

P2C: Self-Supervised Point Cloud Completion from Single Partial Clouds

Ruikai Cui^{1*} Shi Qiu^{1†} Saeed Anwar² Jiawei Liu¹ Chaoyue Xing¹ Jing Zhang¹ Nick Barnes¹

¹Australian National University ²King Fahd University of Petroleum and Minerals

Abstract

Point cloud completion aims to recover the complete shape based on a partial observation. Existing methods require either complete point clouds or multiple partial observations of the same object for learning. In contrast to previous approaches, we present **Partial2Complete (P2C)**, the first self-supervised framework that completes point cloud objects using training samples consisting of only a single incomplete point cloud per object. Specifically, our framework groups incomplete point clouds into local patches as input and predicts masked patches by learning prior information from different partial objects. We also propose *Region-Aware Chamfer Distance* to regularize shape mismatch without limiting completion capability, and devise the *Normal Consistency Constraint* to incorporate a local planarity assumption, encouraging the recovered shape surface to be continuous and complete. In this way, P2C no longer needs multiple observations or complete point clouds as ground truth. Instead, structural cues are learned from a category-specific dataset to complete partial point clouds of objects. We demonstrate the effectiveness of our approach on both synthetic ShapeNet data and real-world ScanNet data, showing that P2C produces comparable results to methods trained with complete shapes, and outperforms methods learned with multiple partial observations. Code is available at <https://github.com/CuiRuikai/Partial2Complete>.

1. Introduction

Point clouds are widely used for 3D shape representation and play a crucial role in a range of applications [18, 30, 28, 29]. However, real-world raw point clouds are collected from sources such as laser scanners [10] and depth cameras [8], and so are often incomplete and noisy due to occlusions and varying lighting conditions. For this reason,

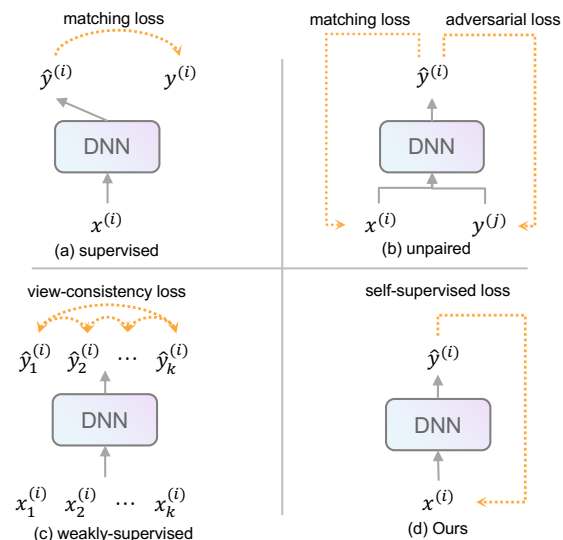


Figure 1. Conceptual comparison of point cloud completion schemes. Let $x_k^{(i)}$ be the k -th incomplete observation of object i , while $\hat{y}^{(i)}$ and $y^{(i)}$ be the corresponding completed prediction and ground truth, respectively. (a) Supervised approaches rely on paired partial-complete samples. (b) Unpaired methods require partial point clouds and complete examples to guide predictions to match the input shape and follow the complete shape distribution. (c) Weakly-supervised models learn completion based on consistency across multi-view partial samples of an object. (d) Our scheme differs from existing settings as only a single partial observation per object instance is available for learning.

point cloud completion (PCC) [40, 9, 34, 14] is studied to obtain complete point clouds from partial ones.

Supervised learning [39, 34, 24] offers a straightforward solution, where both partial point clouds and ground truth completions are required during training. Nevertheless, collecting complete point clouds is challenging. As a result, training data pairs are often obtained by simulating occlusions on 3D model collections like ShapeNet [5]. Due to the distribution gap between real and simulated data, the real-world performance of these approaches is often limited.

Unpaired (or unsupervised) PCC [6] is an alternative to

*Email: ruikai.cui@anu.edu.au

†Corresponding author. Email: shi.qiu@anu.edu.au

supervised PCC, which trains a category-specific network using only partial point clouds and a set of example complete shapes of the same category. This approach enables the use of incomplete shapes from large-scale real scans and virtual 3D object datasets, as the partial points and complete shapes do not need to be paired. However, obtaining a large, complete, and clean 3D point cloud dataset remains challenging, due to factors such as labor cost, equipment expenses, *etc.* Weakly-supervised methods [13, 23, 31] have been proposed by constructing weak supervision cues using multiple unaligned observations from different views of the same object. However, the performance can be significantly affected by alignment errors, and collecting observations from many views is difficult due to hardware limitations or viewing angle restrictions.

To address these challenges, we propose a new self-supervised approach to PCC, where for training, we only require one point cloud observation with unknown incompleteness per object. This novel setting offers several benefits for completion: 1) it eliminates the need for complete samples, thereby reducing the difficulty and expense of annotation; 2) partial objects can be easily collected from the actual world even if only a single viewing angle is available, significantly expanding the scope of training data; 3) by leveraging the unknown incompleteness assumption, partial samples, complete shapes and weakly-supervised cues can be unified in the learning framework to improve completion quality. Fig. 1 illustrates the difference of our proposed setting with existing main schemes.

In this paper, we introduce Partial2Complete (P2C), an effective approach for training a category-specific point cloud completion network using only single partial point clouds. Inspired by He *et al.* [15], P2C groups input points as patches that represent a small but possibly continuous region on the underlying surface, where we expect the network to predict masked patches based on unmasked regions. Our approach assumes that a structural prior can be learned by observing a number of training objects with different missing parts, guiding the reconstruction of severely incomplete point clouds. Furthermore, we develop the cycle constraint [43] from unpaired image translation to propose a latent reconstruction loss to the framework. This regularization ensures that completing different partial regions of the same object leads to the same completed shape.

We also present two new components to address problems that are unique to the self-supervised setting. First, traditional point cloud distance measures [40, 36] lack awareness of complete or missing regions that occur in the completion task, leading to either limited completion capability or mismatching predictions. To address this challenge, we introduce Region-Aware Chamfer Distance (RCD) to estimate point cloud correspondence based on regions centered at dynamically generated skeleton points. By optimising

RCD, possible outlier points can be pulled to the target point set and completion of missing regions will not be restricted. On the other hand, motivated by techniques that use differential geometry-based surface curvature to describe and identify local surface shape [2, 20, 35, 33], we propose the Normal Consistency Constraint (NCC) to encourage generated points to follow the local 2D surface manifold of the incomplete point cloud. The NCC queries the normal direction similarity for nearby points and computes the similarity variance as a regularizer to encourage local planarity.

We apply P2C to synthetic and real-world completion tasks to comprehensively verify its effectiveness. We show that, without any complete shape examples, our approach not only achieves comparable results against methods with access to complete samples, but also outperforms weakly-supervised methods trained with multiple incomplete observations. In summary, our main contributions are:

- We propose, P2C, the first self-supervised framework that is able to complete point clouds with only a single partial point cloud per object for learning.
- We design a novel distance measure, Region-Aware Chamfer Distance, which overcomes problems of restricting completion and insufficient supervision, by constructing local regions around dynamically constructed skeleton points.
- We present the Normal Consistency Constraint to refine shape predictions to follow the local surface manifold by minimizing a novel consistency metric, improving surface continuity and completeness.

2. Related Works

Supervised Point Cloud Completion. Earlier efforts to address point cloud completion can be divided into surface reconstruction and template matching. Surface reconstruction methods [18, 19] attempt to restore missing regions by fitting existing points to an implicit surface based on geometric cues, and then resample new points from the estimated surface. On the other hand, template matching techniques [32, 25] retrieve a template shape from a database and deform it to fit the target shape. However, surface reconstruction-based methods are able to fill holes on the surface but are limited in handling severe geometric incompleteness, while template matching methods are computationally expensive and rely on the availability of a sufficient number of example shapes. Starting with the pioneering work PCN [40], deep learning-based methods [38, 37, 24, 39] have gained significant attention in point cloud completion. However, the supervised training approach requires paired ground truth, which is difficult to obtain for real-world scans. As a result, these methods are

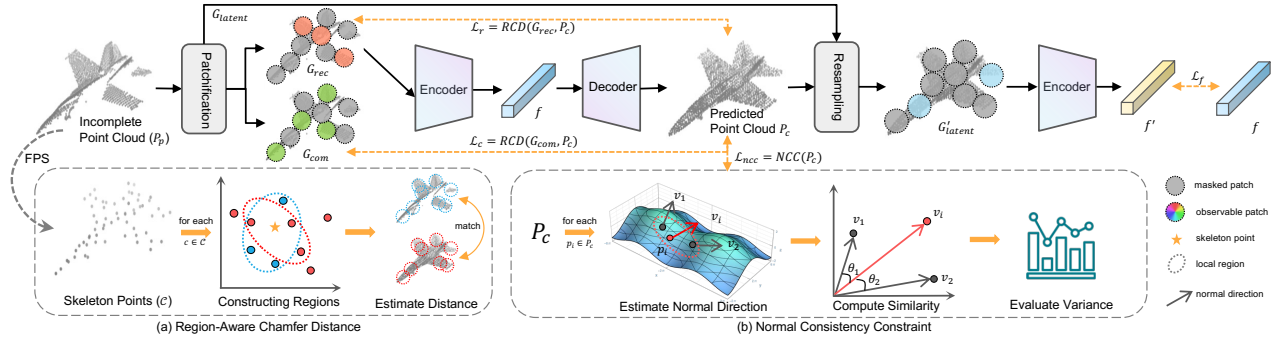


Figure 2. The Pipeline of P2C. Starting from the partial point cloud P_p , we divide it into patches and partition these patches into three groups ($G_{rec}, G_{com}, G_{latent}$). The encoder takes G_{rec} to produce features f then the decoder generates a predicted point cloud P_c based on f . G_{latent} is never observed by the encoder, we resample corresponding regions G'_{latent} in P_c to yield another feature embedding f' . The overall loss has four components. The reconstruction loss \mathcal{L}_r and completion loss \mathcal{L}_c are realized by RCD. The Latent reconstruction loss \mathcal{L}_f and the normal consistency constraint \mathcal{L}_{ncc} are introduced to regularize the inference.

often trained on synthetic datasets, which leads to impressive results on synthetic data but may not generalize well to real-world scans [41].

Unpaired and Weakly-Supervised Completion. To address the issue of data acquisition, Chen *et al.* [6] proposed the first method, Pcl2Pcl, that can be trained without paired partial and complete point sets. This was achieved through a generative adversarial network [11], where the generator transforms a partial shape latent encoding into a representation indistinguishable from the latent variable obtained from real complete shapes by the discriminator. Following Pcl2Pcl, many methods [36, 41, 3, 7] have been proposed to produce more accurate results. Nevertheless, complete shape repositories are still required, and combining unaligned real-world partial scans with complete shapes from other sources may result in poor outcomes due to alignment errors. Different from prior approaches, Gu *et al.* [13] tackle the problem of point cloud completion by using unaligned real-world partial point clouds as their data source. The network is trained with multi-view geometric constraints as weak supervision cues. However, these methods require scans from multiple viewing angles, which are not always feasible to obtain.

Self-Supervised Learning. To mitigate the cost of dataset collection and annotation, self-supervised learning [12] have been proposed. For example, DINO [4] demonstrated improved classification performance using only self-supervised training, without any labels. Self-supervised learning has also gained popularity in point cloud studies. Building upon the work of He *et al.* [15], Liu *et al.* [21] proposed a self-supervised mask discrimination framework for pretraining transformers. For point cloud upsampling, SSPU-Net [42] leverages the consistency between input sparse and generated dense point clouds to train the network using only sparse clouds. Concurrently with our research, Hong *et al.* [16] proposed a related point cloud completion

scheme, but used the same data for training and testing to enable an adaptive closed-loop [1] optimization. In contrast, our approach uses distinct test samples.

3. Method

A complete point cloud can be generated by uniformly sampling an underlying object surface, while an incomplete point cloud is obtained from the surface via biased sampling, *e.g.* due to occlusion. Our proposed self-supervised point cloud completion method aims to predict an object’s complete shape, given only a single incomplete point cloud per object from the same object category during learning. The key motivation of our method is to recover the missing part of one object by observing similar regions of other objects in the same category. Accordingly, even if a large shape collection contains only partial objects, as long as all kinds of parts of a category are exhibited across multiple object instances, the dataset is sufficient for learning to complete partial shapes.

Given only partial observations, our method learns completion via patch-wise self-supervised learning (Sec. 3.1), where patches (Sec. 3.2) of the partial point cloud are generated to achieve both shape augmentation and region-aware regularization (Sec. 3.3). Further, we introduce the Normal Consistency Constraint during training (Sec. 3.4) to enforce the assumption that object surfaces are continuous and closed by leveraging a local planarity along the object surface. The overall pipeline is depicted in Fig. 2.

3.1. Partial2Complete

Let P_p be an incomplete point cloud and P_c a predicted completion of P_p . Our framework takes P_p as input to generate M patches, each of which represents a local region on the surface of the observed shape. The M patches are partitioned into three groups $\{G_{rec}, G_{com}, G_{latent}\}$. G_{rec} is the

observable region for the network, and we force the network to generate a shape prediction P_c that preserves the regions in G_{rec} by introducing the reconstruction loss \mathcal{L}_r . Although \mathcal{L}_r effectively regularizes the predicted shape to match the observed regions in G_{rec} , this loss alone is not enough to guide the network to predict a complete shape. To this end, the completion loss \mathcal{L}_c is used to penalize the network for not predicting the masked group G_{com} . Manually masked parts and those missing from the input are both unseen by the network, hence, are indistinguishable for the network, and so minimizing \mathcal{L}_c guides the network to complete both naturally absent and intentionally masked regions.

The first group G_{rec} is passed through the encoder to obtain a latent feature embedding f , representing an encoding of the corresponding object and serving as input to the decoder to produce a shape prediction P_c . To further regularize the completion, we introduce latent reconstruction loss \mathcal{L}_f to encourage two different sets of local regions of an object sharing the same object latent representation [43]. Particularly, we exploit the third set of patches G_{latent} , which is separate from G_{rec} and not observed by the encoder. By resampling G_{latent} in P_c , we collect the patches at the same spatial location as another group G'_{latent} . Then, we pass G'_{latent} to the encoder, resulting in a latent feature f' , and \mathcal{L}_f is utilized to penalize the difference between f and f' .

3.2. Patchification and Partition

We sample patches from the object surface to provide information about local regions. To achieve this, we use farthest point sampling (FPS) [26] to sample M points as patch centers $C = \{c_i\}_{i=1}^M$ from partial shape P_p . Then, we gather the k -nearest neighbors of each center point based on Euclidean distance to obtain a patch $g_i = \{p | p \in \mathcal{N}_k^{P_p}(c_i)\}$ where $\mathcal{N}_k^{P_p}(c_i)$ denotes the set of k -nearest neighbors for c_i in P_p . Furthermore, the patches are divided into the three partitions: G_{rec} , G_{com} , and G_{latent} , with ratio $r_1 : r_2 : r_3$. Once the decoder produces the predicted shape P_c , we resample G'_{latent} as the regions corresponding to G_{latent} in the prediction by employing the same patch centers used for G_{latent} and searching for the k -nearest neighbors in P_c .

3.3. Region-Aware Chamfer Distance

Chamfer Distance (CD) and Unidirectional Chamfer Distance (UCD) are commonly used to measure the distance between two point clouds that may have different numbers of points [40, 36]. UCD between two point sets S_1 and S_2 is defined as follows:

$$d_{UCD}(S_1, S_2) = \frac{1}{|S_1|} \sum_{x \in S_1} \min_{y \in S_2} \|x - y\|_2. \quad (1)$$

CD takes both directions into account and can be defined through UCD as $d_{CD} = d_{UCD}(S_1, S_2) + d_{UCD}(S_2, S_1)$.




Figure 3. Comparison of pulling direction to minimize different distance measures. (a) CD takes the nearest neighbor for every point in the predicted set, leading to restrictions in completing missing parts; (b) UCD considers the nearest neighbor only for every point in the target set, resulting in no moving directions for noisy points; (c) RCD is aware of observed and unseen regions and thus only evaluates point distance for observed regions, pulling outlier points to the underlying surface while allowing completion of unseen parts.

Let P_p be a partial point cloud of an object with some missing regions, and P_c be a prediction corresponding to a complete but possibly noisy shape of the same object. When applying the two distance measures to self-supervised completion, where we have no access to a complete shape as ground truth, CD is not aware of incompleteness while UCD has no regularization for outliers. For $d_{CD}(P_p, P_c)$, predicted points $p \in P_c$ that correspond to unseen parts in the partial shape are estimated as far away from the underlying surface. Therefore, as shown in Fig. 3 (a), the two points in the blue box are located on the true surface of the object, but they will be displaced to minimize the CD. Thus, CD prevents the network from inferring missing parts. Moreover, $d_{UCD}(P_p, P_c)$ measures the distance by only considering the points in the prediction that are the nearest neighbors of points in P_p . We show the effect on UCD in Fig. 3 (b) that although completion of unseen regions will not be restricted, outlier points bounded by the red box are less likely to be selected as nearest neighbors of points in the target set, leading to no distance measure for outliers in the prediction. As a consequence, the network will not learn to avoid outliers in the prediction when using UCD as the distance measure.

Region-Aware Chamfer Distance (RCD) addresses the problem of seen/unseen region awareness by constructing local regions in both prediction and partial input centered at skeleton points that are dynamically sampled from the partial shape P_p . Specifically, given two point sets, P_p and P_c , m points are sampled from P_p as skeleton points $C = \{c_i\}_{i=1}^m$ through farthest point sampling [27], representing a rough observed shape. Then, the k -nearest neighbors in each point set are gathered for each skeleton point in C , forming two sets that represent the matched regions R_p

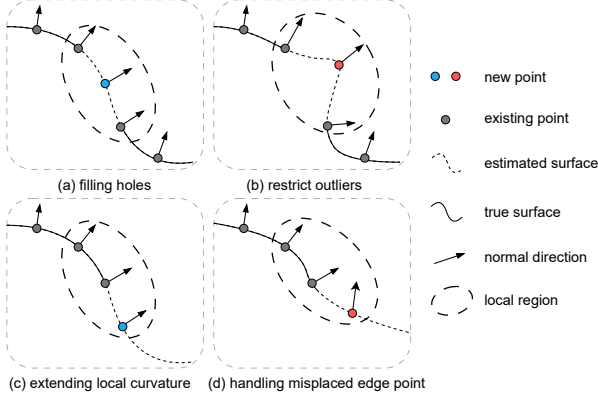


Figure 4. Illustration of the effect of NCC in a 2D case: The variance of normal similarity is lower when the point follows the underlying surface, as shown in (a) and (c), while the variance is larger when the new point results in a surface that diverges from the existing surface curvature, as shown in (b) and (d).

and R_c . Then, RCD can be defined through UCD as:

$$d_{RCD}(P_p, P_c) = d_{UCD}(R_p, R_c) + d_{UCD}(R_c, R_p), \quad (2)$$

where

$$R_p = \bigcup_{i=1}^m \{\mathcal{N}_k^{P_p}(c_i) \mid c_i \in \mathcal{C}\} \text{ for } P_p,$$

$$R_c = \bigcup_{i=1}^m \{\mathcal{N}_k^{P_c}(c_i) \mid c_i \in \mathcal{C}\} \text{ for } P_c,$$

are the union of k -nearest neighbors for all skeleton points in P_p and P_c , respectively.

3.4. Normal Consistency Constraint

To further regularize the completion, we introduce the Normal Consistency Constraint (NCC) to improve surface continuity. Specifically, given a point cloud $P = \{p_i\}_{i=1}^n$, the total least squares estimation of the normal direction [2] of a tangent plane centered at p_i is obtained by eigenvalue decomposition of the covariance matrix Cov of the k -nearest neighbors $\forall q_j \in \mathcal{N}_k^P(p_i)$, defined as:

$$Cov = \frac{1}{k} \sum_{j=1}^k (q_j - \hat{p})(q_j - \hat{p})^T, \quad \hat{p} = \frac{1}{k} \sum_{j=1}^k q_j, \quad (3)$$

where the eigenvector corresponding to the smallest eigenvalue of Cov is the estimated normal direction v_i , and v_i is normalized as $\|v_i\| = 1$. We define the normal consistency of a point p_i as:

$$nc(p_i) = \left(\sum_{j=1}^k (v_i^T v_j - \mu_i)^T (v_i^T v_j - \mu_i) \right)^{1/2}, \quad (4)$$

where a dot product between two normal directions is applied as the similarity measure, and $\mu_i = \frac{1}{k} \sum_{j=1}^k v_i^T v_j$ is the mean of similarities between v_i and v_j . The value of $nc(\cdot)$ represents the variance of the normal similarity, which estimates the local surface curvature. As the local surface approaches piece-wise planar, $nc(\cdot)$ decreases to 0, while $nc(\cdot)$ increases as the curvature increases. Further, NCC is formulated as:

$$NCC(P) = \frac{1}{n} \sum_{i=1}^n nc(p_i). \quad (5)$$

As illustrated in Fig. 4, when a new point added to fill a hole or extend an edge following the local plane, this results in a smaller nc value than if the point diverges from the surface curvature. Therefore, the NCC regularizes the prediction to be more smooth and extends edge points to make the prediction more complete, leading to better shape completion.

3.5. Optimization

The reconstruction loss and completion loss are defined as $\mathcal{L}_r = d_{RCD}(G_{rec}, P_c)$ and $\mathcal{L}_c = d_{RCD}(G_{com}, P_c)$, respectively. We encode G_{latent} in P_c as a latent representation $f' \in \mathbb{R}^d$, and we encourage it to be consistent with the first latent embedding $f \in \mathbb{R}^d$ via latent reconstruction loss:

$$\mathcal{L}_f(f, f') = \frac{1}{d} \sum_{i=1}^d \phi(f_i - f'_i), \quad (6)$$

where $\phi(\cdot)$ is the Huber [17] loss function. Together with the NCC as a loss function $\mathcal{L}_{ncc} = NCC(P_c)$, we have the overall loss defined as:

$$\mathcal{L} = \lambda_{rec} \mathcal{L}_r + \lambda_{com} \mathcal{L}_c + \lambda_{latent} \mathcal{L}_f + \lambda_{ncc} \mathcal{L}_{ncc}, \quad (7)$$

where $\lambda_{rec}, \lambda_{com}, \lambda_{latent}, \lambda_{ncc}$ are weighting parameters.

4. Experiments

4.1. Implementation Details

We employ the encoder from PCN [40] for our method. The decoder is implemented as a multi-layer perceptron with two hidden layers of 2048 dimensions. For the loss functions, we set the weights for the reconstruction loss, completion loss, and latent reconstruction loss to 1, 1, and 0.1, respectively. The weight for the NCC loss is set to 0.1. The number of patches used is 64, each formed by a local region of 32 points. The three groups $G_{rec}, G_{com}, G_{latent}$ each contain 20, 40, and 4 patches, respectively. The network is trained using the AdamW [22] optimizer with a starting learning rate of 10^{-3} and a weight decay of 10^{-3} for 300 epochs.

Table 1. Quantitative comparison result of our method and other methods on the 3D-EPN dataset using $CD-\ell_2 \downarrow (\times 10^4)$.

Method	Data Source	Average	Plane	Cabinet	Car	Chair	Lamp	Couch	Table	Boat
FoldingNet [38]	paired	6.8	2.6	7.6	4.8	8.3	9.7	7.4	8.0	5.8
PCN [40]	paired	7.4	2.5	8.0	4.8	9.0	12.2	8.1	8.9	6.0
TopNet [34]	paired	6.4	2.3	7.5	4.6	7.6	8.9	7.3	7.5	5.2
PoinTr [39]	paired	4.3	1.2	6.5	4.0	5.1	4.5	5.4	5.4	2.6
Pcl2Pcl [6]	unpaired	17.4	4.0	19.0	10.0	20.0	23.0	26.0	26.0	11.0
C4C [36]	unpaired	14.3	3.7	12.6	8.1	14.6	18.2	26.2	22.5	8.7
Inv [41]	complete	23.6	4.3	20.7	11.9	20.6	25.9	54.8	38.0	12.8
Cai <i>et al.</i> [3]	unpaired	13.6	3.5	12.2	9.0	12.1	17.6	26.0	19.8	13.6
P2C*(Ours)	unpaired	10.9	3.7	12.5	7.7	11.3	15.3	13.2	15.2	8.0
Gu <i>et al.</i> [13]	multi-view	21.3	5.9	20.8	9.5	20.4	34.9	27.1	36.7	14.8
PPNet [23]	multi-view	28.1	5.6	46.6	22.4	24.3	46.1	28.4	36.4	15.0
P2C(Ours)	single partial	14.1	4.3	19.4	8.6	13.5	16.3	20.2	18.1	12.0

Table 2. Quantitative comparison result of our method and other methods on the PCN dataset using $CD-\ell_2 \downarrow (\times 10^4)$.

Method	Air	Cab	Car	Cha	Lam	Sof	Tab	Wat	Avg
Folding [38]	2.4	8.4	4.9	9.2	11.5	9.6	8.4	7.4	7.7
PCN [40]	3.0	7.5	5.7	9.7	9.2	9.5	9.2	6.2	7.5
TopNet [34]	2.3	8.2	4.7	8.6	11.0	9.3	7.5	5.2	6.4
C4C [36]	4.1	14.2	9.9	14.6	19.2	27.8	16.8	9.0	14.4
Inv [41]	3.9	17.4	11.0	13.8	14.2	23.0	20.3	9.7	14.1
P2C(Ours)	3.5	11.7	9.0	12.8	16.4	16.2	18.6	9.1	12.2

4.2. Dataset and Evaluation Metrics

Dataset. For a comprehensive comparison, we evaluate our method on synthetic and real-world datasets following state-of-the-art point cloud completion works [39, 36]. We evaluate our method on synthetic datasets 3D-EPN [9] and PCN [40], where the former is usually adopted as an unpaired method benchmark, and the latter is widely used in supervised method evaluation. Moreover, we also extract real-world objects from ScanNet [8]. In particular, 4357 chairs and 1271 tables are extracted as the training set, while the validation set contains 1368 chairs and 350 tables. ScanNet objects are unaligned and have around 800 points on average, creating a more challenging scenario.

Evaluation Metric. We use ℓ_2 Chamfer Distance (CD) as the evaluation metric for synthetic datasets. In the case of real-world datasets, where ground-truth complete shapes are unavailable, we evaluate the prediction in terms of both fidelity and quality. To measure the preservation of observed regions in the prediction, we adopt the Unidirectional Chamfer Distance (UCD), Unidirectional Hausdorff Distance (UHD), and our proposed Region-Aware Chamfer Distance (RCD). To evaluate the quality of the generated shapes, we utilize a complete shape example set extracted from ShapeNet [5] and employ the Minimal Matching Dis-




Figure 5. Visual comparison of point cloud completion results on the 3D-EPN dataset.

tance (MMD) [39] as the quality metric.

4.3. Synthetic Data Evaluation

We compare the performance of our proposed P2C with state-of-the-art methods in the field, including supervised, unpaired (or unsupervised), and weakly-supervised methods. To ensure a fair comparison, we use their open-source implementation and the same hyperparameters except Cai *et*




Figure 6. Visual comparison of point cloud completion results on the ScanNet dataset.

al. [3] where open-source implementation is not available, so we cited their reported result. Since the unsupervised method utilizes unpaired partial and complete samples for training, we provide results of our P2C trained with the same data source, indicated as P2C*. The results on the 3D-EPN dataset are shown in Tab. 1, demonstrating the superiority of our method. P2C outperforms the best unpaired method [3] by 2.7 w.r.t $CD-\ell_2$ without any design to utilize known complete example shapes. Moreover, compared with the best weakly-supervised method [13], our proposed P2C improves the CD score by 7.2 with only single partial observations for training. Although fully supervised methods still show numerical advantages by heavily exploiting complete and paired ground-truth data, our self-supervised framework P2C has significantly reduced the performance gap between the two different learning schemes.

Tab. 2 shows the performance comparison on the PCN dataset. Our method is trained with the same data source as unpaired methods for a fair comparison. On average, we achieve 12.2 $CD-\ell_2$, while other unpaired methods have around 14, showing that our approach attains a much better overall object completion quality. The per-category results demonstrate that our proposed method outperforms the best unpaired model in six out of all eight testing categories.

Fig. 5 presents a qualitative comparison between our method and some recent methods [36, 41, 13, 23], showcasing that our method can successfully complete objects with diverse missing regions even in the absence of complete samples. In particular, our method trained on unpaired data recovers not only realistic geometry, such as the lamp post, but also captures fine-grained details, such as the car’s wheel and the desk’s edges.

Table 3. Shape completion comparison with supervised and unpaired methods on the ScanNet dataset. The numbers shown are RCD \downarrow , UCD- ℓ_2 \downarrow , UHD \downarrow , and MMD \downarrow scaled by 10^3 and 10^4 , 10^2 , and 10^3 , respectively.

Metric	RCD		UCD		UHD		MMD	
Category	Chair	Table	Chair	Table	Chair	Table	Chair	Table
Folding [38]	14.2	11.9	124.6	86.1	23.5	16.9	6.5	8.0
PCN [40]	17.9	14.9	131.8	85.1	24.5	16.8	5.9	7.2
TopNet [34]	20.3	14.1	114.6	82.5	23.0	16.7	5.8	7.6
C2C [36]	16.2	10.1	18.5	14.6	13.0	10.2	9.8	9.1
Inv [41]	18.4	9.5	8.5	7.5	10.0	8.6	15.2	16.2
P2C(Ours)	4.6	6.7	7.7	7.2	8.3	8.2	14.1	8.1

4.4. Real-world Data Evaluation

We evaluate the effectiveness of our method on the ScanNet dataset by training P2C on only partial objects and comparing it with relevant methods pretrained on the ShapeNet dataset. The results are shown in Tab. 3, which indicates that our method outperforms methods trained with complete shape examples in terms of fidelity (RCD, UCD, UHD), including both supervised and unsupervised methods. While we achieve the best result in the table category compared with other unpaired methods using MMD as a quality measure, the unpaired method C2C [36] outperforms ours on the chair category and supervised methods perform better than ours on both categories. Considering the fact that MMD measures the distance of a prediction and its ShapeNet nearest neighbor [39], the MMD scores for the compared methods are usually better since they are all trained on ShapeNet to closely resemble ShapeNet samples.

Table 4. Ablation study on four categories (plane, car, chair, table) of the 3D-EPN dataset. We investigate the impact of RCD, \mathcal{L}_f , and NCC designs. Results reported in $CD-\ell_2 \downarrow$ scaled by 10^4 .

Model	RCD	Latent Recon.	NCC	$CD-\ell_2$
A				18.6
B	✓			13.5
C	✓	✓		12.0
D	✓	✓	✓	11.2




Figure 7. Visualization of per-point distance measures. Our proposed RCD computes point cloud distance by considering the corresponding regions that appear in the target shape. In contrast, the vanilla CD assigns large distances to the unseen parts of the candidate shape, while the UCD lacks a distance measure to outliers near the observed part of the target shape.

4.5. Ablation Study

Model Design Analysis. To examine the effectiveness of our design, we conducted a detailed ablation study on the key components using four main categories in the 3D-EPN dataset. The Chamfer Distance-based evaluation results are summarized in Tab. 4. The baseline model (Model A) is the same framework, employing only reconstruction and completion losses. We then replace the vanilla CD measure used in the baseline model with the proposed RCD in Eq. 2 to form Model B and observe a significant improvement compared to the baseline. This is because the vanilla CD restricts the prediction to overfit existing regions, thereby preventing the model from inferring missing regions. When the latent reconstruction loss (\mathcal{L}_f) is incorporated in Model C, the performance increases by 1.5 compared to Model B, indicating the effectiveness of \mathcal{L}_f . Finally, to retain more completeness but fewer outliers, we further introduce the NCC (Eq. 5) to form our complete P2C framework (Model D), which helps to establish a state-of-the-art result as shown in Tab. 1.

Region-Aware Chamfer Distance. We demonstrate the effect of RCD in Fig. 7, where we visualize different distance measures. Given a partial shape and a noisy complete prediction, vanilla CD computes the distances for all points in the prediction without considering some parts that have no correspondence in the partial observation. This issue causes substantial distance estimations for missing parts. Consequently, as the objective function aims to minimize distance, an overfitted network will reconstruct the exact inputs instead of recovering missing regions. On the other hand,

Table 5. The effect of different schemes for enforcing local planarity in $CD-\ell_2 \downarrow$ scaled by 10^4 .

Method	Average	Plane	Car	Chair	Table
Baseline	13.5	4.7	14.2	14.4	20.8
Mean	12.8	4.4	12.3	14.2	19.8
Min	13.2	4.7	13.5	14.4	20.1
Variance	11.6	4.3	10.1	13.9	18.2

UCD lacks regularization to outlier points as those points are usually further from observed points than valid point predictions. This limitation allows the network to cheat the metric by outputting a shape that fills the whole 3D space, in which case the UCD value will be zero for any possible input shape. Our proposed RCD addresses the above problems by introducing region awareness since we only evaluate distance for points near the observed region, assigning no distance to unseen parts. Therefore, outlier points are constrained in the observed region, while completion capability will not be restricted.

Normal Consistency Constraint. To evaluate the effectiveness of NCC in improving surface continuity, we compare two alternative strategies for calculating the normal consistency (Eq. 4) of a given point. Instead of incorporating the local planarity, we estimate the curvature as the mean of normal vector dot product similarity or minimum of the similarity and compare our method with them. Tab. 5 shows the quantitative results comparing the strategies, where the baseline model is a simplified variant that only utilizes the reconstruction and completion losses. Based on the mean similarity and minimal similarity, we observe incremental improvements compared to the baseline model, where the average $CD-\ell_2$ only drops from 13.5 to 12.8 and 13.2, respectively. In comparison, our proposed NCC that uses the variance can better estimate the local surface curvature and improve the completion quality.

We provide more ablation studies in the supplementary material, including model complexity and efficiency, empirical hyperparameter selection, visualizations, *etc.*

5. Conclusion

In this paper, we propose P2C, the first self-supervised point cloud completion method that only requires a single partial point cloud observation per object for learning. Our method employs a novel Region-Aware Chamfer Distance to measure input-prediction similarity, and we design the Normal Consistency Constraint to enhance prediction completeness. Experimental results demonstrate that P2C exhibits state-of-the-art performance on both synthetic and real-world completion tasks, even outperforming models trained with known complete samples. Overall, our proposed method provides an effective solution for point cloud completion given only partial observation data.

References

- [1] KJ Astrom and B Wittenmark. Adaptive control. courier corporation. *Courier Corporation*, 2013. 3
- [2] Jens Berkmann and Terry Caelli. Computation of surface geometry and segmentation using covariance techniques. *IEEE Trans. Pattern Anal. Mach. Intell.*, 16(11):1114–1116, 1994. 2, 5
- [3] Yingjie Cai, Kwan-Yee Lin, Chao Zhang, Qiang Wang, Xiaogang Wang, and Hongsheng Li. Learning a structured latent space for unsupervised point cloud completion. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 5533–5543. IEEE, 2022. 3, 6, 7
- [4] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 9630–9640. IEEE, 2021. 3
- [5] Angel X. Chang, Thomas A. Funkhouser, Leonidas J. Guibas, Pat Hanrahan, Qi-Xing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. Shapenet: An information-rich 3d model repository. *CoRR*, abs/1512.03012, 2015. 1, 6
- [6] Xuelin Chen, Baoquan Chen, and Niloy J. Mitra. Unpaired point cloud completion on real scans using adversarial training. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. 1, 3, 6
- [7] Ruikai Cui, Shi Qiu, Saeed Anwar, Jing Zhang, and Nick Barnes. Energy-based residual latent transport for unsupervised point cloud completion. In *33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21-24, 2022*, page 48. BMVA Press, 2022. 3
- [8] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas A. Funkhouser, and Matthias Nießner. ScanNet: Richly-annotated 3d reconstructions of indoor scenes. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 2432–2443. IEEE Computer Society, 2017. 1, 6
- [9] Angela Dai, Charles Ruizhongtai Qi, and Matthias Nießner. Shape completion using 3d-encoder-predictor cnns and shape synthesis. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 6545–6554. IEEE Computer Society, 2017. 1, 6
- [10] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The KITTI dataset. *Int. J. Robotics Res.*, 32(11):1231–1237, 2013. 1
- [11] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In Zoubin Ghahramani, Max Welling, Corinna Cortes, Neil D. Lawrence, and Kilian Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 2672–2680, 2014. 3
- [12] Naman Goyal. A survey on self supervised learning approaches for improving multimodal representation learning. *CoRR*, abs/2210.11024, 2022. 3
- [13] Jiayuan Gu, Wei-Chiu Ma, Sivabalan Manivasagam, Wenyuan Zeng, Zihao Wang, Yuwen Xiong, Hao Su, and Raquel Urtasun. Weakly-supervised 3d shape completion in the wild. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part V*, volume 12350, pages 283–299. Springer, 2020. 2, 3, 6, 7
- [14] 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 *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 85–93. IEEE Computer Society, 2017. 1
- [15] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, pages 16000–16009, 2022. 2, 3
- [16] Sangmin Hong, Mohsen Yavartanoo, Reyhaneh Neshatavar, and Kyoung Mu Lee. Acl-spc: Adaptive closed-loop system for self-supervised point cloud completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9435–9444, June 2023. 3
- [17] Peter J Huber. Robust estimation of a location parameter. *Breakthroughs in statistics: Methodology and distribution*, pages 492–518, 1992. 5
- [18] Michael M. Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In Alla Sheffer and Konrad Polthier, editors, *Proceedings of the Fourth Eurographics Symposium on Geometry Processing, Cagliari, Sardinia, Italy, June 26-28, 2006*, volume 256 of *ACM International Conference Proceeding Series*, pages 61–70. Eurographics Association, 2006. 1, 2
- [19] Michael M. Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. *ACM Trans. Graph.*, 32(3):29:1–29:13, 2013. 2
- [20] Ping Liang and John S. Todhunter. Representation and recognition of surface shapes in range images: A differential geometry approach. *Comput. Vis. Graph. Image Process.*, 52(1):78–109, 1990. 2
- [21] Haotian Liu, Mu Cai, and Yong Jae Lee. Masked discrimination for self-supervised learning on point clouds. *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 3
- [22] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 5
- [23] Himangi Mittal, Brian Okorn, Arpit Jangid, and David Held. Self-supervised point cloud completion via inpainting. *arXiv preprint arXiv:2111.10701*, 2021. 2, 6, 7
- [24] Liang Pan, Xinyi Chen, Zhongang Cai, Junzhe Zhang, Haiyu Zhao, Shuai Yi, and Ziwei Liu. Variational relational point completion network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8524–8533, 2021. 1, 2

- [25] Mark Pauly, Niloy J. Mitra, Joachim Giesen, Markus H. Gross, and Leonidas J. Guibas. Example-based 3d scan completion. In Mathieu Desbrun and Helmut Pottmann, editors, *Third Eurographics Symposium on Geometry Processing, Vienna, Austria, July 4-6, 2005*, volume 255 of *ACM International Conference Proceeding Series*, pages 23–32. Eurographics Association, 2005. [2](#)
- [26] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5099–5108, 2017. [4](#)
- [27] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5099–5108, 2017. [4](#)
- [28] Shi Qiu, Saeed Anwar, and Nick Barnes. Geometric back-projection network for point cloud classification. *IEEE Transactions on Multimedia*, 24:1943–1955, 2021. [1](#)
- [29] Shi Qiu, Saeed Anwar, and Nick Barnes. Semantic segmentation for real point cloud scenes via bilateral augmentation and adaptive fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1757–1767, 2021. [1](#)
- [30] Shi Qiu, Saeed Anwar, and Nick Barnes. Pnp-3d: A plug-and-play for 3d point clouds. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(1):1312–1319, 2023. [1](#)
- [31] Yiming Ren, Peishan Cong, Xinge Zhu, and Yuexin Ma. Self-supervised point cloud completion on real traffic scenes via scene-concerned bottom-up mechanism. In *IEEE International Conference on Multimedia and Expo, ICME 2022, Taipei, Taiwan, July 18-22, 2022*, pages 1–6. IEEE, 2022. [2](#)
- [32] Minhyuk Sung, Vladimir G. Kim, Roland Angst, and Leonidas J. Guibas. Data-driven structural priors for shape completion. *ACM Trans. Graph.*, 34(6):175:1–175:11, 2015. [2](#)
- [33] Shuai Tang, Xiaoyu Wang, Xutao Lv, Tony X. Han, James Keller, Zhihai He, Marjorie Skubic, and Shihong Lao. Histogram of oriented normal vectors for object recognition with a depth sensor. In Kyoung Mu Lee, Yasuyuki Matsushita, James M. Rehg, and Zhanyi Hu, editors, *Computer Vision - ACCV 2012, 11th Asian Conference on Computer Vision, Daejeon, Korea, November 5-9, 2012, Revised Selected Papers, Part II*, volume 7725 of *Lecture Notes in Computer Science*, pages 525–538. Springer, 2012. [2](#)
- [34] Lyne P. Tchapmi, Vineet Kosaraju, Hamid Rezafofighi, Ian D. Reid, and Silvio Savarese. Topnet: Structural point cloud decoder. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 383–392. Computer Vision Foundation / IEEE, 2019. [1](#), [6](#), [7](#)
- [35] Federico Tombari, Samuele Salti, and Luigi Di Stefano. Unique signatures of histograms for local surface description. In Kostas Daniilidis, Petros Maragos, and Nikos Paragios, editors, *Computer Vision - ECCV 2010, 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5-11, 2010, Proceedings, Part III*, volume 6313 of *Lecture Notes in Computer Science*, pages 356–369. Springer, 2010. [2](#)
- [36] Xin Wen, Zhizhong Han, Yan-Pei Cao, Pengfei Wan, Wen Zheng, and Yu-Shen Liu. Cycle4completion: Unpaired point cloud completion using cycle transformation with missing region coding. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 13080–13089. Computer Vision Foundation / IEEE, 2021. [2](#), [3](#), [4](#), [6](#), [7](#)
- [37] Peng Xiang, Xin Wen, Yu-Shen Liu, Yan-Pei Cao, Pengfei Wan, Wen Zheng, and Zhizhong Han. Snowflakenet: Point cloud completion by snowflake point deconvolution with skip-transformer. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 5479–5489. IEEE, 2021. [2](#)
- [38] Yaoqing Yang, Chen Feng, Yiru Shen, and Dong Tian. Foldingnet: Point cloud auto-encoder via deep grid deformation. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 206–215. Computer Vision Foundation / IEEE Computer Society, 2018. [2](#), [6](#), [7](#)
- [39] Xumin Yu, Yongming Rao, Ziyi Wang, Zuyan Liu, Jiwen Lu, and Jie Zhou. Pointr: Diverse point cloud completion with geometry-aware transformers. In *ICCV*, 2021. [1](#), [2](#), [6](#), [7](#)
- [40] Wentao Yuan, Tejas Khot, David Held, Christoph Mertz, and Martial Hebert. Pcn: Point completion network. In *2018 International Conference on 3D Vision*, pages 728–737. IEEE Computer Society, 2018. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#)
- [41] Junzhe Zhang, Xinyi Chen, Zhongang Cai, Liang Pan, Haiyu Zhao, Shuai Yi, Chai Kiat Yeo, Bo Dai, and Chen Change Loy. Unsupervised 3d shape completion through GAN inversion. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 1768–1777. Computer Vision Foundation / IEEE, 2021. [3](#), [6](#), [7](#)
- [42] Yifan Zhao, Le Hui, and Jin Xie. Sspu-net: Self-supervised point cloud upsampling via differentiable rendering. In Heng Tao Shen, Yueting Zhuang, John R. Smith, Yang Yang, Pablo César, Florian Metze, and Balakrishnan Prabhakaran, editors, *MM '21: ACM Multimedia Conference, Virtual Event, China, October 20 - 24, 2021*, pages 2214–2223. ACM, 2021. [3](#)
- [43] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 2242–2251. IEEE Computer Society, 2017. [2](#), [4](#)