

Self-supervised Learning of Implicit Shape Representation with Dense Correspondence for Deformable Objects

Baowen Zhang^{1,2} Jiahe Li^{1,2} Xiaoming Deng^{1,2*} Yinda Zhang^{3*}
Cuixia Ma^{1,2} Hongan Wang^{1,2}

¹Institute of Software, Chinese Academy of Sciences ²University of Chinese Academy of Sciences
³Google

Abstract

Learning 3D shape representation with dense correspondence for deformable objects is a fundamental problem in computer vision. Existing approaches often need additional annotations of specific semantic domain, e.g., skeleton poses for human bodies or animals, which require extra annotation effort and suffer from error accumulation, and they are limited to specific domain. In this paper, we propose a novel self-supervised approach to learn neural implicit shape representation for deformable objects, which can represent shapes with a template shape and dense correspondence in 3D. Our method does not require the priors of skeleton and skinning weight, and only requires a collection of shapes represented in signed distance fields. To handle the large deformation, we constrain the learned template shape in the same latent space with the training shapes, design a new formulation of local rigid constraint that enforces rigid transformation in local region and addresses local reflection issue, and present a new hierarchical rigid constraint to reduce the ambiguity due to the joint learning of template shape and correspondences. Extensive experiments show that our model can represent shapes with large deformations. We also show that our shape representation can support two typical applications, such as texture transfer and shape editing, with competitive performance. The code and models are available at <https://iscas3dv.github.io/deformshape>.

1. Introduction

Shape representation with dense correspondence is a fundamental problem in computer vision. It plays a key role in many applications such as shape reconstruction [28, 34, 24], texture mapping [32, 10], and shape editing [10, 9, 36]. Early works often need additional semantic prior or annotations to learn such representation, e.g., SMPL [22] and




Figure 1: We present a self-supervised method to learn neural implicit representation for deformable objects with a collection of shapes. Our method can generate shapes by deforming a learned template and get dense correspondence.

SMAL [42] use registered meshes in certain semantic categories and LEAP [27] and NASA [9] require annotations of skeletons and skinning weights, which limits the use cases and scalability of these representations. On the other hand, with the emerging implicit representation, more 3D assets are encoded in implicit sign distance function (SDF), and a shape representation for deformable objects in SDF is still largely missing in the community.

In this paper, we aim to design a neural representation for deformable objects (Figure 1). Given a target deformable object represented as a set of sign distance field under various deformations, our method learns an implicit representation that is able to reconstruct the 3D shapes, interpolate between the given examples, and provide dense correspondence across shapes, in a fully self-supervised manner without any additional annotation or semantic prior. Following the common idea [10, 31], we formulate the deformable shape as a static shape in a canonical (or template) space, plus the mappings from any target deformation space to the template space for arbitrary locations in the 3D space. However, in the existence of large deformation, such as humans

*indicates corresponding author

and animals, we empirically found the above-mentioned approaches tend to be unstable and can easily get stuck in local optima if learned self-supervised due to unique challenges:

Highly Under-constrained Optimization. The free-form template shape and deformation need optimizing jointly but highly under-constrained. Under the case with large deformation across training examples, the per-location mapping is error-prone, which in return affects the template shape, and there lacks a good constraint to push both back to the correct case. To mitigate the issue, we learn a generative model for the training shapes, where each shape is represented as a code in a latent space. We then enforce a valid shape for the template shape by sampling from this latent space. We found this helps constrain the template shape and benefits the learning of dense correspondence.

Incomplete Local Rigid Constraint. As-rigid-as-possible (ARAP) constraint [1, 36] has been extensively used for discretized surface such as meshes to penalize irregular surface deformation in many previous work [1, 36, 15]. However, defining ARAP equivalent constraint on continuous SDF is non-trivial. Existing works conduct a few attempts but all have their drawbacks. For example, Deng *et al.* [10] use smooth constraints to avoid predicted deformation being large. However, this work cannot model the shapes with large deformation. Park *et al.* [31] propose a local rigid constraint however does not penalize flip mapping in a local region, and as a result a point on the left hand might be incorrectly mapped to the right hand. In contrast, we propose a novel formulation of local rigid constraint that enforces rigid transformation in local region and addresses local flip mapping issue, as illustrated in reflection issue of Figure 3(a). We give theoretical analysis and show that our local rigid constraint is, to the best of our knowledge, the first ARAP equivalent constraint define with implicit representation in infinitesimal scopes.

Insufficient Large-scale Deformation Prior. We found the learned shape representations with only local rigid constraint at each point [31, 10] still suffer from local optima. This is because shapes with large deformation often have rigid deformation in large scope, such as small neighborhood rigid regions (Figure 3(b)) and large rigid regions such as the limb on human body (Figure 3(c)), and the local rigid constraint does not have enough spatial context to take effect. Several rigid constraints on meshes [41] or point clouds [19] have been designed and proved to be effective for shape representation by utilizing connections between surface points. However, it is not straight-forward to extend these rigid constraints to implicit shape representation due to the lack of explicit connections between points. Therefore, previous implicit representation methods neglect large-scale deformation prior. To this end, we design hierarchical rigid constraints for implicit representation to utilize

spatial context of shapes at rigid part level and neighborhood level to reduce the ambiguity to learn shape and correspondence, which effectively constrains rigid motion in large scale and stabilizes the learning of the representation.

We perform extensive experiments to verify that our neural representation, learned with three above mentioned contributions, has superior capability in shape reconstruction, deformation interpolation, and building dense correspondence. We also show that high quality results can be achieved in various applications, including texture transfer and shape editing, using our learned representation.

2. Related Work

Neural Implicit Representation for Rigid Object. Implicit function is widely used in 3D shape representation. Park *et al.* [30] propose an efficient model named DeepSDF to learn SDF to represent shapes. Mescheder *et al.* [26] and Chen *et al.* [7] achieve neural implicit representation assignment by means of a binary classifier. Chibane *et al.* [8] use unsigned distance field to achieve high resolution output of arbitrary shape. Deng *et al.* [10] propose DIF, which learns a template to deform to a class of objects. This method works well for rigid objects, but fails when large deformation occurs, such as moving human body.

Neural Shape Representation with Shape Priors. Several methods present neural shape representation for dynamic shapes using shape priors such as skeleton and skinning for human body. Prior art works on neural representation of human body [9, 27, 6] impose bone transformation to constrain the deformation space. Jiang *et al.* [16] use pose-shape pairs in training data to learn a model to represent moving body. Ma *et al.* [25] project human body on a pre-defined UV map and represent body shape as point cloud. These methods require the pre-defined topology space, and the known or easy-to-learn diffeomorphism to realize human body shape representation. However, none of them can optimize template and correspondence simultaneously. Skeleton provides much prior of deformation, so methods using skeleton are advantageous when ground truth pose are available. In this paper, we focus on deformable shapes without pre-defined skeleton and propose a skeleton-free shape representation with dense correspondence.

Neural Implicit Representation with Dense Correspondence. Training a neural implicit model with dense correspondence is a longstanding task [10]. Oflow [29] can model motion sequence of a deformed shape. This method can model shapes with large deformation, but it requires continuous shape sequence as input. DIF [10] can represent shapes of the same category and generate dense correspondences among shapes. However, DIF cannot model shapes with large deformation, such as human body.

Neural Dense Correspondence. Finding accurate dense




Figure 2: Illustration of our shape representation network.

correspondence among shapes is a fundamental problem in computer vision. Several methods have explored this problem in supervised manner [21, 13, 38] or self-supervised manner [13, 14, 4, 12, 11]. In supervised methods, the corresponding points between the input shapes are required to know. In self-supervised case, some works [13, 14, 12, 11] use topology information to constrain the point location on the surface. However, these self-supervised methods require mesh as input, which is harder to access than point cloud in real world capture. Several works in human mesh registration [4, 39] also predict dense correspondences between shapes. Bharat *et al.* [4] use the shape model SMPL [22] as a prior for training, yet it cannot be generalized to other shape categories without a pre-built shape model. Giovanni *et al.* [39] present a learning approach to register non-rigid 3D point clouds. However, this work requires ground truth correspondences for training. In this paper, we focus on learning dense correspondence with neural implicit representation using self-supervised method.

3. Method

Our method learns a neural representation for a deformable object exhibited in a collection of shapes represented as signed distance fields (SDF). Inspired by DIF [10], we formulate the deformable object as a template shape encoded in an implicit neural network Φ , and dense correspondence fields $D_{i \rightarrow tmpl}$ predicted by a network ($\mathbb{R}^3 \rightarrow \mathbb{R}^3$) from arbitrary deformation S_i (*i.e.*, target space) to the template space $tmpl$. Therefore, to reconstruct the target shape S_i , signed distance values for arbitrary 3D location \mathbf{p} are queried from the template space via the dense correspondence as

$$SDF_i(\mathbf{p}) = \Phi(D_{i \rightarrow tmpl}(\mathbf{p})). \quad (1)$$

With the SDF of a target shape, the 3D mesh can be extracted using surface reconstruction algorithms such as Marching Cubes [23]. Though the overall representation is straightforward, our method focuses on shapes with large deformation, such as moving humans and animals, while DIF [10] can only deal with static categories, such as cars and chairs, we show later in this section effective learning in observations of large deformation is non-trivial.

3.1. Embedded Shapes and Template

In this section, we investigate how to learn a reasonable template field. The previous template field network [10] did not enforce the template shape to be a valid shape of any subject. In practice, this method often generates template with many floating artifacts when dealing with shapes with large deformation (Figure 9). The floating artifacts can mislead the network to find wrong correspondences on template field. However, an ideal template field should have common shape pattern of the target shapes to provide key clues for correspondence.

In order to learn a reasonable template field, we propose to constrain the template shape in the same latent space with the training shape examples. To this end, we extend the neural implicit SDF function Φ to condition on a latent code α , where each training shape and the template shape are mapped to an unique latent code. In this way, the training shape collection naturally forms a strong regularization to ensure a reasonable template shape and effectively prevent flyers. The latent space also serves naturally for the dense deformation fields between target shape and the template shape. As illustrated in Figure 2, we use the latent codes from the target shape α_i and the template shape α_{tmpl} to drive an encoder and a decoder respectively for dense correspondence prediction. The encoder-decoder network is denoted as D . Inspired the key finding by Simeonov *et al.* [35], that the distance from surface is a key clue to learn 3D correspondence, we add the SDF together with the point location as input of encoder to provide geometry clues. Therefore, the target shape can be reconstructed as

$$SDF_i(\mathbf{p}) = \Phi\left(D(\mathbf{p}, \Phi(\mathbf{p}|\alpha_i)|\alpha_i, \alpha_{tmpl})|\alpha_{tmpl}\right). \quad (2)$$

For simplicity, we use $D_{i \rightarrow tmpl}(\mathbf{p})$ for $D(\mathbf{p}, \Phi(\mathbf{p}|\alpha_i)|\alpha_i, \alpha_{tmpl})$ in the following sections.

3.2. Local Rigid Constraint

During the training, Φ , D , and latent code space $\{\alpha\}$ are optimized jointly by minimizing the reconstruction loss on the training shape collection. However, even with the regularization on the template shape, this is still a highly under-constrained optimization problem, and additional regularization is needed. A common option is to assume rigid motion in local infinitesimal scopes, and ARAP [1, 36] is a widely adopted solution on 3D surfaces. In this section, we propose, for the first time, a novel ARAP equivalent constraint defined in implicit representation.

Inspired by Nerfies [31], local deformation can be regularized by constraining the singular values from the Jacobian matrix of the deformation field. In their work, all three singular values are encouraged to be close to 1. However, this does not rule out the reflection as part of the rotation [2], as illustrated in Figure 3 (a), which will deform the shape




Figure 3: Illustration of hierarchical rigid constraints.

inside out or erroneously map the symmetry geometry, *e.g.*, left hand in one shape maps to the right hand in another. While rarely observed in Nerfies [31] since the deformation in their data is relatively small, this drawback becomes vital when deformations are large.

Our ARAP equivalent constraint is also achieved by constraining the singular values of the Jacobian matrix of the deformation field. The analysis on ARAP equivalence can be found in the supplementary material. In theory, preference over local rigid deformation is equivalent as encouraging the Jacobian matrix $J(D_{i \rightarrow tpl})$ to be a rotation matrix [31]. According to Umeyama [40], the closest orthogonal matrix of $J(D_{i \rightarrow tpl})$ in the Froebenius norm is $\mathbf{R} = \mathbf{U}\mathbf{S}\mathbf{V}^T$ with $\mathbf{S} = \text{diag}(1, 1, \det(\mathbf{U}\mathbf{V}^T))$, where \mathbf{U}, \mathbf{V} are obtained via singular value decomposition (SVD), *i.e.* $J(D_{i \rightarrow tpl}) = \mathbf{U}\mathbf{S}\mathbf{V}^T$ (See more detail in the supplementary material). Therefore, denoting $\sigma_1, \sigma_2, \sigma_3$ to be singular values of Jacobian in point \mathbf{p} , *i.e.* $J(D_{i \rightarrow tpl})(\mathbf{p})$, we define the ARAP loss in local region as

$$L_{arap} = \text{smoothL1}(\sigma_1, 1) + \text{smoothL1}(\sigma_2, 1) + \text{smoothL1}(\sigma_3, \det(\mathbf{U}\mathbf{V}^T)). \quad (3)$$

Note that the reflection happens when $\det(\mathbf{U}\mathbf{V}^T) < 0$, which is penalized in our ARAP loss since σ_3 is always positive. To further penalize reflection, we also directly penalize negative $\det(J(D_{i \rightarrow tpl}))$. The overall local rigid loss is defined as

$$L_{lr} = \sum_i \sum_{\mathbf{p} \in S_i^- \cup S_i^0} L_{arap}(\mathbf{p}) + \text{relu}(-\det(J(D_{i \rightarrow tpl})(\mathbf{p}))), \quad (4)$$

where S_i^- is the points from shape interior, *i.e.*, $SDF(p) < 0$, and S_i^0 is the points from shape surface, *i.e.*, $SDF(p) = 0$. We obtain the Jacobian matrix via auto gradient mechanism, and use a smooth L1 loss for stable training.

3.3. Hierarchical Rigid Constraint

Existing implicit learning methods [31, 10] only supervise local rigidity of deformation at each point. However, these methods do not leverage spatial context (*i.e.*, semantic parts, neighborhood distribution) of shapes effectively. To mitigate the correspondence ambiguity with spatial context, we propose a new implicit-based hierarchical rigid constraint that consists of three terms at different levels, *i.e.*, the local rigid constraint in infinitesimal scopes (Sec. 3.2), a neighborhood rigid constraint for nearby region, and a piece-wise rigid constraint for large part.

Neighborhood Rigid Constraint. This constraint is applied on small regions but in larger scale than L_{lr} , constraining the implicit field in each region to remain consistent after transformation. As shown in Figure 3(b), if points in a small neighborhood undergo the rigid transformation, the SDF of each point near the subject’s surface will remain unchanged during deformation.

To this end, we add constraints respecting the above-mentioned property between the template space and each target deformation space. For each point \mathbf{p} on the surface S_i^0 of target deformation, we sample points around it using a Gaussian distribution with σ (set to 0.05 in the experiment). We then estimate the local rotation \mathbf{R} around \mathbf{p} from $J(D_{i \rightarrow tpl})$ similar to Sec. 3.2. Each point sampling $\mathbf{p} + \eta$ ($\eta \in \mathbb{R}^3$) in the neighborhood is then mapped to the template space at $\mathbf{p}_i^{tpl} = D_{i \rightarrow tpl}(\mathbf{p}) + \mathbf{R}\eta$. We propose a neighborhood rigid loss that penalizes inconsistent SDF values sampled from the target and template space as

$$L_{nbr} = \sum_i \sum_{\mathbf{p} \in S_i^0} \mathbb{E}_{\eta \sim \mathcal{N}(\mathbf{0}, \sigma)} \|\Phi(\mathbf{p}_i^{tpl} | \alpha_{tpl}) - \Phi(\mathbf{p} + \eta | \alpha_i)\|_2^2, \quad (5)$$

where \mathbb{E} is the expectation over Gaussian sampling, which is implemented by averaging the deviation of SDF values of sampled points. Taking numerical stability into account, we follow Levinson *et al.* [18] to calculate gradient of \mathbf{R} .

Piece-wise Rigid Constraint. In fact, rigid motion can happen not only locally but also in a much larger semantic scope [41, 19], such as the limb on human body (Figure




Figure 4: Part classification and shapes of two subjects.

3 (c)). These large but rigid structures are often the source of large deformation in 3D space. We thus add a loss term, named piece-wise rigid loss, to favor rigid motion in large scale to help detect the existence of large rigid parts if any.

Our piece-wise rigid loss is enabled via part classification networks that predicts for each 3D point, inside or on the surface of the 3D shape, the probability belonging to each of N_P parts. With the predicted part association, the least square solution of rigid transformation $(\mathbf{R}_h, \mathbf{t}_h)$ for each part h can be obtained. The piece-wise rigid loss then penalizes the sum of minimal rigid transformation error over points of all parts as

$$L_{pr} = \sum_i \sum_{h \in [1, N_P]} \min_{\mathbf{R}_h, \mathbf{t}_h} \sum_{\mathbf{p} \in S_i^0 \cup S_i^-} \mathbf{P}_h(\mathbf{p}) \|(\mathbf{R}_h \mathbf{p} + \mathbf{t}_h) - D_{i \rightarrow \text{tpl}}(\mathbf{p})\|_2^2, \quad (6)$$

where \mathbf{P}_h is the predicted probability belonging to part h . In general, the points from shape interior have strong correlation with surface points. Therefore, we not only constrain surface points, but also inner points. In practice, the part classification network is learned in self-supervised way, jointly with $D_{i \rightarrow \text{tpl}}$, which needs only a pre-defined total number of parts N_P . Because the calculation of per-part rigid transformation $(\mathbf{R}_h, \mathbf{t}_h)$ could be slow, we therefore leverage a closed-form analytical solution to get the minimal rigid transformation error following Sorkine-Hornung *et al.* [37] for efficient and differentiable implementation of Eq. 6. Figure 4 shows that our method can effectively learn part classification in self-supervised manner.

Overall, our hierarchical rigid constraint is defined as

$$L_{rigid} = w_{lr} L_{lr} + w_{nbr} L_{nbr} + w_{pr} L_{pr} \quad (7)$$

where w_{lr} , w_{nbr} and w_{pr} are loss weights.

3.4. Implementation Details

We train our model in an end-to-end manner, and the latent code and parameters of all networks are optimized

together during training.

Besides the hierarchical rigid constraint, we also add a loss term to incorporate directly supervision from the training shapes on the SDF as

$$L_{sdf} = \sum_i \left(w_s \sum_{\mathbf{p} \in S_i} |\Phi(\mathbf{p}|\alpha_i) - \bar{s}| + w_n \sum_{\mathbf{p} \in S_i^0} (1 - S_c(\nabla \Phi(\mathbf{p}|\alpha_i), \bar{\mathbf{n}})) \right) + w_{Eik} \sum_{\mathbf{p} \in S_i} (|\|\nabla \Phi(\mathbf{p}|\alpha_i)\|_2 - 1| + w_\rho \sum_{\mathbf{p} \in S_i \setminus S_i^0} \rho(\Phi(\mathbf{p}|\alpha_i))), \quad (8)$$

where S_c is cosine similarity, \bar{s} is ground truth SDF value, and $\bar{\mathbf{n}}$ is ground truth normal. The first term directly supervises the SDF value, the second term supervises the surface normal, and the third term regularizes the amplitude of SDF gradient to satisfy Eikonal equation. The fourth term refrains from the off-surface points with SDF values close to 0, where $\rho(s) = \exp(-\delta \cdot |s|)$, $s \gg 1$. We use the similar method to supervise queried SDF values (Eq. 2), more details can be found in supplementary material.

We also supervise the surface normal consistency jointly with the deformation field. Specifically, we rotate the ground truth $\bar{\mathbf{n}}$ from the target space to the template space using predicted rotation $J(D_{i \rightarrow \text{tpl}})(\mathbf{p})$, i.e., $J(D_{i \rightarrow \text{tpl}})(\mathbf{p})\bar{\mathbf{n}}$ [17], and then compare it with the normal directly estimated with SDFs of correspondence in the template space, i.e., $\nabla_{D_{i \rightarrow \text{tpl}}} \Phi(D_{i \rightarrow \text{tpl}}(\mathbf{p})|\alpha_{\text{tpl}})$. The normal loss L_{pfn} is defined as

$$L_{pfn} = \sum_i \sum_{\mathbf{p} \in S_i^0} (1 - S_c(\nabla_{D_{i \rightarrow \text{tpl}}} \Phi(D_{i \rightarrow \text{tpl}}(\mathbf{p})|\alpha_{\text{tpl}}), J(D_{i \rightarrow \text{tpl}})(\mathbf{p})\bar{\mathbf{n}})) \quad (9)$$

We also use regularization terms on latent codes $L_{reg} = \sum_i \|\alpha_i\|$ and enforce the latent code of template to be close to its nearest latent code of shape in the training set.

Inspired by 3D-CODED [13] that pre-trains the network to enforce the predicted correspondence of a point in input shape to close to the input point, we use a loss term L_{recon} to enforce the self-correspondence $D_{i \rightarrow i}(\mathbf{p})$ of the target shape to close to \mathbf{p} using L_2 loss, which enables good initial correspondence. The hyper-parameters such as loss weights are fixed over all the experiments.

4. Experiment

4.1. Dataset

We use the human dataset D-FAUST [5] and animal dataset DeformingThings4D [20] for evaluation. D-FAUST contains 5 males and 5 females. Each person performs various movements, such as punching and waving arms. We use the same data split as Atzmon *et al.* [3]. Shapes from several

	D-FAUST[5]			DeformingThings4D[20]																	
				bear			rabbit			elephant			whale			deer			average		
	CD↓	IoU↑	corr↓	CD↓	IoU↑	corr↓	CD↓	IoU↑	corr↓	CD↓	IoU↑	corr↓	CD↓	IoU↑	corr↓	CD↓	IoU↑	corr↓	CD↓	IoU↑	corr↓
DIF	11.936	0.647	0.0901	16.579	0.636	0.2062	14.005	0.566	0.1515	187.662	0.431	0.0612	22.146	0.492	0.0812	24.772	0.489	0.1948	44.053	0.547	0.1303
3D-CODED	3.450	0.592	0.1038	1.327	0.824	0.1654	1.303	0.797	0.1271	0.832	0.888	0.0377	2.689	0.692	0.0710	2.683	0.656	0.1440	1.658	0.784	0.0688
Our	1.594	0.881	0.0304	0.439	0.940	0.0700	1.731	0.897	0.0925	0.587	0.910	0.0198	1.754	0.908	0.0175	2.121	0.870	0.0786	1.165	0.912	0.0557

Table 1: Capacity evaluation on D-FAUST and DeformingThings4D.

	MANO[33]		
	CD↓	IoU↑	corr↓
DIF	8.137	0.824	0.0557
3D-CODED	0.833	0.879	0.0241
Our	0.150	0.935	0.0024

Table 2: Capacity evaluation on generated MANO[33] hands.




Figure 5: Comparison of model representation capability from our method with DIF and 3D-CODED. Our method outperforms DIF and 3D-CODED by a large margin on the representation capability.




Figure 6: Results on datasets with synthesized noise. We can observe that our method can generate reasonable results with synthesized noise.

sequences are randomly split. Although pose is available in the dataset, it is not used in our experiments. DeformingThings4D contains various animals. In our experiment, we select 5 animals with very different shapes and skeleton structures, including bear, bunny, whale, elephant and deer.

4.2. Representation Capability

We evaluate the representation capability of our method by comparing reconstructed shapes in the training set with SOTA shape representation methods, DIF [10] and 3D-CODED [13]. To get the results with DIF, we only use the

queried SDF value from template field via correspondence without per-point correction to reconstruct shapes, because shapes with correction term are not just a deformed template, their geometry mainly depends on correction instead of correspondence. Table 1 shows quantitative comparison on D-FAUST and DeformingThings4D, measured with Chamfer distance ($CD \times 1000$) and Intersection over Union (IoU). On D-FAUST, our method outperforms 3D-CODED and DIF. On DeformingThing4D, our method also outperforms DIF and achieves better IoU than 3D-CODED with comparable average performance on CD. We find that in our task IoU is a more stable evaluation metric than CD, because CD is sensitive to small floating components. In the supplementary material, we show several failure cases where floating components make CD increase. We also test our method on synthesized MANO [33] dataset. As shown in Table 2 and Figure 5, our model can deal with challenging hand shapes and outperforms 3D-CODED and DIF. We further demonstrate the robustness of our method against noise. We follow DeepSDF [30] to apply synthesized noise on the depth maps where $\sigma = 0.01$, before calculating SDFs from these depths. Although suffering from heavy noises, our model still performs well. Figure 6 demonstrates the qualitative experiment, and for the quantitative experiment: Chamfer=0.936, IoU=0.830.

4.3. Shape Interpolation

In this section, we demonstrate that our model can also represent shapes similar to the shapes in training set. We evaluate shape representation capability through model fitting from full observation or partial point cloud rendered from D-FAUST and DeformingThings4D.

In experiments, we obtain each partial point cloud from single depth image, while obtaining each full observation from 20 depth images captured from multiple views. During evaluation, we use the shapes from training sequences but not involved in training. With trained models, we conduct model fitting from partial point clouds or full observations by optimizing the latent code α and a global transformation with the following function

$$L = w_s L_{sdf} + w_s L_{pbs} + w_n L_{pbn} + w_n L_{pfn} + w_{reg} L_{reg}, \quad (10)$$

where L_{pbs} and L_{pbn} are constraints to supervise queried SDF values from the template and normal of queried SDF values. Details can be found in supplementary material.

	D-FAUST[5]			DeformingThings4D[20]																	
				bear			rabbit			elephant			whale			deer			average		
	CD↓	IoU↑	corr↓	CD↓	IoU↑	corr↓	CD↓	IoU↑	corr↓	CD↓	IoU↑	corr↓	CD↓	IoU↑	corr↓	CD↓	IoU↑	corr↓	CD↓	IoU↑	corr↓
DIF	11.790	0.636	0.0917	17.010	0.629	0.1959	15.057	0.548	0.1648	192.317	0.425	0.0617	22.247	0.486	0.0851	28.907	0.426	0.2360	45.672	0.531	0.1487
3D-CODED	3.389	0.597	0.1068	1.276	0.826	0.1720	1.467	0.791	0.1326	0.807	0.889	0.0360	2.580	0.697	0.0799	2.823	0.657	0.1631	1.669	0.784	0.0742
Our	1.480	0.890	0.0307	0.609	0.940	0.0789	1.716	0.891	0.0831	0.720	0.909	0.0216	1.613	0.914	0.0241	2.917	0.867	0.1003	1.321	0.911	0.0616

Table 3: Shape reconstruction from full observation.

	D-FAUST[5]			DeformingThings4D[20]																	
				bear			rabbit			elephant			whale			deer			average		
	CD↓	IoU↑	corr↓	CD↓	IoU↑	corr↓	CD↓	IoU↑	corr↓	CD↓	IoU↑	corr↓	CD↓	IoU↑	corr↓	CD↓	IoU↑	corr↓	CD↓	IoU↑	corr↓
DIF	11.787	0.632	0.0924	17.177	0.621	0.2228	15.238	0.543	0.1963	190.240	0.428	0.0878	22.255	0.495	0.1001	27.268	0.467	0.2293	45.239	0.534	0.1673
Our	1.689	0.881	0.0324	1.757	0.910	0.0922	2.407	0.868	0.0928	2.196	0.873	0.0359	1.685	0.885	0.0237	3.403	0.833	0.0960	2.156	0.882	0.0681

Table 4: Shape reconstruction from partial point clouds.

For full observation input, our method outperforms 3D-CODED on 10 subjects of D-FAUST (Table 3). As shown in Figure 7, the reconstructed shapes with 3D-CODED may distort to reduce CD. For shape reconstruction from partial observation, we compare our method with the SoTA method DIF [10]. As shown in Figure 8 and Table 4, our method outperforms DIF by a large margin for partial point clouds. Since 3D-CODED fails to work well on partial observation, we do not show its results in Table 4.




Figure 7: Comparison of shape reconstruction from full point clouds with DIF and 3D-CODED. Our method outperforms DIF and 3D-CODED by a large margin on the reconstruction results.




Figure 8: Comparison of shape reconstruction from partial point clouds with DIF and our method. Our method outperforms DIF by a large margin on the reconstruction results.

4.4. Correspondence

We evaluate the accuracy of correspondence for training shapes (Table 1), which shows the capacity of our

method, and for shapes reconstructed from partial observation (Table 4) and full observation (Table 3) of unseen data. We compare our method with unsupervised correspondence learning methods: 3D-CODED [13] and DIF [10]. Our method achieves better correspondence performance, *i.e.*, *corr*, than DIF and 3D-CODED. The correspondence metric *corr* is calculated in the following manner. Given the reconstructed shapes S_{f_1} and S_{f_2} , we optimize Eq. 10 to get deformation field, and then warp S_{f_1} and S_{f_2} to template space. For each point \mathbf{p}_{f_1} on S_{f_1} warped to template, we find the nearest point warped from a point (*e.g.*, \mathbf{p}_{f_2}) on S_{f_2} to template, then we set \mathbf{p}_{f_2} on S_{f_2} as the correspondence of \mathbf{p}_{f_1} on S_{f_1} . We calculate the geodesic distance between ground truth corresponding point $\bar{\mathbf{p}}_{f_2}$ and predicted corresponding point \mathbf{p}_{f_2} as error. For each subject, we randomly select 100 shapes from the training set and the testing set, and evaluate the geodesic distance error between each pair. The correspondence metric *corr* is calculated by averaging the geodesic distance error of each pair of shapes. We also test model trained on noise depths (Figure. 6) and *corr*=0.0233.

4.5. Template Visualization

Figure 9 visualizes the template fields generated by our method and DIF. The template built by our method is in the manifold of human shape, while the template by DIF does not follow human shape, which shows that the same points in differently posed shapes may have different correspondences in template, such as points on arms. Therefore, our method is effective to get reasonable template.

4.6. Ablation Study

We first investigate the effect of our contributed constraints on representation capability. We conduct comparisons by removing L_{recon} , L_{pr} , L_{nbr} , L_{pfn} , L_{lr} and L_{arap} , respectively, and we also compare the results if the L_{lr} only has L_{arap} ('only L_{arap} ') or the L_{lr} is replaced by the elastic loss [31] ('replace L_{lr} '). As shown in Table 5, all our constraints are useful in our method. Our local rigid constraint is more effective than the elastic loss [31], and our




Figure 9: Comparison of zero-level set of the learned template field by our method and DIF [10].

local rigid constraint could resolve reflection issue of the elastic loss. Figure 10 shows visual results of our rigid loss terms. Different colors indicate correspondence. Without L_{pr} , the model suffers from heavy artifacts. The loss L_{nbr} is effective to reconstruct flexible regions such as front legs. For 'only L_{arap} ', we do not use the second term of L_{lr} to further penalize the reflection issue on top of L_{arap} , and we see inside-out flip at the bear's right front paw. For 'wo L_{arap} ', we penalize the second term of L_{lr} (Eq. 4). Figure 11 shows that this term not only alleviates the floating artifacts, but also helps to find correct correspondences. So, the term L_{arap} can reduce irregular shapes, and the second term of L_{lr} helps eliminate floating artifacts.




Figure 10: Qualitative experiments of the rigid loss terms.




Figure 11: Qualitative experiments of wo L_{arap} terms.

We also investigate the effect of our embedded template space. We replace our embedded template module with a separate network that has the same structure as Φ to predict template field. As shown in Figure 12, the separate template representation learns unreasonable template, and further reduces the ability of our method to represent shapes.

To investigate the effect of different scope of rigid constraints, we change the number of parts in piece-wise rigid constraint, and the standard deviation σ of the random sample in Eq. 5 for the neighborhood rigid constraint, and show the quantitative results in Table 6 and Table 7. In terms of the physic scale, L_{lr} works in infinitesimal scope. The mean volume, proportional to cube of σ , of convex hull of sampled points in neighborhood constraint is roughly 5×10^{-4} when σ is 0.05, and the mean volume of parts is 0.0768, 0.0334, 0.02175 for part number 5, 10, 20 respectively. We find that $\sigma = 0.05$ corresponds to preferred scale while smaller and larger scale lead to worse performance. With relatively small scope, piece-wise rigid constraint is flexibly applied on shapes and performs well.

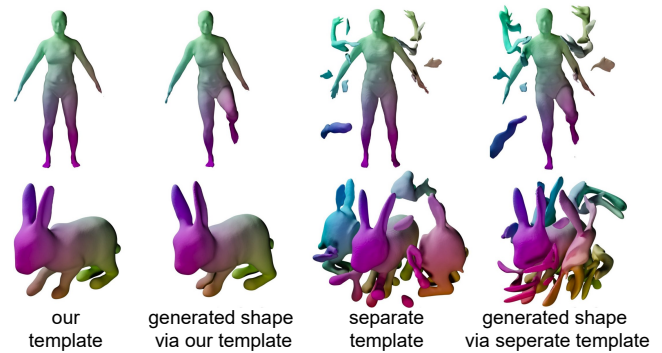


Figure 12: Qualitative comparison of our embedded template shape and separate template for shape representation.

4.7. Applications

Texture Transfer. We show a texture transfer application of our method. Shapes S_{f_1} and S_{f_2} are represented by our model. We apply texture to S_{f_1} , sample points on S_{f_1} and S_{f_2} , and transfer these points to the template space. Points on the surface of S_{f_2} query color from the nearest point on the surface of S_{f_1} in template space. As shown in Figure 13, the textures of the source shapes are well transferred to the correct regions of the target shapes under various poses.




Figure 13: Results of texture transfer. Textures on the source shapes are transferred to the target shapes.

Shape Editing. We describe how to achieve shape editing

	wo L_{recon}	wo L_{arap}	wo L_{pr}	wo L_{nbr}	wo L_{pfn}	wo L_{lr}	only L_{arap}	replace L_{lr}	full model
CD ↓	358.129	0.767	0.736	2.045	0.547	0.735	0.609	0.815	0.439
corr ↓	0.0867	0.0772	0.0858	0.0789	0.0772	0.0709	0.0761	0.0701	0.0700

Table 5: Ablation study on bear from DeformingThings4D [20]. We conduct the experiment on training data.

	5 parts	10 parts	20 parts
CD ↓	2.169	2.645	0.687
IoU ↑	0.876	0.865	0.885
corr ↓	0.0448	0.0375	0.0141

Table 6: Results of different number of parts in piece-wise rigid constraint.

	$\sigma=0.01$	$\sigma=0.03$	$\sigma=0.1$	$\sigma=0.05$
CD ↓	1.109	1.421	1.803	0.687
IoU ↑	0.877	0.872	0.852	0.885
corr ↓	0.0139	0.0226	0.0520	0.0141

Table 7: Results of different standard deviations for the neighborhood rigid constraint.

with our method. Given a set of surface points \mathbf{p}_1 on template space and corresponding target points \mathbf{p}_2 , we follow DIF [10] to optimize a latent code α_{opt} so that the target positions lie on the surface of the generated shape as

$$L = \sum_{\mathbf{p}_1, \mathbf{p}_2} w_1 (|\Phi(\mathbf{p}_2 | \alpha_{opt})| + |\Phi(D_{opt \rightarrow tpl}(\mathbf{p}_2) | \alpha_{tpl})|) + w_2 \|D_{opt \rightarrow tpl}(\mathbf{p}_2) - \mathbf{p}_1\|_2^2 + w_3 L_{pr} + w_4 \|\alpha_{opt} - \alpha_{tpl}\|_2^2. \quad (11)$$

The first term constrains the SDFs of target points in target shape and their correspondences in template, the second term enforces correspondence consistency, the third term encourages piece-wise rigid motion, and the fourth term constrains the latent code α_{opt} close to α_{tpl} . To calculate L_{pr} , we directly sample points in the bounding box of the template shape and the target points \mathbf{p}_2 . As shown in Figure 14, the template shapes are well deformed to target points. At the top right of the figure, raising the hands can cause the left leg to move in an undesired manner, and the resulting average shape after shape editing may become unreasonable. This issue is out of the scope of this paper and we leave it as the future work.

5. Limitations and Future Works

Although our method can deal with shapes with large deformation and achieve promising results, there are some problems left for future works. First, our method cannot represent out-of-distribution data. our method acquires knowledge about shape deformation from training set and does not use any prior of skeleton, so it cannot represent shapes that are out of distribution of training set. This prob-




Figure 14: Shape editing results. Shapes on the left are warped to shapes on the right. Red points are selected surface points and green points are the target positions. Arrows are used to illustrate the direction of deformations.

lem limits the generalization of shape reconstruction and degrades shape editing (Sec. 4.7). Second, our method cannot represent shape collections with topology changes because the key features (such as SDF values and normals) used in our method cannot be correctly calculated under this data setting. For example, if two body parts contact with each other, we can not get accurate SDF values due to the invisible contact surface from any camera view. In fact, this is a common issue for other methods using implicit representation. Third, the latent code space lacks smooth constraint. Our method is capable of optimizing the latent code to get a valid interpolation, but cannot guarantee valid shapes when receiving random codes. In fact, most codes in the latent space is invalid, which may be tackled by proper smooth constraint in latent. In the future work, we plan to address these challenges and propose effective shape representation approaches for more general deformable shapes.

6. Conclusions

We present a model to represent shape with dense correspondence in a self-supervised manner. Even for subjects with large deformation, our method can learn good shape and dense correspondences. We also show two typical applications of shape representation, and our method can achieve competitive performance.

Acknowledgments. This work was supported in part by National Key R&D Program of China (2022ZD0117900).

References

- [1] Marc Alexa, Daniel Cohen-Or, and David Levin. As-rigid-as-possible shape interpolation. In *Siggraph*, pages 157–164, 2000.
- [2] K Somani Arun, Thomas S Huang, and Steven D Blostein. Least-squares fitting of two 3-d point sets. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (5):698–700, 1987.
- [3] Matan Atzmon and Yaron Lipman. Sal: Sign agnostic learning of shapes from raw data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2565–2574, 2020.
- [4] Bharat Lal Bhatnagar, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. Loopreg: Self-supervised learning of implicit surface correspondences, pose and shape for 3d human mesh registration. *Advances in Neural Information Processing Systems*, 33:12909–12922, 2020.
- [5] Federica Bogo, Javier Romero, Gerard Pons-Moll, and Michael J. Black. Dynamic FAUST: Registering human bodies in motion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [6] Xu Chen, Yufeng Zheng, Michael J Black, Otmar Hilliges, and Andreas Geiger. Snarf: Differentiable forward skinning for animating non-rigid neural implicit shapes. *arXiv preprint arXiv:2104.03953*, 2021.
- [7] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5939–5948, 2019.
- [8] Julian Chibane, Gerard Pons-Moll, et al. Neural unsigned distance fields for implicit function learning. *Advances in Neural Information Processing Systems*, 33:21638–21652, 2020.
- [9] Boyang Deng, JP Lewis, Timothy Jeruzalski, Gerard Pons-Moll, Geoffrey Hinton, Mohammad Norouzi, and Andrea Tagliasacchi. Neural articulated shape approximation. In *The European Conference on Computer Vision (ECCV)*. Springer, August 2020.
- [10] Yu Deng, Jiaolong Yang, and Xin Tong. Deformed implicit field: Modeling 3d shapes with learned dense correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10286–10296, June 2021.
- [11] Marvin Eisenberger, David Novotny, Gael Kerchenbaum, Patrick Labatut, Natalia Neverova, Daniel Cremers, and Andrea Vedaldi. Neuromorph: Unsupervised shape interpolation and correspondence in one go. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7473–7483, 2021.
- [12] Marvin Eisenberger, Aysim Toker, Laura Leal-Taixé, and Daniel Cremers. Deep shells: Unsupervised shape correspondence with optimal transport. *Advances in Neural information processing systems*, 33:10491–10502, 2020.
- [13] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan Russell, and Mathieu Aubry. 3d-coded : 3d correspondences by deep deformation. In *The European Conference on Computer Vision (ECCV)*, 2018.
- [14] Oshri Halimi, Or Litany, Emanuele Rodola, Alex M Bronstein, and Ron Kimmel. Unsupervised learning of dense shape correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4370–4379, 2019.
- [15] Takeo Igarashi, Tomer Moscovich, and John F Hughes. As-rigid-as-possible shape manipulation. *ACM Transactions on Graphics (TOG)*, 24(3):1134–1141, 2005.
- [16] Boyan Jiang, Yinda Zhang, Xingkui Wei, Xiangyang Xue, and Yanwei Fu. Learning compositional representation for 4d captures with neural ode. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5340–5350, 2021.
- [17] John M Lee. Smooth manifolds. In *Introduction to smooth manifolds*, pages 1–31. Springer, 2013.
- [18] Jake Levinson, Carlos Esteves, Kefan Chen, Noah Snavely, Angjoo Kanazawa, Afshin Rostamizadeh, and Ameesh Makadia. An analysis of svd for deep rotation estimation. *Advances in Neural Information Processing Systems*, 33:22554–22565, 2020.
- [19] Chun-Liang Li, Tomas Simon, Jason Saragih, Barnabás Póczos, and Yaser Sheikh. Lbs autoencoder: Self-supervised fitting of articulated meshes to point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11967–11976, 2019.
- [20] Yang Li, Hikari Takehara, Takafumi Taketomi, Bo Zheng, and Matthias Nießner. 4dcomplete: Non-rigid motion estimation beyond the observable surface. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [21] Or Litany, Tal Remez, Emanuele Rodol’a, Alex Bronstein, and Michael Bronstein. Deep functional maps: Structured prediction for dense shape correspondence. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2017.
- [22] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM Transactions on Graphics (TOG)*, 34(6):1–16, 2015.
- [23] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. *Siggraph*, 21(4):163–169, 1987.
- [24] Qianli Ma, Jinlong Yang, Michael J. Black, and Siyu Tang. Neural point-based shape modeling of humans in challenging clothing. In *International Conference on 3D Vision (3DV) 2022*, Sept. 2022.
- [25] Qianli Ma, Jinlong Yang, Siyu Tang, and Michael J. Black. The power of points for modeling humans in clothing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10974–10984, October 2021.
- [26] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4460–4470, 2019.

- [27] Marko Mihajlovic, Yan Zhang, Michael J. Black, and Siyu Tang. Leap: Learning articulated occupancy of people. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10461–10471, June 2021.
- [28] Richard A Newcombe, Dieter Fox, and Steven M Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 343–352, 2015.
- [29] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Occupancy flow: 4d reconstruction by learning particle dynamics. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5379–5389, 2019.
- [30] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 165–174, 2019.
- [31] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [32] Amit Raj, Julian Tanke, James Hays, Minh Vo, Carsten Stoll, and Christoph Lassner. Anr: Articulated neural rendering for virtual avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3722–3731, 2021.
- [33] Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: Modeling and capturing hands and bodies together. *arXiv preprint arXiv:2201.02610*, 2022.
- [34] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2304–2314, 2019.
- [35] Anthony Simeonov, Yilun Du, Andrea Tagliasacchi, Joshua B Tenenbaum, Alberto Rodriguez, Pulkit Agrawal, and Vincent Sitzmann. Neural descriptor fields: Se (3)-equivariant object representations for manipulation. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 6394–6400. IEEE, 2022.
- [36] Olga Sorkine and Marc Alexa. As-rigid-as-possible surface modeling. In *Symposium on Geometry Processing*, volume 4, pages 109–116, 2007.
- [37] Olga Sorkine-Hornung and Michael Rabinovich. Least-squares rigid motion using svd. *Computing*, 1(1):1–5, 2017.
- [38] Ramana Sundararaman, Gautam Pai, and Maks Ovsjanikov. Implicit field supervision for robust non-rigid shape matching. *The European Conference on Computer Vision (ECCV)*, 2022.
- [39] Giovanni Trappolini, Luca Cosmo, Luca Moschella, Riccardo Marin, Simone Melzi, and Emanuele Rodolà. Shape registration in the time of transformers. *Advances in Neural Information Processing Systems*, 34:5731–5744, 2021.
- [40] Shinji Umeyama. Least-squares estimation of transformation parameters between two point patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(04):376–380, 1991.
- [41] Silvia Zuffi and Michael J Black. The stitched puppet: A graphical model of 3d human shape and pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3537–3546, 2015.
- [42] Silvia Zuffi, Angjoo Kanazawa, David Jacobs, and Michael J. Black. 3D menagerie: Modeling the 3D shape and pose of animals. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.