

Pixel2Mesh++: Multi-View 3D Mesh Generation via Deformation

Chao Wen^{1*} Yinda Zhang^{2*} Zhuwen Li^{3*} Yanwei Fu^{1†}
¹Fudan University ²Google LLC ³Nuro, Inc.

Abstract

We study the problem of shape generation in 3D mesh representation from a few color images with known camera poses. While many previous works learn to hallucinate the shape directly from priors, we resort to further improving the shape quality by leveraging cross-view information with a graph convolutional network. Instead of building a direct mapping function from images to 3D shape, our model learns to predict series of deformations to improve a coarse shape iteratively. Inspired by traditional multiple view geometry methods, our network samples nearby area around the initial mesh’s vertex locations and reasons an optimal deformation using perceptual feature statistics built from multiple input images. Extensive experiments show that our model produces accurate 3D shape that are not only visually plausible from the input perspectives, but also well aligned to arbitrary viewpoints. With the help of physically driven architecture, our model also exhibits generalization capability across different semantic categories, number of input images, and quality of mesh initialization.

1. Introduction

3D shape generation has become a popular research topic recently. With the astonishing capability of deep learning, lots of works have been demonstrated to successfully generate the 3D shape from merely a single color image. However, due to limited visual evidence from only one viewpoint, single image based approaches usually produce rough geometry in the occluded area and do not perform well when generalized to test cases from domains other than training, e.g. cross semantic categories.

Adding a few more images (e.g. 3-5) of the object is an effective way to provide the shape generation system with more information about the 3D shape. On one hand, multi-view images provide more visual appearance information, and thus grant the system with more chance to build the connection between 3D shape and image priors. On the other hand, it is well-known that traditional multi-view ge-

*indicates equal contributions.

†indicates corresponding author. This work is supported by the STCSM project (19ZR1471800), and Eastern Scholar (TP2017006).

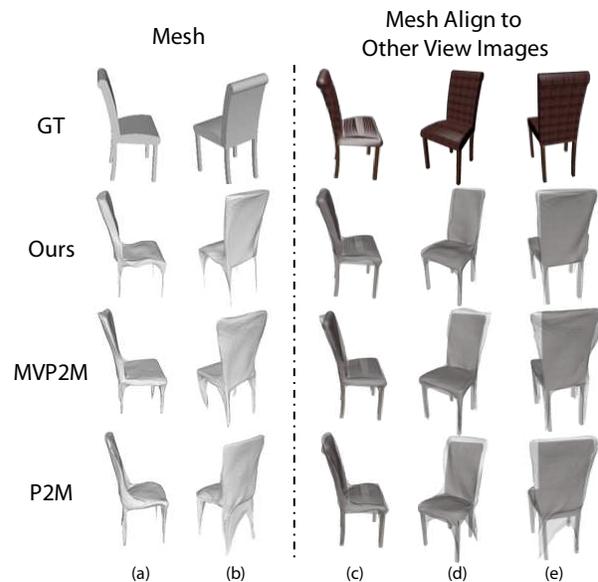


Figure 1. **Multi-View Shape Generation.** From multiple input images, we produce shapes aligning well to input (c and d) and arbitrary random (e) camera viewpoint. Single view based approach, e.g. Pixel2Mesh (P2M) [41], usually generates shape looking good from the input viewpoint (c) but significantly worse from others. Naive extension with multiple views (MVP2M, Sec. 4.2) does not effectively improve the quality.

ometry methods [12] accurately infer 3D shape from correspondences across views, which is analytically well defined and less vulnerable to the generalization problem. However, these methods typically suffer from other problems, like large baselines and poorly textured regions. Though typical multi-view methods are likely to break down with very limited input images (e.g. less than 5), the cross-view connections might be implicitly encoded and learned by a deep model. While well-motivated, there are very few works in the literature exploiting in this direction, and a naive multi-view extension of single image based model does not work well as shown in Fig. 1.

In this work, we propose a deep learning model to generate the object shape from multiple color images. Especially, we focus on endowing the deep model with the capacity of improving shapes using cross-view information. We resort to designing a new network architecture, named Multi-View Deformation Network (MDN), which works in

conjunction with the Graph Convolutional Network (GCN) architecture proposed in Pixel2Mesh [41] to generate accurate 3D geometry shape in the desirable mesh representation. In Pixel2Mesh, a GCN is trained to deform an initial shape to the target using features from a single image, which often produces plausible shapes but lack of accuracy (Fig. 1 P2M). We inherit this characteristic of “generation via deformation” and further deform the mesh in MDN using features carefully pooled from multiple images. Instead of learning to hallucinate via shape priors like in Pixel2Mesh, MDN reasons shapes according to correlations across different views through a physically driven architecture inspired by classic multi-view geometry methods. In particular, MDN proposes hypothesis deformations for each vertex and move it to the optimal location that best explains features pooled from multiple views. By imitating correspondences search rather than learning priors, MDN generalizes well in various aspects, such as cross semantic category, number of input views, and the mesh initialization.

Besides the above-mentioned advantages, MDN is in addition featured with several good properties. First, it can be trained end-to-end. Note that it is non-trivial since MDN searches deformation from hypotheses, which requires a non-differentiable argmax/min. Inspired by [20], we apply a differentiable 3D soft argmax, which takes a weighted sum of the sampled hypotheses as the vertex deformation. Second, it works with varying number of input views in a single forward pass. This requires the feature dimension to be invariant with the number of inputs, which is typically broken when aggregating features from multiple images (e.g. when using concatenation). We achieve the input number invariance by concatenating the statistics (e.g. mean, max, and standard deviation) of the pooled feature, which further maintains input order invariance. We find this statistics feature encoding explicitly provides the network cross-view information, and encourages it to automatically utilize image evidence when more are available. Last but not least, the nature of “generation via deformation” allows an iterative refinement. In particular, the model output can be taken as the input, and quality of the 3D shape is gradually improved throughout iterations. With these desiring features, our model achieves the state-of-the-art performance on ShapeNet for shape generation from multiple images under standard evaluation metrics.

To summarize, we propose a GCN framework that produces 3D shape in mesh representation from a few observations of the object in different viewpoints. The core component is a physically driven architecture that searches optimal deformation to improve a coarse mesh using perceptual feature statistics built from multiple images, which produces accurate 3D shape and generalizes well across different semantic categories, numbers of input images, and the quality of coarse meshes.

2. Related Work

3D Shape Representations Since 3D CNN is readily applicable to 3D volumes, the volume representation has been well-exploited for 3D shape analysis and generation [4, 42]. With the debut of PointNet [30], the point cloud representation has been adopted in many works [7, 29]. Most recently, the mesh representation [19, 41] has become competitive due to its compactness and nice surface properties. Some other representations have been proposed, such as geometry images [33], depth images [36, 31], classification boundaries [26, 3], signed distance function [28], etc., and most of them require post-processing to get the final 3D shape. Consequently, the shape accuracy may vary and the inference take extra time.

Single view shape generation Classic single view shape reasoning can be traced back to shape from shading [6, 45], texture [25], and de-focus [8], which only reason the visible parts of objects. With deep learning, many works leverage the data prior to hallucinate the invisible parts, and directly produce shape in 3D volume [4, 9, 43, 11, 32, 37, 16], point cloud [7], mesh models [19], or as an assembling of shape primitive [40, 27]. Alternatively, 3D shape can be also generated by deforming an initialization, which is more related to our work. Tulsiani *et al.* [39] and Kanazawa *et al.* [17] learn a category-specific 3D deformable model and reasons the shape deformations in different images. Wang *et al.* [41] learn to deform an initial ellipsoid to the desired shape in a coarse to fine fashion. Combining deformation and assembly, Huang *et al.* [14] and Su *et al.* [34] retrieve shape components from a large dataset and deform the assembled shape to fit the observed image. Kuryenkov *et al.* [22] learns free-form deformations to refine shape. Even though with impressive success, most deep models adopt an encoder-decoder framework, and it is arguable if they perform shape generation or shape retrieval [38].

Multi-view shape generation Recovering 3D geometry from multiple views has been well studied. Traditional multi-view stereo (MVS) [12] relies on correspondences built via photo-consistency and thus it is vulnerable to large baselines, occlusions, and texture-less regions. Most recently, deep learning based MVS models have drawn attention, and most of these approaches [44, 13, 15, 46] rely on a cost volume built from depth hypotheses or plane sweeps. However, these approaches usually generate depth maps, and it is non-trivial to fuse a full 3D shape from them. On the other hand, direct multi-view shape generation uses fewer input views with large baselines, which is more challenging and has been less addressed. Choy *et al.* [4] propose a unified framework for single and multi-view object generation reading images sequentially. Kar *et al.* [18] learn a multi-view stereo machine via recurrent feature fusion. Gwak *et al.* [10] learns shapes from multi-view silhouettes

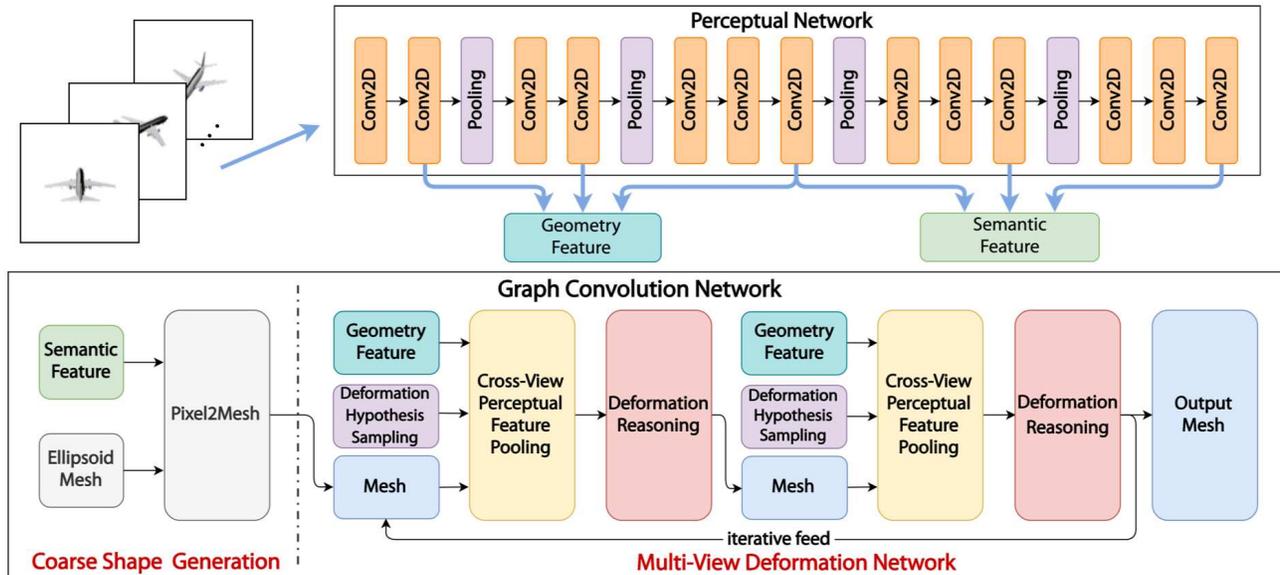


Figure 2. **System Pipeline.** Our whole system consists of a 2D CNN extracting image features and a GCN deforming an ellipsoid to target shape. A coarse shape is generated from Pixel2Mesh and refined iteratively in Multi-View Deformation Network. To leverage cross-view information, our network pools perceptual features from multiple input images for hypothesis locations in the area around each vertex and predicts the optimal deformation.

by ray-tracing pooling and further constrains the ill-posed problem using GAN. Our approach belongs to this category but is fundamentally different from the existing methods. Rather than sequentially feeding in images, our method learns a GCN to deform the mesh using features pooled from all input images at once.

3. Method

Our model receives multiple color images of an object captured from different viewpoints (with known camera poses) and produces a 3D mesh model in the world coordinate. The whole framework adopts the strategy of coarse-to-fine (Fig. 2), in which a plausible but rough shape is generated first, and details are added later. Realizing that existing 3D shape generators usually produce reasonable shape even from a single image, we simply use Pixel2Mesh [41] trained either from single or multiple views to produce the coarse shape, which is taken as input to our Multi-View Deformation Network (MDN) for further improvement. In MDN, each vertex first samples a set of deformation hypotheses from its surrounding area (Fig. 3 (a)). Each hypothesis then pools cross-view perceptual feature from early layers of a perceptual network, where the feature resolution is high and contains more low-level geometry information (Fig. 3 (b)). These features are further leveraged by the network to reason the best deformation to move the vertex. It is worth noting that our MDN can be applied iteratively for multiple times to gradually improve shapes.

3.1. Multi-View Deformation Network

In this section, we introduce Multi-View Deformation Network, which is the core of our system to enable the network exploiting cross-view information for shape generation. It first generates deformation hypotheses for each vertex and learns to reason an optimum using feature pooled from inputs. Our model is essentially a GCN, and can be jointly trained with other GCN based models like Pixel2Mesh. We refer reader to [1, 21] for details about GCN, and Pixel2Mesh [41] for graph residual block which will be used in our model.

3.1.1 Deformation Hypothesis Sampling

The first step is to propose deformation hypotheses for each vertex. This is equivalent as sampling a set of target locations in 3D space where the vertex can be possibly moved to. To uniformly explore the nearby area, we sample from a level-1 icosahedron centered on the vertex with a scale of 0.02, which results in 42 hypothesis positions (Fig. 3 (a), left). We then build a local graph with edges on the icosahedron surface and additional edges between the hypotheses to the vertex in the center, which forms a graph with 43 nodes and $120 + 42 = 162$ edges. Such local graph is built for all the vertices, and then fed into a GCN to predict vertex movements (Fig. 3 (a), right).

3.1.2 Cross-View Perceptual Feature Pooling

The second step is to assign each node (in the local GCN) features from the multiple input color images. Inspired by

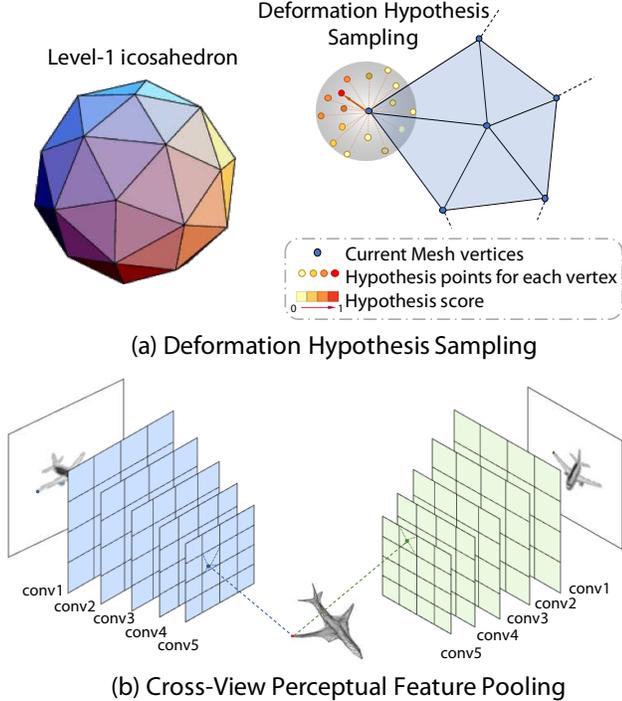


Figure 3. **Deformation Hypothesis and Perceptual Feature Pooling.** (a) Deformation Hypothesis Sampling. We sample 42 deformation hypotheses from a level-1 icosahedron and build a GCN among hypotheses and the vertex. (b) Cross-View Perceptual Feature Pooling. The 3D vertex coordinates are projected to multiple 2D image planes using camera intrinsics and extrinsics. Perceptual features are pooled using bilinear interpolation, and feature statistics are kept on each hypothesis.

Pixel2Mesh, we use the prevalent VGG-16 architecture to extract perceptual features. Since we assume known camera poses, each vertex and hypothesis can find their projections in all input color image planes using known camera intrinsics and extrinsics and pool features from four neighboring feature blocks using bilinear interpolation (Fig. 3 (b)). Different from Pixel2Mesh where high level features from later layers of the VGG (i.e. ‘conv3_3’, ‘conv4_3’, and ‘conv5_3’) are pooled to better learn shape priors, MDN pools features from early layers (i.e. ‘conv1_2’, ‘conv2_2’, and ‘conv3_3’), which are in high spatial resolution and considered maintaining more detailed information.

To combine multiple features, concatenation has been widely used as a loss-less way, however ends up with total dimension changing with respect to (w.r.t.) the number of input images. Statistics feature has been proposed for multi-view shape recognition [35] to handle this problem. Inspired by this, we concatenate some statistics (*mean*, *max*, and *std*) of the features pooled from all views for each vertex, which makes our network naturally adaptive to variable input views and behave invariant to different input orders. This also encourages the network to learn from cross-view

feature correlations rather than each individual feature vector. In addition to image features, we also concatenate the 3-dimensional vertex coordinate into the feature vector. In total, we compute for each vertex and hypothesis a 1347 dimension feature vector.

3.1.3 Deformation Reasoning

The next step is to reason an optimal deformation for each vertex from the hypotheses using pooled cross-view perceptual features. Note that picking the best hypothesis of all needs an argmax operation, which requires stochastic optimization and usually is not optimal. Instead, we design a differentiable network component to produce desirable deformation through soft-argmax of the 3D deformation hypotheses, which is illustrated in Fig. 4. Specifically, we first feed the cross-view perceptual feature P into a scoring network, consisting of 6 graph residual convolution layers [41] plus ReLU, to predict a scalar weight c_i for each hypothesis. All the weights are then fed into a softmax layer and normalized to scores s_i , with $\sum_{i=1}^{43} s_i = 1$. The vertex location is then updated as the weighted sum of all the hypotheses, i.e. $v = \sum_{i=1}^{43} s_i * h_i$, where h_i is the location of each deformation hypothesis including the vertex itself. This deformation reasoning unit runs on all local GCN built upon every vertex with shared weights, as we expect all the vertices leveraging multi-view feature in a similar fashion.

3.2. Loss

We train our model fully supervised using ground truth 3D CAD models. Our loss function includes all terms from Pixel2Mesh, but extends the Chamfer distance loss to a re-sampled version. Chamfer distance measures “distance” between two point clouds, which can be problematic when points are not uniformly distributed on the surface. We propose to randomly re-sample the predicted mesh when calculating Chamfer loss using the re-parameterization trick proposed in Ladický *et al.* [23]. Specifically, given a triangle defined by 3 vertices $\{v_1, v_2, v_3\} \in \mathbb{R}^3$, a uniform sampling can be achieved by:

$$s = (1 - \sqrt{r_1}) v_1 + (1 - r_2) \sqrt{r_1} v_2 + \sqrt{r_1} r_2 v_3,$$

where s is a point inside the triangle, and $r_1, r_2 \sim U[0, 1]$. Knowing this, when calculating the loss, we uniformly sample our generated mesh for 4000 points, with the number of points per triangle proportional to its area. We find this is empirically sufficient to produce a uniform sampling on our output mesh with 2466 vertices, and calculating Chamfer loss on the re-sampled point cloud, containing 6466 in total, helps to remove artifacts in the results.

3.3. Implementation Details

For initialization, we use Pixel2Mesh to generate a coarse shape with 2466 vertices. To improve the quality of

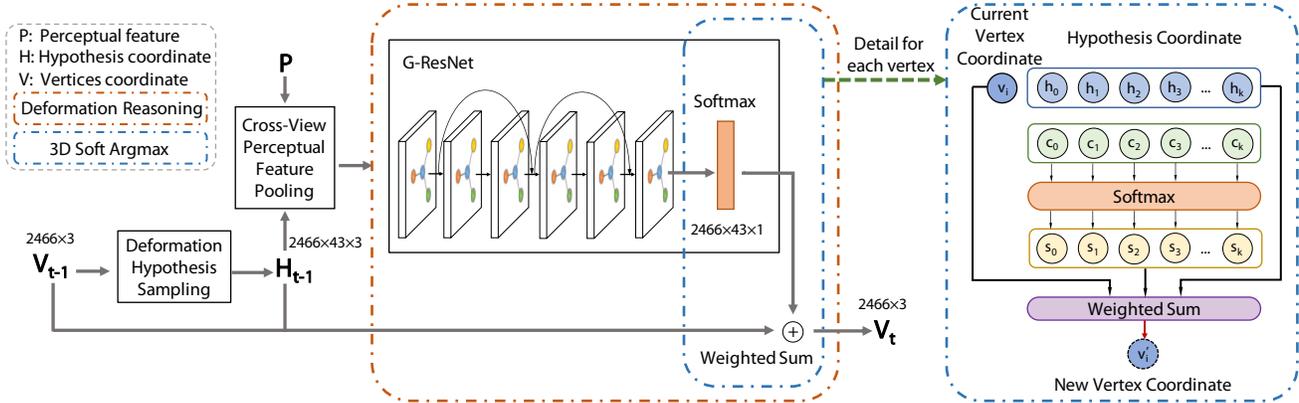


Figure 4. **Deformation Reasoning.** The goal is to reason a good deformation from the hypotheses and pooled features. We first estimate a weight (green circle) for each hypothesis using a GCN. The weights are normalized by a softmax layer (yellow circle), and the output deformation is the weighted sum of all the deformation hypotheses.

initial mesh, we equip the Pixel2Mesh with our cross-view perceptual feature pooling layer, which allows it to extract features from multiple views.

The network is implemented in Tensorflow and optimized using Adam with weight decay as $1e-5$ and mini-batch size as 1. The model is trained for 50 epochs in total. For the first 30 epochs, we only train the multi-view Pixel2Mesh for initialization with learning rate $1e-5$. Then, we make the whole model trainable, including the VGG for perceptual feature extraction, for another 20 epoch with the learning rate as $1e-6$. The whole model is trained on NVIDIA Titan Xp for 96 hours. During training, we randomly pick three images for a mesh as input. During testing, it takes 0.32s to generate a mesh.

4. Experiments

In this section, we perform extensive evaluation of our model for multi-view shape generation. We compare to state-of-the-art methods, as well as conduct controlled experiments w.r.t. various aspects, *e.g.* cross category generalization, quantity of inputs, etc.

4.1. Experimental setup

Dataset We adopt the dataset provided by Choy *et al.* [4] as it is widely used by many existing 3D shape generation works. The dataset is created using a subset of ShapeNet[2] containing 50k 3D CAD models from 13 categories. Each model is rendered from 24 randomly chosen camera view-points, and the camera intrinsic and extrinsic parameters are given. For fair comparison, we use the same training/testing split as in Choy *et al.* [4] for all our experiments.

Evaluation Metric We use standard evaluation metrics for 3D shape generation. Following Fan *et al.* [7], we calculate Chamfer Distance(CD) between points clouds uni-

formly sampled from the ground truth and our prediction to measure the surface accuracy. We also use F-score following Wang *et al.* [41] to measure the completeness and precision of generated shapes. For CD, the smaller is better. For F-score, the larger is better.

4.2. Comparison to Multi-view Shape Generation

We compare to previous works for multi-view shape generation and show effectiveness of MDN in improving shape quality. While most shape generation methods take only a single image, we find Choy *et al.* [4] and Kar *et al.* [18] work in the same setting with us. We also build two competitive baselines using Pixel2Mesh. In the first baseline (Tab.1, P2M-M), we directly run single-view Pixel2Mesh on each of the input image and fuse multiple results [5, 24]. In the second baseline (Tab.1, MVP2M), we replace the perceptual feature pooling to our cross-view version to enable Pixel2Mesh for the multi-view scenario (more details in supplementary materials).

Tab. 1 shows quantitative comparison in F-score. As can be seen, our baselines already outperform other methods, which shows the advantage of mesh representation in finding surface details. Moreover, directly equipping Pixel2Mesh with multi-view features does not improve too much (even slightly worse than the average of multiple runs of single-view Pixel2Mesh), which shows dedicate architecture is required to efficiently learn from multi-view features. In contrast, our Multi-View Deformation Network significantly further improves the results from the MVP2M baseline (*i.e.* our coarse shape initialization).

More qualitative results are shown in Fig. 8. We show results from different methods aligned with one input view (left) and a random view (right). As can be seen, Choy *et al.* [4] (3D-R2N2) and Kar *et al.* [18] (LSM) produce 3D volume, which lose thin structures and surface details. Pixel2Mesh (P2M) produces mesh models but shows obvi-

Category	F-score(τ) \uparrow					F-score(2τ) \uparrow				
	3DR2N2 \dagger	LSM	MVP2M	P2M-M	Ours	3DR2N2 \dagger	LSM	MVP2M	P2M-M	Ours
Couch	45.47	43.02	53.17	53.70	57.56	59.97	55.49	73.24	72.04	75.33
Cabinet	54.08	50.80	56.85	63.55	65.72	64.42	60.72	76.58	79.93	81.57
Bench	44.56	49.33	60.37	61.14	66.24	62.47	65.92	75.69	75.66	79.67
Chair	37.62	48.55	54.19	55.89	62.05	54.26	64.95	72.36	72.36	77.68
Monitor	36.33	43.65	53.41	54.50	60.00	48.65	56.33	70.63	70.51	75.42
Firearm	55.72	56.14	79.67	74.85	80.74	76.79	73.89	89.08	84.82	89.29
Speaker	41.48	45.21	48.90	51.61	54.88	52.29	56.65	68.29	68.53	71.46
Lamp	32.25	45.58	50.82	51.00	62.56	49.38	64.76	65.72	64.72	74.00
Cellphone	58.09	60.11	66.07	70.88	74.36	69.66	71.39	82.31	84.09	86.16
Plane	47.81	55.60	75.16	72.36	76.79	70.49	76.39	86.38	82.74	86.62
Table	48.78	48.61	65.95	67.89	71.89	62.67	62.22	79.96	81.04	84.19
Car	59.86	51.91	67.27	67.29	68.45	78.31	68.20	84.64	84.39	85.19
Watercraft	40.72	47.96	61.85	57.72	62.99	63.59	66.95	77.49	72.96	77.32
Mean	46.37	49.73	61.05	61.72	66.48	62.53	64.91	77.10	76.45	80.30

Table 1. **Comparison to Multi-view Shape Generation Methods.** We show F-score on each semantic category. Our model significantly outperforms previous methods, i.e. 3DR2N2 [4] and LSM [18], and competitive baselines derived from Pixel2Mesh [41]. Please see supplementary materials for Chamfer Distance. The notation \dagger indicates the methods which does not require camera extrinsics.

ous artifacts when visualized in viewpoint other than the input. In comparison, our results contain better surface details and more accurate geometry learned from multiple views.

4.3. Generalization Capability

Our MDN is inspired by multi-view geometry methods, where 3D location is reasoned via cross-view information. In this section, we investigate the generalization capability of MDN in many aspects to improve the initialization mesh. For all the experiments in this section, we fix the coarse stage and train/test MDN under different settings.

4.3.1 Semantic Category

We first verify how our network generalizes across semantic categories. We fix the initial MVP2M and train MDN with 12 out 13 categories and test on the one left out, and the improvements upon initialization are shown in Fig. 5 (a). As can be seen, the performance is only slightly lower when the testing category is removed from the training set compared to the model trained using all categories. To make it more challenging, we also train MDN on only one category and test on all the others. Surprisingly, MDN still generalizes well between most of the categories as shown in Fig. 5 (b). Strong generalizing categories (e.g. chair, table, lamp) tend to have relatively complex geometry, thus the model has better chance to learn from cross-view information. On the other hand, categories with super simple geometry (e.g. speaker, cellphone) do not help to improve other categories, even not for themselves. On the whole, MDN shows good generalization capability across semantic categories.

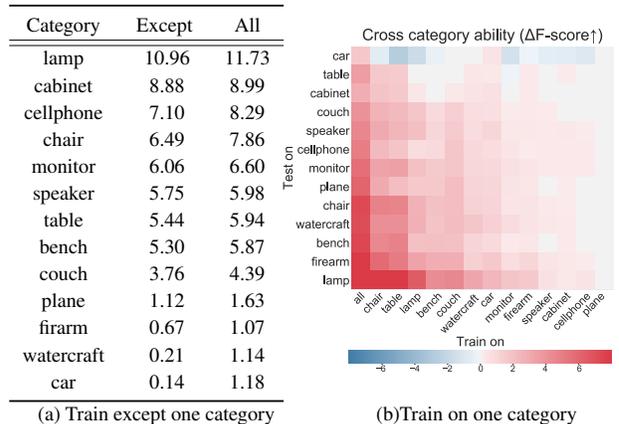


Figure 5. **Cross-Category Generalization.** (a) MDN trained on 12 out of 13 categories and tested on the one left. (b) MDN trained on 1 category and tested on the other. Each block represents the experiment with MDN trained on horizontal category and tested on vertical category. Both (a) and (b) show improvements of F-score(τ) upon MVP2M through MDN.

4.3.2 Number of Views

We then test how MDN performs w.r.t. the number of input views. In Tab. 2, we see that MDN consistently performs better when more input views are available, even though the number of view is fixed as 3 for efficiency during the training. This indicates that features from multiple views are well encoded in the statistics, and MDN is able to exploit additional information when seeing more images. For reference, we train five MDNs with the input view number

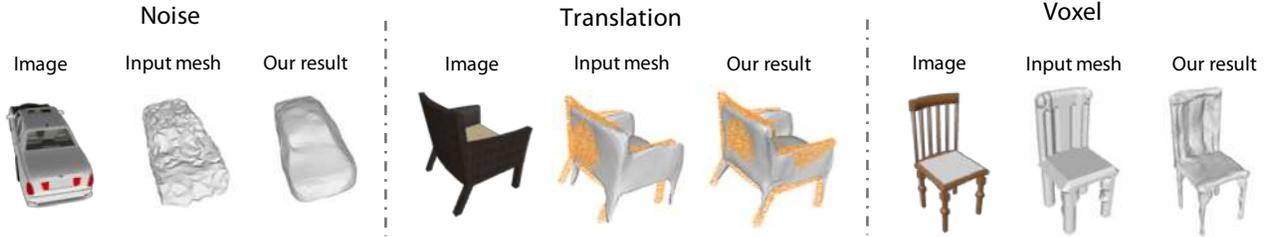


Figure 6. **Robustness to Initialization.** Our model is robust to added noise, shift, and input mesh from other sources.

fixed at 2 to 5 respectively. As shown in Tab. 2 “Resp.,” the 3-view MDN performs very close to models trained with more views (e.g. 4 and 5), which shows the model learns efficiently from fewer number of views during the training. The 3-view MDN also outperform models trained with less views (e.g. 2), which indicates additional information provided during the training can be effectively activated during the test even when observation is limited. Overall, MDN generalizes well to different number of inputs.

#train	#test	2	3	4	5
3	F-score(τ) \uparrow	64.48	66.44	67.66	68.29
	F-score(2τ) \uparrow	78.74	80.33	81.36	81.97
	CD \downarrow	0.515	0.484	0.468	0.459
Resp.	F-score(τ) \uparrow	64.11	66.44	68.54	68.82
	F-score(2τ) \uparrow	78.34	80.33	81.56	81.99
	CD \downarrow	0.527	0.484	0.467	0.452

Table 2. **Performance w.r.t. Input View Numbers.** Our MDN performs consistently better when more view is given, even trained using only 3 views.

4.3.3 Initialization

Lastly, we test if the model overfits to the input initialization, i.e. the MVP2M. To this end, we add translation and random noise to the rough shape from MVP2M. We also take as input the mesh converted from 3DR2N2 using marching cube [24]. As shown in Fig. 6, MDN successfully removes the noise, aligns the input with ground truth, and adds significant geometry details. This shows that MDN is tolerant to input variance.

4.4. Ablation Study

In this section, we verify the qualitative and quantitative improvements from statistic feature pooling, re-sampled Chamfer distance, and iterative refinement.

4.4.1 Statistical Feature

We first check the importance of using feature statistics. We train MDN with the ordinary concatenation. This maintains

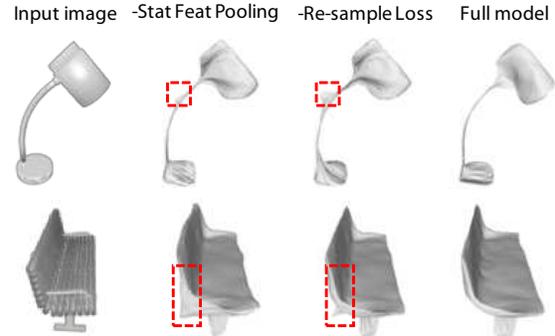


Figure 7. **Qualitative Ablation Study.** We show meshes from the MDN with statistics feature or re-sampling loss disabled.

Metrics	F-score(τ) \uparrow	F-score(2τ) \uparrow	CD \downarrow
-Feat Stat	65.26	79.13	0.511
-Re-sample Loss	66.26	80.04	0.496
Full Model	66.48	80.30	0.486

Table 3. **Quantitative Ablation Study.** We show the metrics of the MDN with statistics feature or re-sampling loss disabled.

all the features loss-less to potentially produce better geometry, but does not support variable number of inputs any more. Surprisingly, our model with feature statistics (Tab. 3, “Full Model”) still outperforms the one with concatenation (Tab. 3, “-Feat Stat”). This is probably because that our feature statistics is invariant to the input order, such that the network learns more efficiently during the training. It also explicitly encodes cross-view feature correlations, which can be directly leveraged by the network.

4.4.2 Re-sampled Chamfer Distance

We then investigate the impact of the re-sampled Chamfer loss. We train our model using the traditional Chamfer loss only on mesh vertices as defined in Pixel2Mesh, and all metrics drop consistently (Tab. 3, “-Re-sample Loss”). Intuitively, our re-sampling loss is especially helpful for places with sparse vertices and irregular faces, such as the elongated lamp neck as shown in Fig. 7, 3rd column. It also

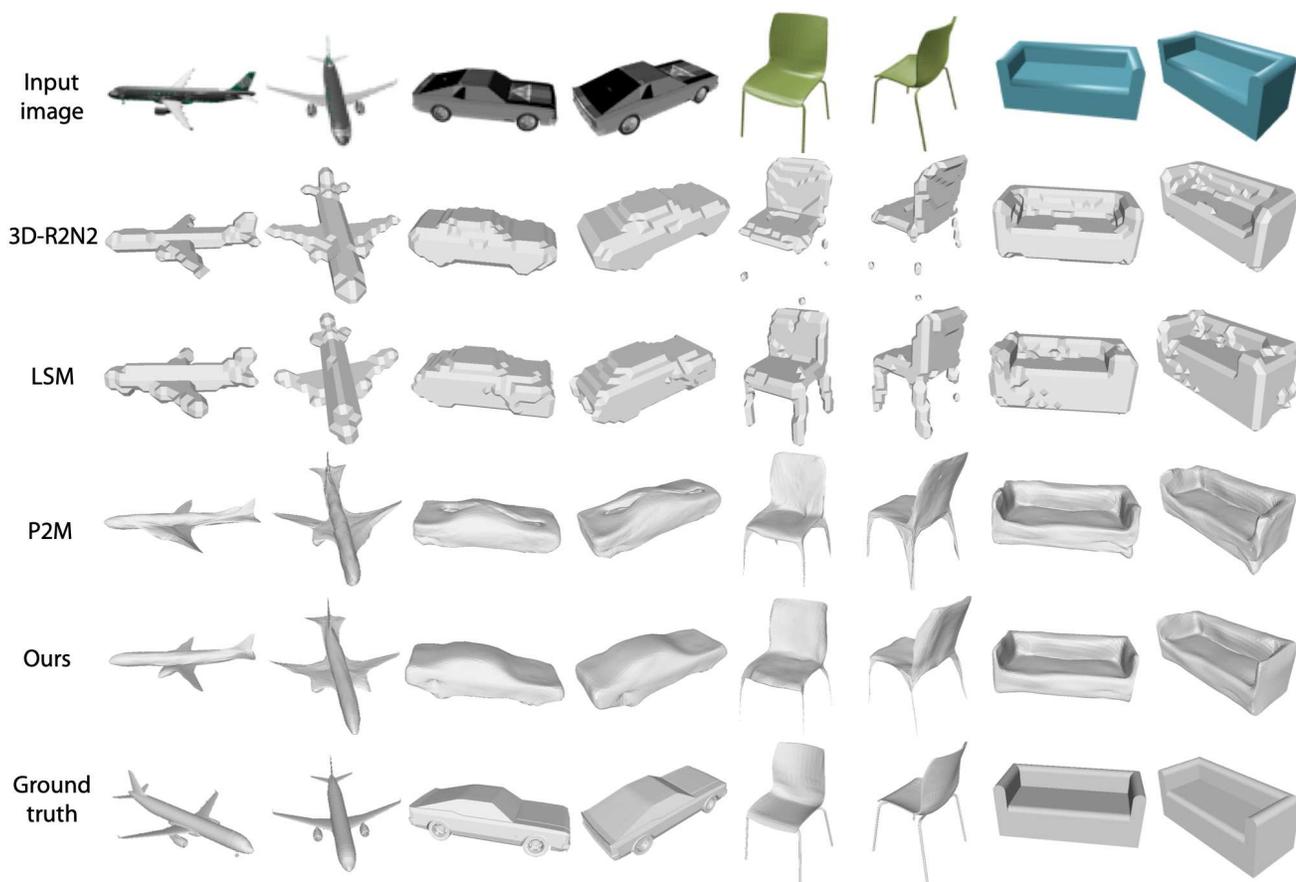


Figure 8. **Qualitative Evaluation.** From top to bottom, we show in each row: two camera views, results of 3DR2N2, LSM, multi-view Pixel2Mesh, ours, and the ground truth. Our predicts maintain good details and align well with different camera views. Please see supplementary materials for more results.

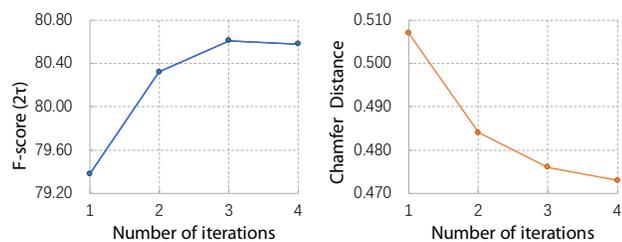


Figure 9. **Performance with Different Iterations.** The performance keeps improving with more iterations and saturate at three.

prevents big mistakes from happening on a single vertex, e.g. the spike on bench, where our loss penalizes a lot of sampled points on wrong faces caused by the vertex but the standard Chamfer loss only penalizes one point.

4.4.3 Number of Iteration

Figure 9 shows that the performance of our model keeps improving with more iterations, and is roughly saturated at three. Therefore we choose to run three iterations during the inference even though marginal improvements can be further obtained from more iterations.

5. Conclusion

We propose a graph convolutional framework to produce 3D mesh model from multiple images. Our model learns to exploit cross-view information and generates vertex deformation iteratively to improve the mesh produced from the direct prediction methods, e.g. Pixel2Mesh and its multi-view extension. Inspired by multi-view geometry methods, our model searches in the nearby area around each vertex for an optimal place to relocate it. Compared to previous works, our model achieves the state-of-the-art performance, produces shapes containing accurate surface details rather than merely visually plausible from input views, and shows good generalization capability in many aspects. For future work, combining with efficient shape retrieval for initialization, integrating with multi-view stereo models for explicit photometric consistency, and extending to scene scales are some of the practical directions to explore. On a high level, how to integrating the similar idea in emerging new representations, such as part based model with shape basis and learned function [28] are interesting for further study.

References

- [1] Michael M. Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: Going beyond euclidean data. *IEEE Signal Process. Mag.*, 34(4):18–42, 2017. 3
- [2] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 5
- [3] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *CVPR*, pages 5939–5948, 2019. 2
- [4] Christopher B Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *ECCV*, 2016. 2, 5, 6
- [5] Brian Curless and Marc Levoy. A volumetric method for building complex models from range images. 1996. 5
- [6] Jean-Denis Durou, Maurizio Falcone, and Manuela Sagona. Numerical methods for shape-from-shading: A new survey with benchmarks. *Computer Vision and Image Understanding*, 109(1):22–43, 2008. 2
- [7] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3d object reconstruction from a single image. In *CVPR*, 2017. 2, 5
- [8] Paolo Favaro and Stefano Soatto. A geometric approach to shape from defocus. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(3):406–417, 2005. 2
- [9] Rohit Girdhar, David F. Fouhey, Mikel Rodriguez, and Abhinav Gupta. Learning a predictable and generative vector representation for objects. In *ECCV*, 2016. 2
- [10] JunYoung Gwak, Christopher B. Choy, Manmohan Chandraker, Animesh Garg, and Silvio Savarese. Weakly supervised 3d reconstruction with adversarial constraint. In *3DV*, 2017. 2
- [11] Christian Hane, Shubham Tulsiani, and Jitendra Malik. Hierarchical surface prediction for 3d object reconstruction. In *3DV*, 2017. 2
- [12] Andrew Hartley and Andrew Zisserman. *Multiple view geometry in computer vision (2. ed.)*. Cambridge University Press, 2006. 1, 2
- [13] Po-Han Huang, Kevin Matzen, Johannes Kopf, Narendra Ahuja, and Jia-Bin Huang. Deepmvs: Learning multi-view stereopsis. In *CVPR*, 2018. 2
- [14] Qixing Huang, Hai Wang, and Vladlen Koltun. Single-view reconstruction via joint analysis of image and shape collections. *ACM Trans. Graph.*, 34(4):87:1–87:10, 2015. 2
- [15] Sunghoon Im, Hae-Gon Jeon, Stephen Lin, and In So Kweon. Dpsnet: End-to-end deep plane sweep stereo. In *ICLR*, 2018. 2
- [16] Adrian Johnston, Ravi Garg, Gustavo Carneiro, and Ian D. Reid. Scaling cnns for high resolution volumetric reconstruction from a single image. In *ICCV*, 2017. 2
- [17] Angjoo Kanazawa, Shubham Tulsiani, Alexei A. Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *ECCV*, 2018. 2
- [18] Abhishek Kar, Christian Häne, and Jitendra Malik. Learning a multi-view stereo machine. In *NeurIPS*, pages 365–376, 2017. 2, 5, 6
- [19] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3d mesh renderer. In *CVPR*, 2018. 2
- [20] Alex Kendall, Hayk Martirosyan, Saumitro Dasgupta, and Peter Henry. End-to-end learning of geometry and context for deep stereo regression. In *ICCV*, 2017. 2
- [21] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2016. 3
- [22] Andrey Kurenkov, Jingwei Ji, Animesh Garg, Viraj Mehta, JunYoung Gwak, Christopher Choy, and Silvio Savarese. Deformnet: Free-form deformation network for 3d shape reconstruction from a single image. In *WACV*, pages 858–866, 2018. 2
- [23] Lubor Ladicky, Olivier Saurer, SoHyeon Jeong, Fabio Maninchedda, and Marc Pollefeys. From point clouds to mesh using regression. In *ICCV*, pages 3893–3902, 2017. 4
- [24] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *SIGGRAPH*, 1987. 5, 7
- [25] Constantinos Marinos and Andrew Blake. Shape from texture: the homogeneity hypothesis. In *ICCV*, pages 350–353, 1990. 2
- [26] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *CVPR*, pages 4460–4470, 2019. 2
- [27] Chengjie Niu, Jun Li, and Kai Xu. Im2struct: Recovering 3d shape structure from a single RGB image. In *CVPR*, 2018. 2
- [28] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *CVPR*, June 2019. 2, 8
- [29] Charles R. Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J. Guibas. Frustum pointnets for 3d object detection from RGB-D data. In *CVPR*, 2018. 2
- [30] Charles Ruizhongtai Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017. 2
- [31] Stephan R. Richter and Stefan Roth. Matryoshka networks: Predicting 3d geometry via nested shape layers. In *CVPR*, 2018. 2
- [32] Gernot Riegler, Ali Osman Ulusoy, and Andreas Geiger. Octnet: Learning deep 3d representations at high resolutions. In *CVPR*, 2017. 2
- [33] Ayan Sinha, Asim Unmesh, Qixing Huang, and Karthik Ramani. Surfnet: Generating 3d shape surfaces using deep residual networks. In *CVPR*, 2017. 2
- [34] Hao Su, Qixing Huang, Niloy J. Mitra, Yangyan Li, and Leonidas J. Guibas. Estimating image depth using shape collections. *ACM Trans. Graph.*, 33(4):37:1–37:11, 2014. 2
- [35] Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik G. Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *ICCV*, 2015. 4

- [36] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Multi-view 3d models from single images with a convolutional network. In *ECCV*, 2016. 2
- [37] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Octree generating networks: Efficient convolutional architectures for high-resolution 3d outputs. In *ICCV*, 2017. 2
- [38] Maxim Tatarchenko, Stephan Richter, René Ranftl, Zhuwen Li, Vladlen Koltun, and Thomas Brox. What do single-view 3d reconstruction networks learn? In *CVPR*, 2019. 2
- [39] Shubham Tulsiani, Abhishek Kar, João Carreira, and Jitendra Malik. Learning category-specific deformable 3d models for object reconstruction. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(4):719–731, 2017. 2
- [40] Shubham Tulsiani, Hao Su, Leonidas J. Guibas, Alexei A. Efros, and Jitendra Malik. Learning shape abstractions by assembling volumetric primitives. In *CVPR*, 2017. 2
- [41] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *ECCV*, 2018. 1, 2, 3, 4, 5, 6
- [42] Peng-Shuai Wang, Yang Liu, Yu-Xiao Guo, Chun-Yu Sun, and Xin Tong. O-cnn: Octree-based convolutional neural networks for 3d shape analysis. *ACM Transactions on Graphics (TOG)*, 36(4):72, 2017. 2
- [43] Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Josh Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In *Neurips*, 2016. 2
- [44] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In *ECCV*, 2018. 2
- [45] Ruo Zhang, Ping-Sing Tsai, James Edwin Cryer, and Mubarak Shah. Shape from shading: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.*, 21(8):690–706, 1999. 2
- [46] Yinda Zhang, Sameh Khamis, Christoph Rhemann, Julien Valentin, Adarsh Kowdle, Vladimir Tankovich, Michael Schoenberg, Shahram Izadi, Thomas Funkhouser, and Sean Fanello. Activestereonet: End-to-end self-supervised learning for active stereo systems. In *ECCV*, pages 784–801, 2018. 2