Self-supervised Shape Completion via Involution and Implicit Correspondences

Mengya Liu¹, Ajad Chhatkuli^{1,3}, Janis Postels¹, Luc Van Gool^{1,3}, and Federico Tombari²

¹ Computer Vision Lab, ETH Zurich ² Google, TU Munich ³ INSAIT, Sofia University {mengya.liu, ajad.chhatkuli, jpostels, vangool}@vision.ee.ethz.ch, tombari@google.com

Abstract. 3D shape completion is traditionally solved using supervised training or by distribution learning on complete shape examples. Recently self-supervised learning approaches that do not require any complete 3D shape examples have gained more interests. In this paper, we propose a non-adversarial self-supervised approach for the shape completion task. Our first finding is that completion problems can be formulated as an involutory function trivially, which implies a special constraint on the completion function \mathcal{G} , such that $\mathcal{G} \circ \mathcal{G}(X) = X$. Our second constraint on self-supervised shape completion relies on the fact that shape completion becomes easier to solve with correspondences and similarly, completion can simplify the correspondences problem. We formulate a consistency measure in the canonical space in order to supervise the completion function. We efficiently optimize the completion and correspondence modules using "freeze and alternate" strategy. The overall approach performs well for rigid shapes in a category as well as dynamic non-rigid shapes. We ablate our design choices and compare our solution against state-of-the-art methods, showing remarkable accuracy approaching supervised accuracy in some cases.

1 Introduction

Shape completion for partial shape scans of a category has only recently been tackled as a self-supervised learning task [4,8,18,29,34], among which some [4,34] use a training set of complete examples. Despite not having complete shape examples due to self-occlusions, human vision can reliably understand complete shapes of objects [31], possibly due to recognition and shape understanding. Our goal is to explore a similar ability of explicit low-level shape understanding in a neural network for shape completion.

State-of-the-art approaches for shape completion either rely on full supervision [41, 43, 48] or on adversarial learning through unpaired examples of partial and complete shapes [4,34,37]. The challenges of learning shape completion without complete examples are significant, as this requires a thorough exploration of the relationship between different 3D shapes within a given category. Recent

work [29] explores the distribution of local surface patches in order to devise an adversarial learning problem without complete shape examples. Even more recently [8] explores local similarities between shape patches in order to devise self-supervised shape completion. We argue that local patch distributions may be insufficient for accurate completion and explore novel directions for the task.

In this paper, we study the problem of self-supervised partial shape completion via correspondence estimation. Similar to the state-of-the-art methods [4, 8, 18, 29, 34, we perform completion of rigidly aligned partial shapes. In order to use correspondences for completion, we parameterize correspondences via a deformation warp from the input shape into a common template or canonical space [11, 19, 47]. Such canonical pasteurization of shapes in a category allows shape reconstruction as a function of correspondences. We base the network architecture on the hypothesis that better completion can induce better correspondences and vice-versa. The quality of the correspondences in this case can be measured by the deformation consistency of many shapes onto the common template space [19,21]. In other words, completion can be guided by the objective of having better consistency of shape deformations to a common space. We supervise the deformation warp using regularization and reconstruction accuracy similar to [47]. Our completion module uses a point generator (the completion function) and an upsampler [39]. Moreover, we observe that completion functions, regardless of their domains, can be formulated to be an exact involution. An involution is a function which is an inverse of itself. In contrast to e.g., standard cyclic constraints, involutions are highly restrictive constraints but rarely applicable. Consequently, involution allows us to devise a strong self-supervised loss without using any external augmentation or external data prior. We propose representing both incomplete and complete shapes using Implicit Neural Representation [24,27,47]. As partial shapes are often non-watertight, we use an Unsigned Distance Field (UDF) representation [5, 15, 49, 50].

Despite the simple design, the joint optimization of completion and correspondence estimation is ill-posed. Taking inspiration from Generative Adversarial Networks (GANs), we opt for a batch-wise alternate optimization of the losses, where each iteration optimizes only one of the two sets of parameters for completion and correspondences, leading to the desired convergence. We evaluate and ablate our methodology using partial shapes of both synthetic data and real scans. Results show that our method performs superior to other selfsupervised methods and the unpaired adversarial approaches. Overall, our main contributions can be summarized as:

- 1. We propose to formulate the completion task as an involution, which implies a special property of the completion function \mathcal{G} , such that $\mathcal{G} \circ \mathcal{G}(X') = X'$ where X' is the partial point set input.
- 2. We supervise completion using the consistency of the canonical space, which is a proxy measure for correspondence quality. We supervise correspondences using the partial shape reconstruction loss and deformation regularizers. We show these losses are enough for the self-supervision of shape completion.

- 3. We use a batch-wise training strategy to alternately optimize the completion module and the template-based INR module.
- 4. We provide a partial shape UDF dataset that contains the dense partial surface and the UDF field for each instance.

2 Related works

Supervised Shape Completion. Shape completion is largely tackled as a supervised learning problem with paired data [7, 10, 16, 22, 25, 30, 35, 39–41, 43, 48]. 3D-EPN [10] trains an autoencoder to complete the shape in voxelized representation. PatchComplete [30] uses the multi-resolution local patch prior and completes the SDF grid. DiffComplete [7] learns a diffusion model conditioned on the incomplete shape and outputs the occupancy representation. In favor of fine-grained details, [36, 39, 45, 48] adopt a coarse-to-fine strategy. They first generate the coarse points, and upsample them step by step. Additionally, the use of self-attention via transformers [28, 39, 41, 42, 46, 48] show improved performance. Snowflakenet [39] progressively upsamples points where the child points are recursively generated from parent points using a skip-transformer. Recently [48] introduced the patch seeds and an upsample transformer for better local details. **Unpaired Shape Completion.** Although supervised approaches show their strong capability in shape completion, they are limited by the availability of paired data, for example in real scans. Such paired data may not be available in all categories. Consequently, many works take advantage of unpaired examples of complete and partial shapes. cGan [4] first proposes to use the unpaired data to achieve shape completion through adversarial loss. However, it needs to first train two auto-encoders with the complete shape examples and partial shapes respectively. The follow-up work [37] created a variational autoencoder to learn multiple shapes with one partial input. Later many works [2, 34, 44] have designed the partial-to-full latent code transformation using the adversarial loss. ShapeInversion [44] proposes GAN inversion to learn the shape removal to go from full shape to partial shape with a pre-trained GAN from the full shape. The latest DDIT [20] pre-trains a deep implicit template that holds the deformation from various shapes to the template shape. Later the shapes are completed by deforming the template with the latent code conditioned on the partial shapes. SCoDA [38] at the same time, solves the domain gap between synthetic data completion and real data by transferring knowledge with the cross-domain feature fusion method.

Self-supervised Shape Completion. There exist very few truly self-supervised shape completion approaches. They do so using the local patch or local shape priors [6, 8, 18, 29]. [6] uses test-time optimization to directly learn contextual information from the known regions to complete a shape based on the neural tangent kernel. [18] on the other hand, takes advantage of the poses from multiviews to integrate the complete point cloud. [8] generates patches of the partial point cloud in order to achieve both shape augmentation and region-aware regularization. [29] develop an adversarial framework by using a discriminator network to reject incomplete shapes via a loss function that separately assesses local

sub-regions of the generated example and accepts only regions with sufficiently high point count in a voxelized representation. On the other hand, we solve the task of shape completion by observing similarities between completed shapes in a category. Similar to DDIT [20], but without a pre-trained template shape, our method learns to generate a consistent template space gradually from only the partial inputs and generates the dense meshes.

3 Background and Method Architectures

In this section, we outline our architectural design for self-supervised shape completion and correspondences. We illustrate the architecture in Figure 1.

3.1 Overview

The method consists of two modules, the completion module and the templatebased INR module. The completion module consists of a completion function (or a generator) \mathcal{G} and an umsampler function \mathcal{U} . We first generate the missing part using \mathcal{G} and get the detailed shape after the upsampler \mathcal{U} . Later we feed all points to the template based INR \mathcal{T} in order to obtain the completed shape implicit field. The two modules are supervising each other during training. During the inference, we obtain the final completed shape using Marching Cubes (MC) [15] on the final implicit field. Next, we describe each module in necessary details.

3.2 Completion Module (\mathcal{G} and \mathcal{U})

In contrast to unsupervised shape correspondences, where multiple paradigms exist for solving the problem (including ones with template-based INR [19,47]), self-supervised shape completion is rather ill-posed and has a very few classes of solutions. Furthermore, we want to devise a completion module that does not need full shape examples. Nonetheless, we re-purpose the completion module proposed in Snowflakenet [39] to predict sample points on the missing parts of the shape guided by self-supervised losses that will be described in Section 4.

Formally, let X' denote the set of points on the partial shape \mathcal{X}' . Then the completion function \mathcal{G} takes as input X' and outputs the set of points Y which contain missing points, i.e., $\mathcal{G}(X') = Y$. The concatenation of the points Y and the initial points X' thus provides the complete coarse shape points X_c . As we will see further on, this definition of the completion function provides a natural self-supervision loss. The architecture of the completion function \mathcal{G} consists of an encoder \mathcal{E} providing a single feature vector and a decoder G resulting in the missing points.

$$\mathcal{G}(X';\theta_{\mathcal{G}}) = G \circ \mathcal{E} = Y \tag{1}$$

 $\theta_{\mathcal{G}}$ denotes the network parameters of the completion function. Our completion module has an additional purpose: the encoder also provides a latent code describing the given set of points. Mathematically,

$$\mathcal{E}(X') = c', \quad \mathcal{E}(Y) = c_Y \quad \text{and} \quad \mathcal{E}(X_c) = c.$$
 (2)



Fig. 1: Architecture of our method. We show the detailed implementation of the completion module (\mathcal{G} and \mathcal{U}), and the template-based INR module \mathcal{T} . The completion function \mathcal{G} starts with partial input X', where an encoder \mathcal{E} is used to extract a shape code c' and is fed to the decoder G to produce Y. The involution is implemented by feeding Y to \mathcal{G} that should ideally result in the input partial points. The concatenated coarse but complete points X_c go through upsampler \mathcal{U} and generates detailed shape points X and provides conditioning on \mathcal{T} . \mathcal{L} denotes the losses applied on each module.

The latent codes provide additional context for the modules in the forward path.

The upsampler \mathcal{U} obtains finer shape by increasing the number of points by splitting each parent point into multiple child points. We refer the reader to [39] for detailed model architecture. Thus, the upsampler consists of the encoder \mathcal{E} and the upsampler decoder U to get the detailed shape.

$$\mathcal{U}(X_c;\theta_{\mathcal{U}}) = U \circ \mathcal{E} = X \tag{3}$$

Here, X denotes the final set of points from the completion module. U is the decoder that provides child points from X_c guided by the shape code c extracted from the encoder \mathcal{E} (see in Eq. (2)). $\theta_{\mathcal{U}}$ represents the network parameters of the upsampler. The upsampled points further provide better latent code to condition the warp function D in the template-based INR module \mathcal{T} .

3.3 Template-based INR Module \mathcal{T}

As partial shapes cannot be represented using the Signed Distance Fields (SDF), we sample them using the Unsigned Distance Fields (UDF). Instead of directly fitting an INR on the partial shape similar to [27], we use template-based INR [11, 19–21, 47] as it encodes the shape correspondences in a category. At any point $x \in \mathbb{R}^3$ that is around the shape surface, the INR can be written as:

$$\mathcal{T}(x; c_X, \theta_{\mathcal{T}}) = T \circ D = d. \tag{4}$$

We collectively denote the network parameters as $\theta_{\mathcal{T}}$ for the Template INR \mathcal{T} . Essentially, the warp D transforms an input point x near the shape surface to a consistent template space before passing it to the INR decoder T. Thus,

 $D(x; c_X)$ is a function of a point in space conditioned on the latent code of the shape instance c_X extracted with the encoder \mathcal{E} . Finally d denotes the UDF distance for the point x to its closest surface point.

State-of-the-art methods train the template-based INR with complete shape instances of a category, using the reconstruction loss on the output of T. In such cases, Eq. (4) can lead to a consistent template space, which thus results in valid correspondences across the input shapes [11, 19, 21, 47]. In case of only partial shape examples, the warp D cannot deform to a consistent template. Therefore, we provide the completed (although naturally with errors) points Xto the warp D. In the next section, we describe the training of the completion (θ_G) , upsampler (θ_U) and the template-based INR (θ_T) parameters.

4 Self-supervised Constraints and Training

The self-supervised training of the completion module and the template-based INR module on incomplete shapes form the core of our approach. We introduce the proposed constraints and describe their use in training as follows.

4.1 Completion Self-constraint

The definition of the completion function $\mathcal{G}(X') = Y$ means that any partial shape points X' as input results in the missing points Y. Conversely, if Y is used as the initial points, by definition, the completion function $\mathcal{G}(Y)$ should result in the points X'. Thus, the completion function \mathcal{G} by definition is an exact *involution*. Consequently, the ideal completion function must satisfy the following constraint:

$$\mathcal{G} \circ \mathcal{G}(X') = \mathcal{G}(Y) = X'.$$
(5)

Although Eq. (5) is defined for shape completion, it is straight-forward to note that Eq. (5) is true for any local representation of X' (*e.g.*, point clouds, pixels, voxels in contrast to non-local *e.g.* SDF or UDF). We represent point sets as matrices for implementation, but in fact Eq. (5) are equalities of sets. Thus we use the following loss on \mathcal{G} :

$$\mathcal{L}_{invo} = \operatorname{Chamfer}(\mathcal{G}(Y) - X').$$
(6)

Here Chamfer(.) is the operator for the bi-directional Chamfer distance. Eq. (6) is the self-loss of the completion function, derived from its definition. Note that the identity function and $\mathcal{G}(X) = 1/X$ are trivial solutions for involution. However, using further constraints, specifically, the template shape consistency loss, we guide the optimization away from the trivial solutions. Nonetheless, the involution restrains the function \mathcal{G} to be a completion function.

Involution and the cyclic constraints. Cyclic constraints [14,17,34,51] in computer vision problems have been known for their efficacy as well as elegance. They are also associated to a fundamental conservation principle in Kirchoff's second (voltage) law [12] and in transport or graph flow problems [23]. A cyclic

constraint is defined by two different transformations/functions f, g such that $g \circ f(x) = x$. Thus, the involution constraint in Eq. (5) is a special cyclic constraint with g = f, which is a stricter condition. Although involution itself has many interesting properties, for the purpose of this work, it is only important to note that the solution space of $f = f^{-1}$ is much smaller than that of $f = g^{-1}$. One more note on involution in machine learning is that it can be interpreted as self-augmentation. The second use of the function f in f(f(x)) is thus on the augmented data obtained via prediction f(x).

4.2 Completion Module Supervision Loss

As highlighted in Section 3, the template INR requires complete shape examples in a category for obtaining consistent template space. This essentially provides a key supervisory constraint for us to optimize the completion module. We thus use the template space consistency similar to [21] but in order to supervise the completion module rather than the warp. We first generate the template shape P by directly predicting the zero level set in the warp space, i.e., by taking the zero level set of T(x). For more details, we refer the reader to [47]. Thus the template consistency loss is designed to constrain each warped shape (with generated points) to be close to the template shape P. Given X as the shape points that are output from the completion module, we obtain the following loss for the completion module parameters $\theta_{\mathcal{G}}, \theta_{\mathcal{U}}$:

$$\mathcal{L}_{\mathcal{G}} = \operatorname{Chamfer}(\mathcal{D}(X_c; c), P), \ \mathcal{L}_{\mathcal{U}} = \operatorname{Chamfer}(\mathcal{D}(X; c_X), P)$$
(7)

Here, Chamfer(.,.) is the operator that measures the bidirectional Chamfer distance between two sets of points. Moreover, we follow [39] to use a partial matching loss \mathcal{L}_{part} which forces the output point cloud X from the upsampler \mathcal{U} to partially match the input X_c . Therefore, we use a single-side Chamfer Distance to constrain the loss.

$$\mathcal{L}_{part} = \operatorname{Chamfer}_{X_c \to X}(X_c, X) \tag{8}$$

We use the following loss in order to optimize the completion network:

$$\mathcal{L}_{\theta_{\mathcal{G}}} = \mathcal{L}_{\mathcal{G}} + l_1 \mathcal{L}_{invo}
\mathcal{L}_{\theta_{\mathcal{U}}} = \mathcal{L}_{\mathcal{U}} + l_2 \mathcal{L}_{nart}$$
(9)

Here, $l_1 = 1$ and $l_2 = 1$ for optimizing the completion module using a mini-batch.

4.3 Template-based INR Module Loss

While the ground-truth complete shapes are unknown for any shape instance, the partial shape is always available for supervising the INR \mathcal{T} . We thus use the partial shape reconstruction loss in order to guide the INR. Collectively we can define the constraint as follows:

$$\mathcal{L}_{\mathcal{T}} = |\mathcal{T}([X]; c_X) - [d]|.$$
(10)

[X] denotes the set of neighboring points in matrix form with the UDF distances [d] from the initial points X' on the partial shape surface, as well as the generated points, where the UDF are assigned to be 0. The loss formulates the base of the template implicit field. The operator [] emphasizes the matrix/vector representation.

Moreover, it is clear that the deformation warp D is central to the optimization of INR as well as the completion module. Thus, with strong non-linearities in D, the overall solution may be pushed towards incorrect correspondences and completion despite the other losses. In order to counteract that, similar to [21,47] we use an identity regularization \mathcal{L}_{pw} and a smoothness regularization \mathcal{L}_{pp} as:

$$\mathcal{L}_{pw} = \sum_{x \in [X']} h(\|D(x; c_X) - \mathsf{x}\|_2)$$

$$\mathcal{L}_{pp} = \sum_{x, y \in [X']} \sum_{x \neq y} \max(\frac{||\Delta x - \Delta y||_2}{||x - y||_2}, 0)$$
 (11)

Here, h(.) is the Huber function and $\triangle x = D(x; c_X) - x$ is the position shift of a point $x \in [X']$. The loss \mathcal{L}_{pw} prevents excessive deviation of a point, and \mathcal{L}_{pp} minimizes the relative space distortions. We use the following loss for the INR module.

$$\mathcal{L}_{\theta_{\mathcal{T}}} = \mathcal{L}_{\mathcal{T}} + l_3 \mathcal{L}_{pw} + l_4 \mathcal{L}_{pp}.$$
 (12)

We choose $l_3 = 0.0005, l_4 = 0.0001$.

4.4 Batch-wise Alternate Training

Now that we have defined the complete set of losses, we describe how we use them in order to alternately optimize the parameters $\theta_{\mathcal{T}}$, $\theta_{\mathcal{G}}$ and $\theta_{\mathcal{U}}$. We observe that optimizing both parameters jointly does not result in the desired convergence. It is ill-posed to optimize both modules in one batch. Taking inspiration from GANs, we use the batch-wise alternate training strategy. To be specific, in the first batch, we train the template INR module with its losses while freezing the completion module parameters. Later we optimize the completion module parameters while freezing the template INR module. We exchange the two modules training alternatively, until the training converges.

We optimize the losses $\mathcal{L}_{\theta_{\mathcal{T}}}$ and $\mathcal{L}_{\theta_{\mathcal{G}},\theta_{\mathcal{U}}}$ for a total of 50K iterations alternately. More precisely, we freeze the parameters $\theta_{\mathcal{G}}, \theta_{\mathcal{U}}$ while optimizing $\theta_{\mathcal{T}}$ in a mini-batch and vice-versa.

5 Experimental Results

5.1 Implementation Details

We use three layers of set abstraction (SA) modules and two point transformers to learn the shape code and generate 128 seeds of the missing part in the completion function \mathcal{G} . The seeds are further merged with the input partial points to

Ssp. Shape Completion via Involution and Correspondences

		F1						CD					
	type	car	plane	sofa	chair	table	avg.	car	plane	sofa	chair	table	avg.
SeedF $[48]$	sp.	80.85	97.08	71.30	80.65	83.82	82.74	0.82	0.30	1.24	1.07	0.84	0.69
cGan [4]	unp.	49.87	78.08	32.25	38.95	44.91	48.81	2.62	1.77	4.69	6.11	5.38	4.11
SInv [44]	unp.	54.87	78.68	n/a	55.70	62.31	62.89	4.00	1.64	n/a	5.57	4.74	3.99
P2C [8]	ssp.	39.43	68.56	38.09	48.74	57.67	50.50	7.54	3.37	13.45	7.92	5.21	7.50
Ours	ssp.	68.63	81.49	53.73	56.02	65.68	65.11	1.75	2.12	9.58	8.19	11.39	$\underline{6.61}$

 Table 1: Quantitative results of shape completion on our dataset. Type "sp." and "ssp."

 represent supervised and self-supervised methods, "unp." implies the unpaired method.

 We write the best performance among all unpaired and self-supervised methods in bold and underline the better performance in self-supervised methods.

create a coarse point set. From this point set, 512 points are extracted in total by farthest point sampling (FPS). The Upsampler is designed from [39], which outputs 2048 points. \mathcal{T} contains a warp function D with 10 layers of MLP and LSTM. The UDF decoder T consists of 4 MLP layers and a hyperbolic tangent function as the final activation, followed by the absolute operator to predict the UDF. We use the Adam optimizer with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. The initial learning rate is 5e-4 and the continuous decay rate is 0.5 every 500 epochs. The model is trained for 2500 epochs on 4 NVIDIA RTX GPUs with a batch size of 24. More network architecture and training details can be found in the supplementary. The code and dataset can be found in https://github.com/lmy1001/Selfsupervised-Shape-Completion.

5.2 Partial Shape UDF (PartialUDF) Dataset

There exist few datasets that provide partial shapes with implicit representations. 3D-EPN [10] proposes a 3D partial shape dataset containing a truncated signed distance field (TSDF) not satisfying our UDF representation. PatchComplete [30] provides the same representation based on ScanNet dataset [9]. Most partial shape datasets such as MVP dataset [26], and PCN dataset [32] provide only the point cloud representation. These datasets typically offer a limited number of points (2048/4096/8192/16384), which are not enough to generate dense surface and UDF samples. To address this limitation, we develop a partial shape dataset containing dense partial surfaces and UDF samples for each instance.

Dataset preparation. We develop our dataset based on ShapenetCore.v2 [3] that provides a wide range of water-tight objects across different categories. Our PartialUDF-Shapenet dataset contains diverse categories (cars/lamps/sofas/planes/tables/chairs/bathtubs/beds). For each object instance, we carefully position 30 cameras in a uniform distribution on a unit sphere which creates 30 partial observations, similar to MVP dataset [26]. For each partial shape, we densely sample 30,000 points on the seen surface and compute UDF around them following DeepSDF [27]. We use the same train/test splits as in the DeepSDF [27]. For each object instance, we randomly select 6 partial observations to collect the

Methods	type	F1	CD	Corr $\ell 2$	Methods	type	Fidelity	MMD
SeedF [48] cGan [4]	sp. unp.	$81.18 \\ 57.90$	$2.82 \\ 7.34$	-	SeedF [48] GRNet [40]	sp. sp.	$2.11 \\ 12.97$	$1.22 \\ 1.17$
P2C [8]	SSD.	42.27	7.50	-	SInv [44]	unp.	1.88	1.11
INR_only	ssp.	80.29	2.84	0.34	P2C [8]	ssp.	15.32	1.20
Ours	ssp.	84.57	1.76	0.28	Ours	ssp.	4.61	0.08

 Table 2: Quantitative results on DFaust Table 3: Quantitative results on KITTI.

 dataset. Our method performs the best.
 We achieve the best MMD metrics.

training set, in order to make the training set large. During inference time, we randomly test 2 partial inputs for each instance. Note that, we make sure that each mini-batch during training/inference does not contain the partial observation from the same shape instance, therefore, avoiding multi-view knowledge during training and test. Furthermore, we extend the PartialUDF dataset to the non-rigid DFaust [1] dataset (PartialUDF-DFaust) which provides the groundtruth (GT) correspondences. We use 7 sequences for training set and 2 sequences for test set. More details are in the supplementary.

5.3 Shape Completion on PartialUDF-Shapenet

Quantitative results. We evaluate the performance of our method on the PartialUDF-Shapenet dataset across cars, planes, sofas, chairs, and tables categories. We compare our method with the supervised SeedFormer [48], the unpaired cGan [4], the pre-trained GAN-based method ShapeInversion [44], and the self-supervised method P2C [8]. We train SeedFormer, cGan, and P2C models on our train set and directly use the provided trained model from ShapeInversion and test on our test set. We measure F1 score ($\tau = 0.03$) and CD errors, results are displayed in Table 1. Although our method cannot compete with the supervised SeedFormer, we get the second-best F1 score in all categories, and the second-lowest CD error in the cars category. Moreover, we beat the self-supervised P2C in most categories. However, it is worth noting that for complex categories like chairs and tables, our method performs worse in CD errors. Our method has no prior on whether the completed shape contains additional parts (e.g., table compartments), thus it tends to complete shapes aggressively with more parts than that exist in the shape instance.

Qualitative results. We present some qualitative results in Figure 2. Given the partial observations, our method shows impressive completion performance. cGan [4] always completes noisy shapes and ShapeInversion [44] sometimes fails to keep local input parts. P2C [8] completion is noisy and fails sometimes, especially when there is large missing part in the input shape, see the first three rows in Figure 2, P2C almost fails to generate the part while our method complete the mesh. Conversely, SeedFormer generates detailed shapes closely resembling the ground truth. Our method generates dense meshes and preserves the input points by design. Note that, due to the use of a common template space, our

11



Fig. 2: Qualitative results of shape comple-**Fig. 3:** Qualitative results on DFaust. tion on Shapenet various categories, com- We show generated points (in red) pared with SeedFormer [48], cGan [4], from the generator (completion func-ShapeInversion [44], P2C [8], our method tion) and the upsampler, and meshes completes dense meshes.

method may generate non-existent but valid parts, as seen from the last row in Figure 2. The table instance without the lower compartment generates it while completing the missing leg. This behavior results in worse CD errors.

5.4 Shape Completion and Correspondences on PartialUDF-DFaust

Quantitative results. We train our method on the non-rigid dataset and test its performance. We compare the method with the supervised SeedFormer [48], unpaired cGan [4] and self-supervised P2C [8] methods. Moreover, we additionally compare with the INR only method which consists of only the templatebased INR module. Meanwhile, we measure the "Corr $\ell 2$ " error that denotes the difference between the predicted correspondences (by finding the nearest neighbor in the template filed) and the GT correspondences provided by the DFaust dataset [1] following the protocol in [33]. Results are presented in Table 2. It shows great improvement with our method in both shape completion and dense correspondences. More importantly, the PartialUDF-DFaust dataset has an average missing rate of around 72%, and our method still achieves a high F1-score. We are 4 times smaller in CD errors compared with the supervised SeedFormer [48]. At the same time, we also improve the dense correspondences compared with INR only by gradually completing the template space, which proves that the completion module is also helping the template-based INR module to learn the template shape well. One reason for the excellent performance on DFaust [1] is that the common template space exactly satisfies the problem at hand. Thus together with the involution constraint, the completion function quickly learns to fill in the missing parts correctly. P2C, however, performs even worse than INR only due to its weak ability in non-rigid shapes and the large missing rates. More analysis of correspondences is displayed in supplementary.

Qualitative results. We present qualitative results in Figure 3. The figure showcases 4 partial inputs with various gestures, alongside the completion



Fig. 4: Qualitative results on KITTI dataset. Our Fig. 5: Visualization results of method completes the partial data even when there ablation study on sub-modules are extremely few points available. with different settings.

process of our method, INR_only, cGan [4], P2C [8], SeedFormer [48], and GT meshes respectively. Our method completes the water-tight meshes that resemble the ground-truth meshes faithfully. While cGan, P2C and SeedFormer fail to deal with large non-rigid deformations, *e.g.*, the arms. Furthermore, we showcase the generated template and the dense correspondences in Figure 7. The correspondences are presented across the same colors. Our method generates the template shape more like a mean shape of the category with good correspondences while the INR only method results in the wrong correspondences.

5.5 Shape Completion on Real-World Scans

Quantitative results. We investigate the completion capability further in realworld scans extracted from KITTI [13] dataset. Compared with the synthesized partial datasets, real-world data is noisier and sparser, increasing the difficulty of completion. For the test on KITTI dataset, we follow the protocol of GRNet [40] to prepare the test data. For all the methods, we use the trained models from the ShapeNet cars category and test the performance on KITTI. We evaluate our method following SeedFormer [48] with two metrics: Fidelity Distance, which measures the distance from every partial scan to its completed shape, denoting how well the input is reserved. The other metric is the Minimal Matching Distance (MMD), measuring the Chamfer Distance between the completed shape to its closest shape in the ShapeNet cars category, it represents the similarity between the completed shape to the synthetic shape. Results are presented in Table 3. Although our method is slightly worse in Fidelity, we achieve the lowest MMD, even better than the supervised method, showing that our method generates more car-like shapes.

Qualitative results. We further present the qualitative results on KITTI scans in Figure 4. Our method generates the complete car even with extremely few points while SeedFormer fails and GRNet raises many noisy points, P2C on the other hand, always generates fewer points in the missing area. Moreover, when there are outliers in the input scans, *e.g.*, the points circled out in the last

Methods	type	F1	CD	Methods	type	F1	CD
INR_only Completion_only	ssp.	51.54	24.52	Ours w/o. Invo Ours	ssp. ssp.	16.31 <u>53.73</u>	$\frac{24.01}{9.58}$
Ours	ssp.	53.73	9.58	sp. Ours w/o. Invo sp. Ours	sp. sp.	55.58 55.31	7.19 6.67

Table 4: Ablation on sub-modules. We ablate the effectiveness of each module by separately training each part. All ablation studies are performed in the sofas category.

row (cars should not contain points that are on the ground), while ShapeInversion and P2C still preserve the points, our method ignores the outliers, which results in slightly worse Fidelity.

5.6 Ablation Study

We analyze the effectiveness of each module of our network with different setups. First of all, we test the effectiveness of the completion module by training the template-based INR module only (INR_only) and vice versa (Completion_only). We then measure the capability of involution loss by training the network without self-involution (Ours w/o. Invo). We further train the network in a supervised way (sp. Ours) by supervising the completion module with the complete point clouds instead of the learned template shape, and the supervised way without involution loss (sp. Ours w/o. Invo). All the experiments are trained and tested on the sofas category. Results are presented in Table 4. More visualizations are presented in Figure 5.

Effet of Completion module. Compared with INR_only, the CD error decreases by 61% after adding the completion module. Moreover, see in Table 2, the correspondences are improved, meaning that the generated parts are helping with the build of the consistent template shape.

Effect of template-based INR module. In Figure 5, Completion_only generates only the input shape. we show the great importance of the template-based INR module providing the supervision on the completion module.

Effect of involution. In the self-supervised setting, the involution constraint plays an essential role in training the completion function. The network performs even worse than INR_only if trained without involution loss. From Figure 5, we find the method without involution loss tends to complete the same noisy mesh in the field. Furthermore, in the supervised setting, involution is still beneficial in decreasing CD errors, as it encourages the completion function towards completing the missing parts only.

Compared with supervised setting. We evaluate our method in a supervised setting. We observe an obvious decrease in CD errors, as well as an improvement in the F1 score. The supervised setting is powerful in completing accurate point clouds and thus is beneficial in generating the dense complete mesh, resulting in both lower CD errors and better F1 scores.

Train with smaller dataset. Although most partial shape dataset contains more than one partial observation for each instance, for example, the popular PCN [32] provides 8 random views, 3D-EPN [10] uses 6 views, and we know P2C can work with only one partial observation for each instance. Therefore, we experiment our method with much smaller dataset, consisting of 1/2/4/6partial observations respectively, and test on the same test set. The results are presented in Table 6. We test on cars and sofas categories. Note that, 6 views are the common settings we use for all the methods. Our method still performs better than cGan and P2C with a single view on cars, as the involution and consistent correspondences are relatively strong priors. And all methods are not performing well having only one partial observation in each instance.

views	SeedF	F cGan	1 P2C	Ours	SeedF	Cl cGan	D P2C	R	
	sp.	unp.	ssp.	ssp.	sp.	unp.	ssp.	ssp.	
1	76.51	47.53	38.19	<u>56.39</u>	2.30	3.93	6.61	<u>2.59</u>	
2	78.53	49.16	37.44	$\underline{62.29}$	2.20	4.25	6.29	2.12	
4	81.52	50.84	37.34	66.63	2.07	3.68	6.42	1.73	
6	80.85	49.87	39.43	<u>68.63</u>	0.82	2.62	7.54	1.63	Ours
1	68.17	34.90	27.16	41.18	3.07	6.32	15.48	14.23	
6	71.30	32.25	38.09	<u>53.73</u>	1.24	<u>4.69</u>	13.45	9.58	12

Fig. 6: Quantitative results of shape completion Fig. 7: Visualization of dense in cars (the above sub-table) and sofas (the be- correspondences and the generlow sub-table) categories. The training data contains ated template shape from our 1/2/4/6 views of partial shapes for each instance re- methods and the INR_only spectively. The best performance is shown bold, and method. Points with the same the second best performance is shown underlined. color are corresponding points.

6 Conclusions

In this paper, we presented an approach for partial shape completion based on the geometric prior of deformation warp consistency and the involution self-loss of completion. We verified two important hypotheses. First, correspondences or a common template space can be used as a supervisory signal for the completion task. Second, a completion function can be formulated as an involution, if it is constructed to predict the missing part. Our experiments showed that both contributions can play important roles in guiding the completion module. Surprisingly, in some cases such as the human body completion, our method achieves impressive performance. Similarly, we observe that our method, despite not using any complete examples, out-competes several unpaired completion approaches in the F1 score, in the rigid category instances.

Acknowledgements

The research is partially funded by VIVO collaboration project on real-time scene reconstruction.

References

- 1. Bogo, F., Romero, J., Pons-Moll, G., Black, M.J.: Dynamic faust: Registering human bodies in motion. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 6233–6242 (2017)
- Cao, Z., Zhang, W., Wen, X., Dong, Z., Liu, Y.S., Yang, B.: Mfm-net: Unpaired shape completion network with multi-stage feature matching. arXiv preprint arXiv:2111.11976 (2021)
- Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al.: Shapenet: An information-rich 3d model repository. arXiv preprint arXiv:1512.03012 (2015)
- 4. Chen, X., Chen, B., Mitra, N.J.: Unpaired point cloud completion on real scans using adversarial training. arXiv preprint arXiv:1904.00069 (2019)
- Chibane, J., Pons-Moll, G., et al.: Neural unsigned distance fields for implicit function learning. Advances in Neural Information Processing Systems 33, 21638– 21652 (2020)
- Chu, L., Pan, H., Wang, W.: Unsupervised shape completion via deep prior in the neural tangent kernel perspective. ACM Transactions on Graphics (TOG) 40(3), 1–17 (2021)
- Chu, R., Xie, E., Mo, S., Li, Z., Nießner, M., Fu, C.W., Jia, J.: Diffcomplete: Diffusion-based generative 3d shape completion. arXiv preprint arXiv:2306.16329 (2023)
- Cui, R., Qiu, S., Anwar, S., Liu, J., Xing, C., Zhang, J., Barnes, N.: P2c: Selfsupervised point cloud completion from single partial clouds. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14351–14360 (2023)
- Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: Scannet: Richly-annotated 3d reconstructions of indoor scenes. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5828–5839 (2017)
- Dai, A., Ruizhongtai Qi, C., Nießner, M.: Shape completion using 3d-encoderpredictor cnns and shape synthesis. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5868–5877 (2017)
- 11. Deng, Y., Yang, J., Tong, X.: Deformed implicit field: Modeling 3d shapes with learned dense correspondence. In: CVPR (2021)
- 12. Feynman, R.P.: The feynman lectures on physics. (No Title) 1, 46 (1963)
- Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The kitti dataset. The International Journal of Robotics Research 32(11), 1231–1237 (2013)
- 14. Ginzburg, D., Raviv, D.: Cyclic functional mapping: Self-supervised correspondence between non-isometric deformable shapes. In: ECCV (2020)
- Guillard, B., Stella, F., Fua, P.: Meshudf: Fast and differentiable meshing of unsigned distance field networks. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part III. pp. 576– 592. Springer (2022)

- 16 M. Liu, A. Chhatkuli et al.
- Huang, Z., Yu, Y., Xu, J., Ni, F., Le, X.: Pf-net: Point fractal network for 3d point cloud completion. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 7662–7670 (2020)
- Kalal, Z., Mikolajczyk, K., Matas, J.: Forward-backward error: Automatic detection of tracking failures. In: 2010 20th international conference on pattern recognition. pp. 2756–2759. IEEE (2010)
- Kim, J., Kweon, H., Yang, Y., Yoon, K.J.: Learning point cloud completion without complete point clouds: A pose-aware approach. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14203–14213 (2023)
- Lei, J., Daniilidis, K.: Cadex: Learning canonical deformation coordinate space for dynamic surface representation via neural homeomorphism. arXiv preprint arXiv:2203.16529 (2022)
- Li, H., Dong, J., Wen, B., Gao, M., Huang, T., Liu, Y.H., Cremers, D.: Ddit: Semantic scene completion via deformable deep implicit templates. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 21894–21904 (2023)
- Liu, M., Chhatkuli, A., Postels, J., Gool, L.V., Tombari, F.: Unsupervised template warp consistency for implicit surface correspondences. In: Computer Graphics Forum. vol. 42, pp. 77–87. Wiley Online Library (2023)
- Liu, M., Sheng, L., Yang, S., Shao, J., Hu, S.M.: Morphing and sampling network for dense point cloud completion. In: Proceedings of the AAAI conference on artificial intelligence. vol. 34, pp. 11596–11603 (2020)
- Masch, V.V.: The cyclic method of solving the transhipment problem with an additional linear constraint. Networks 10(1), 17–31 (1980)
- 24. Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space. In: CVPR (2019)
- 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: CVPR (2019)
- Pan, L., Chen, X., Cai, Z., Zhang, J., Zhao, H., Yi, S., Liu, Z.: Variational relational point completion network. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8524–8533 (2021)
- Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous signed distance functions for shape representation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 165– 174 (2019)
- Parmar, N., Vaswani, A., Uszkoreit, J., Kaiser, L., Shazeer, N., Ku, A., Tran, D.: Image transformer. In: International conference on machine learning. pp. 4055– 4064. PMLR (2018)
- Peters, T., Schindler, K., Brenner, C.: Self-supervised adversarial shape completion. ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences 2, 143–150 (2022)
- Rao, Y., Nie, Y., Dai, A.: Patchcomplete: Learning multi-resolution patch priors for 3d shape completion on unseen categories. Advances in Neural Information Processing Systems 35, 34436–34450 (2022)
- Sekuler, A., Gold, J., Murray, R.: Visual completion of partly occluded objects: Insights from behavioral studies (2000)
- 32. Wang, S., Cheng, J., Liu, H., Tang, M.: Pcn: Part and context information for pedestrian detection with cnns. arXiv preprint arXiv:1804.04483 (2018)

- Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M., Solomon, J.M.: Dynamic graph cnn for learning on point clouds. Acm Transactions On Graphics (tog) 38(5), 1–12 (2019)
- Wen, X., Han, Z., Cao, Y.P., Wan, P., Zheng, W., Liu, Y.S.: Cycle4completion: Unpaired point cloud completion using cycle transformation with missing region coding. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 13080–13089 (2021)
- Wen, X., Li, T., Han, Z., Liu, Y.S.: Point cloud completion by skip-attention network with hierarchical folding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1939–1948 (2020)
- Wen, X., Xiang, P., Han, Z., Cao, Y.P., Wan, P., Zheng, W., Liu, Y.S.: Pmp-net: Point cloud completion by learning multi-step point moving paths. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 7443–7452 (2021)
- Wu, R., Chen, X., Zhuang, Y., Chen, B.: Multimodal shape completion via conditional generative adversarial networks. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16. pp. 281–296. Springer (2020)
- Wu, Y., Yan, Z., Chen, C., Wei, L., Li, X., Li, G., Li, Y., Cui, S., Han, X.: Scoda: Domain adaptive shape completion for real scans. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17630–17641 (2023)
- 39. Xiang, P., Wen, X., Liu, Y.S., Cao, Y.P., Wan, P., Zheng, W., Han, Z.: Snowflakenet: Point cloud completion by snowflake point deconvolution with skiptransformer. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 5499–5509 (2021)
- Xie, H., Yao, H., Zhou, S., Mao, J., Zhang, S., Sun, W.: Grnet: Gridding residual network for dense point cloud completion. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IX. pp. 365–381. Springer (2020)
- Yan, X., Lin, L., Mitra, N.J., Lischinski, D., Cohen-Or, D., Huang, H.: Shapeformer: Transformer-based shape completion via sparse representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6239–6249 (2022)
- 42. Yu, X., Rao, Y., Wang, Z., Liu, Z., Lu, J., Zhou, J.: Pointr: Diverse point cloud completion with geometry-aware transformers. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 12498–12507 (2021)
- Yuan, W., Khot, T., Held, D., Mertz, C., Hebert, M.: Pcn: Point completion network. In: 2018 international conference on 3D vision (3DV). pp. 728–737. IEEE (2018)
- 44. Zhang, J., Chen, X., Cai, Z., Pan, L., Zhao, H., Yi, S., Yeo, C.K., Dai, B., Loy, C.C.: Unsupervised 3d shape completion through gan inversion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1768–1777 (2021)
- Zhang, W., Yan, Q., Xiao, C.: Detail preserved point cloud completion via separated feature aggregation. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXV 16. pp. 512– 528. Springer (2020)
- Zhao, H., Jiang, L., Jia, J., Torr, P.H., Koltun, V.: Point transformer. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 16259– 16268 (2021)

- 18 M. Liu, A. Chhatkuli et al.
- Zheng, Z., Yu, T., Dai, Q., Liu, Y.: Deep implicit templates for 3d shape representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1429–1439 (2021)
- Zhou, H., Cao, Y., Chu, W., Zhu, J., Lu, T., Tai, Y., Wang, C.: Seedformer: Patch seeds based point cloud completion with upsample transformer. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part III. pp. 416–432. Springer (2022)
- 49. Zhou, J., Ma, B., Li, S., Liu, Y.S., Han, Z.: Learning a more continuous zero level set in unsigned distance fields through level set projection. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 3181–3192 (2023)
- Zhou, J., Ma, B., Liu, Y.S., Fang, Y., Han, Z.: Learning consistency-aware unsigned distance functions progressively from raw point clouds. Advances in Neural Information Processing Systems 35, 16481–16494 (2022)
- Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE international conference on computer vision. pp. 2223–2232 (2017)