

Multimodal Shape Completion via Implicit Maximum Likelihood Estimation

Himanshu Arora Saurabh Mishra Shichong Peng Ke Li Ali Mahdavi-Amiri
Simon Fraser University
Burnaby, CA

{haa48, skm24, spa176, keli, amahdavi}@sfu.ca

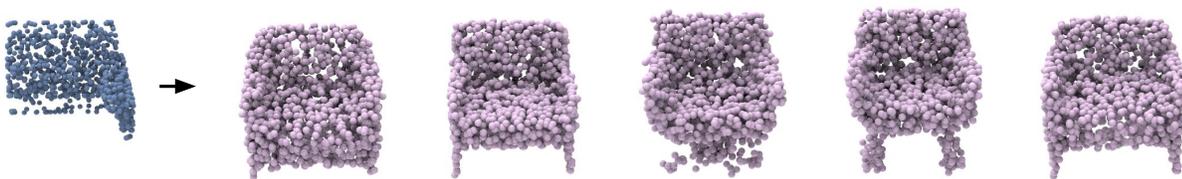


Figure 1. Having a partial point cloud as input, our method is capable of producing *diverse* complete outputs.

Abstract

Shape completion is the problem of completing partial input shapes such as partial scans. This problem finds important applications in computer vision and robotics due to issues such as occlusion or sparsity in real-world data. However, most of the existing research related to shape completion has been focused on completing shapes by learning a one-to-one mapping which limits the diversity and creativity of the produced results. We propose a novel multimodal shape completion technique that is effectively able to learn a one-to-many mapping and generates diverse complete shapes. Our approach is based on the conditional Implicit Maximum Likelihood Estimation (IMLE) technique wherein we condition our inputs on partial 3D point clouds. We extensively evaluate our approach by comparing it to various baselines both quantitatively and qualitatively. We show that our method is superior to alternatives in terms of completeness and diversity of shapes.

1. Introduction

As 3D scanning devices are now widely available, scenes and objects in the form of point clouds are rapidly becoming more and more prevalent. However, many such objects are not fully scanned due to occlusion; as a result, the scanned shapes may be missing substantial portions of the complete shape. Therefore, the ability to complete such shapes by

generating the missing portions is desired. More precisely, the objective of shape completion is to meaningfully complete a partial shape and obtain a complete shape without violating the geometry of the partial shape.

Early work has focused on generating a single complete shape; the mapping from the input and output is one-to-one. However, shape completion is ill-posed by nature as the target shape is not unique – given a partial shape, there are multiple ways to generate the missing portions, especially when the input point cloud is noisy and coarse, which is the case in real-world scanned data. Therefore, given a partial shape, it is desired to generate a *set* of many possible complete shapes, all of which respecting the given input partial shape. Not only is multimodality a way to resolve ill-posedness, but it also adds more *diversity* and *creativity* to the final outputs.

Recently, Wu et al. [37] proposed the task of multimodal shape completion, where the goal is to generate various complete shapes for a single partial shape. This is challenging, because the training data only contains one ground truth complete shape for each partial shape. In [37], a conditional GAN-based model was proposed to learn a one-to-many mapping from partial shapes to complete shapes.

While [37] took a step towards multimodal shape completion, the different shapes it generates for a given partial point cloud tend to be similar in structure. For example, for an incomplete chair, the different complete planes it generates are of a similar style and vary mostly in width and

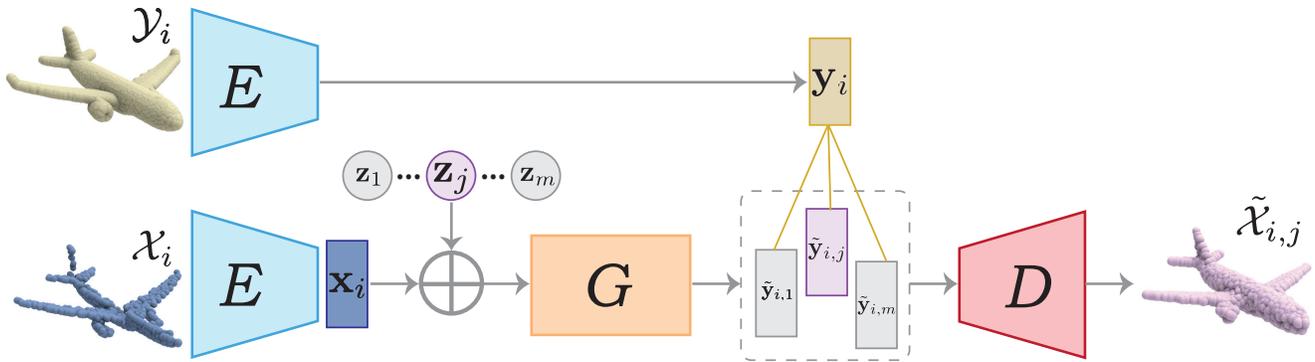


Figure 2. Our network receives a partial point cloud \mathcal{X}_i and a complete point cloud (in training time) \mathcal{Y}_i . The latent code generated from the incomplete shape, \mathbf{x}_i along with a noise, \mathbf{z}_j is given to generator G to be mapped to $\tilde{\mathbf{y}}_{i,j}$. Closest $\tilde{\mathbf{y}}_{i,j}$ to the latent code \mathbf{y}_i of the complete point cloud is found and passed to decoder D to produce a novel and complete shape $\tilde{\mathcal{X}}_{i,j}$.

height, as shown in Figure 5. In this paper, our goal is to generate *structurally diverse* complete shapes for the same incomplete shape. In particular, consider the case where a part of the object is entirely missing, the different generated complete shapes should ideally vary in the style of the missing part while remaining compatible with the other parts that are present.

Our key insight is that structural similarity between the generated complete shapes comes from the tendency of GANs to collapse modes. To achieve structural diversity, we must therefore get around mode collapse. To this end, we propose a novel method by devising a model formulation that allows for training with the method of Implicit Maximum Likelihood Estimation (IMLE). Prior to our work, IMLE has never been successfully used for point cloud generation. In this paper, we show how to do so, which results in a simple and elegant model with fewer parameters than [37]. Yet, despite this simplicity, we demonstrate that our method is capable of generating structurally diverse results. Quantitatively, our method achieves superior performance than prior methods in terms of Total Mutual Difference, and comparable or superior performance than prior methods in terms of Unidirectional Hausdorff Distance (UHD). Moreover, we demonstrate the capability of our method to generalize to new kinds of incomplete shapes. Despite not having part-level supervision, our method can produce complete shapes for incomplete shapes with an entire part missing.

2. Related Work

Since the missing regions of an incomplete point cloud are unknown and might be geometrically complex, simple surface repair algorithms are not fully effective in reproducing parts and geometric features. In fact, missing regions should be better *generated* after observing many relevant shapes to learn their geometric properties. Therefore, here, we mostly provide a review of deep generative models

that are reasonable solutions for shape completion. In addition, since our input is a point cloud, we provide a review of deep learning models designed for point clouds, and we particularly discuss point cloud completion methods. Point cloud completion may have similarities with other problems, including point upsampling, 3D part assembly, 3D reconstruction, and single view reconstruction. However, since the focus of our paper is multimodal point cloud completion, we refrain from discussing works related to those areas and refer the interested reader to the available survey papers such as [1, 3].

2.1. Generative Models

Generating new images or 3D models is an important area in computer vision and computer graphics. One of the most well-known techniques to generate new shapes is *variational autoencoder* [16] or VAE that tries to fit a known distribution (e.g., Gaussian) to the latent space of an autoencoder so that by sampling the known distribution, the decoder can produce novel data. VAEs have been used in many applications, including generating new 3D datasets and shape completion [3, 30].

Generative Adversarial Networks (GANs) [9] are alternative generative models composed of a generator and a discriminator, where the generator tries to produce novel data that can fool the discriminator, and the discriminator tries to distinguish the produced results (i.e., fake) from real samples. There are various forms of GAN that have been designed to tackle different problems, including conditional GANs that try to generate data based on a condition such as an incomplete image for the purpose of inpainting [13, 43] or a partial shape for shape completion [5]. In fact, GAN-based networks have been successful in many 3D tasks and applications [3].

Despite the popularity and success of GANs, they have limitations including mode collapse. This problem is more evident when you need diversity in generated data (such as

in the case of multimodal shape completion). This has been the primary motivation behind IMLE, which is to avoid mode collapse and it has been applied to various multimodal image generation problems including image super-resolution [17], image synthesis from semantic layout [18] and image decompression [24].

2.2. Deep Learning on Point Clouds

One of the possibilities is to feed a neural network by an array of points (i.e., point clouds) capturing the geometry of shapes. One of the earliest networks to deal with point clouds is PointNet [25] that tries to learn feature points to perform different tasks including segmentation or classification. As PointNet only learns global features, PointNet++ [26] is introduced to learn multi-scale features and better abstract local patterns. Due to the effectiveness of PointNet and PointNet++, they have been used as the main component of many networks with various applications including point to point translation [4, 40, 41], edge detection on point clouds [35], learning implicit surfaces [8], denoising [27] and many more [10].

2.3. Shape Completion

Classical methods. Shape completion has been studied in the area of mesh processing even before the popularity of deep learning. For example, the missing region of a mesh can be repaired by minimizing an energy function via Laplacian or Poisson equation [14, 15, 21, 29, 46], benefiting from shapes' skeletal structures [38, 42], or retrieving a template from a database and deforming it into the incomplete input point cloud [23, 28]. Interested readers can refer to [1] for a thorough study on shape reconstruction and completion for point clouds.

Deep shape completion. The missing region of a 3D shape, however, might be so geometrically complex that it can only be *generated* through a deep network trained on a dataset with variety. As a result, deep learning techniques have been proposed to leverage available 3D shape datasets to learn how to reproduce the missing regions successfully.

In [7], an autoencoder structure is used to first predict a complete distance field of the voxelized incomplete input shape. The distance field is then contrasted against a shape database to obtain the final shape through a multi-resolution patch-based 3D shape synthesis. Stutz and Geiger [30] provide a weakly supervised VAE-based method for shape completion. They first train a denoising variational autoencoder on ShapeNet in a supervised manner. They then fix and reuse the decoder to train a new encoder on incomplete shapes where alignment of the predicted result with the incomplete input shape is respected by applying a maximum likelihood (ML) loss. ASFM-Net [39] maps the partial and complete input point clouds into a common latent space to capture detailed shape priors and better respect the input.

A series of methods have utilized a multi-stage approach to complete a shape. Inspired by patch-match technique in image processing, Han et al. [11] employed two networks to produce missing parts of an incomplete shape, where one network tries to respect the overall structure of the shape and the other one produces local patches to repair the given incomplete shape. PCN [44] is a multi-stage encoder-decoder network where the decoder first produces a coarse complete shape, and then for each point of the coarse shape, a local patch is generated and deformed to reproduce details. Liu et al. [19] also introduced a multi-stage method where a coarse shape is first approximated by a set of parametric surfaces and then it is refined by learning point-wise residuals. Similarly, Wang et al. [34] designed a network that performs the shape completion task in two stages (i.e., coarse and dense reconstruction). They have also designed a patch discriminator to guarantee that local patches resemble the same pattern as those in the ground truth. In addition to the coarse/fine point completion techniques, TopNet [32] has been designed as a tree structure where each branch represents a subset of the point cloud.

Detailed shape completion. There are recent works that have focused on completing partial input shapes that attain a great amount of detail. These works try to replicate existing features in the partial shape in the proper location of the generated complete shape. Wen et al. [36] have defined a skip-attention mechanism to avoid using a global representation that may suffer from altering the initial geometry of the input shape. PF-Net [12] employs a multi-resolution encoding for the shape and tries to only generate the missing parts of the partial shape to respect the input's features. Most recently, VRCNet [22] benefits from the multi-stage completion (i.e., coarse and fine stages) and is capable of replicating fine details when they are available. Although these works produce nice results when the inputs are detailed, our work is naturally different from these methods. Here, we focus on completing input shapes that are relatively coarse, and therefore multiple interpretations are possible. Only under this setting, the multimodality is meaningful since when the partial input shape is dense and detailed, the complete shape will be limited to only one ground truth shape with the same features as the input.

GANs and multimodality. GAN-based techniques have also been employed for shape completion. Wang et al. [33] use a GAN in the form of a convolutional autoencoder to produce a coarse shape which is then sliced into images and passed to a recurrent convolutional network for refinement. To setup an unpaired shape completion, two autoencoders are initially trained on complete and incomplete shapes in [5]. The latent code of an incomplete shape is then mapped to a complete shape's latent code via a generator, and two discriminators are used to evaluate complete shapes and their latent codes. The other generative method

is proposed by Wu et al. [37] where they combine VAE and conditional GAN to achieve multimodality. Later in this paper, we show that although [37] is successful in point cloud completion due to mode collapse, our IMLE-based approach is more successful in producing shapes with more variety. A series of experiments have been performed to examine the effectiveness of our method, which is discussed in Section 4. In addition to GAN-based approaches that are capable of providing multimodality, a voxel-based generative model (GCA) is introduced by Zhang et al. [45] that produce results by gradually completing voxels. Although we have provided a qualitative comparison with GCA, it might not be directly relevant to our work as GCA is designed for voxels instead of point clouds, and also it does not use GAN in its structure while we try to provide an alternative to GAN.

3. Method

Overview. In the shape completion problem, we want to predict a complete point cloud based on a given partial point cloud with missing regions or parts. In the multimodal version of this problem, a set of complete shapes is desired under only one given partial point cloud. Formally, given an input \mathbf{x} in a one-to-many problem setting, the goal is to learn $p(\mathbf{y}|\mathbf{x})$ where \mathbf{y} is a prediction. Here, \mathbf{x} is our partial point cloud and \mathbf{y} is a predicted complete point cloud. Since there are many possible complete shapes that are consistent with \mathbf{x} , we want to generate non-deterministic predictions $p(\mathbf{y}|\mathbf{x})$ instead of a single prediction. The simplest way to generate non-deterministic predictions is to add a random code input $\mathbf{z} \sim \mathcal{N}(0, \mathbf{I})$ to a deterministic function T_θ where $\mathbf{y} := T_\theta(\mathbf{x}, \mathbf{z})$. This is also known as an implicit generative model [20]. One way to train this implicit generative model is to use a conditional GAN, in this context, where $T_\theta(\mathbf{x}, \mathbf{z})$ is often known as *generator*. However, the generator may fail to produce diverse outputs due to mode collapse. Therefore, we use conditional Implicit Maximum Likelihood Estimation (IMLE) to train the generative model to overcome the mode collapse issue.

IMLE reverses the direction of how generated samples are matched to the target data: instead of making the generated samples similar to some target data, IMLE ensures that every target data has a similar sample being generated. In addition, IMLE does not require an adversarial objective which simplifies training. If we denote the generator as $T_\theta(\cdot, \cdot)$, an input point cloud as \mathbf{x}_i , the corresponding ground truth as \mathbf{y}_i , and a latent code as $\mathbf{z}_{i,k}$, then IMLE would optimize the following equation:

$$\min_{\theta} \mathbb{E}_{\mathbf{Z}} \left[\sum_{i=1}^n \min_{k \in \{1, \dots, m\}} d(T_\theta(\mathbf{x}_i, \mathbf{z}_{i,k}), \mathbf{y}_i) \right] \quad (1)$$

where $\mathbf{Z} = \mathbf{z}_{1,1}, \dots, \mathbf{z}_{n,m} \sim \mathcal{N}(0, \mathbf{I})$, $d(\cdot, \cdot)$ is a distance

metric and, m is a hyperparameter. Algorithm 1 shows the conditional IMLE training procedure. In our case, IMLE is conditioned on the given partial shape.

Algorithm 1 Conditional IMLE Training Procedure

Require: The set of inputs $\{\mathbf{x}_i\}_{i=1}^n$ and the set of corresponding observed outputs $\{\mathbf{y}_i\}_{i=1}^n$
Initialize the parameters θ of the generator T_θ
for $p = 1$ **to** N **do**
 Pick a random batch $S \subseteq \{1, \dots, n\}$
 for $i \in S$ **do**
 Randomly generate i.i.d. m latent codes
 $\mathbf{z}_1, \dots, \mathbf{z}_m$
 $\tilde{\mathbf{y}}_{i,k} \leftarrow T_\theta(\mathbf{x}_i, \mathbf{z}_k) \forall k \in [m]$
 $j \leftarrow \arg \min_k d(\mathbf{y}_i, \tilde{\mathbf{y}}_{i,k}) \forall k \in [m]$
 for $q = 1$ **to** M **do**
 Pick a random mini-batch $\tilde{S} \subseteq S$
 $\theta \leftarrow \theta - \eta \nabla_{\theta} (\sum_{i \in \tilde{S}} d(\mathbf{y}_i, \tilde{\mathbf{y}}_{i,j})) / |\tilde{S}|$
return θ

Shape completion network’s overall structure. The structure of our network is composed of an auto-encoder (i.e., encoder E and decoder D) along with generator G (see Figure 2). The auto-encoder encodes partial shape \mathcal{X}_i and complete shape \mathcal{Y}_i via its encoder and generates a novel and complete shape $\tilde{\mathcal{X}}_i$ via its decoder. Generator G plays the role of T_θ in Algorithm 1 by mapping the latent code \mathbf{x}_i of the partial shape, aggregated by a noise vector \mathbf{z}_j sampled from normal distribution, to a new code $\tilde{\mathbf{y}}_{i,j}$. $\tilde{\mathbf{y}}_{i,j}$ is then passed to decoder D to produce a new shape $\tilde{\mathcal{X}}_{i,j}$. Therefore, each partial shape \mathcal{X}_i can be mapped to many shapes $\tilde{\mathcal{X}}_{i,j}$.

Multi-modality and training. To train the network on the entire data, for each input \mathbf{x}_i , m forward passes with different random vectors ($\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_m$) are processed. Therefore, the output from the generator is m latent vectors ($\tilde{\mathbf{y}}_{i,1}, \tilde{\mathbf{y}}_{i,2}, \dots, \tilde{\mathbf{y}}_{i,m}$) where each latent vector corresponds to a different complete shape or *mode*. As also described in Algorithm 1, we then use a nearest neighbour algorithm to select $\tilde{\mathbf{y}}_{i,j}$ which is closest to the latent vector encoded from the corresponding complete shape \mathbf{y}_i . We compute an L_1 loss between $\tilde{\mathbf{y}}_{i,j}$ and \mathbf{y}_i to bring $\tilde{\mathbf{y}}_{i,j}$ closer to the selected mode (see Figure 3):

$$L_{L_1} = d(\mathbf{y}_i, \tilde{\mathbf{y}}_{i,j}) \quad (2)$$

Enforcing that every complete shape in our training data has a *similar* shape in generated samples ensures that mode collapse does not occur and we have diversity in generated shapes. Finally, we decode $\tilde{\mathbf{y}}_{i,j}$ through decoder D and produce a novel complete shape. We use Unidirectional Hausdorff Distance (UHD) loss between the given partial shape

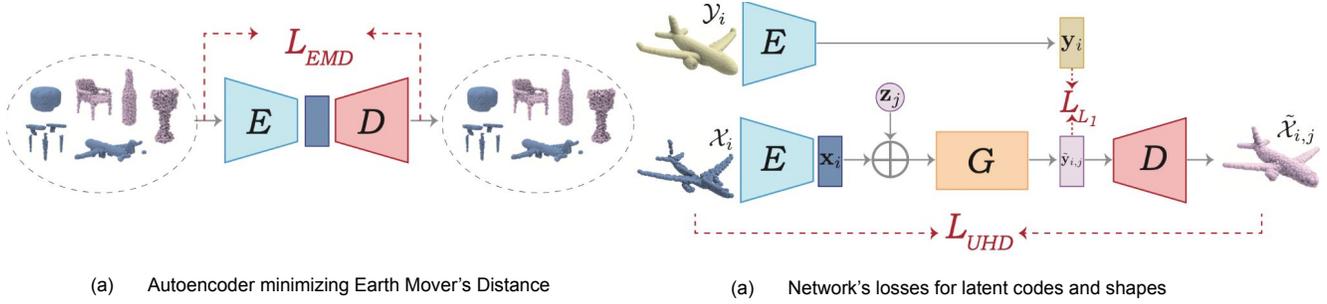


Figure 3. Autoencoder is trained on both partial (blue) and complete (purple) shapes using Earth mover’s distance (a). Two losses are used on shapes and latent codes in our network (b). Losses are highlighted in red.

and the generated complete shape to ensure that the generated complete shape encloses the partial shape effectively. This loss can be defined as:

$$L_{UHD} = \max_{x \in \mathcal{X}_i} \min_{y \in \mathcal{X}_{i,j}} \|x - y\|, \quad (3)$$

where x and y are single points belonging to partial and the generated complete point clouds.

Note that our auto-encoder is trained for reconstruction task on a point cloud dataset which includes both partial and complete shapes. We employed Earth Mover’s Distance (EMD) loss between input and output point clouds \mathbf{p}_{in} and \mathbf{p}_{out} to train our auto-encoder:

$$L_{EMD}(\mathbf{p}_{in}, \mathbf{p}_{out}) = \min_{\xi: \mathbf{p}_{in} \rightarrow \mathbf{p}_{out}} \sum_{\alpha \in \mathbf{p}_{in}} \|\alpha - \xi(\alpha)\| \quad (4)$$

Testing. During testing, our model does not require any accompanied complete shape \mathcal{Y}_i because of which there is no nearest neighbour selection involved. We simply pass the partial shapes along with different randomly sampled noise vectors to the network and decode all the produced latent codes into complete generated shapes.

4. Experiments and Results

In this section, we provide various experiments and compare our work both qualitatively and quantitatively with baseline methods. In the end, we are providing some statistics about timing along with qualitative results to show how much our method respects the partial input shape. We have also provided some complementary experiments on baselines and real-world scan data

Baselines. We compare the result of our point completion tasks against two multimodal methods: MSC [37], and kNN-latent (a baseline introduced by [37]). kNN-latent is a simple baseline to look for most similar complete shapes in dataset using encoded latent vector of partial shape. MSC uses a conditional GAN architecture to generate multiple shapes. We perform quantitative and qualitative comparisons with these baselines in the following.

Datasets. We use three datasets to evaluate our work in different settings. First, we use 3Depn dataset [6] which is generated from ShapeNet using virtual scans. The dataset has varied proportions of incompleteness in the samples. We evaluate our method on three classes of this dataset - Chair, Plane, and Table. Second, we evaluate Partnet dataset [2] in which parts are randomly removed from a complete shape while training/testing. In Partnet, we perform evaluation on Chair, Table, and Lamp because of high number of samples for these classes. Finally, we use the recently released RobustPointSet [31] dataset for evaluating transferability of our network. Shapes in this dataset have missing regions synthetically created on various objects.

Evaluation. Here we provide quantitative and qualitative results to evaluate our method against baselines. In Table 1 we compare our result using Total Mutual Difference on 3Depn (left) and Partnet (right) datasets. Total Mutual Difference (TMD) quantifies the diversity among the generated samples. This is calculated using Chamfer distance between every pair among generated samples. We show that we have significant gains in terms of diversity in generated samples in comparison with multimodal baselines. The difference can be explained by the fact that our work tends to use all modes in the dataset. Similarly, in Table 2 we compare our result using Unidirectional Hausdorff Distance (UHD) on 3Depn (left) and Partnet (right) datasets. UHD is a measure of how well the partial shape is enclosed by the generated complete shape. We either outperform the baselines or perform relatively similarly. To provide qualitative results, we show our generated samples for various shape categories in Figure 4. It can be observed that in the generated shapes, our method tends to produce various geometric and shape attributes.

Figure 5 contrasts our results to conditional GAN [37] and shows that our model has more diversity. Note that our method respects the partial input very well as is apparent in Figure 6. To make a better contrast, we have overlapped the partial input highlighted in green over the generated complete shape highlighted in blue.



Figure 4. Our completion results for multiple object categories (e.g., Chair, Lamp, Plane and Bottle). It can be observed that the generated complete shapes are diverse for every object and respects partial shapes.

Methods	Chair	Plane	Table
kNN-latent	2.84	1.13	3.25
Multimodal-GAN	2.56	2.03	4.49
Ours	2.93	2.31	4.92

Methods	Chair	Lamp	Table
kNN-latent	2.28	4.18	2.85
Multimodal-GAN	2.56	3.31	3.88
Ours	2.76	5.49	4.45

Table 1. Comparison between our methods and baselines in terms of Total Mutual Difference (TMD) or *diversity* on two datasets - 3Depn (top) and PartNet (bottom). Higher TMD values indicate more diversity. TMD values in the table are multiplied by a factor of 10^2 .

Methods	Chair	Plane	Table
kNN-latent	8.94	9.54	10.04
Multimodal-GAN	9.14	9.59	8.74
Ours	8.51	9.55	8.52

Methods	Chair	Plane	Table
kNN-latent	8.58	8.47	7.61
Multimodal-GAN	6.65	5.40	5.38
Ours	6.17	5.58	5.16

Table 2. Comparison between our method and baselines in terms of Unidirectional Hausdorff Distance (UHD) or *completeness* on two datasets - 3Depn (top) and PartNet (bottom). Lower UHD values indicate better completion. UHD values in the table are multiplied by a factor of 10^2 .

Part-based completion. In Figure 7, we demonstrate the ability of our method to generalize to unseen types of partial shapes as input. We remove a part entirely and visualize different completions generated by our method. Even



Figure 5. Qualitative Comparison of our work (left) and conditional GAN [37] (right). As it is apparent, our method produces more diverse objects.

though no part labels were used during training and so the model has never seen an entire part taken out, the model can generate reasonable and diverse completions. Moreover, the model can generate a diverse range of parts in place of the missing part that is functionally similar but structurally different, which suggest several interesting properties: (1) the model has learned which parts have the same function and are interchangeable, suggesting it has acquired an implicit understanding of the semantics of the part in question even though it has never seen part labels during training. (2) The model does not necessarily generate the particular shape of the part that was taken out in all completions, suggesting it did not collapse to a single mode. (3) The model is able to generate structurally diverse shapes for missing parts in dif-



Figure 6. Partial shapes (green) and generated complete shapes (blue) are overlapped to illustrate how well our complete shape respects the partial shape.



Figure 7. Three partial portions are taken from complete shapes for two categories (Chair and Lamp) and completed by our technique.

ferent completions, suggesting that it has learned the space of structural variations for a part. Moreover, the variations can be localized to a part even though global noise is injected into the architecture.

Transfer Learning. To examine whether our model can be used on other datasets, possibly with more non-uniform missing parts and noises, we perform an experiment on Robustpointset dataset [31]. Robustpointset contains various artifacts including missing regions, rotations, noise, etc. We use our model trained on 3Depn dataset and fine-tune it on this dataset. Our model successfully fills in or completes the missing regions and parts. In Figure 8, we show that our model can be successfully transferred to other datasets. It is worth noting that these samples do not have much scope for a diverse generation, so our model naturally converges to the actual shape.

Timing Comparison. Since our network does not have a discriminator, its training and inference times are lower than conditional GAN. Our method needs 13 ms per shape for training and 2 ms for inference, while conditional GAN needs 20ms per second and 2ms for inference. [37]. On a Geforce RTX 2080 GPU, the overall training time for conditional GAN work for the main class Chair (3Depn) is 10 hours for the autoencoder, 10 hours for the VAE and 5 hours for the generator/discriminator network. In comparison, for the same case, our work takes total 13.5 hours (10 hours for the autoencoder and 3.5 hours for the generator training).

Noise Tolerance. We perform noise tolerance analysis of our work by adding Gaussian noise to input point clouds. We perform this on 3Depn dataset for Chair category. As it can be seen in Figure 9, we are able to achieve reasonable completion despite having noise in partial shapes. We achieve 2.71×10^{-2} as Total Mutual Difference value and 8.84×10^{-2} as Unidirectional Hausdorff Distance value with noisy input point clouds.

Real-World Scans Here we show our results on ScanNet dataset shared by [5]. This dataset comprises of real-world scans of chairs. The scanned samples are more noisy and sparse as compared to other datasets which makes them a difficult dataset to evaluate on. We use the model we trained on 3Depn dataset to evaluate on real-world scans and we demonstrate successful completion with decent variety (see Figure 11).

Comparison with GCA [45] We have also provided some qualitative results with GCA [45] that is a probabilistic model capable of performing multimodal shape completion. As illustrated in Figure 10, both our method and GCA are capable of producing diverse results. Note that our method can produce various plausible bases for the two provided tops in comparison with GCA and cGAN. It is also worth mentioning that due to the simplicity of our network, our inference time is only 0.02 seconds while GCA takes 0.7 seconds to complete a single inference execution since they run their network for 100 times to generate a single out-



Figure 8. Generator trained on 3Depn dataset is fine-tuned on RobustPointSet dataset and results are evaluated on Chair class.

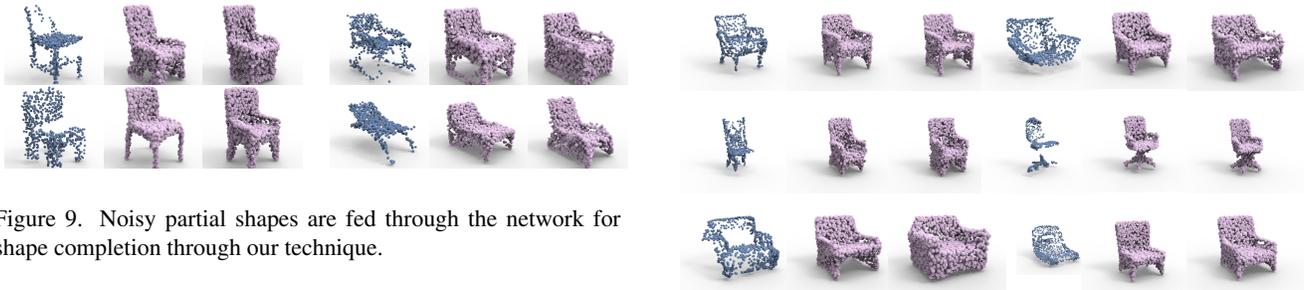


Figure 9. Noisy partial shapes are fed through the network for shape completion through our technique.

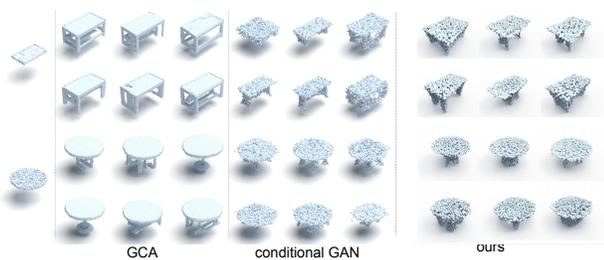


Figure 10. Comparative results to highlight the difference between GCA [45], conditional GAN [37], and our method. GCA results are directly taken from the paper [45].

put. Quantitative results/code are not provided as the GCA’s code is not publicly available yet.

Implementation details. We tune our hyperparameters through experiments to have a standard set of parameters for all classes. For the encoder and decoder model, we use 1D convolutional layers coupled with batch-normalization layers. We train this auto-encoder network for 1000 epochs with initial learning rate of $5e - 4$. Similarly, we use linear layers for the generator and train it for 500 epochs with the same initial learning rate. During testing, we set $m = 10$ for the number of samples we generate for every partial shape.

Limitations. Our current network cannot generate very detailed features in a multimodal manner. This is primarily due to the network’s simple architecture and can be resolved by adding refinement networks.

5. Conclusion and Future Work

In this paper, we tackled the problem of multimodal shape completion problem, in which a given partial shape can be completed in several different forms. In fact, when a partial shape is missing a region or part, there is not a single and fixed complete shape and the result can be subject to interpretation. As a result, a method capable of produc-

Figure 11. Qualitative results on real-world scans from ScanNet dataset. The model trained on 3depn dataset is used to perform multimodal completion for real-world partial scans.

ing several plausible results is desired to expand the creativity and diversity of shape completion techniques. Here, we used IMLE that tends to avoid mode collapse as opposed to conditional GAN to make sure that we can deliver diverse completed shapes. Through several experiments, we have shown that our method is superior in respecting the input partial shapes’ geometry while producing more diverse results. We have also shown that without direct supervision, the network learns to respect geometric properties of shapes such as symmetry and it learns to generate missing parts in the right location with accurate geometry. The results of our shape completion are also comparable with the state-of-the-art uni-modal shape completion techniques, which shows that IMLE is capable of respecting features when they are present.

Several avenues can be explored related to shape completion and also IMLE. Our shape completion only adds parts and missing pieces. However, it would be interesting to study how we can add styles and fine details in a multimodal manner. IMLE can also be used in other tasks where a one-to-many mapping is needed. An example can be single view reconstruction, where several objects can be produced according to a single input image.

References

- [1] Matthew Berger, Andrea Tagliasacchi, Lee Seversky, Pierre Alliez, Joshua Levine, Andrei Sharf, and Claudio Silva. State of the art in surface reconstruction from point clouds. In *Eurographics 2014-State of the Art Reports*, volume 1, pages 161–185, 2014. 2, 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] Siddhartha Chaudhuri, Daniel Ritchie, Jiajun Wu, Kai Xu, and Hao Zhang. Learning generative models of 3d structures. *Computer Graphics Forum (Eurographics STAR)*, 2020. 2
- [4] Qimin Chen, Johannes Merz, Aditya Sanghi, Hooman Shayani, Ali Mahdavi-Amiri, and Hao Zhang. UNIST: unpaired neural implicit shape translation network. *CoRR*, abs/2112.05381, 2021. 3
- [5] Xuelin Chen, Baoquan Chen, and Niloy J Mitra. Unpaired point cloud completion on real scans using adversarial training. *arXiv preprint arXiv:1904.00069*, 2019. 2, 3, 7
- [6] Angela Dai, Charles Ruizhongtai Qi, and Matthias Nießner. Shape completion using 3d-encoder-predictor cnns and shape synthesis. In *CVPR*, 2017. 5
- [7] Angela Dai, Charles Ruizhongtai Qi, and Matthias Nießner. Shape completion using 3d-encoder-predictor cnns and shape synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5868–5877, 2017. 3
- [8] Philipp Erler, Paul Guerrero, Stefan Ohrhallinger, Niloy J Mitra, and Michael Wimmer. Points2surf learning implicit surfaces from point clouds. In *ECCV*, pages 108–124. Springer, 2020. 3
- [9] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *arXiv preprint arXiv:1406.2661*, 2014. 2
- [10] Yulan Guo, Hanyun Wang, Qingyong Hu, Hao Liu, Li Liu, and Mohammed Bennamoun. Deep learning for 3d point clouds: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 2020. 3
- [11] Xiaoguang Han, Zhen Li, Haibin Huang, Evangelos Kalogerakis, and Yizhou Yu. High-resolution shape completion using deep neural networks for global structure and local geometry inference. In *Proceedings of the IEEE international conference on computer vision*, pages 85–93, 2017. 3
- [12] Zitian Huang, Yikuan Yu, Jiawen Xu, Feng Ni, and Xinyi Le. Pf-net: Point fractal network for 3d point cloud completion. In *CVPR*, pages 7662–7670, 2020. 3
- [13] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *ICCV*, pages 1125–1134, 2017. 2
- [14] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, volume 7, 2006. 3
- [15] Michael Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. *ACM Transactions on Graphics (ToG)*, 32(3):1–13, 2013. 3
- [16] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 2
- [17] Ke Li*, Shichong Peng*, Tianhao Zhang*, and Jitendra Malik. Multimodal image synthesis with conditional implicit maximum likelihood estimation. *International Journal of Computer Vision*, 2020. 3
- [18] K. Li, T. Zhang, and Jitendra Malik. Diverse image synthesis from semantic layouts via conditional imle. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4219–4228, 2019. 3
- [19] Minghua Liu, Lu Sheng, Sheng Yang, Jing Shao, and Shi-Min Hu. Morphing and sampling network for dense point cloud completion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11596–11603, 2020. 3
- [20] Shakir Mohamed and Balaji Lakshminarayanan. Learning in implicit generative models. *arXiv preprint arXiv:1610.03483*, 2016. 4
- [21] Andrew Nealen, Takeo Igarashi, Olga Sorkine, and Marc Alexa. Laplacian mesh optimization. In *Proceedings of the 4th international conference on Computer graphics and interactive techniques in Australasia and Southeast Asia*, pages 381–389, 2006. 3
- [22] Liang Pan, Xinyi Chen, Zhongang Cai, Junzhe Zhang, Haiyu Zhao, Shuai Yi, and Ziwei Liu. Variational relational point completion network. *arXiv preprint arXiv:2104.10154*, 2021. 3
- [23] Mark Pauly, Niloy J Mitra, Joachim Giesen, Markus H Gross, and Leonidas J Guibas. Example-based 3d scan completion. In *Symposium on Geometry Processing*, number CONF, pages 23–32, 2005. 3
- [24] Shichong Peng and Ke Li. Generating unobserved alternatives. *ArXiv*, abs/2011.01926, 2020. 3
- [25] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, pages 652–660, 2017. 3
- [26] Charles R Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *arXiv preprint arXiv:1706.02413*, 2017. 3
- [27] Marie-Julie Rakotosaona, Vittorio La Barbera, Paul Guerrero, Niloy J Mitra, and Maks Ovsjanikov. Pointcleannet: Learning to denoise and remove outliers from dense point clouds. In *CGF*, volume 39, pages 185–203. Wiley Online Library, 2020. 3
- [28] Jason Rock, Tanmay Gupta, Justin Thorsen, JunYoung Gwak, Daeyun Shin, and Derek Hoiem. Completing 3d object shape from one depth image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2484–2493, 2015. 3
- [29] Olga Sorkine and Daniel Cohen-Or. Least-squares meshes. In *Proceedings Shape Modeling Applications, 2004.*, pages 191–199. IEEE, 2004. 3
- [30] David Stutz and Andreas Geiger. Learning 3d shape completion under weak supervision. *International Journal of Computer Vision*, 128(5):1162–1181, 2020. 2, 3
- [31] Saeid Asgari Taghanaki, Jieliang Luo, Ran Zhang, Ye Wang, Pradeep Kumar Jayaraman, and Krishna Murthy Jatavallabhula. Robustpointset: A dataset for benchmarking robustness of point cloud classifiers. *arXiv preprint arXiv:2011.11572*, 2020. 5, 7

- [32] Lyne P Tchammi, Vineet Kosaraju, Hamid Rezatofighi, Ian Reid, and Silvio Savarese. Topnet: Structural point cloud decoder. In *CVPR*, pages 383–392, 2019. 3
- [33] Weiyue Wang, Qiangui Huang, Suya You, Chao Yang, and Ulrich Neumann. Shape inpainting using 3d generative adversarial network and recurrent convolutional networks. In *ICCV*, pages 2298–2306, 2017. 3
- [34] Xiaogang Wang, Marcelo H Ang Jr, and Gim Hee Lee. Cascaded refinement network for point cloud completion. In *CVPR*, pages 790–799, 2020. 3
- [35] Xiaogang Wang, Yuelang Xu, Kai Xu, Andrea Tagliasacchi, Bin Zhou, Ali Mahdavi-Amiri, and Hao Zhang. Pie-net: Parametric inference of point cloud edges, 2020. 3
- [36] Xin Wen, Tianyang Li, Zhizhong Han, and Yu-Shen Liu. Point cloud completion by skip-attention network with hierarchical folding. In *CVPR*, pages 1939–1948, 2020. 3
- [37] Rundi Wu, Xuelin Chen, Yixin Zhuang, and Baoquan Chen. Multimodal shape completion via conditional generative adversarial networks. In *ECCV*, August 2020. 1, 2, 4, 5, 6, 7, 8
- [38] Shihao Wu, , Hui Huang, Minglun Gong, Matthias Zwicker, and Daniel Cohen-Or. Deep points consolidation. *ACM Transactions on Graphics (Proc. of SIGGRAPH Asia 2015)*, 34(6):176:1–176:13, 2015. 3
- [39] Yaqi Xia, Yan Xia, Wei Li, Rui Song, Kailang Cao, and Uwe Stilla. Asfm-net: Asymmetrical siamese feature matching network for point completion. *arXiv preprint arXiv:2104.09587*, 2021. 3
- [40] Kangxue Yin, Zhiqin Chen, Hui Huang, Daniel Cohen-Or, and Hao Zhang. LOGAN: Unpaired shape transform in latent overcomplete space. *TOG*, 38(6):Article 198, 2019. 3
- [41] Kangxue Yin, Hui Huang, Daniel Cohen-Or, and Hao Zhang. P2p-net: Bidirectional point displacement net for shape transform. *TOG*, 37(4):Article 152, 2018. 3
- [42] Kangxue Yin, Hui Huang, Hao Zhang, Minglun Gong, Daniel Cohen-Or, and Baoquan Chen. Morfit: interactive surface reconstruction from incomplete point clouds with curve-driven topology and geometry control. *ACM Trans. Graph.*, 33(6):202–1, 2014. 3
- [43] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Generative image inpainting with contextual attention. In *CVPR*, pages 5505–5514, 2018. 2
- [44] Wentao Yuan, Tejas Khot, David Held, Christoph Mertz, and Martial Hebert. PCN: Point completion network. In *2018 International Conference on 3D Vision (3DV)*, pages 728–737. IEEE, 2018. 3
- [45] Dongsu Zhang, Changwoon Choi, Jeonghwan Kim, and Young Min Kim. Learning to generate 3D shapes with generative cellular automata. *arXiv preprint arXiv:2103.04130*, 2021. 4, 7, 8
- [46] Wei Zhao, Shuming Gao, and Hongwei Lin. A robust hole-filling algorithm for triangular mesh. *The Visual Computer*, 23(12):987–997, 2007. 3