

# Render4Completion: Synthesizing Multi-View Depth Maps for 3D Shape Completion

Tao Hu, Zhizhong Han, Abhinav Shrivastava, Matthias Zwicker  
University of Maryland, College Park

taohu@cs.umd.edu, h312h@umd.edu, abhinav@cs.umd.edu, zwicker@cs.umd.edu

## Abstract

*We propose a novel approach for 3D shape completion by synthesizing multi-view depth maps. While previous work for shape completion relies on volumetric representations, meshes, or point clouds, we propose to use multi-view depth maps from a set of fixed viewing angles as our shape representation. This allows us to be free of the memory limitations of volumetric representations and point clouds by casting shape completion into an image-to-image translation problem. Specifically, we render depth maps of the incomplete shape from a fixed set of viewpoints, and perform depth map completion in each view. Different from image-to-image translation networks that process each view separately, our novel multi-view completion net (MVCN) leverages information from all views of a 3D shape to help the completion of each single view. This enables MVCN to leverage more information from different depth views to achieve high accuracy in single depth view completion, and improve the consistency among the completed depth images in different views. Benefiting from the multi-view representation and novel network structure, MVCN significantly improves the accuracy of 3D shape completion in large-scale benchmarks compared to the state of the art.*

## 1. Introduction

Shape completion is an important challenge in 3D shape analysis, serving as a building block in applications such as 3D scanning in robotics, autonomous driving, or 3D modeling and fabrication. While learning-based methods that leverage large shape databases have achieved significant advances recently, choosing a suitable 3D representation for such tasks remains a difficult problem. On the one hand, volumetric approaches such as binary voxel grids or distance functions have the advantage that convolutional neural networks can readily be applied, but including a third dimension increases the memory requirements and limits the resolutions that can be handled. On the other hand,

point-based techniques provide a more parsimonious shape representation, and recently there has been much progress in generalizing convolutional networks to such irregularly sampled data. However, most generative techniques for 3D point clouds involve fully connected layers that limit the number of points and level of shape detail that can be obtained [1, 4, 31].

In this paper, we propose to use a shape representation that is based on multi-view depth maps for shape completion. The representation consists of a fixed number of depth images taken from a set of pre-determined viewpoints. Each pixel is a 3D point, and the union of points over all depth images yields the 3D point cloud of a shape. This has the advantage that we can use several recent advances in neural networks that operate on images, like U-Net [24] and 2D convolutional networks. In addition, the number of points is not fixed and the point density can easily be increased by using higher resolution depth images, or more viewpoints.

Here we leverage this representation for shape completion. Our key idea to perform shape completion is to render multiple depth images of an incomplete shape from a set of pre-defined viewpoints, and then to complete each depth map using image-to-image translation networks. To improve the completion accuracy, we further propose a novel multi-view completion net (MVCN) that leverages information from all depth views of a 3D shape to achieve high accuracy for single depth view completion. In summary, our contributions are as follows:

- A strategy to address shape completion by re-rendering multi-view depth maps to represent the incomplete shape, and performing image translation of these rendered views.
- A multi-view completion architecture that leverages information from all rendered views and outperforms separate depth image completion for each view.
- More accurate 3D shape completion results than previous state of the art methods.

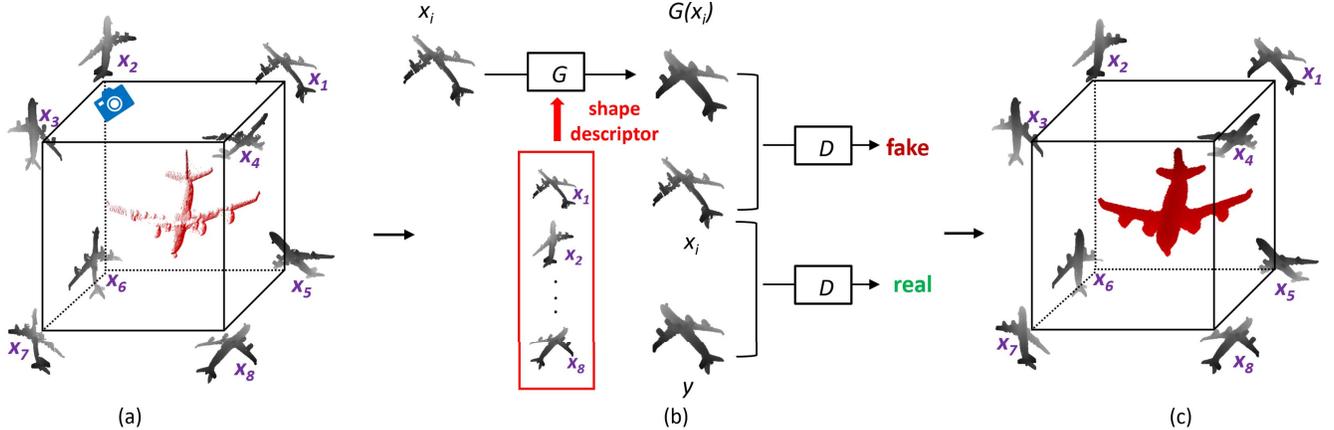


Figure 1: Overview of our approach. (a) We render 8 depth maps of an incomplete shape (shown in red) from 8 viewpoints on the corners of a cube; (b) These rendered 8 depth maps are passed through a multi-view completion net including an adversarial loss, which generates 8 completed depth maps; (c) We back-project the 8 depth maps into a completed 3D model.

## 2. Related Work

**Deep learning on 3D shapes.** Pioneering work on deep learning for 3D shapes has relied on volumetric representations [16, 30], which allow the straightforward application of convolutional neural networks. To avoid the computation and memory costs of 3D convolutions and 3D voxel grids, multi-view convolutional neural networks have also been proposed early for shape analysis [22, 27] such as recognition and classification. But these techniques cannot address shape completion. In addition to volumetric and multi-view representations, point clouds have also been popular for deep learning on 3D shapes. Groundbreaking work in this area includes PointNet and its extension [21, 23].

**3D shape completion.** Shape completion can be performed using volumetric grids, as proposed by Dai et al. [3] and Han et al. [6], which are convenient for CNNs, like 3D-Encoder-Predictor CNNs for [3] and encoder-decoder CNN for patch-level geometry refinement in [6]. However, when represented with volumetric grids, data size grows cubically as the size of the space increases, which severely limits resolution and application. To address this problem, point based shape completion methods were presented, like [1, 31, 32]. The point completion network (PCN) [32] is the state-of-the-art approach that extends the PointNet architecture [21] to provide an encoder, followed by a multi-stage decoder that uses both fully connected [1] and folding layers [31]. They show that their decoder leads to better results than using a fully connected [1] or folding based [31] decoder separately. However, for these voxels or points based shape completion methods, the numbers of input and output voxels or points are still fixed. For example, the input should be voxelized on a  $32^3$  grid [6] and the output point cloud size is 2048 [31], however, which can lead to loss of detail in many scenarios.

**3D reconstruction from images.** The problem of 3D shape

reconstruction from single RGB images shares similarities with 3D shape completion, but is arguably even harder. While a complete survey of these techniques is beyond the scope of this paper, our work shares some similarities with the approach by Lin et al. [15]. They use a multi-view depth map representation for shape reconstruction from single RGB images using a differentiable renderer. In contrast to their technique, we address shape completion, and our approach allows us to solve the problem directly using image-to-image translation. Soltani et al. [26] do shape synthesis and reconstruction from multi-view depth images, which are generated by a variational autoencoder [14]. However, they do not consider the relations between the multi-view depth images of the same model in their generative net.

**Image translation and completion.** A key advantage of our approach is that it allows us to leverage powerful image-to-image translation architectures to address the shape completion problem, including techniques based on generative adversarial networks (GAN) [5], and U-Net structures [24]. Based on conditional GANs, image-to-image translation networks can be applied on a variety of tasks [11]. Satoshi et al. [9] and Portenier et al. [20] propose to use conditional GANs for image completion or editing. However, each image is completed individually in their networks. We propose a network that can combine information from other related images to help the completion of one single image.

## 3. Method

### 3.1. Multi-view Representation

As discussed above, high resolution completion is difficult to achieve by existing methods that operate on voxels or point clouds due to memory limitations or fully connected network structures. In contrast, multi-view representations of 3D shapes [22, 27, 7, 8] can circumvent these obstacles to achieve high resolution and dense completion. As

shown in Fig. 1 (a), given an incomplete point cloud, our method starts from rendering 8 depth maps for this point cloud. Specifically, the renderings are generated by placing 8 virtual cameras at the 8 vertices of a cube enclosing the shape, all pointing towards the centroid of the shape. We also render 8 depth maps from the ground truth point cloud, and then we use these image pairs to train our network.

With this multi-view representation, the shape completion problem can be formulated as image-to-image translation, i.e., translating an incomplete depth map to a corresponding complete depth map, for which we can take full advantage of several recent advances in net structures that operate successfully on images, like U-Net architectures and GANs. After the completion net shown in Fig. 1(b), we get 8 completed depth maps in Fig. 1(c), which can be back-projected into a completed point cloud.

### 3.2. Multi-view Depth Maps Completion

In the completion problem, we learn a mapping from an incomplete depth map  $x_i$  to a completed depth map  $G(x_i)$ , where  $x_i$  is rendered from a partial 3D shape  $S$ ,  $i \in [1, V]$ . We render  $V$  views for each shape and expect to complete each depth map  $x_i$  of  $S$  as similar as possible to the corresponding depth map  $y_i$  of the ground truth 3D shape  $S_1$ .

Although completing each of the  $V$  depth maps of a 3D shape separately would be straightforward, there are two disadvantages. First, we cannot encourage consistency among the completed depth maps from the same 3D shape, which affects the accuracy of the resulting 3D shapes obtained by back-projecting the completed depth maps. Second, we cannot leverage information from other depth maps of the same 3D shape while completing one single depth map. This limits the accuracy of completing a single depth image, since views of the same 3D model share some common information that could be exploited, like global shape and local parts as seen from different viewpoints.

To resolve these issues, we propose a multi-view completion net (MVCN) architecture to complete one single depth image by jointly considering the global 3D shape information. In order to complete a depth image  $x_i$  as similar as possible to the ground truth  $y_i$  in terms of both low-frequency correctness and high-frequency structure, MVCN is designed based on a conditional GAN [5], which is formed by an image-to-image net  $G$  and a discriminator  $D$ . In addition, we introduce a shape descriptor  $d$  for each 3D shape  $S$  to contribute to the completion of each depth image  $x_i$  from  $S$ , where  $d$  holds global information of shape  $S$ . The shape descriptor  $d$  is learned along with the other parameters in MVCN, and it is updated dynamically with the completion of all the depth images  $x_i$  of shape  $S$ .

### 3.3. MVCN Architecture

We use a U-Net based structure [24] as our image-to-image net  $G$ , which has shown its effectiveness over encoder-decoder nets in many tasks including image-to-image translation [11]. Including our shape descriptor, we propose an end-to-end architecture as illustrated in Fig. 2.

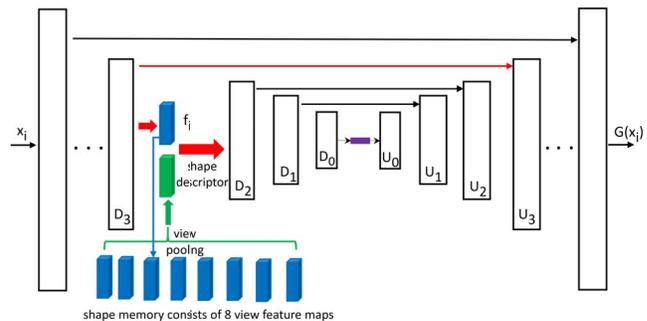


Figure 2: Architecture of MVCN. The shape descriptor represents the information of a 3D shape, which contributes to the completion of each single depth map from the 3D shape.

We adopt the generator and discriminator architecture of [11]. MVCN consists of 8 U-Net modules with an input resolution of  $256 \times 256$ , and each U-Net module has two submodules, DOWN and UP. DOWN (e.g.  $D_3$ ) consists of the Convolution-BatchNorm-ReLU layers [10, 18], and UP (e.g.  $U_3$ ) consists of UpReLU-UpConv-UpNorm layers. More details can be found in [11].

In MVCN, DOWN modules are used to extract a view feature  $f_i$  of each depth image  $x_i$ . For each 3D shape  $S$ , we learn a shape descriptor  $d$  by aggregating all  $V$  view features  $f_i$  through a view-pooling layer. Since not all the features are necessary to represent the global shape, we use max pooling to extract the maximum activation in each dimension of all  $f_i$  to form the shape descriptor, as illustrated in Fig. 2. In addition, the shape descriptor  $d$  is applied to contribute to the completion of each depth image  $x_i$ .

Specifically, for an input  $x_i$  we employ the output of DOWN module  $D_3$  as the view feature  $f_i$ , and insert the view-pooling layer after  $D_3$ . For each shape  $S$  we use a shape memory to store all its  $V$  view features  $f_i$  as shown in Fig. 2. When we get  $f_i$ , we first use it to update the corresponding feature map in shape memory. For example, if  $i = 3$ , the third feature map in shape memory will be replaced with  $f_3$ . Then we obtain the shape descriptor of  $S$  in the current iteration by a view-pooling layer (max pooling all feature maps in the shape memory of  $S$ ). This strategy dynamically keeps the best view features in all training iterations, as illustrated in Fig. 2. Subsequently, we use shape descriptor  $d$  to contribute to the completion of depth map  $x_i$  by concatenating  $d$  with view feature  $f_i$  as the input of module  $D_2$ . This concatenated feature is also forwarded to module  $U_3$  via a skip connection.

### 3.4. Loss Function

The objective of our conditional GAN is similar to image-to-image translation [11],

$$\mathcal{L}_{cGAN}(G, D) = \mathbb{E}_{x,y}[\log D(x, y)] + \mathbb{E}_x[\log(1 - D(x, G(x)))]. \quad (1)$$

In our completion problem, we expect the completion net (G) could not only deceive the discriminator but also produce a completion result near the ground truth. Hence we combine the GAN objective with a traditional pixel-wise loss, such as L1 or L2 distance, which is consistent with previous approaches [11, 19]. Since L1 is less prone to blurring than L2, and considering Eq. 4, there is a linear mapping from a pixel in a depth image to a 3D point, we want to push the generated image to be near the ground truth in L1 sense rather than L2. Therefore, the loss of the completion net is

$$\mathcal{L}_{L1}(G) = \mathbb{E}_{x,y}[\|y - G(x)\|_1]. \quad (2)$$

Our final object in training is then

$$G^* = \arg \min_G \max_D \mathcal{L}_{cGAN}(G, D) + \lambda \mathcal{L}_{L1}(G), \quad (3)$$

where  $\lambda$  is a balance parameter that controls the contributions of the two terms.

### 3.5. Optimization and Inference

Unlike some approaches that focus on image generation [25], our method does not generate images from noise, which also makes our training stable, as mentioned in [9]. Similar to pix2pix [11], we only provide noise in the form of dropout in our network.

To optimize our net, we follow the standard approach [5, 11]. The training of  $D$  and  $G$  is alternated, one gradient descent step on  $D$ , then one step on  $G$ . Minibatch SGD and the Adam solver [13] are applied, with a learning rate of 0.0006 for  $G$  and 0.000006 for  $D$ , which slows down the rate at which  $D$  learns relative to  $G$ . Momentum parameters are  $\beta_1 = 0.5$ ,  $\beta_2 = 0.999$ , and the batch size is 32.

During inference, we first run MVCN with all the 8 rendered views of an incomplete 3D shape to build the shape memory and extract the shape descriptor. Then we run the net again for the second time to complete each view leveraging the learned shape descriptor.

Our final target is 3D shape completion. Given a generated depth image  $G(x_i)$ , for each pixel  $p$  at location  $(x_p, y_p)$  with depth value  $d_p$ , we can back-project  $p$  to a 3D point  $P$  through an inverse perspective transformation,

$$P = R^{-1}(K^{-1}[x_p, y_p, d_p]^T - t), \quad (4)$$

where  $K$ ,  $R$ , and  $t$  are the camera intrinsic matrix, rotation matrix, and translation vector respectively. Note that  $K$ ,  $R$ ,

and  $t$  are always known since these are the parameters of the 8 virtual cameras placed on the corners of a cube. The final shape is the union of the completed, back-projected point clouds from all 8 virtual views.

## 4. Experiments

In this section, we first describe the creation of a multi-category dataset to train our model, and then we illustrate the effectiveness of our method and the improvement of MVCN over a single view completion net (VCN) used as a baseline, where each view is completed individually without shape descriptor. Finally, we analyze the performance of our method, and make comparisons with existing methods. By default, we conduct the training of MVCN under the MVCN-Airplane600 (trained with the first 600 shapes of airplane in ShapeNet [2]), and test it under the same 150 models involved in [32]).

### 4.1. Data Generation and Evaluation Metrics

We use synthetic CAD models from ShapeNet to create a dataset to train our model. Specifically, we take models from 8 categories: airplane, cabinet, car, chair, lamp, sofa, table, and vessel. Our inputs are partial point clouds. For each model, we extract one partial point cloud by back-projecting a 2.5D depth map (from a random viewpoint) into 3D, and render this partial point cloud into  $V = 8$  depth maps of resolution  $256 \times 256$  as training samples. The reason why we use back-projected depth maps as partial point clouds instead of subsets of the complete point cloud is that our training samples are closer to real-world sensor data in this way. In addition, similar to other works, we choose to use a synthetic dataset to generate training data because it contains detailed 3D shapes, which are not available in real-world datasets. In the same way, we also render  $V = 8$  depth maps from the ground truth point clouds as the ground truth depth maps. More details of rendering and back-projecting depth maps are in the supplementary.

Similar to [32], here we also use the symmetric version of Chamfer Distance (CD) [4] to calculate the average closest point distance between the ground truth shape and the completed shape.

### 4.2. Analysis of the Objective Function

We conduct ablation studies to justify the effectiveness of our objective function for the completion problem. Table 1(a) shows the quantitative effects of these variations, and Fig. 3 shows the qualitative effects. The cGAN alone (bottom left, setting  $\lambda = 0$  in Eq. (3)) gives very noisy results. L2+cGAN (bottom middle) leads to reasonable but blurry results. L1 alone (top right) also produces reasonable results, but we can find some visual defects, like some unfilled holes as marked, which makes the final CD distance higher than that of L1+cGAN. These visual defects can be

reduced when including both L1 and cGAN in the loss function (bottom right). As shown by the example in Fig. 5, the combination of L1 and cGAN can complete the depth images with high accuracy. We further explore the importance of the two components of the objective function for point cloud completion by using different weights ( $\lambda$  in Eq. (3)) of the L1 loss. In Table 1(b), the best completion result is achieved when  $\lambda = 1$ . We set  $\lambda = 1$  in our experiments.

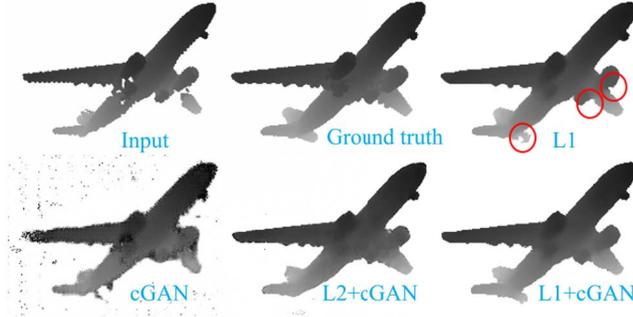


Figure 3: Completion results for different losses.

Loss	Avg CD	$\lambda$ in Eq. (3)	Avg CD
cGAN	0.010729	$\lambda = 50$	0.005748
L1	0.005672	$\lambda = 10$	0.005665
L2 + cGAN	0.006467	$\lambda = 1$	<b>0.005512</b>
L1 + cGAN	<b>0.005512</b>	$\lambda = 0.5$	0.005541

(a)

(b)

Table 1: Analysis of the objective function: average CD for different losses (a), and different  $\lambda$  (b).

### 4.3. Analysis of the View-pooling Layer

**Pooling methods.** We also study different view-pooling methods to construct the shape descriptor, including element-wise max-pooling and mean-pooling. According to our experiments, mean-pooling is not as effective as max-pooling to extract the shape descriptor for image completion, which is similar to the recognition problem [27]. The average CD is 0.005926 for mean-pooling, but that of max-pooling is 0.005512, so max-pooling is used.

Position	Avg L1 distance	Avg CD
$D_2$	<b>3.376642</b>	<b>0.005512</b>
$D_1$	3.433185	0.005604
$D_0$	3.500945	0.005919
Code	3.477186	0.005836

Table 2: Completion results for different positions of view-pooling layer

**Position of the view-pooling layer.** Here we insert the view-pooling layer into different positions to extract the

shape descriptor and further evaluate its effectiveness, including  $D_2$ ,  $D_1$ , and  $D_0$ , which are marked in Fig. 2. Intuitively, the shape descriptor would have the biggest impact on the original network if we place the view-pooling layer before  $D_2$ , and the experimental results illustrate this in Table 2, where both average L1 distance and CD are the lowest. We also try to do view pooling after  $D_0$  and concatenate the shape descriptor with the latent code (marked in purple in Fig. 2) and then pass them through a fully connected layer, but experiments show that the shape descriptor will be ignored since both the average L1 distance and CD do not decrease compared with single view completion net (average L1 distance is 3.473643 and CD is 0.005839 in Table 5).

Model Name	Avg L1 Distance
MVCN-V3	3.794273
MVCN-V8-3	3.616740
MVCN-V5	3.564278
MVCN-V8-5	3.397558

Table 3: Average L1 distance for different numbers of views in view-pooling.

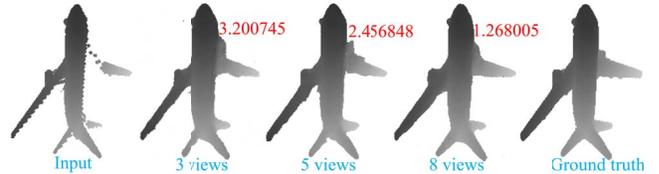


Figure 4: Completion results for different numbers of views in view-pooling.

**Number of views in view-pooling.** We also analyze the effect of the number of views used in view-pooling. In Table 3, MVCN-V3 was trained with 3 depth images (No.1, 3, 5) of the 8 depth images of each 3D model, and MVCN-V5 was trained with 5 depth images (No. 1, 3, 5, 6, 8). MVCN-V8-3 and MVCN-V8-5 were trained with all the 8 depth images, but were tested with 3 views and 5 views respectively. In order to make fair comparisons, we took the 1st, 3rd, and 5th view images to test MVCN-V8-3 and MVCN-V3, and 1st, 3rd, 5th, 6th, 8th to test MVCN-V8-5 and MVCN-V5. The results show that the completion of one single view will be better when we increase the number of views, which means other views are helpful for the completion of one single view, and the more the views, the higher the completion accuracy. Fig. 4 shows an example of the completion. As we increase the number of views in view-pooling, the completion results are improved.

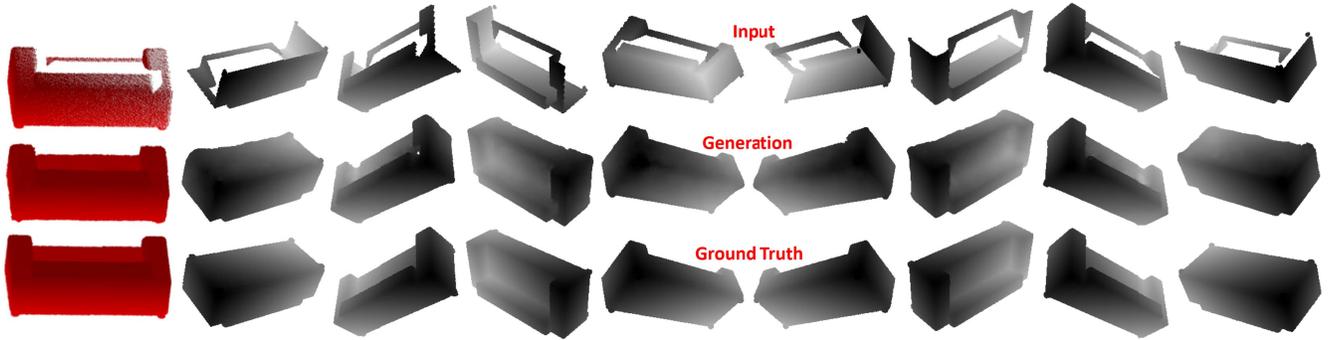


Figure 5: An example of the completion of sofa. The 1st row: incomplete point cloud and 8 depth maps of it; The 2nd row: generated point cloud and related 8 depth maps; The 3rd row: ground truth point cloud and its 8 depth maps.

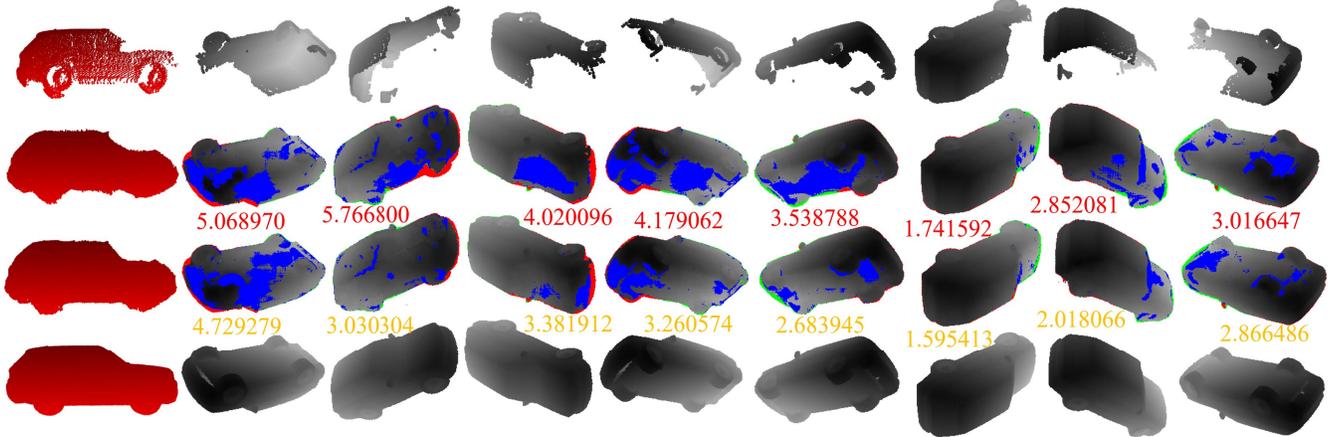


Figure 6: Visual comparison between VCN and MVCN. Starting from the partial point cloud in the first row, VCN and MVCN perform completions of depth maps in the second and third row, respectively, where the completed point clouds are also shown. We use colormaps (from blue to green to red) to highlight the pixels with bigger errors than 10 in terms of L1 distance. Ground truth data is in the last row. MVCN achieves lower L1 distance on all the 8 depth maps.

#### 4.4. Improvements over Single View Completion

**Pervasive improvements on L1 distance and CD.** From Table 5, we find significant and pervasive improvements over single view completion net (VCN) on both average L1 distance and CD on different categories. Nets in Table 5 were trained with 600 3D models for airplane, 1600 for lamp, and 1000 for other categories. We use 150 models of each category to evaluate our network, the same test dataset in [32]. We further conduct visual comparison with VCN in Fig. 6, where we can see MVCN can achieve higher completion accuracy with the help of the shape descriptor.

**Better generalization capability.** Table 4 shows that we can improve the performance of VCN and MVCN while increasing the number of training samples. We find that the performance differences between MVCN-Lamp1000 and VCN-Lamp1000 are not obvious. The reason is that there are relatively large individual differences among lamp models in ShapeNet, and the completion results are bad in several unusual lamp models in the test set. For these models, the comparisons between VCN and MVCN are less

Model	Avg L1 Distance	Avg CD
MVCN-Airplane600	3.376642	0.005512
MVCN-Airplane1200	3.156200	0.005273
MVCN-Lamp1000	6.660511	0.012012
MVCN-Lamp1600	6.245297	0.010576
VCN-Lamp1000	6.763339	0.012091
VCN-Lamp1600	6.430318	0.012007

Table 4: Improvements while increasing training samples.

meaningful, so the improvement is not obvious. But this can be solved when we add another 600 training samples in training. MVCN-Lamp1600 has a bigger improvement than VCN-Lamp1600 on average L1 distance and CD, which indicates a better generalization capability of MVCN.

#### 4.5. Comparisons with the State-of-the-art

**Baselines.** Some previous completion methods need prior knowledge of the shape [28], or assume more complete inputs [12], so they are not directly comparable to our method.

Model	Average L1 Distance								
	Avg	Airplane	Cabinet	Car	Chair	Lamp	Sofa	Table	Vessel
VCN	5.431036	3.473643	4.304635	3.858853	7.644824	6.430318	5.716992	7.572865	4.44616
MVCN	<b>5.102478</b>	<b>3.376642</b>	<b>3.991407</b>	<b>3.609639</b>	<b>7.143200</b>	<b>6.245297</b>	<b>5.284686</b>	<b>7.155616</b>	<b>4.013339</b>
Model	Mean Chamfer Distance per point								
	Avg	Airplane	Cabinet	Car	Chair	Lamp	Sofa	Table	Vessel
VCN	0.008800	0.005839	0.007297	0.006589	0.010398	0.012007	0.009565	0.009371	0.009334
MVCN	<b>0.008328</b>	<b>0.005512</b>	<b>0.007154</b>	<b>0.006322</b>	<b>0.010077</b>	<b>0.010576</b>	<b>0.009174</b>	<b>0.009020</b>	<b>0.008790</b>

Table 5: Comparison of average L1 Distance and mean Chamfer Distance between VCN and MVCN.

Here we compare MVCN with several strong baselines. **PCN-CD** [32] trained with point completion net with CD as loss function, is the state of the art while this work was developed. **PCN-EMD** uses Earth Mover’s Distance (EMD) [4] as loss function, but it is intractable for dense completion due to the calculation complexity of EMD. The encoders of **FC** [1], **Folding** [31] are the same with PCN-CD, but decoders are different, a 3-layer fully-connected network for FC, and folding-based layer for Folding. **PN2** uses the same decoder, but the encoder is PointNet++ [23]. **3D-EPN** [3] is a representative of the class of volumetric completion methods. For fair comparison, the distance field outputs of 3D-EPN are converted into point clouds as mentioned in [32]. TopNet [29] is a recent point-based method, but it can only generate sparse point clouds because their decoder mostly consists of multilayer perceptron networks, which limits the number of points they can process.

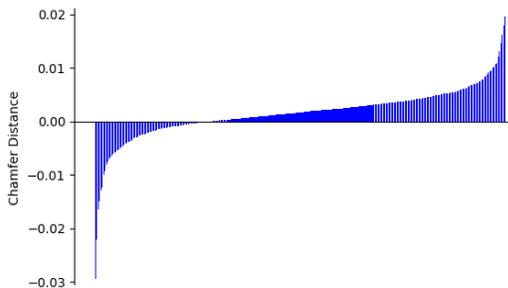


Figure 7: Comparison between MVCN and PCN-CD.

**Comparisons.** As done in [32], we use the symmetric version of CD to calculate the average closest point distance, where ground truth point clouds and generated point clouds are not required to be the same size, which is different from EMD [4]. For point-based methods like PCN [32], the input is sampled and the output size is fixed, which makes the calculation of EMD relatively easy. Different from these methods, the number of output points of our approach is not fixed, which would require resampling our output to compute the EMD. CD is more suitable for a fair comparison among different techniques. Table 6 shows the quantitative results, where the completion results of other methods are from [32]. Our method achieves the lowest CD across

almost all object categories. A more detailed comparison with PCN-CD is in Fig. 7, where the height of the blue bar indicates the amounts of improvement of our method over PCN-CD on each object. Our model outperforms PCN on most objects. Fig. 9 shows the qualitative results. Our completions are denser, and we recover more details in the results. Another obvious advantage is that our method can complete shapes with complex geometry, like the 2nd to 4th objects, but other methods fail to recover these shapes.

#### 4.6. Completion Results on KITTI

Our goal is to obtain high quality and high resolution shape completion from data similar to individual range scans focused on individual objects. Hence we obtain incomplete data using synthetic depth images, which is similar to data from RGB-D cameras. However, for data like KITTI, which is extremely sparse and does not contain ground truth, the usual objective is to obtain rough not high resolution completion. Our method performs reasonably well on KITTI data, as shown in Fig. 8. More completion results on noisy, sparse, and occluded data can be found in the supplementary material.

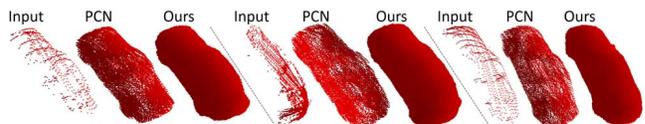


Figure 8: Completion results on KITTI.

## 5. Conclusion

We have presented a method for shape completion by rendering multi-view depth maps of incomplete shapes, and then performing image completion of these rendered views. Our multi-view completion net shows significant improvements over a baseline single view completion net across multiple object categories. Experiments show that our view based representation and novel network structure can achieve better results with less training samples, perform better on objects with complex geometry, and generate higher resolution results than previous methods.

Model	Mean Chamfer Distance per point								
	Avg	Airplane	Cabinet	Car	Chair	Lamp	Sofa	Table	Vessel
3D-EPN	2.0147	1.3161	2.1803	2.0306	1.8813	2.5746	2.1089	2.1716	1.8543
FC	0.9799	0.5698	1.1023	0.8775	1.0969	1.1131	1.1756	0.9320	0.9720
Folding	1.0074	0.5965	1.0831	0.9272	1.1245	1.2172	1.1630	0.9453	1.0027
PN2	1.3999	1.0300	1.4735	1.2187	1.5775	1.7615	1.6183	1.1676	1.3521
PCN-CD	0.9636	0.5502	1.0625	0.8696	1.0998	1.1339	1.1676	<b>0.8590</b>	0.9665
PCN-EMD	1.0021	0.5849	1.0685	0.9080	1.1580	1.1961	1.2206	0.9014	0.9789
MVCN	<b>0.8298</b>	<b>0.5273</b>	<b>0.7154</b>	<b>0.6322</b>	<b>1.0077</b>	<b>1.0576</b>	<b>0.9174</b>	0.9020	<b>0.8790</b>

Table 6: Comparison with the state-of-the-art in terms of mean CD (multiplied by 100) per point over multiple categories.

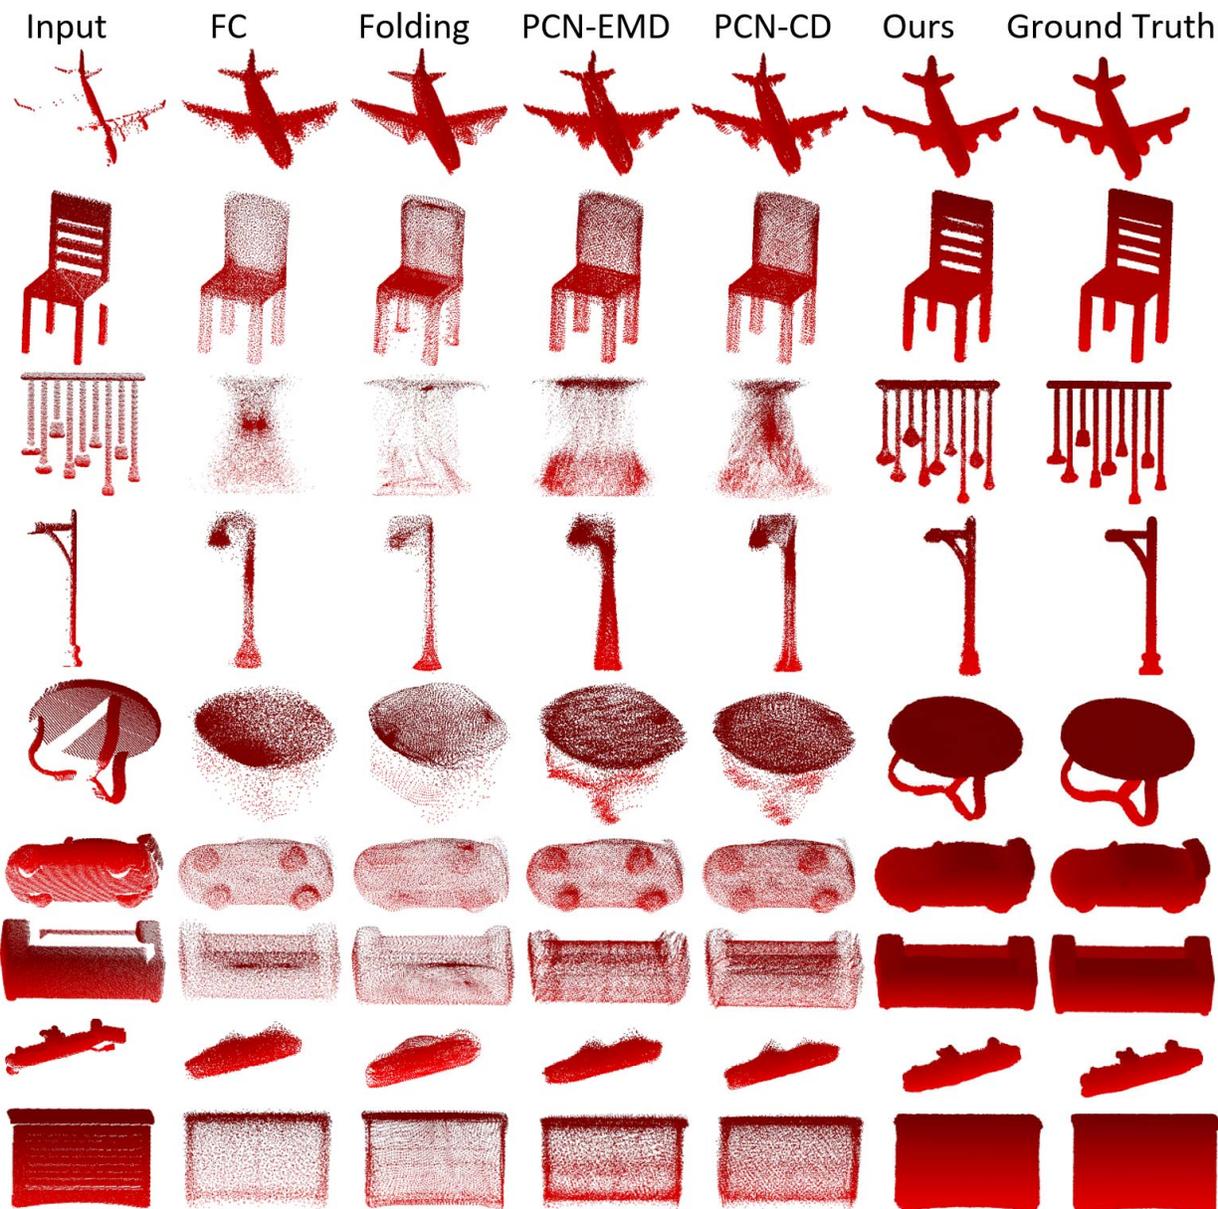


Figure 9: Qualitative completion on ShapeNet, where MVCN can complete complex shapes with high resolution.

## References

- [1] P. Achlioptas, O. Diamanti, I. Mitliagkas, and L. J. Guibas. Learning representations and generative models for 3d point clouds. In *ICML*, 2018.
- [2] A. X. Chang, T. A. Funkhouser, L. J. Guibas, P. Hanrahan, Q.-X. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu. Shapenet: An information-rich 3d model repository. *CoRR*, abs/1512.03012, 2015.
- [3] A. Dai, C. R. Qi, and M. Nießner. Shape completion using 3d-encoder-predictor cnns and shape synthesis. *CVPR*, pages 6545–6554, 2017.
- [4] H. Fan, H. Su, and L. J. Guibas. A point set generation network for 3d object reconstruction from a single image. *CVPR*, pages 2463–2471, 2017.
- [5] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014.
- [6] X. Han, Z. Li, H. Huang, E. Kalogerakis, and Y. Yu. High-resolution shape completion using deep neural networks for global structure and local geometry inference. In *ICCV*, Oct 2017.
- [7] Z. Han, M. Shang, Y.-S. Liu, and M. Zwicker. View Inter-Prediction GAN: Unsupervised representation learning for 3D shapes by learning global shape memories to support local view predictions. In *AAAI*, 2019.
- [8] Z. Han, M. Shang, X. Wang, Y.-S. Liu, and M. Zwicker. Y2Seq2Seq: Cross-modal representation learning for 3D shape and text by joint reconstruction and prediction of view and word sequences. In *AAAI*, pages 126–133, 2019.
- [9] S. Iizuka, E. Simo-Serra, and H. Ishikawa. Globally and locally consistent image completion. *ACM Trans. Graph.*, 36(4):107:1–107:14, July 2017.
- [10] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.
- [11] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. *CVPR*, pages 5967–5976, 2017.
- [12] M. Kazhdan and H. Hoppe. Screened poisson surface reconstruction. *ACM Trans. Graph.*, 32(3):29:1–29:13, July 2013.
- [13] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 12 2014.
- [14] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *CoRR*, abs/1312.6114, 2014.
- [15] C.-H. Lin, C. Kong, and S. Lucey. Learning efficient point cloud generation for dense 3d object reconstruction. In *AAAI*, 2018.
- [16] D. Maturana and S. Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 922–928, Sep. 2015.
- [17] L. M. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger. Occupancy networks: Learning 3d reconstruction in function space. *ArXiv*, abs/1812.03828, 2018.
- [18] V. Nair and G. E. Hinton. Rectified linear units improve restricted boltzmann machines. In *ICML*, 2010.
- [19] D. Pathak, P. Krähenbühl, J. Donahue, T. Darrell, and A. A. Efros. Context encoders: Feature learning by inpainting. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2536–2544, 2016.
- [20] T. Portenier, Q. Hu, A. Szabó, S. A. Bigdeli, P. Favaro, and M. Zwicker. Faceshop: deep sketch-based face image editing. *ACM Trans. Graph.*, 37:99:1–99:13, 2018.
- [21] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 77–85, 2017.
- [22] C. R. Qi, H. Su, M. Nießner, A. Dai, M. Yan, and L. J. Guibas. Volumetric and multi-view cnns for object classification on 3d data. In *CVPR*, pages 5648–5656, June 2016.
- [23] C. R. Qi, L. Yi, H. Su, and L. J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NIPS*, 2017.
- [24] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, volume 9351 of *LNCS*, pages 234–241. Springer, 2015. (available on arXiv:1505.04597 [cs.CV]).
- [25] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, X. Chen, and X. Chen. Improved techniques for training gans. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 2234–2242. Curran Associates, Inc., 2016.
- [26] A. A. Soltani, H. Huang, J. Wu, T. D. Kulkarni, and J. B. Tenenbaum. Synthesizing 3d shapes via modeling multi-view depth maps and silhouettes with deep generative networks. *CVPR*, pages 2511–2519, 2017.
- [27] H. Su, S. Maji, E. Kalogerakis, and E. G. Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *Proc. ICCV*, 2015.
- [28] M. Sung, V. G. Kim, R. Angst, and L. Guibas. Data-driven structural priors for shape completion. *ACM Trans. Graph.*, 34(6):175:1–175:11, Oct. 2015.
- [29] L. P. Tchammi, V. Kosaraju, H. Rezatofighi, I. Reid, and S. Savarese. Topnet: Structural point cloud decoder. In *CVPR*, June 2019.
- [30] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao. 3d shapenets: A deep representation for volumetric shapes. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1912–1920, 2015.
- [31] Y. Yang, C. Feng, Y. Shen, and D. Tian. Foldingnet: Interpretable unsupervised learning on 3d point clouds. *CoRR*, abs/1712.07262, 2017.
- [32] W. Yuan, T. Khot, D. Held, C. Mertz, and M. Hebert. Pcn: Point completion network. *2018 International Conference on 3D Vision (3DV)*, pages 728–737, 2018.