Mesh [13], Mesh R-CNN [16], AtlasNet [12], TMNet [14], SDFNet [19], and Ours, respectively.

performance on unseen novel classes. We dive into this result and argue that: 1) elaborated mechanisms usually negatively affect generalization; 2) the global feature based reconstruction may function by recognition instead of generation, as pointed out in [15], leading to significant performance decline on data out of distribution, *i.e.*, novel classes. The performance of SDFNet [19] also proves the weaknesses of global features, which is only slightly better than AtlasNet and TMNet on novel classes, although it is designed for generalized 3D reconstruction. In contrast, Pixel2Mesh, Mesh R-CNN and our ZeroMesh take local features as the guidance of deformation to emphasize detailed geometry information. This strategy enforces the network to pay more attention to the shape structure. Also, some local structures are shared across objects, *e.g.*, legs of sofas and tables, tires of cars and buses, and triggers of rifles and pistols, so the model could adapt well to unseen objects. Note that compared with Pixel2Mesh using only 2D local features, our ZeroMesh uses both 2D and 3D local features, and the benefit of this strategy could be observed from Tab. III in the ablation study. Moreover, our intermediate point cloud representation and multi-view silhouette loss also guarantee the generalization and robustness on novel classes, so that our ZeroMesh outperforms existing methods with a significant margin under both coordinates.

Tab. II compares the model generalization from the synthetic ShapeNet dataset to the real Pix3D dataset. Although the large domain gap leads to serious performance decline, our approach still outperforms the state-of-the-art methods on all metrics, especially under CD our approach outperforms the second best methods Mesh R-CNN nearly 24×10^{-3} . Such generalization

ability verified the merits of our proposed contributions.

B. Qualitative Results.

Fig. 5 shows six visual examples of reconstruction results from ZeroMesh and previous models, Pixel2Mesh [13], Mesh R-CNN [16], AtlasNet [12], TMNet [14] and SDFNet [19]. For the base classes (the top two rows in Fig. 5), all methods in comparison are able to generate accurate and smooth shapes. Our ZeroMesh can achieve on-par performance with competitive works. For novel classes (the 3rd to 6th rows in Fig. 5), if the target object is similar to the base classes (the 3rd and 4th rows), *e.g.*, the bench and pistol similar to the chair and rifle, the outputs of AtlasNet and TMNet are not too bad but with visible reconstruction differences. If the input objects are very different from the base classes (the 5th and 6th rows), both AtlasNet and TMNet totally fail and even misidentify the watercraft as the plane (the 5th row), implying they may function by recognition. Only Pixel2Mesh, Mesh R-CNN, SDFNet and our ZeroMesh can generate reasonable results, while ZeroMesh achieves the most faithful reconstruction. Note that although Mesh R-CNN outperforms Pixel2Mesh quantitatively, the former shows uneven surface and leads to worse visual results possibly due to its relatively weak mesh-based supervision.

Furthermore, to qualitatively evaluate the generalization ability on the real images, we test our ZeroMesh and other methods on the Pix3D dataset [27] by using the model trained on the ShapeNet [26] without any further training or refinement. This is a challenging task as the images from Pix3D have remarkable differences with those from ShapeNet on

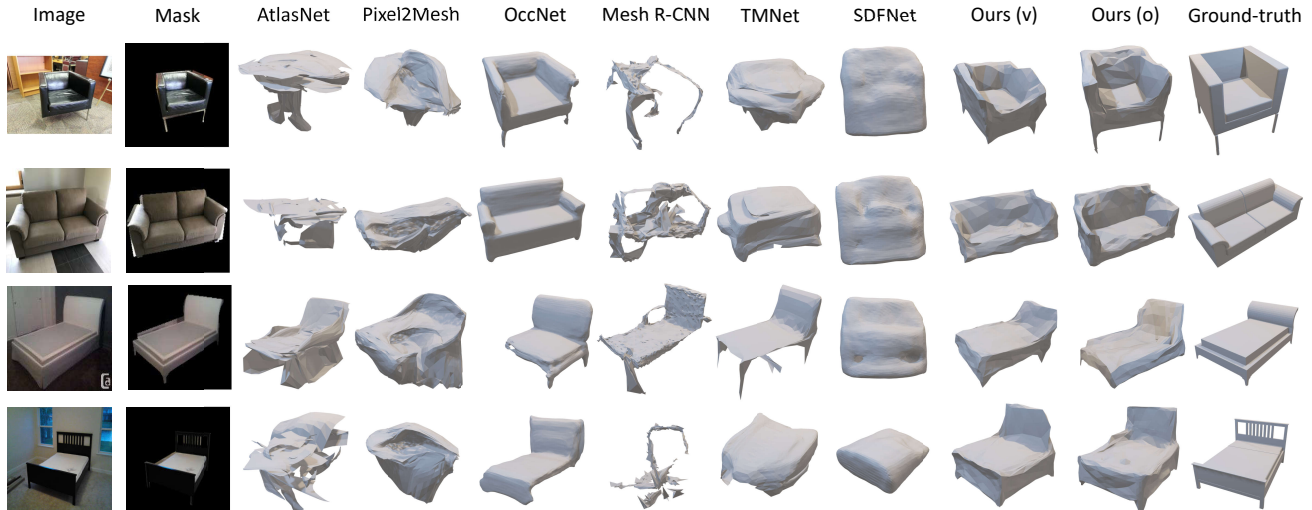


Fig. 6. Visual reconstruction examples on real images from Pix3D [27] dataset. The objects in the images are segmented by masks provided by Pix3D. All models are trained on ShapeNet and directly applied on Pix3D without any further training or refinement.

both the style and the category. Fig. 6 shows the results from ours and the comparing methods, confirming the superiority of our method. ZeroMesh is still able to reconstruct a variety of objects on real images faithfully, while most of the comparing methods fail in such a case.

V. DISCUSSION

A. Ablation Study

We conduct ablative studies to single out the improvement from each of our claimed contributions. Firstly, to demonstrate the local features on generalization, we compare the performance of our model by using either local features or global features. Secondly, we re-design a one-stage version of our proposed network to demonstrate the benefits from the intermediate point cloud representation on 3D mesh reconstruction. Thirdly, we compare our model performance with or without the proposed multi-view silhouette loss. Finally, we provide a vanilla baseline “Pixel2Mesh⁺” to verify the improvement collectively. This is a modified Pixel2Mesh model, where the graph convolution used in Pixel2Mesh is replaced by MLPs and better training recipe is applied, consistent with our ZeroMesh. Compared with our “Full model”, “Pixel2Mesh⁺” only employs the 2D local features like Pixel2Mesh, without any of other proposed components.

Local vs Global. To use global features, we pool the feature maps or feature groups from the 2D or 3D encoders into feature vectors, and concatenate them with vertex coordinates for the following deformation, same as AtlasNet and TMNet. The results show that global features yield obviously poor performance on novel classes, which is comparable to AtlasNet and TMNet in Tab. III, verifying the necessity of sampling local features. The dramatic performance decline of ZeroMesh without 2D local features is due to the unreasonable intermediate representation, which invalidates the 3D feature sampling at the second stage and further degrades the local feature model into the global feature model.

Two-stage vs one-stage. To verify the two-stage architecture, we develop a one-stage variant of our model without using the

intermediate point cloud generation. In particular, we sample the local features from the image plane according to the vertex coordinates of the mesh template, and concatenate the local features with the vertex coordinates to get the per-vertex feature representation. Lacking the intermediate point clouds, we cannot sample 3D features. By comparing “one-stage” and “Full Model” in Tab. III, the performance of the one-stage model significantly decreases on all criteria, verifying the advantages of the proposed two-stage reconstruction for model generalization.

Multi-view Silhouette Loss. We investigate the effect of our proposed multi-view silhouette loss, by comparing our full framework with a variant without using the IoU loss. By comparing “w/o IoU loss” with “Full Model”, and “Pixel2Mesh⁺” with “one-stage” in Tab. III, the employment of IoU loss for training not only yields improvement on the IoU metric for evaluation, but also reduces CD and EMD, on both the base and the novel classes.

Comparison with the baseline “Pixel2Mesh⁺”. As mentioned, “Pixel2Mesh⁺” is our vanilla baseline, similar to the original Pixel2Mesh [13]. With the Multi-view Silhouette Loss, 3D local features and intermediate point cloud generation, our “Full Model” outperforms “Pixel2Mesh⁺” under all evaluation protocols, further demonstrating the merits of our collective design.

B. Sensitivity to Object Scales.

To investigate the effect of object scales to reconstruction quality, we re-render the same objects with same parameters except depth. Specifically, we scale the depth with 1.01, 1.05, and 1.10, respectively. We evaluate our model on the three test sets directly without any further training to verify its robustness to object scales. The quantitative comparison in Tab. IV shows that novel classes are more sensitive to depth variation than base classes, while our method remains reasonably robust to the scale of 2D objects in the images.

TABLE III
 ABLATION STUDY ABOUT THE EFFECTS OF DIFFERENT MODULES ON SHAPENET. *Pixel2Mesh*⁺ IS OUR BASELINE MODIFIED FROM PIXEL2MESH. ALL MODELS RECONSTRUCT UNDER OBJECT-CENTERED COORDINATES.

Methods	Base Classes					Novel Classes				
	CD↓ ×10 ⁻³	EMD↓ ×10 ⁻²	IoU↑ ×10 ⁻²	NC↓ ×10 ⁻¹	F↑ ×10 ⁻²	CD↓ ×10 ⁻³	EMD↓ ×10 ⁻²	IoU↑ ×10 ⁻²	NC↓ ×10 ⁻¹	F↑ ×10 ⁻²
<i>Pixel2Mesh</i> ⁺	4.47	5.69	82.69	4.70	56.24	8.01	5.96	78.04	6.14	45.87
w/o 2D local	6.58	5.29	79.68	4.99	50.97	37.25	10.71	66.87	7.24	35.48
w/o 3D local	4.09	5.49	85.44	4.55	58.71	7.07	6.00	80.80	5.93	48.38
one-stage	4.34	5.58	84.64	4.62	57.23	7.91	6.24	79.45	6.03	46.67
w/o IoU loss	4.11	5.59	83.77	4.58	58.62	6.94	6.06	80.06	5.88	49.33
Full Model	3.96	5.36	85.85	4.51	59.57	6.69	5.96	81.50	5.80	50.09

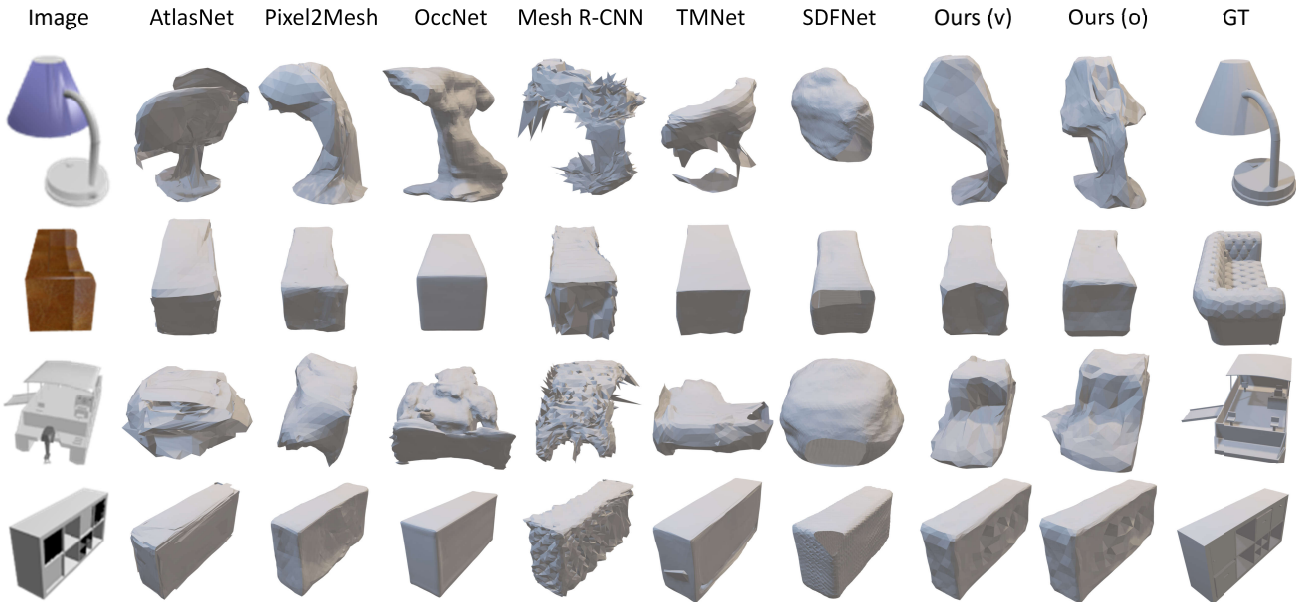


Fig. 7. Four failure examples from ShapeNet [26]. Both existing methods and our ZeroMesh occasionally fail to reconstruct thin lamp pole (Row 1), blur boundaries (Row 2), and complex topology (Row 3 and Row 4).

TABLE IV
 SENSITIVITY TO OBJECT SCALES ON SHAPENET [26].

Methods		CD↓	EMD↓	IoU↑	NC↓	F-score↑
Base	original	3.96	5.36	85.85	4.51	59.57
	depth +1%	3.95	5.36	85.85	4.51	59.72
	depth +5%	3.99	5.36	85.83	4.51	59.59
	depth +10%	4.02	5.35	85.75	4.51	59.52
Novel	original	6.69	5.96	81.50	5.80	50.09
	depth +1%	6.75	5.92	81.46	5.80	50.03
	depth +5%	6.81	5.95	81.46	5.79	50.00
	depth +10%	7.04	6.03	81.30	5.91	49.74

C. Limitation

Examples of failure cases are provided in Fig. 7. As shown, existing methods and our ZeroMesh may occasionally fail on thin structures (e.g., lamp pole), blur boundaries (e.g., sofa armrest), and complex topology (e.g., boat, cabinet), which are common challenges for the community. First, the failure on thin structures is mainly caused by two reasons: 1) point number imbalance in training the loss function, and 2) deformation difficulty. These can be mitigated by raising the importance of points from thin structures to force the model to pay more attention to thin structures during deformation.

Second, the blur boundaries are introduced by the datasets. It’s an ill-posed problem for single-view 3D reconstruction and can only be mitigated by learned priors. Third, by using the sphere template, our ZeroMesh could keep the surface closed, but this also restricts the deformation of objects with holes (e.g., cabinet) and complex topology (e.g., boat). This issue could be handled by a post-processing step to treat holes specially in the literature [14]. We would like to point out that such a restriction could also be possibly relaxed by using multiple spheres to form the template meshes. This could not only generate more complex structures but also keep the surface closed, and will be explored in our future study. Note that existing methods that could avoid the topology problem theoretically (i.e., AtlasNet [12], OccNet [2], Mesh R-CNN [16], TMNet [14], SDFNet [19]) also fail on these cases.

VI. CONCLUSION

In this paper, we put forward a learning framework, ZeroMesh, to solve zero-shot single-view 3D mesh reconstruction. We propose three strategies to improve the model generalization ability on novel classes and prevent overfitting, namely, learning intermediate point cloud representation, employing

local features, and introducing multi-view silhouette loss for model regularization. Our model demonstrates a promising capacity for cross-category object reconstruction and generalizes to unseen object classes well.

VII. REFERENCES SECTION

REFERENCES

- [1] C. B. Choy, D. Xu, J. Gwak, K. Chen, and S. Savarese, “3d-r2n2: A unified approach for single and multi-view 3d object reconstruction,” in *European conference on computer vision*. Springer, 2016, pp. 628–644.
- [2] L. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger, “Occupancy networks: Learning 3d reconstruction in function space,” in *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [3] S. R. Richter and S. Roth, “Matryoshka networks: Predicting 3d geometry via nested shape layers,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 1936–1944.
- [4] S. Popov, P. Bauszat, and V. Ferrari, “Corenet: Coherent 3d scene reconstruction from a single rgb image,” in *European Conference on Computer Vision*. Springer, 2020, pp. 366–383.
- [5] J. Wu, C. Zhang, T. Xue, B. Freeman, and J. Tenenbaum, “Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling,” *Advances in neural information processing systems*, vol. 29, 2016.
- [6] H. Fan, H. Su, and L. 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*, 2017, pp. 605–613.
- [7] Z. Han, C. Chen, Y.-S. Liu, and M. Zwicker, “DRWR: A differentiable renderer without rendering for unsupervised 3D structure learning from silhouette images,” in *International Conference on Machine Learning*, 2020.
- [8] C. Chen, Z. Han, Y. shen Liu, and M. Zwicker, “Unsupervised learning of fine structure generation for 3d point clouds by 2d projection matching,” in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021.
- [9] Z. Chen and H. Zhang, “Learning implicit fields for generative shape modeling,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 5939–5948.
- [10] S. Saito, Z. Huang, R. Natsume, S. Morishima, H. Li, and A. Kanazawa, “Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization,” *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, Oct 2019. [Online]. Available: <http://dx.doi.org/10.1109/iccv.2019.00239>
- [11] S. Saito, T. Simon, J. Saragih, and H. Joo, “Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, June 2020.
- [12] T. Groueix, M. Fisher, V. G. Kim, B. C. Russell, and M. Aubry, “A papier-mâché approach to learning 3d surface generation,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 216–224.
- [13] N. Wang, Y. Zhang, Z. Li, Y. Fu, W. Liu, and Y.-G. Jiang, “Pixel2mesh: Generating 3d mesh models from single rgb images,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 52–67.
- [14] J. Pan, X. Han, W. Chen, J. Tang, and K. Jia, “Deep mesh reconstruction from single rgb images via topology modification networks,” *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, Oct 2019. [Online]. Available: <http://dx.doi.org/10.1109/iccv.2019.01006>
- [15] M. Tatarchenko, S. R. Richter, R. Ranftl, Z. Li, V. Koltun, and T. Brox, “What do single-view 3d reconstruction networks learn?” *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun 2019. [Online]. Available: <http://dx.doi.org/10.1109/cvpr.2019.00352>
- [16] G. Gkioxari, J. Malik, and J. Johnson, “Mesh r-cnn,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 9785–9795.
- [17] X. Zhang, Z. Zhang, C. Zhang, J. B. Tenenbaum, W. T. Freeman, and J. Wu, “Learning to Reconstruct Shapes From Unseen Classes,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [18] J. Wang and Z. Fang, “Gsir: Generalizable 3d shape interpretation and reconstruction,” in *European Conference on Computer Vision*. Springer, 2020, pp. 498–514.
- [19] A. Thai, S. Stojanov, V. Upadhyaya, and J. M. Rehg, “3d reconstruction of novel object shapes from single images,” in *2021 International Conference on 3D Vision (3DV)*. IEEE, 2021, pp. 85–95.
- [20] B. Wallace and B. Hariharan, “Few-shot generalization for single-image 3d reconstruction via priors,” *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, Oct 2019. [Online]. Available: <http://dx.doi.org/10.1109/iccv.2019.00392>
- [21] M. Michalkiewicz, S. Parisot, S. Tsogkas, M. Baktashmotlagh, A. Eriksson, and E. Belilovsky, “Few-shot single-view 3-d object reconstruction with compositional priors,” in *European Conference on Computer Vision*. Springer, 2020, pp. 614–630.
- [22] N. Ravi, J. Reizenstein, D. Novotny, T. Gordon, W.-Y. Lo, J. Johnson, and G. Gkioxari, “Accelerating 3d deep learning with pytorch3d,” 2020. [Online]. Available: <https://arxiv.org/abs/2007.08501>
- [23] S. Liu, T. Li, W. Chen, and H. Li, “Soft rasterizer: A differentiable renderer for image-based 3d reasoning,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 7708–7717.
- [24] X. Li, S. Liu, K. Kim, S. D. Mello, V. Jampani, M.-H. Yang, and J. Kautz, “Self-supervised single-view 3d reconstruction via semantic consistency,” in *European Conference on Computer Vision*. Springer, 2020, pp. 677–693.
- [25] A. Kanazawa, S. Tulsiani, A. A. Efros, and J. Malik, “Learning category-specific mesh reconstruction from image collections,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 371–386.
- [26] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su et al., “Shapenet: An information-rich 3d model repository,” *arXiv preprint arXiv:1512.03012*, 2015.
- [27] X. Sun, J. Wu, X. Zhang, Z. Zhang, C. Zhang, T. Xue, J. B. Tenenbaum, and W. T. Freeman, “Pix3d: Dataset and methods for single-image 3d shape modeling,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [28] K. Ikeuchi and B. K. Horn, “Numerical shape from shading and occluding boundaries,” *Artif. Intell.*, vol. 17, no. 1–3, p. 141–184, Aug. 1981. [Online]. Available: [https://doi.org/10.1016/0004-3702\(81\)90023-0](https://doi.org/10.1016/0004-3702(81)90023-0)
- [29] B. K. Horn, “Shape from shading: A method for obtaining the shape of a smooth opaque object from one view,” 1970.
- [30] A. P. Pentland, “Local shading analysis,” *IEEE Transactions on Pattern Analysis & Machine Intelligence*, vol. 6, no. 02, pp. 170–187, mar 1984.
- [31] A. Pentland, “Shape information from shading: A theory about human perception,” [1988 Proceedings] *Second International Conference on Computer Vision*, pp. 404–413, 1988.
- [32] H. Kato, Y. Ushiku, and T. Harada, “Neural 3d mesh renderer,” in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3907–3916.
- [33] W. Chen, H. Ling, J. Gao, E. Smith, J. Lehtinen, A. Jacobson, and S. Fidler, “Learning to predict 3d objects with an interpolation-based differentiable renderer,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [34] J. K. Pontes, C. Kong, S. Sridharan, S. Lucey, A. Eriksson, and C. Fookes, “Image2mesh: A learning framework for single image 3d reconstruction,” in *Asian Conference on Computer Vision*. Springer, 2018, pp. 365–381.
- [35] Y. Ye, S. Tulsiani, and A. Gupta, “Shelf-supervised mesh prediction in the wild,” in *Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [36] Y. Nie, X. Han, S. Guo, Y. Zheng, J. Chang, and J. J. Zhang, “Total3dunderstanding: Joint layout, object pose and mesh reconstruction for indoor scenes from a single image,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [37] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *International Conference on Learning Representations*, 2015.
- [38] M. A. Bautista, W. Talbott, S. Zhai, N. Srivastava, and J. M. Susskind, “On the generalization of learning-based 3d reconstruction,” *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, Jan 2021. [Online]. Available: <http://dx.doi.org/10.1109/wacv48630.2021.00223>
- [39] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, “Pointnet++: Deep hierarchical feature learning on point sets in a metric space,” *Advances in neural information processing systems*, vol. 30, 2017.
- [40] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.