

Multimodal Shape Completion via Conditional Generative Adversarial Networks

Rundi Wu^{1*}, Xuelin Chen^{2*}, Yixin Zhuang¹, and Baoquan Chen¹

¹ Center on Frontiers of Computing Studies, Peking University

² Shandong University

Abstract. Several deep learning methods have been proposed for completing partial data from shape acquisition setups, i.e., filling the regions that were missing in the shape. These methods, however, only complete the partial shape with a single output, ignoring the ambiguity when reasoning the missing geometry. Hence, we pose a *multi-modal* shape completion problem, in which we seek to complete the partial shape with multiple outputs by learning a one-to-many mapping. We develop the first multimodal shape completion method that completes the partial shape via conditional generative modeling, without requiring paired training data. Our approach distills the ambiguity by conditioning the completion on a learned multimodal distribution of possible results. We extensively evaluate the approach on several datasets that contain varying forms of shape incompleteness, and compare among several baseline methods and variants of our methods qualitatively and quantitatively, demonstrating the merit of our method in completing partial shapes with both diversity and quality.

Keywords: Shape completion, multimodal mapping, conditional generative adversarial network

1 Introduction

Shape completion, which seeks to reason the geometry of the missing regions in incomplete shapes, is a fundamental problem in the field of computer vision, computer graphics and robotics. A variety of solutions now exist for efficient shape acquisition. The acquired shapes, however, are often incomplete, e.g., incomplete work in the user modeling interface, and incomplete scans resulted from occlusion. The power of shape completion enables the use of these incomplete data in downstream applications, e.g., virtual walk-through, intelligent shape modeling, path planning.

With the rapid progress made in deep learning, many data-driven methods have been proposed and demonstrated effective in shape completion [5, 7, 11, 12, 18, 24, 26, 28, 30–33, 36, 37]. However, most approaches in this topic have focused on completing the partial shape with a *single* result, learning a one-to-one

* Equal contribution

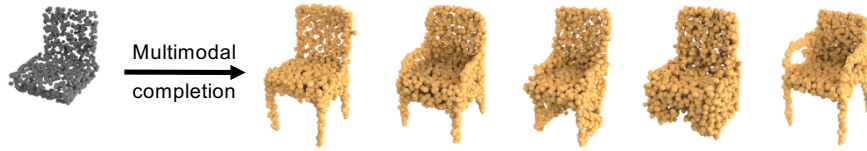


Fig. 1. We present a point-based shape completion network that can complete the partial shape with multiple plausible results. Here we show a sampling of our results, where a partial chair is completed with different types of chairs.

mapping for shape completion. In contrast, we model a distribution of potential completion results, as the shape completion problem is *multimodal* in nature, especially when the incompleteness causes significant ambiguity. For example, as shown in Fig. 1, a partial chair can be completed with different types of chairs. Hence, we pose the multimodal shape completion problem, which seeks to associate each incomplete shape with multiple different complete shapes.

In this work, we propose a first point-based multimodal shape completion method, in which the multimodality of potential outputs is distilled in a low-dimensional latent space, enabling random sampling of completion results at inference time. The challenge to achieve this is two-fold. First, modeling the multimodality in the high-dimensional shape space and mapping it into a low-dimensional latent space is challenging. A common problem is *mode collapse*, where only a subset of the modes are represented in the low-dimensional space. Second, the ground-truth supervision data, on which most learning-based methods rely for success, is extremely hard to acquire in our problem. Without the availability of particular supervision data (i.e., for each training incomplete shape, multiple corresponding complete shapes are required), it is challenging to learn the multimodal mapping for shape completion in an unsupervised manner *without* any paired data.

We address the challenge by completing the partial shape in a conditional generative modeling setting. We design a conditional generative adversarial network (cGAN) wherein a generator learns to map incomplete training data, combined with a latent vector sampled from a learned multimodal shape distribution, to a suitable latent representation such that a discriminator cannot differentiate between the mapped latent variables and the latent variables obtained from complete training data (i.e., complete shape models). An encoder is introduced to encode mode latent vectors from complete shapes, learning the multimodal distribution of all possible outputs. We further apply this encoder to the completion output to extract and recover the input latent vector, forcing the bijective mapping between the latent space and the output modes. The mode encoder is trained to encode the multimodality in an explicit manner (Section 3.3), alleviating the aforementioned mode collapse issue. By conditioning the generation of completion results on the learned multimodal shape distribution, we achieve multimodal shape completion.

We extensively evaluate our method on several datasets that contain varying forms of shape incompleteness. We compare our method against several baseline methods and variants of our method, and rate the different competing methods using a combination of established metrics. Our experiments demonstrate the superiority of our method compared to other alternatives, producing completion results with both high diversity and quality, all the while remaining faithful to the partial input shape.

2 Related work

Shape completion With the advancement of deep learning in 3D domain, many deep learning methods have been proposed to address the shape completion challenge. Following the success of CNN-based 2D image completion networks, 3D convolutional neural networks applied on voxelized inputs have been widely adopted for 3D shape completion task [7, 12, 24, 26, 28, 30–33]. To avoid geometric information loss resulted from quantizing shapes into voxel grids, several approaches [5, 11, 18, 36, 37] have been develop to work directly on point sets to reason the missing geometry. While all these methods resort to learning a parameterized model (i.e., neural networks) as a mapping from incomplete shapes to completed shapes, the learned mapping function remains injective. Consequently, these methods can only complete the partial shape with a single deterministic result, ignoring the ambiguity of the missing regions.

Generative modeling The core of generative modeling is parametric modeling of the data distribution. Several classical approaches exist for tackling the generative modeling problem, e.g., restricted Boltzmann machines [25], and autoencoders [13, 29]. Since the introduction of Generative Adversarial Networks (GANs) [8], it has been widely adopted for a variety of generative tasks. In 2D image domain, researchers have utilized GANs in tasks ranging from image generation [19, 2, 3], image super-resolution [15], to image inpainting in 2D domain. In the context of 3D, a lot of effort has also been put on the task of content generation [1, 6, 9, 21]. The idea of using GANs to generatively model the missing regions for shape completion has also been explored in the pioneering works [5, 11]. In this work, we achieve multimodal shape completion by utilizing GANs to reason the missing geometry in a conditional generative modeling setting.

Deep learning on points Our method is built upon recent success in deep neural networks for point cloud representation learning. Although many improvements to PointNet [22] have been proposed [16, 17, 23, 27, 38], the simplicity and effectiveness of PointNet and its extension PointNet++ make them popular in many other analysis tasks [10, 34–36]. To achieve point cloud generation, [1] proposed to train GANs in the latent space produced by a PointNet-based autoencoder, reporting significant performance gain in point cloud generation. Similar to [5], we also leverage the power of such point cloud generation model to complete shapes, but conditioning the generation on a learned distribution of potential completion outputs for multimodal shape completion.

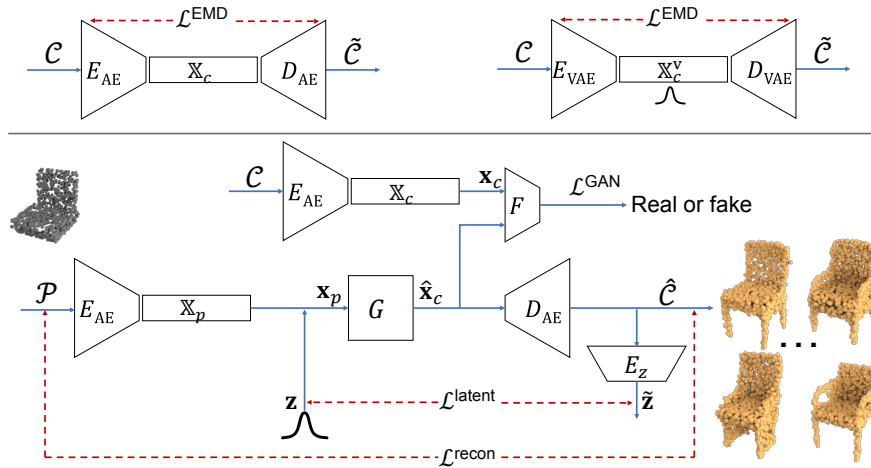


Fig. 2. The proposed network architecture for learning multimodal shape completion. We use the encoder E_{VAE} of a shape variational autoencoder as the mode encoder E_z to encode the multimodality from shapes explicitly.

3 Method

Given a partial shape domain $\mathcal{P} \subset \mathbb{R}^{K \times 3}$, we seek to learn a multimodal mapping from \mathcal{P} to the complete shape domain $\mathcal{C} \subset \mathbb{R}^{N \times 3}$, achieving the goal of multimodal shape completion. Unlike existing methods, which are primarily limited to producing a single deterministic completion output $\hat{\mathbf{C}}$ from a partial shape $\mathbf{P} \in \mathcal{P}$, our method learns a mapping that could sample the completion $\hat{\mathbf{C}}$ from the conditional distribution of possible completions, producing diverse completion results. During training, our method only requires access to a set of partial point sets, and a set of complete point sets. It is important to note that there is no any paired completion instances in \mathcal{C} for point sets in \mathcal{P} .

Following the spirit of [5], without paired training data, we address the multimodal shape completion problem via adversarial training on the latent space learned from point sets, while introducing a low-dimensional code $\mathbf{z} \in \mathbb{R}^Z$, of which the latent space is learned by modeling the multimodality presented in possible outputs, as a conditional input in addition to the partial shape. To enable stochastic sampling, we desire \mathbf{z} to be drawn from a prior distribution $p(\mathbf{z})$; a standard Gaussian distribution $\mathcal{N}(0, \mathcal{I})$ is used in this work.

More specifically, we learn two class-specific point set manifolds, \mathbb{X}_p for the partial shapes, and \mathbb{X}_c for the complete shapes. Solving the multimodal shape completion problem then amounts to learning a mapping $\mathbb{X}_p \rightarrow \mathbb{X}_c$ in a conditional generative modeling setting between respective latent spaces. Hence, we train a generator $G: (\mathbb{X}_p, p(\mathbf{z})) \rightarrow \mathbb{X}_c$, to perform the multimodal mapping with the latent code \mathbf{z} as a condition input. In absence of paired training data, we opt to use adversarial training to complete shapes in a generative modeling set-

ting. Furthermore, to force the generator to use the latent code \mathbf{z} , we introduce an encoder E_z to recover \mathbf{z} from the completion output, forcing the connection between the latent space and the shape space to be invertible. Fig. 2 shows the setup of the proposed multimodal shape completion network. All network modules are detailed next.

3.1 Learning latent spaces for point sets

The latent space of a given set of point sets is obtained by training an autoencoder, which encodes the given input to a low-dimension latent feature and then decodes to reconstruct the original input.

For point sets coming from the complete point sets \mathcal{C} , we learn an encoder network E_{AE} that maps \mathbf{C} from the original parameter space $\mathbb{R}^{N \times 3}$, defined by concatenating the coordinates of the N points, to a lower-dimensional latent space \mathbb{X}_c . A decoder network D_{AE} performs the inverse transformation back to $\mathbb{R}^{N \times 3}$ giving us a reconstructed point set $\hat{\mathbf{C}}$ with also N points. The encoder-decoders are trained with reconstruction loss:

$$\mathcal{L}^{\text{EMD}} = \mathbb{E}_{\mathbf{C} \sim p(\mathbf{C})} d^{\text{EMD}}(\mathbf{C}, D_{\text{AE}}(E_{\text{AE}}(\mathbf{C}))), \quad (1)$$

where $\mathbf{C} \sim p(\mathbf{C})$ denotes point set samples drawn from the set of complete point sets, $d^{\text{EMD}}(X_1, X_2)$ is the Earth Mover’s Distance (EMD) between point sets X_1, X_2 . Once trained, the network weights are held fixed and the *complete* latent code $\mathbf{x}_c = E_{\text{AE}}(\mathbf{C})$, $\mathbf{x}_c \in \mathbb{X}_c$ for a complete point set \mathbf{C} provides a compact representation for subsequent training and implicitly captures the manifold of complete data. As for the point set coming from the partial point sets \mathcal{P} , instead of training another autoencoder for its latent parameterization, we directly feed the partial point sets to E_{AE} obtained above for producing *partial* latent space \mathbb{X}_p , which in [5] is proved to yield better performance in subsequent adversarial training. Note that, to obtain $\mathbf{x}_p \in \mathbb{X}_p$, we duplicate the partial point set of K points to align with the number of complete point set before feed it to E_{AE} .

3.2 Learning multimodal mapping for shape completion

Next, we setup a min-max game between the generator and the discriminator to perform the multimodal mapping between the latent spaces. The generator is trained to fool the discriminator such that the discriminator fails to reliably tell if the latent variable comes from original \mathbb{X}_c or the remapped \mathbb{X}_p . The mode encoder is trained to model the multimodal distribution of possible complete point sets, and is further applied to the completion output to encode and recover the input latent vector, encouraging the use of conditional mode input.

Formally, the latent representation of the input partial shape $\mathbf{x}_p = E_{\text{AE}}(\mathbf{P})$, along with a Gaussian-sampled condition \mathbf{z} , is mapped by the generator to $\hat{\mathbf{x}}_c = G(\mathbf{x}_p, \mathbf{z})$. Then, the task of the discriminator F is to distinguish between latent representations $\hat{\mathbf{x}}_c$ and $\mathbf{x}_c = E_{\text{AE}}(\mathbf{C})$. The mode encoder E_z will encode the completion point set, which can be decoded from the completion latent code $\hat{\mathbf{C}} =$

$D_{\text{AE}}(\hat{\mathbf{x}}_c)$, to reconstruct the conditional input $\tilde{\mathbf{z}} = E_z(\hat{\mathbf{C}})$. We train the mapping function using a GAN. Given training examples of complete latent variables \mathbf{x}_c , remapped partial latent variables $\hat{\mathbf{x}}_c$, and Gaussian samples \mathbf{z} , we seek to optimize the following training losses over the generator G , the discriminator F , and the encoder E_z :

Adversarial loss. We add the adversarial loss to train the generator and discriminator. In our implementation, we use least square GAN [19] for stabilizing the training. Hence, the adversarial losses minimized for the generator and the discriminator are defined as:

$$\mathcal{L}_F^{\text{GAN}} = \mathbb{E}_{\mathbf{C} \sim p(\mathbf{C})} [F(E_{\text{AE}}(\mathbf{C})) - 1]^2 + \mathbb{E}_{\mathbf{P} \sim p(\mathbf{P}), \mathbf{z} \sim p(\mathbf{z})} [F(G(E_{\text{AE}}(\mathbf{P}), \mathbf{z}))]^2 \quad (2)$$

$$\mathcal{L}_G^{\text{GAN}} = \mathbb{E}_{\mathbf{P} \sim p(\mathbf{P}), \mathbf{z} \sim p(\mathbf{z})} [F(G(E_{\text{AE}}(\mathbf{P}), \mathbf{z})) - 1]^2, \quad (3)$$

where $\mathbf{C} \sim p(\mathbf{C})$, $\mathbf{P} \sim p(\mathbf{P})$ and $\mathbf{z} \sim p(\mathbf{z})$ denotes samples drawn from, respectively, the set of complete point sets, the set of partial point sets, and $\mathcal{N}(0, \mathcal{I})$.

Partial reconstruction loss. Similar to previous work [5], we add a reconstruction loss to encourage the generator to *partially* reconstruct the partial input, so that the completion output is faithful to the input partial:

$$\mathcal{L}_G^{\text{recon}} = \mathbb{E}_{\mathbf{P} \sim p(\mathbf{P}), \mathbf{z} \sim p(\mathbf{z})} [d^{\text{HL}}(\mathbf{P}, D_{\text{AE}}(G(E_{\text{AE}}(\mathbf{P}), \mathbf{z})))]^2, \quad (4)$$

where d^{HL} denotes the unidirectional Hausdorff distance from the partial point set to the completion point set.

Latent space reconstruction. A reconstruction loss on the \mathbf{z} latent space is also added to encourage G to use the conditional mode vector \mathbf{z} :

$$\mathcal{L}_{G, E_z}^{\text{latent}} = \mathbb{E}_{\mathbf{P} \sim p(\mathbf{P}), \mathbf{z} \sim p(\mathbf{z})} [\|\mathbf{z}, E_z(D_{\text{AE}}(G(E_{\text{AE}}(\mathbf{P}), \mathbf{z})))\|_1], \quad (5)$$

Hence, our full objective function for training the multimodal shape completion network is described as:

$$\underset{(G, E_z)}{\operatorname{argmin}} \underset{F}{\operatorname{argmax}} \mathcal{L}_F^{\text{GAN}} + \mathcal{L}_G^{\text{GAN}} + \alpha \mathcal{L}_G^{\text{recon}} + \beta \mathcal{L}_{G, E_z}^{\text{latent}}, \quad (6)$$

where α and β are importance weights for the partial reconstruction loss and the latent space reconstruction loss, respectively.

3.3 Explicitly-encoded multimodality

To model the multimodality of possible completion outputs, we resort to an explicit multimodality encoding strategy, in which E_z is trained as a part to explicitly reconstruct the complete shapes. More precisely, a variational autoencoder ($E_{\text{VAE}}, D_{\text{VAE}}$) is pre-trained to encode complete shapes into a standard Gaussian distribution $\mathcal{N}(0, \mathcal{I})$ and then decode to reconstruct the original shapes. Once trained, E_{VAE} can encode complete point sets as $\mathbf{x}_c^v \in \mathbb{X}_c^v$, and can be used to recover the conditional input \mathbf{z} from the completion output. Hence, the mode encoder is set to $E_z = E_{\text{VAE}}$ and held fixed during the GAN training.

Another strategy is to implicitly encode the multimodality, in which the E_z is jointly trained to map complete shapes into a latent space without being trained as a part of explicit reconstruction. Although it has been shown effective to improve the diversity in the multimodal mapping learning [39, 40], we demonstrate that using explicitly-encoded multimodality in our problem yields better performance. In Section 4.2, we present the comparison against variants of using implicit multimodality encoding.

3.4 Implementation details

In our experiments, a partial shape is represented by $K = 1024$ points and a complete shape by $N = 2048$ points. The the point set (variational) autoencoder follows [1, 5]: using a PointNet[22] as the encoder and a 3-layer MLP as the decoder. The autoencoder encodes a point set into a latent vector of fixed dimension $|\mathbf{x}| = 128$. Similar to [1, 5], we use a 3-layer MLP for both generator G and discriminator F . The E_z also uses the PointNet to map a point set into a latent vector \mathbf{z} , of which the length we set to $|\mathbf{z}| = 64$. Unless specified, the trade-off parameters α and β in Eq. 6 are set to 6 and 7.5, respectively, in our experiments. For training the point set (variational) autoencoder, we use the Adam optimizer[14] with an initial learning rate 0.0005, $\beta_1 = 0.9$ and train 2000 epochs with a batch size of 200. To train the GAN, we use the Adam optimizer with an initial learning rate 0.0005, $\beta_1 = 0.5$ and train for a maximum of 1000 epochs with a batch size of 50. More details about each network module are presented in the supplementary material.

4 Experiments

In this section, we present results produced from our method on multimodal shape completion, and both quantitative and qualitative comparisons against several baseline methods and variants of our method, along with a set of experiments for evaluating different aspects of our method.

Datasets. Three datasets are derived to evaluate our method under different forms of shape incompleteness: (A) **PartNet** dataset simulates part-level incompleteness in the user modeling interface. With the semantic segmentation provided in the original PartNet dataset [20], for each provided point set, we remove points of randomly selected parts to create a partial point set with at least one part. (B) **PartNet-Scan** dataset resembles the scenario where the partial scan suffers from part-level incompleteness. For each shape in [20], we randomly remove parts and virtually scan residual parts to obtain a partial scan with part-level incompleteness. (C) **3D-EPN** dataset [7] is derived from ShapeNet [4] and provides simulated partial scans with arbitrary incompleteness. Scans are represented as Signed Distance Field but we only use the provided point cloud representations. Last, the complete point sets provided in PartNet [20] serve as the complete training data for the first two datasets, while for **3D-EPN** dataset, we use the complete virtual scan of ShapeNet objects as the complete training




Fig. 3. Our multimodal shape completion results. We show result examples, where the input partial shape is colored in grey and is followed by five different completions in yellow. From top to bottom: **PartNet** (rows 1-3), **PartNet-Scan** (rows 4-6), and **3D-EPN** (rows 7-9).

data. We use Chair, Table and Lamp categories for **PartNet** and **PartNet-scan**, and use Chair, Airplane and Table categories for **3D-EPN**. In all our experiments, we train separate networks for each category in each dataset. More details about data processing can be found in the supplementary material.

Evaluation measures. For each partial shape in the test set, we generate $k = 10$ completion results and adopt the following measures for quantitative evaluation:

- Minimal Matching Distance (**MMD**) measures the *quality* of the completed shape. We calculate the Minimal Matching Distance (as described in [1]) between the set of completion shapes and the set of test shapes.
- Total Mutual Difference (**TMD**) measures the completion *diversity* for a partial input shape, by summing up all the difference among the k completion shapes of the same partial input. For each shape i in the k generated shapes, we calculate its average Chamfer distance d_i^{CD} to the other $k - 1$ shapes. The diversity is then calculated as the sum $\sum_{i=1}^k d_i^{\text{CD}}$.
- Unidirectional Hausdorff Distance (**UHD**) measures the completion *fidelity* to the input partial. We calculate the average Hausdorff distance from the input partial shape to each of the k completion results.



Fig. 4. Shape completion guided by reference shapes. The completion result varies accordingly when the reference shape changes.

More in-depth description can be found in the supplementary material.

4.1 Multimodal completion results

We first present qualitative results of our method on multimodal shape completion, by using randomly sampled \mathbf{z} from the standard Gaussian distribution. Figure 3 shows a collection of our multimodal shape completion results on the aforementioned datasets. More visual examples can be found in supplementary material.

To allow more explicit control over the modes in completion results, the mode condition \mathbf{z} can also be encoded from a user-specified shape. As shown in Figure 4, this enables us to complete the partial shape under the guidance of a given reference shape. The quantitative evaluation of our results, along with the comparison results, is presented next.

4.2 Comparison results

We present both qualitative and quantitative comparisons against baseline methods and variants of our method:

- `pc12pc1` [5], which also uses GANs to complete via generative modeling without paired training data. Without the conditional input, this method, however, cannot complete with diverse shapes for a single partial input.
- `KNN-latent`, which retrieves a desired number of best candidates from the latent space formed by our complete point set autoencoder, using k-nearest neighbor search algorithm.
- `Ours-im-12z`, a variant of our method as described in Section 3.3, jointly trains the E_z to implicitly model the multimodality by mapping complete data into a low-dimensional space. The E_z can either take input as complete latent codes (denoted by `Ours-im-12z`) or complete point clouds (denoted by `Ours-im-pc2z`) to map to \mathbf{z} space.

PartNet	MMD (lower is better)				TMD (higher is better)				UHD (lower is better)			
Method	Chair	Lamp	Table	Avg.	Chair	Lamp	Table	Avg.	Chair	Lamp	Table	Avg.
pcl2pcl	1.90	2.50	1.90	2.10	0.00	0.00	0.00	0.00	4.88	4.64	4.78	4.77
KNN-latent	1.39	1.72	1.30	1.47	2.28	4.18	2.36	2.94	8.58	8.47	7.61	8.22
Ours-im-12z	1.74	2.36	1.68	1.93	3.74	2.68	3.59	3.34	8.41	6.37	7.21	7.33
Ours-im-pc2z	1.90	2.55	1.54	2.00	1.01	0.56	0.51	0.69	6.65	5.40	5.38	5.81
Ours	1.52	1.97	1.46	1.65	2.75	3.31	3.30	3.12	6.89	5.72	5.56	6.06

PartNet-Scan	MMD (lower is better)				TMD (higher is better)				UHD (lower is better)			
Method	Chair	Lamp	Table	Avg.	Chair	Lamp	Table	Avg.	Chair	Lamp	Table	Avg.
pcl2pcl	1.96	2.36	2.09	2.14	0.00	0.00	0.00	0.00	5.20	5.34	4.73	5.09
KNN-latent	1.40	1.80	1.39	1.53	3.09	4.47	2.85	3.47	8.79	8.41	7.50	8.23
Ours-im-12z	1.79	2.58	1.92	2.10	3.85	3.18	4.75	3.93	7.88	6.39	7.40	7.22
Ours-im-pc2z	1.65	2.75	1.84	2.08	1.91	0.50	1.86	1.42	7.50	5.36	5.68	6.18
Ours	1.53	2.15	1.58	1.75	2.91	4.16	3.88	3.65	6.93	5.74	6.24	6.30

3D-EPN	MMD (lower is better)				TMD (higher is better)				UHD (lower is better)			
Method	Chair	Plane	Table	Avg.	Chair	Plane	Table	Avg.	Chair	Plane	Table	Avg.
pcl2pcl	1.81	1.01	3.12	1.98	0.00	0.00	0.00	0.00	5.31	9.71	9.03	8.02
KNN-latent	1.45	0.93	2.25	1.54	2.24	1.13	3.25	2.21	8.94	9.54	12.70	10.40
Ours-im-12z	1.91	0.86	2.78	1.80	3.84	2.17	4.27	3.43	9.53	10.60	9.36	9.83
Ours-im-pc2z	1.61	0.91	3.19	1.90	1.51	0.82	1.67	1.33	8.18	9.55	8.50	8.74
Ours	1.61	0.82	2.57	1.67	2.56	2.03	4.49	3.03	8.33	9.59	9.03	8.98

Table 1. Quantitative comparison results on PartNet (top), PartNet-Scan (middle) and 3D-EPN (bottom). Top two methods on each measure are highlighted. Note that MMD (quality), TMD (diversity) and UHD (fidelity) presented in the tables are multiplied by 10^3 , 10^2 and 10^2 , respectively.

- Ours-im-pc2z, in which, as stated above, the E_z takes complete point clouds as input to encode the multimodality in an implicit manner.

More details of the above methods can be found in the supplementary material.

We present quantitative comparison results in Table 1. We can see that, by rating these methods using the combination of the established metrics, our method outperforms other alternatives with high completion quality (low MMD) and diversity (high TMD), while remains faithful to the partial input. More specifically, pcl2pcl has the best fidelity (lowest UHD) to the input partial shape, the completion results, however, present no diversity; Ours-im-12z presents the best completion diversity, but fails to produce high partial matching between the partial input and the completion results; Ours-im-pc2z suffers from severe mode collapse; our method, by adopting the explicit multimodality encoding, can complete the partial shape with high completion fidelity, quality and diversity.

To better understand the position of our method among those competing methods when rating with the established metrics, we present Fig. 6 to visualize the performance of each method in comparisons using combinations of the metrics. Note that the percentages in Fig. 6 are obtained by taking the performance

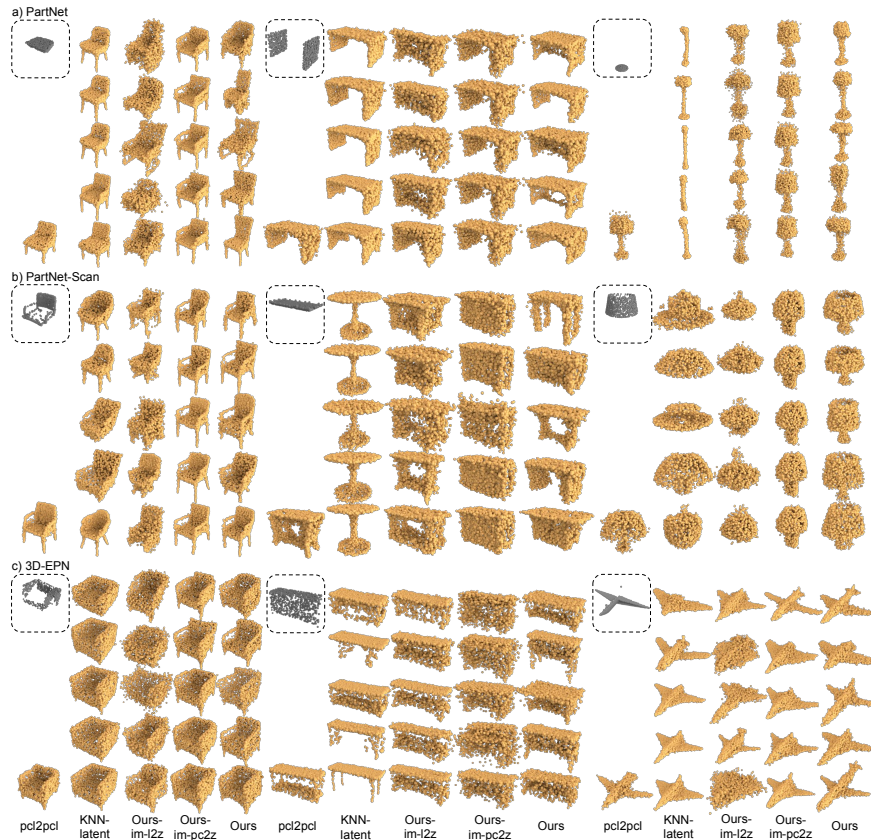


Fig. 5. Qualitative comparison on each three categories of **PartNet** (top), **PartNet-Scan** (middle) and **3D-EPN** (bottom) dataset. Our method produces results that are both diverse and plausible.

of our method as reference and calculating the relative performance of other methods to ours.

We also show qualitative comparison results in Fig. 5. Compared to other methods, our method shows the superiority on multimodal shape completion: high completion diversity and quality while still remains faithful to the input, as consistently exhibited in the quantitative comparisons. Note that **pc12pc1** cannot complete the partial shape with multiple outputs, thus only a single completion result is shown in the figure.

4.3 Results on real scans

The nature of unpaired training data setting enables our method to be directly trained on real scan data with ease. We have trained and tested our method on the real-world scans provided in **pc12pc1**, which contains around 550 chairs. We

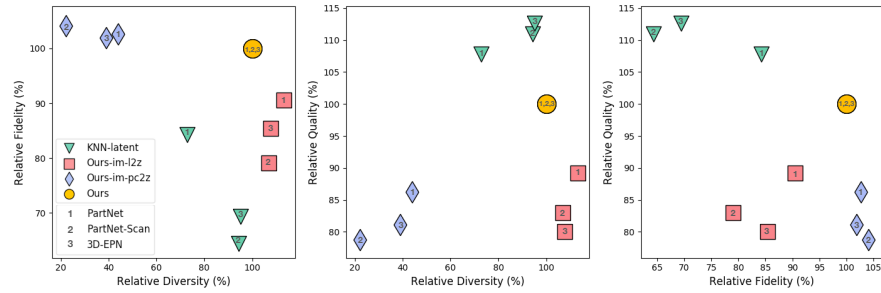


Fig. 6. Comparisons using metrics combinations. Our results present high diversity, quality and fidelity in comparisons using combinations of metrics. *Relative* performance is plotted, and pc12p1c is excluded as it fails to present completion diversity.

randomly picked 50 of them for testing and used the rest for training. Qualitative results are shown in Fig. 7. For quantitative evaluation, our method achieved MMD of 2.42×10^{-3} , TMD of 3.17×10^{-2} and UHD of 8.60×10^{-2} .

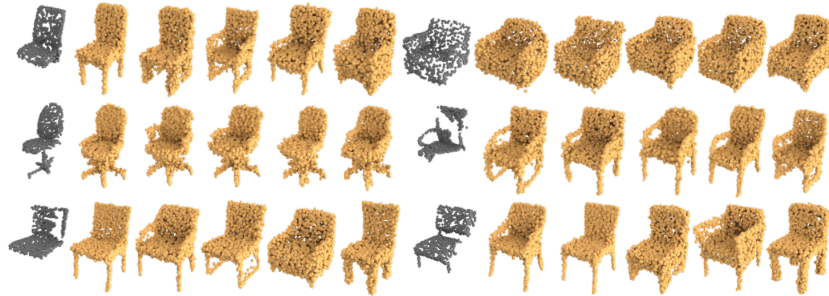


Fig. 7. Examples of our multimodal shape completion results on real scan data.

4.4 More experiments

We also present more experiments conducted to evaluate different aspects of our method.

Effect of trade-off parameters. Minimizing the \mathcal{L}^{GAN} , $\mathcal{L}^{\text{recon}}$ and $\mathcal{L}^{\text{latent}}$ in Eq. 6 corresponds to the improvement of completion quality, fidelity, and diversity, respectively. However, conflict exists when simultaneously optimizing these three terms, e.g., maximizing the completion fidelity to the partial input would potentially compromise the diversity presented in the completion results, as part of the completion point set is desired to be fixed. Hence, we conduct experiments and show in Fig. 8 to see how the diversity and fidelity vary with respect to the change of trade-off parameters α and β .

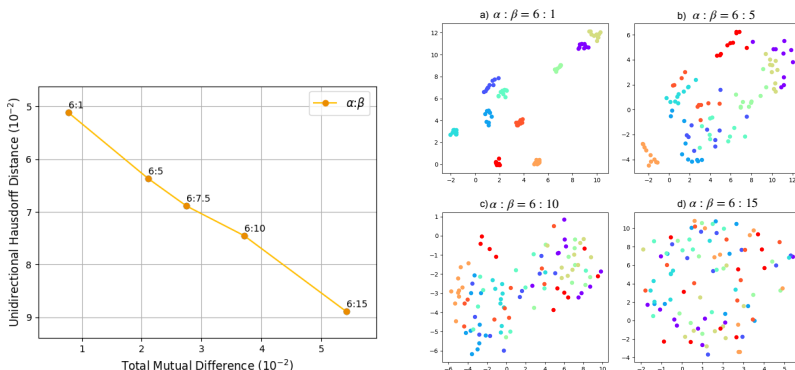


Fig. 8. Effect of trade-off parameters. Left: the completion fidelity (UHD) decreases and the completion diversity (TMD) increases as we set β to be larger. Right: t-SNE visualization of completion latent vectors under different parameter settings. Dots with identical color indicates completion latent vectors resulted from the same partial shape. a) setting a larger weight for $\mathcal{L}^{\text{recon}}$ to encourage completion fidelity leads to mode collapse within the completion results of each partial input; b)-d) setting larger weights for $\mathcal{L}^{\text{latent}}$ to encourage completion diversity results in more modes.

Effect of input incompleteness. The shape completion problem has more ambiguity when the input incompleteness increases. Thus, it is desired that the model can complete with more diverse results under increasing incompleteness of the input. To this end, we test our method on **PartNet** dataset, where the incompleteness can be controlled by the number of missing parts. Figure 9 shows how diversity in our completion results changes with respect to the number of missing parts in the partial input.

5 Conclusion

We present a point-based shape completion framework that can produce multiple completion results for the incomplete shape. At the heart of our approach lies a generative network that completes the partial shape in a conditional generative modeling setting. The generation of the completions is conditioned on a learned mode distribution that is explicitly distilled from complete shapes. We extensively evaluate our method on several datasets containing different forms of incompleteness, demonstrating that our method consistently outperforms other alternative methods by producing completion results with both high diversity and quality for a single partial input.

The work provides a novel approach to indoor scene modeling suggestions and large-scale scene scan completions, rather than object-oriented completions. While the completion diversity in our method is already demonstrated, the explicit multimodality encoding module is, nonetheless, suboptimal and could be potentially improved. Our method shares the same limitations as many of the

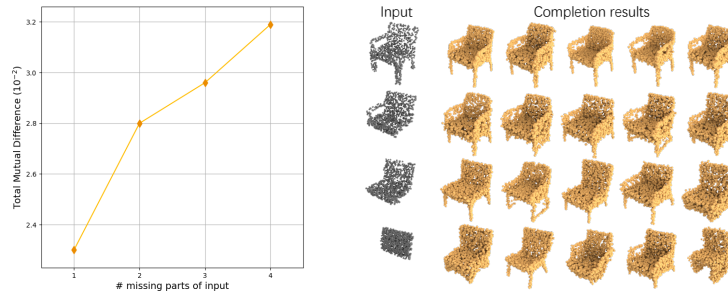


Fig. 9. Effect of input incompleteness. Completion results tend to show more variation as the input incompleteness increases. Left: the completion diversity (TMD) increases as the number of missing parts in the partial input rises. Right: we show an example of completion results for each input incompleteness.

counterparts: not producing shapes with fine-scale details and requiring the input to be canonically oriented. A promising future direction would be to investigate the possibility of guiding the mode encoder network to concentrate on the missing regions for multimodality extraction, alleviating the compromise between the completion diversity and the faithfulness to the input.

Acknowledgement

We thank the anonymous reviewers for their valuable comments. This work was supported in part by National Key R&D Program of China (2019YFF0302902) and NSFC (61902007).

References

1. Achlioptas, P., Diamanti, O., Mitliagkas, I., Guibas, L.: Learning representations and generative models for 3d point clouds. In: International Conference on Machine Learning (ICML). pp. 40–49 (2018)
2. Alec, R., Luke, M., Soumith, C.: Unsupervised representation learning with deep convolutional generative adversarial networks. In: International Conference on Learning Representations (ICLR) (2016)
3. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein generative adversarial networks. In: International Conference on Machine Learning (ICML). pp. 214–223 (2017)
4. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., Xiao, J., Yi, L., Yu, F.: ShapeNet: An Information-Rich 3D Model Repository. Tech. Rep. arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago (2015)
5. Chen, X., Chen, B., Mitra, N.J.: Unpaired point cloud completion on real scans using adversarial training. In: International Conference on Learning Representations (ICLR) (2020)

6. Chen, Z., Zhang, H.: Learning implicit fields for generative shape modeling. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5939–5948 (2019)
7. Dai, A., Ruizhongtai Qi, C., Nießner, M.: Shape completion using 3d-encoder-predictor cnns and shape synthesis. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5868–5877 (2017)
8. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Advances in Neural Information Processing Systems (NeurIPS). pp. 2672–2680 (2014)
9. Groueix, T., Fisher, M., Kim, V.G., Russell, B., Aubry, M.: AtlasNet: A Papier-Mâché Approach to Learning 3D Surface Generation. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
10. Guerrero, P., Kleiman, Y., Ovsjanikov, M., Mitra, N.J.: Pepnet learning local shape properties from raw point clouds. In: Computer Graphics Forum. vol. 37, pp. 75–85 (2018)
11. Gurusurthy, S., Agrawal, S.: High fidelity semantic shape completion for point clouds using latent optimization. pp. 1099–1108. IEEE (2019)
12. Han, X., Li, Z., Huang, H., Kalogerakis, E., Yu, Y.: High-resolution shape completion using deep neural networks for global structure and local geometry inference. In: International Conference on Computer Vision (ICCV). pp. 85–93 (2017)
13. Hinton, G.E., Salakhutdinov, R.R.: Reducing the dimensionality of data with neural networks. *science* **313**(5786), 504–507 (2006)
14. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization (2014)
15. Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., et al.: Photo-realistic single image super-resolution using a generative adversarial network. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4681–4690 (2017)
16. Li, J., Chen, B.M., Hee Lee, G.: So-net: Self-organizing network for point cloud analysis. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9397–9406 (2018)
17. Li, Y., Bu, R., Sun, M., Wu, W., Di, X., Chen, B.: Pointcnn: Convolution on x-transformed points. In: Advances in Neural Information Processing Systems (NeurIPS). pp. 820–830 (2018)
18. Liu, M., Sheng, L., Yang, S., Shao, J., Hu, S.M.: Morphing and sampling network for dense point cloud completion. In: Association for the Advancement of Artificial Intelligence (AAAI) (2019)
19. Mao, X., Li, Q., Xie, H., Lau, R.Y., Wang, Z., Paul Smolley, S.: Least squares generative adversarial networks. In: International Conference on Computer Vision (ICCV). pp. 2794–2802 (2017)
20. Mo, K., Zhu, S., Chang, A.X., Yi, L., Tripathi, S., Guibas, L.J., Su, H.: PartNet: A large-scale benchmark for fine-grained and hierarchical part-level 3D object understanding. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
21. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning continuous signed distance functions for shape representation. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 165–174 (2019)
22. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 652–660 (2017)

23. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In: *Advances in Neural Information Processing Systems (NeurIPS)*. pp. 5099–5108 (2017)
24. Sharma, A., Grau, O., Fritz, M.: Vconv-dae: Deep volumetric shape learning without object labels. In: *European Conference on Computer Vision (ECCV)*. pp. 236–250 (2016)
25. Smolensky, P.: *Information processing in dynamical systems: Foundations of harmony theory*. Tech. rep., Colorado Univ at Boulder Dept of Computer Science (1986)
26. Stutz, D., Geiger, A.: Learning 3d shape completion under weak supervision. *International Journal of Computer Vision (IJCV)* pp. 1–20 (2018)
27. Su, H., Jampani, V., Sun, D., Maji, S., Kalogerakis, E., Yang, M.H., Kautz, J.: Splatnet: Sparse lattice networks for point cloud processing. In: *Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 2530–2539 (2018)
28. Thanh Nguyen, D., Hua, B.S., Tran, K., Pham, Q.H., Yeung, S.K.: A field model for repairing 3d shapes. In: *Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 5676–5684 (2016)
29. Vincent, P., Larochelle, H., Bengio, Y., Manzagol, P.A.: Extracting and composing robust features with denoising autoencoders. In: *International Conference on Machine Learning (ICML)*. pp. 1096–1103 (2008)
30. Wang, W., Huang, Q., You, S., Yang, C., Neumann, U.: Shape inpainting using 3d generative adversarial network and recurrent convolutional networks. In: *International Conference on Computer Vision (ICCV)*. pp. 2298–2306 (2017)
31. Wu, R., Zhuang, Y., Xu, K., Zhang, H., Chen, B.: Pq-net: A generative part seq2seq network for 3d shapes. *arXiv preprint arXiv:1911.10949* (2019)
32. Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3d shapenets: A deep representation for volumetric shapes. In: *Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1912–1920 (2015)
33. Yang, B., Rosa, S., Markham, A., Trigoni, N., Wen, H.: 3d object dense reconstruction from a single depth view. *arXiv preprint arXiv:1802.00411* **1**(2), 6 (2018)
34. Yin, K., Huang, H., Cohen-Or, D., Zhang, H.: P2p-net: Bidirectional point displacement net for shape transform. *ACM Transactions on Graphics (TOG)* **37**(4), 1–13 (2018)
35. Yu, L., Li, X., Fu, C.W., Cohen-Or, D., Heng, P.A.: Ec-net: an edge-aware point set consolidation network. In: *European Conference on Computer Vision (ECCV)*. pp. 386–402 (2018)
36. Yu, L., Li, X., Fu, C.W., Cohen-Or, D., Heng, P.A.: Pu-net: Point cloud upsampling network. In: *Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 2790–2799 (2018)
37. Yuan, W., Khot, T., Held, D., Mertz, C., Hebert, M.: Pcn: Point completion network. In: *International Conference on 3D Vision (3DV)*. pp. 728–737 (2018)
38. Zaheer, M., Kottur, S., Ravanbakhsh, S., Póczos, B., Salakhutdinov, R.R., Smola, A.J.: Deep sets. In: *Advances in Neural Information Processing Systems (NeurIPS)*. pp. 3391–3401 (2017)
39. Zhu, J.Y., Zhang, R., Pathak, D., Darrell, T., Efros, A.A., Wang, O., Shechtman, E.: Toward multimodal image-to-image translation. In: *Advances in Neural Information Processing Systems (NeurIPS)*. pp. 465–476 (2017)
40. Zhu, J.Y., Zhang, Z., Zhang, C., Wu, J., Torralba, A., Tenenbaum, J., Freeman, B.: Visual object networks: image generation with disentangled 3d representations. In: *Advances in Neural Information Processing Systems (NeurIPS)*. pp. 118–129 (2018)