# DynoSurf: Neural Deformation-based Temporally Consistent Dynamic Surface Reconstruction

Yuxin Yao<sup>1</sup><sup>(6)</sup>, Siyu Ren<sup>1</sup><sup>(6)</sup>, Junhui Hou<sup>1</sup><sup>(6)</sup>\*, Zhi Deng<sup>2</sup>, Juyong Zhang<sup>3</sup><sup>(6)</sup>, and Wenping Wang<sup>4</sup>

<sup>1</sup>City University of Hong Kong <sup>2</sup>Tencent Games <sup>3</sup>University of Science and Technology of China <sup>4</sup>Texas A&M University



Fig. 1: Temporally-consistent dynamic meshes (i.e., vertices are corresponded and connections are identical over time) reconstructed by our DynoSurf from continuous dynamic 3D point cloud sequences **without** using any shape-prior, ground-truth surface, and ground-truth temporal correspondence. The color and texture map are used to illustrate correspondence across reconstructed mesh frames.

Abstract. This paper explores the problem of reconstructing temporally consistent surfaces from a 3D point cloud sequence without correspondence. To address this challenging task, we propose DynoSurf, an unsupervised learning framework integrating a template surface representation with a learnable deformation field. Specifically, we design a coarseto-fine strategy for learning the template surface based on the deformable tetrahedron representation. Furthermore, we propose a learnable deformation representation based on the learnable control points and blending weights, which can deform the template surface non-rigidly while maintaining the consistency of the local shape. Experimental results demonstrate the significant superiority of DynoSurf over current stateof-the-art approaches, showcasing its potential as a powerful tool for

<sup>\*</sup> Corresponding author. Email: jh.hou@cityu.edu.hk.

dynamic mesh reconstruction. The code is publicly available at https: //github.com/yaoyx689/DynoSurf.

Keywords: Dynamic Surface Reconstruction  $\cdot$  Temporally Consistent Meshes  $\cdot$  Low-Dimensional Deformation  $\cdot$  Point Cloud Sequences

# 1 Introduction

The reconstruction of temporally consistent geometric surfaces from a continuous dynamic 3D point cloud sequence without correspondences presents a valuable yet formidable challenge. These geometric surfaces play a crucial role in diverse domains, including film, gaming, animation, virtual reality (VR), augmented reality (AR), and robotics. They facilitate various applications such as motion object editing, texture transfer, and shape analysis, thereby contributing to the creation of realistic visual effects and immersive experiences.

With the development of the static surface reconstruction for a single object from a point cloud [15, 18, 40, 45], a natural idea to model the temporally consistent dynamic surfaces is to select a keyframe from the point cloud sequences and reconstruct it as a template surface through static reconstruction technology. Then the non-rigid deformation for the template surface is estimated by aligning it with other frames. High requirements are placed on the selection of the keyframe point cloud with fewer missing parts or occlusions and a similar shape to other frames. Some optimization-based dynamic reconstruction frameworks [10,34,47] continuously update the template shape during frame-by-frame non-rigid tracking. These approaches are effective and pragmatic, but most of them use depth as input and cannot handle the sparse point clouds without any camera information. Some learning-based method [23, 36, 49] no longer define or pre-establish the template shape, but jointly learn dynamic reconstruction. This way may achieve better results but increases the difficulty of decoupling the deformation field and the template space, which makes deformation have less practical significance. Moreover, these methods use ground truth to supervise training, have poor generalization, and can usually only process sequences of fixed length. To address these problems, we first design an adaptive keyframe selection strategy and learn the template surface representation, then combine it with a deformation field to form a dynamic reconstruction framework.

Furthermore, the design of the deformation field is particularly important. Optimization-based methods [24,25,54] often employ the embedded deformation graph [48] for non-rigid deformation, but they depend on the spatial position and mesh quality of the template surface. Some methods [39, 50, 59] directly use a network to predict the deformation for each point, lacking explicit control of the local structure. [6] proposes a neural deformation graph to reconstruct dynamic objects. However, they have some restrictions on the location and connection relationships of nodes, which necessitates the use of multiple loss functions. Therefore, we propose a deformation field based on explicit learnable control points, which improves the explicit controllability of the deformation field based on implicit representation. We also learn the blending weights between the points on

the template surface and the control points. This allows us to adaptively learn the deformation consistency of surface points based on motion information.

In this paper, we aim to reconstruct a temporally consistent surface from continuous dynamic 3D point cloud sequences with not-fixed length. We propose a reconstruction framework combining a template surface representation and a learnable deformation field based on the learnable control points and blending weight. Extensive experiments on three benchmark datasets demonstrate that our proposed temporally-consistent surface reconstruction outperforms the stateof-the-art methods to a large extent.

To summarize, our contributions are as follows:

- We propose a new learning framework for reconstructing temporally-consistent dynamic surfaces from continuous dynamic point cloud sequences that can be trained without requiring any shape prior, ground-truth surface, and groundtruth temporal correspondence information as supervision.
- We propose a coarse-to-fine learning strategy for unsupervised static surface reconstruction based on the deformable tetrahedron representation.
- We propose a deformation representation based on learnable control points and blending weights that can better maintain local structure while being independent of object category.

#### $\mathbf{2}$ **Related Work**

We briefly review the most relevant works, including 3D shape representation, non-rigid shape deformation, and sequential shape reconstruction.

**3D** Shape Representation. Generally, 3D shape representation could be divided into two main categories, explicit and implicit, and these can be transformed into one another. The point cloud and triangle mesh are the two most commonly utilized explicit representations for 3D shapes in various applications [13, 14, 28, 42, 57]. Implicit representation utilizes an isosurface of a field to depict a surface. The Binary Occupancy Field (BOF) [5,19,32,41] and Signed Distance Field (SDF) [21,29,38] are two prevalent implicit fields for representing shapes, both dividing the entire space into two regions, inside and outside the shape. The implicit fields could be approximated by a neural network, such as DeepSDF [38], and SAL [1].

Given the shapes explicitly represented as point clouds or triangle meshes. diverse methods exist for their conversion into implicit representations [19, 21, 22, 29, 36]. Conversely, when transforming implicit shapes into explicit forms, Marching Cubes (MC) [31] and Marching Tetrahedrons (MT) [35] stand out as the two widely used methods. Considering the detailed parts of shapes, DMTet [45] and FlexiCubes [46] introduce offset for each tetrahedron or cube vertex, allowing their methodologies to effectively manage detailed shape components. **Non-rigid Shape Deformation.** Deforming non-rigid shapes presents a challenge due to the high degree of freedom involved in controlling the shape deformation. To mitigate this, [24, 25] employ an embedded deformation graph [48]

with local rigid properties for modeling deformation. During optimization, various geometric regularization terms, such as as-conformal-as-possible deformations [52] and as-rigid-as-possible deformations [53], are incorporated to impart meaningful deformation. Recently, [54,55] have developed optimization methods to accelerate the optimization process. With the advancement of deep learning in 3D vision, numerous learning-based deformation methods have been introduced. CorrNet3D [58] designs a symmetric pipeline to compute deformations between point clouds. RMA-Net [12] introduces a new deformation representation with a point-wise combination of several rigid transformations. NDP [26] proposes a non-rigid motion representation using a pyramid architecture.

Sequential Shape Reconstruction. Several 3D parametric models, such as SMPL [30] and SMAL [60], have been developed to represent general shapes within specific categories. These parametric models allow for straightforward adjustments in pose and shape by modifying their respective parameters. Subsequently, the combination of neural implicit functions and parametric models enhances the expressive power for reconstruction [7, 16, 17, 37]. However, these methods are constrained to handling sequences within specific categories, limiting their broad applicability.

Recent studies [6, 34, 36, 49] have expanded the scope of reconstructing dynamic shapes by predicting deformation fields between different frames. Approaches like Dynamic Fusion [34] and NDG [6] establish a deformation graph to capture the deformation fields. O-Flow [36] extends ONet [32] from 3D to 4D, representing deformation fields with a Neural-ODE [8], while LPDC [49] employs MLPs to model correspondences in parallel. Both O-Flow and LPDC reconstruct the reference shape based on the initial frame, with subsequent frames having their query points mapped to the chosen reference frame. In contrast, I3DMM [56] learns a canonical reference shape for sequences of heads, while CASPR [44] and Garment Nets [9] use ground truth canonical coordinates for network training, which is often impractical. Cadex [23] employs canonical map factorization to represent deformation without requiring correspondences during training, but it may struggle with sequences featuring substantial motions. Recently, FDN [2] proposes a new computational model that leverages fluid simulation priors to handle dynamic 3D scenes with topological changes.

# 3 Proposed Method

## 3.1 Problem Statement and Method Overview

Let  $\{\mathbf{P}_1, ..., \mathbf{P}_K\}$  be a continuous dynamic 3D point cloud sequence capturing the geometric changes or motion of scenes/objects over time, where the correspondence across frames is unknown and  $\mathbf{P}_k \in \mathbb{R}^{N_k \times 3}$  is the k-th  $(k \in [1, K])$  frame containing  $N_k$  points. We aim to reconstruct from them a *temporally-consistent* mesh sequence  $\{\mathcal{M}_1, ..., \mathcal{M}_k\}$  where the k-th mesh surface  $\mathcal{M}_k := \{\mathbf{V}_k, \mathbf{F}\}$  with  $\mathbf{V}_k$  and  $\mathbf{F}$  being the vertex set and face set, respectively, the vertices across meshes are corresponded, and all meshes share the same  $\mathbf{F}$ .

To address this challenging task, we propose an unsupervised learning framework, named DynoSurf, which requires no ground-truth surface and correspondence information as supervision. As illustrated in Fig. 2, DynoSurf mainly consists of two stages: template surface construction (Sec. 3.2) and deformationbased temporal reconstruction (Sec. 3.3). Generally, we first select a keyframe from the input point clouds, denoted as  $\mathbf{P}_{k^*}$ , and reconstruct its surface based on the deformable tetrahedron representation through a coarse-to-fine learning process. Then, we deform the enhanced template surface to simultaneously align with all point clouds by learning a control points blending-based deformation field through jointly optimizing the template surface and deformation field. In what follows, we will detail the framework.



Fig. 2: Illustration of the proposed DynoSurf, which can reconstruct from continuous dynamic point cloud sequences temporally-consistent dynamic surfaces without requiring any ground-truth surface and temporal correspondence information.

## 3.2 Template Surface Representation via Deformable Tetrahedron

In this stage, we aim to construct a template surface, i.e., a mesh fitting the selected keyframe point cloud. As illustrated in Fig. 3, we initially select a keyframe  $\mathbf{P}_{k^*}$ , and construct its convex hull as a triangle mesh, referred to as  $\mathcal{C}$ . Subsequently, we process  $\mathcal{C}$  to generate a tetrahedron mesh, denoted as  $\mathcal{T}$ , where the signed distance field (SDF) of the surface underlying  $\mathbf{P}_{k^*}$  is defined and learned through a coarse-to-fine learning process. The SDF enables us to extract a mesh, denoted as  $\overline{\mathcal{M}} := \{\overline{\mathbf{V}}, \mathbf{F}\}$  with  $\overline{\mathbf{V}}$  being the vertex set.

Keyframe selection and tetrahedron mesh construction. Since we will deform the reconstructed template surface from  $\mathbf{P}_{k^*}$  for temporal reconstruction, intuitively,  $\mathbf{P}_{k^*}$  should be relatively average to facilitate the deformation process. Thus, we determine  $k^* = \arg \min_{k \in [1,K]} \sum_l CD_{\ell_2}(\mathbf{P}_k, \mathbf{P}_l)$ , where  $CD_{\ell_2}(\cdot, \cdot)$  computes the  $\ell_2$ -norm-based Chamfer distance between two point clouds.



Fig. 3: Illustration of the pipeline of learning template surface via deformable tetrahedron.

We then adopt QuickHull [3] to construct the convex hull  $\mathcal{C}$  of  $\mathbf{P}_{k^*}$ , which is further dilated and re-meshed to obtain a uniform triangle mesh  $\widehat{\mathcal{C}}$  surrounding  $\mathcal{C}$ . By tetrahedralizing  $\widehat{\mathcal{C}}$ , we obtain a tetrahedron mesh  $\mathcal{T} = \{\mathbf{Q}, \mathbf{T}\}$  with  $\mathbf{Q}$  and  $\mathbf{T}$  being the vertex and tetrahedron sets, respectively. Note that the vertices of  $\mathbf{Q}$  are distributed both inside and outside of  $\mathcal{C}$  due to the dilation and remeshing operations.

Coarse-to-fine learning of deformable tetrahedron-based template surface. We learn the deformable tetrahedron representation [45] to construct a template surface  $\overline{\mathcal{M}}$  that approximates the point cloud  $\mathbf{P}_{k^*}$ .

$$[s(\mathbf{q}), \delta(\mathbf{q})] = F_{\mathbf{\Phi}}(\Gamma(\mathbf{q})), \tag{1}$$

where  $\Gamma(\cdot)$  is the Position Encoding operator. Then, based on the deformed tetrahedral grid points  $\{\mathbf{q}+\delta(\mathbf{q})\}_{\mathbf{q}\in\mathbf{Q}}$  and SDF values  $\{s(\mathbf{q})\}_{\mathbf{q}\in\mathbf{Q}}$ , we can employ the differentiable marching tetrahedron algorithm [45], denoted as  $\mathsf{DMT}(\cdot, \cdot)$ , to extract  $\overline{\mathcal{M}}$ , i.e.,  $\overline{\mathcal{M}} = \mathsf{DMT}(\{\mathbf{q}+\delta(\mathbf{q})\}, \{s(\mathbf{q})\})$ .

To train  $F_{\Phi}(\cdot)$  to encode the SDF of the surface  $\overline{\mathcal{M}}$ , we adopt the following coarse-to-fine strategy to make the learning process easier and faster: (1) We initialize  $F_{\Phi}(\cdot)$  by learning the SDF of  $\mathcal{C}$ , i.e.,

$$\min_{\mathbf{\Phi}} L_{\text{init1}} = \sum_{\mathbf{q} \in \mathbf{Q}} \|s(\mathbf{q}) - \text{SDF}(\mathbf{q}, \mathcal{C})\|^2,$$
(2)

where  $s(\mathbf{q})$  is the predicted SDF value of  $F_{\mathbf{\Phi}}(\cdot)$  and  $\text{SDF}(\mathbf{q}, \mathcal{C})$  denotes the groundtruth SDF value of convex hull surface  $\mathcal{C}$  at  $\mathbf{q}$ . (2) We then refine  $F_{\mathbf{\Phi}}(\cdot)$  by constraining  $\overline{\mathcal{M}}$  and  $\mathbf{P}_{k^*}$  to be as close as possible, i.e.,

$$\min_{\mathbf{\sigma}} L_{\text{init2}} = \tilde{w}_1 L_{\text{CD}} + \tilde{w}_2 L_{\text{Norm}}^* + \tilde{w}_3 L_{\text{SDF}},\tag{3}$$

where  $L_{\text{CD}} := \text{CD}_{\ell_1}(\overline{\mathcal{M}}, \mathbf{P}_{k^*})$  denotes the  $\ell_1$ -norm-based Chamfer distance between  $\overline{\mathcal{M}}$  and  $\mathbf{P}_{k^*}$ ;  $L^*_{\text{Norm}} := \text{NC}_{\ell_1}(\overline{\mathcal{M}}, \mathbf{P}_{k^*})$  measures the consistency of the corresponding normals (see the *Supplementary Material* for details);  $L_{\text{SDF}}$  is employed to promote the accuracy of the learned SDF, defined as

$$L_{\text{SDF}} = \frac{1}{|\mathbf{Q}|} \sum_{\mathbf{q} \in \mathbf{Q}} |s(\mathbf{q}) - \text{SDF}_{\text{IMLS}}(\mathbf{q}, \mathbf{P}_{k^*})|, \qquad (4)$$

where  $\text{SDF}_{\text{IMLS}}(\mathbf{q}, \mathbf{P}_{k^*})$  is the SDF value at  $\mathbf{q}$  approximated through implicit moving least-squares [21] (refer to the *Supplementary Material*).

**Remark.** It is worth noting that the resulting template surface  $\overline{\mathcal{M}}$  will be *adaptively* updated/enhanced during the following deformation stage to facilitate the deformation stage, thus promoting the quality of reconstructed surfaces of the remaining point cloud frames.

#### 3.3 Temporal Reconstruction via Learnable Deformation Field

In this stage, as illustrated in Fig. 4, we propose a learnable deformation field based on control points blending. This deformation field enables us to deform the *adaptively-enhanced* template surface to achieve temporal reconstruction. By employing this approach, we ensure the temporal consistency of the reconstructed surface sequence.



**Fig. 4:** Illustration of the proposed control points blending-based learnable deformation stage for temporal reconstruction. Note that the adaptively enhanced template surface will be deformed to all frames.

Control points blending-based learnable deformation field. We first build a set of control point  $\mathbf{U} := {\mathbf{u}_1, ..., \mathbf{u}_{|\mathbf{U}|}}$ , which are obtained using the farthest point sampling (FPS) algorithm [11] on the vertex set of the tetrahedron mesh  $\mathcal{T}$  (i.e.,  $\mathbf{Q}$ ) such that the control points  $\mathbf{u}_r$ ,  $r = 1, 2, ..., |\mathbf{U}|$  are uniformly distributed in  $\mathcal{T}$ . Note that the positions of initial control points are also adaptively optimized during training to make the deformation more flexible. We associate each  $\mathbf{u}_r$  with a rotation matrix  $\mathbf{R}_r^k \in SO(3)$  and a translation  $\mathbf{t}_r^k \in \mathbb{R}^3$  for each frame  $\mathbf{P}_k$ , k = 1, 2, ..., K. Then for each vertex of the template surface  $\overline{\mathbf{v}} \in \overline{\mathbf{V}}$ , which is the vertex set of the surface  $\overline{\mathcal{M}}$ . The deformed vertex at the k-th frame is represented as

$$\mathbf{v}^{k} = \sum_{r=1}^{|\mathbf{U}|} \omega(\overline{\mathbf{v}}, \mathbf{u}_{r}) \cdot (\mathbf{R}_{r}^{k} \overline{\mathbf{v}} + \mathbf{t}_{r}^{k}),$$
(5)

with the weight  $\omega(\overline{\mathbf{v}}, \mathbf{u}_r)$  learned via

$$\omega(\overline{\mathbf{v}}, \mathbf{u}_r) = \frac{\exp(F_{\mathbf{\Omega}}([\overline{\mathbf{v}}||(\overline{\mathbf{v}} - \mathbf{u}_r)]))}{\sum_{r=1}^{|\mathbf{U}|} \exp(F_{\mathbf{\Omega}}([\overline{\mathbf{v}}||(\overline{\mathbf{v}} - \mathbf{u}_r)]))},$$
(6)

where  $F_{\Omega}(\cdot)$  is an MLP parameterized with  $\Omega$ , and  $[\cdot||\cdot]$  is the concatenation operator. Before training, we pre-optimize  $\Omega$  so that  $\omega(\overline{\mathbf{v}}, \mathbf{u}_r) = \tilde{\omega}(\overline{\mathbf{v}}, \mathbf{u}_r)$ , where

$$\tilde{\omega}(\overline{\mathbf{v}}, \mathbf{u}_r) = \frac{\exp(-\|\overline{\mathbf{v}} - \mathbf{u}_r\|^2 / 2\eta^2)}{\sum_{r=1}^{|\mathbf{U}|} \exp(-\|\overline{\mathbf{v}} - \mathbf{u}_r\|^2 / 2\eta^2)},\tag{7}$$

The reconstructed surface of  $\mathbf{P}_k$  can be obtained as  $\mathcal{M}_k = {\mathbf{V}_k, \mathbf{F}}$  with  $\mathbf{V}_k$  being the collection of  $\mathbf{v}^k$ .

We also utilize an MLP  $F_{\Xi}(\cdot)$  parametrized with  $\Xi$  to learn the deformation transformations of all frames, i.e.,  $[\xi_r^k, \mathbf{t}_r^k] = F_{\Xi}(\Gamma([\mathbf{u}_r||k]))$ , where  $\xi_r^k \in \mathfrak{so}(3)$ is the Lie algebra format of  $\mathbf{R}_r^k$ . Owing to the spectral bias property of neural networks [43], the learned transformations would vary smoothly in both spatial and temporal domains to ensure the spatial and temporal smoothness of the constructed mesh sequence.

Learning of deformable field with enhanced template surface. Instead of solely focusing on training the deformation stage, i.e.,  $F_{\Omega}(\cdot)$ ,  $F_{\Xi}(\cdot)$  and U, we jointly train the deformation stage and the template surface representation, i.e.,  $F_{\Phi}(\cdot)$  has been warmed-up in the first stage. Specifically, at each iteration, we extract the template surface using DMT $(\cdot, \cdot)$  and then deform it to all frames for temporally consistent reconstruction. Through joint optimization, we aim to enhance the adaptability of the template surface to the deformation process, ultimately leading to improved temporal reconstruction quality. To achieve this goal, we consider the following loss function:

$$\min_{\mathbf{\Phi},\mathbf{\Omega},\mathbf{\Xi},\mathbf{U}} L = w_1 L_{\text{R-CD}} + w_2 L_{\text{Norm}} + w_3 L_{\text{R-SDF}} + w_4 L_{\text{Smo}} + w_5 L_{\text{Shape}}.$$
 (8)

Specifically,  $L_{\text{R-CD}}$  stands for the proposed robust Chamfer distance to cope with point clouds with holes and outliers, defined as  $L_{\text{R-CD}} = \frac{1}{K} \sum_{k} L_{\text{R-CD}}^{k}$  with

$$L_{\text{R-CD}}^{k} = \frac{1}{|\widetilde{\mathbf{V}}_{k}|} \sum_{\mathbf{v}_{i} \in \widetilde{\mathbf{V}}_{k}} \phi_{\alpha}(\mathbf{v}_{i}, \mathbf{p}_{\rho_{i}}) \cdot \|\mathbf{v}_{i} - \mathbf{p}_{\rho_{i}}\|^{2} + \frac{1}{|\mathbf{P}_{k}|} \sum_{\mathbf{p}_{j} \in \mathbf{P}_{k}} \phi_{\alpha}(\mathbf{p}_{j}, \mathbf{v}_{\tau_{j}}) \cdot \|\mathbf{p}_{j} - \mathbf{v}_{\tau_{j}}\|^{2},$$
(9)

where  $\widetilde{\mathbf{V}}_k$  is the set of points uniformly sampled from  $\mathcal{M}_k$ , and  $\phi_{\alpha}(x, y) = \exp(-\alpha \cdot ||x - y||^2)$ .  $\mathbf{p}_{\rho_i}$  and  $\mathbf{v}_{\tau_j}$  are the closest points for  $\mathbf{v}_i$  and  $\mathbf{p}_j$  respectively. The normal loss term  $L_{\text{Norm}} = \sum_k \text{NC}_{\ell_1}(\widetilde{\mathbf{V}}_k, \mathbf{P}_k)/K$  constrains the consistency of normals. The proposed robust SDF loss term  $L_{\text{R-SDF}}$  regularizes the orientation of the deformed mesh, defined as

$$L_{\text{R-SDF}} = \frac{1}{K \cdot |\mathbf{Q}_s|} \sum_{\mathbf{q} \in \mathbf{Q}_s} \sum_k \overline{\psi}_{\beta}^k(\mathbf{q}) \cdot |f_{\gamma}(s(\widehat{\mathbf{q}}^k)) - f_{\gamma}(\text{SDF}_{\text{IMLS}}(\widehat{\mathbf{q}}^k, \mathbf{P}_k))|, \quad (10)$$

where  $\widehat{\mathbf{q}}^k$  is the deformed position of  $\mathbf{q}$  according to Eq. (5);  $\overline{\psi}^k_{\beta}(\mathbf{q}) = \frac{\psi^k_{\beta}(\mathbf{q})}{\sum_k \psi^k_{\beta}(\mathbf{q})}$ with  $\psi^k_{\beta}(\mathbf{q}) = \exp(-\beta \cdot |\text{SDF}_{\text{IMLS}}(\widehat{\mathbf{q}}^k, \mathbf{P}_k)|^2)$  is to measure the confidence of each term. That is, when the deformed point  $\widehat{\mathbf{q}}^k$  has a smaller value in the SDF approximated by the target point cloud, it is closer to the surface, and we assume that it is more accurate;  $\mathbf{Q}_s$  is a subset randomly sampled from  $\mathbf{Q}$ .  $f_{\gamma}(x) = \frac{1}{1 + \exp(-\gamma \cdot x)}$  is used to scale the amplified SDF value to (0, 1) because we pay more attention to whether the point  $\widehat{\mathbf{q}}^k$  is inside or outside the surface.

The loss term  $L_{\rm Smo}$  ensures the smoothness of the deformation, defined as

$$L_{\text{Smo}} = \frac{1}{2K|\mathcal{E}|} \sum_{k} \sum_{l \in \mathcal{I}(k)} \sum_{(\mathbf{v}_i, \mathbf{v}_j) \in \mathcal{E}} \|(\mathbf{v}_i^k - \mathbf{v}_i^l) - (\mathbf{v}_j^k - \mathbf{v}_j^l)\|^2, \quad (11)$$

where  $\mathcal{E}$  is the edge set derived by face set  $\mathbf{F}$ .  $\mathcal{I}(k) = \{k-1, k+1\}$ . When k = 1 or k = K, we only use one neighbor point cloud frame. Finally, it is crucial to ensure that the enhanced template surface maintains a reasonable shape throughout the joint optimization process. Thus, we constrain the deformation field at keyframe as close to the identity transformation as possible with a loss term  $L_{\text{Shape}}$ :

$$L_{\text{Shape}} = \sum_{\mathbf{u}_r \in \mathbf{U}} \| [\xi_r^{k^*}, \mathbf{t}_r^{k^*}] \|^2,$$
(12)

where  $[\xi_r^{k^*}, \mathbf{t}_r^{k^*}] \in \mathbb{R}^6$  denotes the rotation matrix represented by Lie algebra format and translation vector of the keyframe.

## 4 Experiments

#### 4.1 Experimental Settings

**Implementation details.** In our approach, we employed MLPs consisting of 5 linear layers with a feature dimension of 128 to implement  $F_{\Phi}(\cdot)$ ,  $F_{\Xi}(\cdot)$ , and  $F_{\Omega}(\cdot)$ . Training was conducted using the ADAM optimizer [20] with a learning rate of  $1 \times 10^{-4}$ . The *Supplementary Material* provides more details. We performed all experiments on a single NVIDIA A6000 GPU.

**Datasets.** We utilized three well-established benchmark datasets, namely DFA-UST [4], DT4D [27], and AMA [51], for our evaluation. From motion sequences, we selected 17 consecutive frames to form a set of input point clouds. Randomly sampling 5000 points from the ground truth mesh, we ensured diversity in our point cloud selection. Adopting the train and test data splits provided by [23], we used 109 sets and 89 sets from the test data of the DFAUST dataset and DT4D-animal dataset, respectively. The AMA dataset comprises 10 motion sequences; for supervised learning comparison methods, we employed 1176 sets from 7 sequences as the train set and 32 sets of point clouds from the remaining 3 sequences ("crane", "march1" and "samba") as the test set. To maintain consistency across sequences, we normalized the diagonal length of the bounding



Fig. 5: Comparison of visual results by different methods on the DFAUST dataset.

box in the first frame to 1 and centered it at the origin of the coordinates. Subsequent point clouds underwent transformations based on the scale and translation of the first frame. We utilized the algorithm in pymeshlab [33] to estimate the normals of the point cloud and adjusted the normal directions to face outward.

Dataset	Method	CD (×10 <sup>-4</sup> ) $\downarrow$	$\mathrm{NC}\uparrow$	F-0.5% $\uparrow$	F-1% $\uparrow$	Corr. (×10 <sup>-2</sup> ) $\downarrow$
DFAUST [4]	LPDC [49] Cadex [23] Ours	$2.430 \\ 1.062 \\ 0.688$	$\begin{array}{c} 0.929 \\ 0.941 \\ 0.953 \end{array}$	$0.299 \\ 0.519 \\ 0.894$	$\begin{array}{c} 0.633 \\ 0.823 \\ 0.985 \end{array}$	$1.46 \\ 1.28 \\ 1.02$
DT4D [27]	Cadex [23] Ours	22.377 1.036	$\begin{array}{c} 0.868 \\ 0.933 \end{array}$	$0.269 \\ 0.582$	$0.532 \\ 0.870$	4.99 3.23
AMA [51]	LPDC [49] Cadex [23] Ours	108.822 71.041 0.320	$0.666 \\ 0.663 \\ 0.918$	0.045 0.052 0.683	0.101 0.119 0.943	$14.1 \\ 13.6 \\ 4.44$

Table 1: Quantitative comparisons of different methods on three datasets.

Methods under comparison. We conducted a comparative analysis involving our proposed method and two cutting-edge techniques specialized in dynamic surface reconstruction from sequences of point clouds: LPDC [49] and Cadex [23]. *Note* that both Cadex and LPDC necessitate ground-truth occupancy values and ground-truth correspondences for supervision during training. For Cadex, we employed the officially released pre-trained model provided by the authors for testing on DFAUST and DT4D datasets. As for LPDC, the authors exclusively performed experiments on the DFAUST dataset and released the pre-trained model. To ensure a comprehensive comparison, we added the comparison on the



Fig. 6: Comparison of visual results by different methods on the AMA dataset.

AMA dataset, a human body dataset characterized by *significantly larger* deformation and movement ranges compared to DFAUST. Furthermore, we retrained Cadex and LPDC on the AMA dataset.

**Evaluation metrics.** Following previous works [23, 32, 40], we scaled the test data to its original size and quantitatively assessed various methods using  $\ell_2$ -norm-based Chamfer Distance (CD), Normal Consistency (NC), F-score with thresholds of 0.5% (F-0.5%) and 1% (F-1%), and Correspondences Error (Corr.) to evaluate different methods quantitatively. For the visual results, we assigned the same colors to corresponding vertices to indicate the accuracy of learned correspondences by different methods. In addition, we also show the error map to visualize the pointwise distance error from the ground-truth surface.

#### 4.2 Comparisons with State-of-the-Art Methods

**Results of clean and complete data.** Table 1 shows the comparison of our method against state-of-the-art methods on the three datasets. The results demonstrate that our method attains the highest accuracy. The visual comparisons in Figs. 5 and 6 further validate the significant superiority of our method. Notably, Cadex and LPDC exhibit notably poor performance on the AMA dataset. This discrepancy could be attributed to the smaller data volume of the AMA dataset compared to DFAUST and DT4D, shedding light on the limitations of these supervised learning methods. We also refer readers to the *Supplementary Material* and *Video Demo* for more visual results.

**Results of noisy and partially missing data.** We also evaluate the performance of our method on noisy data. We add 0.5% Gaussian noise for each point in a point cloud sequence in DFAUST dataset [4] and show the results in Fig. 7. Since point cloud data collected by real-depth cameras often contain missing parts, we tested the situation when there are some holes in the surface.

For a sequence in "swing" motion sequences of AMA dataset [51], we placed two virtual cameras in front of and behind the mesh model for simulated acquisition and merged the point clouds scanned by the two cameras. Fig. 7 shows the performance of our method and the variants. We can see that our method can reconstruct partially missing meshes. More results are presented in the *Supplementary Material*.



Fig. 7: Results by our method on noisy data (top) and partially missing data (bottom).

#### 4.3 Ablation Study

In this section, we conducted thorough ablation studies targeting key components of our framework on the AMA dataset, where we took 1 set from each sequence and 10 sets in total for data diversity. We focused on the following aspects:

1) Learnable deformation field. To evaluate the effectiveness of the proposed deformation field, we conducted experiments by maintaining the other components and substituting the deformation field:

- Pointwise deformation based on an MLP: For any point  $\overline{\mathbf{v}} \in \mathbb{R}^3$  in the template surface and frame index k, we established a deformation function  $F_{\Psi}(\Gamma([\overline{\mathbf{v}}||k]))$  and obtained the deformed position  $\mathbf{v}^k = \overline{\mathbf{v}} + F_{\Psi}(\Gamma([\overline{\mathbf{v}}||k]))$ . Here,  $F_{\Psi}(\cdot)$  is an MLP with 5 linear layers and a feature dimension of 128.
- Fixed blending weight formula: We replaced the learnable blending weights in Eq. (5) by setting it to  $\tilde{\omega}(\bar{\mathbf{v}}, \mathbf{u}_r)$  in Eq. (7) throughout the learning process.
- Fixed control point positions: We maintained the positions  $\mathbf{u}_r$  of the control points at their initial values.

From Tab. 2 and Fig. 8, the three variants produce worse quantitative and visual results than Ours, verifying the advantage of our proposed deformation field.



**Fig. 8:** Comparison of visual results by different variants of the deformation field in our method. (a) Input; (b) Point-wise deformation based on an MLP; (c) Fixed blending weight formula; (d) Fixed control point positions; (e)  $|\mathbf{U}| = 10$ ; (f)  $|\mathbf{U}| = 100$ ; (g) Ours.

Fig. 9 also visualizes the learned weights via Eq. (6) and those computed via Eq. (7) for blending, where we can see that the learned weights can capture the motion information and well handle the deformations for the parts with close spatial locations but inconsistent motions.

**2)** Size of  $|\mathbf{U}|$ . From Tab. 2 and Fig. 8, where we set different numbers of control points i.e., 10, 30, and 100, we can observe that too few control points result in insufficient degrees of freedom and suboptimal alignment, while too many control points can weaken local consistency. We set  $|\mathbf{U}| = 30$  as a balanced trade-off.

3) Loss functions during temporal reconstruction. To assess the effectiveness of various loss terms used in the temporal reconstruction stage, we individually excluded the robust SDF loss  $L_{\text{R-SDF}}$ , deformation smooth loss  $L_{\text{Smo}}$ , and shape preservation loss  $L_{\text{Shape}}$ , while keeping other aspects unchanged. The results presented in Tab. 2 demonstrate their necessity and contributions.

4) Keyframe selection, template surface learning and arbitrary length of input sequences. The performances and discussions can be found in the Supplementary Material.

# 5 Conclusion and Discussion

We introduced a new unsupervised temporally-consistent dynamic surface reconstruction framework for time-varying point cloud sequences. Technically, we integrated template surface representation based on the deformable tetrahedron and a learnable deformation field. We first proposed a coarse-to-fine learning strategy for constructing the template surface. Then we designed a learnable deformation field using the learnable control points and blending weights for

Variants	$\big \operatorname{CD}(\times 10^{-4})\downarrow$	$\mathrm{NC}\uparrow$	F-0.5% $\uparrow$	$\text{F-1\%}\uparrow$	Corr. (×10 <sup>-2</sup> ) $\downarrow$
MLP-Deform	0.412	0.909	0.618	0.917	5.91
w. $\tilde{\omega}(\overline{\mathbf{v}}, \mathbf{u}_r)$	0.475	0.911	0.564	0.890	3.26
Fixed $\{\mathbf{u}_r\}$	0.372	0.911	0.637	0.927	3.23
w.o. $L_{\text{R-SDF}}$	0.330	0.918	0.649	0.937	3.14
w.o. $L_{\text{Norm}}$	0.656	0.900	0.582	0.884	3.30
w.o. $L_{\rm Smo}$	0.378	0.914	0.616	0.923	3.30
w.o. $L_{\text{Shape}}$	0.335	0.919	0.636	0.936	<u>2.91</u>
U  = 10	0.437	0.910	0.594	0.910	3.30
U  = 30 (Ours)	0.324	0.919	0.649	0.939	2.89
U  = 100	0.315	0.919	0.651	0.943	3.16

**Table 2:** Ablation studies of our method on the AMA dataset [51]. The best and second best results are highlighted in **bold** and <u>underlined</u>.



**Fig. 9:** Visual learnable control points and blending weights. We assigned a color  $\mathbf{c}_r$  to each control point  $\mathbf{u}_r$  based on its spatial position, and obtained the color for each  $\overline{\mathbf{v}}$  by  $\sum_{r=1}^{|\mathbf{U}|} \omega(\overline{\mathbf{v}}, \mathbf{u}_r) \cdot \mathbf{c}_r$ . Closer colors indicate more similar deformation fields.

temporally reconstruction. We jointly optimized the enhanced template surface and the learnable deformation field. Experiments demonstrate that our proposed method performs even much better than SOTA supervised learning methods.

Although our method has achieved good performance on many examples, it may encounter limitations in cases where the point cloud in certain areas is sparse or the local structure is elongated but exhibits significant non-rigid deformations (as depicted on the right). This is due to our utilization of Chamfer distance as the alignment supervision, which restricts the establishment of correspondences between adjacent frames solely



based on spatial proximity. Incorporating more robust metrics for shape alignment would enhance the effectiveness of our approach.

# Acknowledgement

This project was supported by the Hong Kong Research Grants Council under Grants 11219324, 11219422, and 11202320.

# References

- Atzmon, M., Lipman, Y.: Sal: Sign agnostic learning of shapes from raw data. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 2565–2574 (2020)
- 2. Baieri, D., Esposito, S., Maggioli, F., Rodolà, E.: Fluid dynamics network: Topology-agnostic 4d reconstruction via fluid dynamics priors. In: arXiv.org (2023)
- Barber, C.B., Dobkin, D.P., Huhdanpaa, H.: The quickhull algorithm for convex hulls. ACM Transactions on Mathematical Software (TOMS) 22(4), 469–483 (1996)
- Bogo, F., Romero, J., Pons-Moll, G., Black, M.J.: Dynamic faust: Registering human bodies in motion. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 6233–6242 (2017)
- Boulch, A., Marlet, R.: Poco: Point convolution for surface reconstruction. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 6302–6314 (2022)
- Božič, A., Palafox, P., Zollhöfer, M., Thies, J., Dai, A., Nießner, M.: Neural deformation graphs for globally-consistent non-rigid reconstruction. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 1450–1459 (2021)
- 7. Chen, H., Yao, Y., Zhang, J.: Neural-abc: Neural parametric models for articulated body with clothes. IEEE Trans. Vis. Comput. Graph. (2024)
- Chen, R.T., Rubanova, Y., Bettencourt, J., Duvenaud, D.K.: Neural ordinary differential equations. Adv. Neural Inform. Process. Syst. 31 (2018)
- Chi, C., Song, S.: Garmentnets: Category-level pose estimation for garments via canonical space shape completion. In: Int. Conf. Comput. Vis. pp. 3324–3333 (2021)
- Dou, M., Khamis, S., Degtyarev, Y., Davidson, P., Fanello, S.R., Kowdle, A., Escolano, S.O., Rhemann, C., Kim, D., Taylor, J., et al.: Fusion4d: Real-time performance capture of challenging scenes. ACM Trans. Graph. 35(4), 1–13 (2016)
- Eldar, Y., Lindenbaum, M., Porat, M., Zeevi, Y.Y.: The farthest point strategy for progressive image sampling. IEEE Trans. Image Process. 6(9), 1305–1315 (1997)
- Feng, W., Zhang, J., Cai, H., Xu, H., Hou, J., Bao, H.: Recurrent multi-view alignment network for unsupervised surface registration. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 10297–10307 (2021)
- Feng, Y., Feng, Y., You, H., Zhao, X., Gao, Y.: Meshnet: Mesh neural network for 3d shape representation. In: AAAI. vol. 33, pp. 8279–8286 (2019)
- 14. Hanocka, R., Hertz, A., Fish, N., Giryes, R., Fleishman, S., Cohen-Or, D.: Meshcnn: a network with an edge. ACM Trans. Graph. **38**(4), 1–12 (2019)
- Huang, J., Gojcic, Z., Atzmon, M., Litany, O., Fidler, S., Williams, F.: Neural kernel surface reconstruction. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 4369–4379 (2023)
- Jiang, B., Ren, X., Dou, M., Xue, X., Fu, Y., Zhang, Y.: Lord: Local 4d implicit representation for high-fidelity dynamic human modeling. In: Eur. Conf. Comput. Vis. pp. 307–326. Springer (2022)
- Jiang, B., Zhang, Y., Wei, X., Xue, X., Fu, Y.: H4d: human 4d modeling by learning neural compositional representation. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 19355–19365 (2022)

- 16 Y. Yao, S. Ren et al.
- Kazhdan, M., Bolitho, M., Hoppe, H.: Poisson surface reconstruction. In: Proceedings of the fourth Eurographics symposium on Geometry processing. vol. 7, p. 0 (2006)
- Kazhdan, M., Hoppe, H.: Screened poisson surface reconstruction. ACM Trans. Graph. 32(3), 1–13 (2013)
- Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Int. Conf. Learn. Represent. (2015)
- 21. Kolluri, R.: Provably good moving least squares. ACM Transactions on Algorithms (TALG) 4(2), 1–25 (2008)
- Koneputugodage, C.H., Ben-Shabat, Y., Gould, S.: Octree guided unoriented surface reconstruction. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 16717–16726 (2023)
- Lei, J., Daniilidis, K.: Cadex: Learning canonical deformation coordinate space for dynamic surface representation via neural homeomorphism. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 6624–6634 (2022)
- Li, H., Adams, B., Guibas, L.J., Pauly, M.: Robust single-view geometry and motion reconstruction. ACM Trans. Graph. 28(5), 1–10 (2009)
- Li, H., Sumner, R.W., Pauly, M.: Global correspondence optimization for non-rigid registration of depth scans. In: Computer graphics forum. vol. 27, pp. 1421–1430. Wiley Online Library (2008)
- Li, Y., Harada, T.: Non-rigid point cloud registration with neural deformation pyramid. Adv. Neural Inform. Process. Syst. 35, 27757–27768 (2022)
- Li, Y., Takehara, H., Taketomi, T., Zheng, B., Nießner, M.: 4dcomplete: Non-rigid motion estimation beyond the observable surface. Int. Conf. Comput. Vis. (2021)
- Li, Y., Bu, R., Sun, M., Wu, W., Di, X., Chen, B.: Pointcnn: Convolution on x-transformed points. Adv. Neural Inform. Process. Syst. 31 (2018)
- Liu, S.L., Guo, H.X., Pan, H., Wang, P.S., Tong, X., Liu, Y.: Deep implicit moving least-squares functions for 3d reconstruction. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 1788–1797 (2021)
- Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: A skinned multi-person linear model. ACM Trans. Graphics (Proc. SIGGRAPH Asia) 34(6), 248:1–248:16 (Oct 2015)
- Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3d surface construction algorithm. ACM siggraph computer graphics 21(4), 163–169 (1987)
- Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 4460–4470 (2019)
- Muntoni, A., Cignoni, P.: PyMeshLab (Jan 2021). https://doi.org/10.5281/ zenodo.4438750
- Newcombe, R.A., Fox, D., Seitz, S.M.: Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 343–352 (2015)
- Nielson, G.M.: Dual marching tetrahedra: Contouring in the tetrahedronal environment. In: International Symposium on Visual Computing. pp. 183–194. Springer (2008)
- Niemeyer, M., Mescheder, L., Oechsle, M., Geiger, A.: Occupancy flow: 4d reconstruction by learning particle dynamics. In: Int. Conf. Comput. Vis. pp. 5379–5389 (2019)
- Palafox, P., Božič, A., Thies, J., Nießner, M., Dai, A.: Npms: Neural parametric models for 3d deformable shapes. In: Int. Conf. Comput. Vis. pp. 12695–12705 (2021)

- Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous signed distance functions for shape representation. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 165–174 (2019)
- Park, K., Sinha, U., Barron, J.T., Bouaziz, S., Goldman, D.B., Seitz, S.M., Martin-Brualla, R.: Nerfies: Deformable neural radiance fields. In: Int. Conf. Comput. Vis. pp. 5865–5874 (2021)
- Peng, S., Jiang, C., Liao, Y., Niemeyer, M., Pollefeys, M., Geiger, A.: Shape as points: A differentiable poisson solver. Adv. Neural Inform. Process. Syst. 34, 13032–13044 (2021)
- Peng, S., Niemeyer, M., Mescheder, L., Pollefeys, M., Geiger, A.: Convolutional occupancy networks. In: Eur. Conf. Comput. Vis. pp. 523–540. Springer (2020)
- Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 652–660 (2017)
- Rahaman, N., Baratin, A., Arpit, D., Draxler, F., Lin, M., Hamprecht, F., Bengio, Y., Courville, A.: On the spectral bias of neural networks. In: Int. Conf. Mach. Learn. pp. 5301–5310 (2019)
- Rempe, D., Birdal, T., Zhao, Y., Gojcic, Z., Sridhar, S., Guibas, L.J.: Caspr: Learning canonical spatiotemporal point cloud representations. Adv. Neural Inform. Process. Syst. 33, 13688–13701 (2020)
- 45. Shen, T., Gao, J., Yin, K., Liu, M.Y., Fidler, S.: Deep marching tetrahedra: a hybrid representation for high-resolution 3d shape synthesis. Adv. Neural Inform. Process. Syst. 34, 6087–6101 (2021)
- 46. Shen, T., Munkberg, J., Hasselgren, J., Yin, K., Wang, Z., Chen, W., Gojcic, Z., Fidler, S., Sharp, N., Gao, J.: Flexible isosurface extraction for gradient-based mesh optimization. ACM Trans. Graph. 42(4), 1–16 (2023)
- Slavcheva, M., Baust, M., Cremers, D., Ilic, S.: Killingfusion: Non-rigid 3d reconstruction without correspondences. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 1386–1395 (2017)
- Sumner, R.W., Schmid, J., Pauly, M.: Embedded deformation for shape manipulation. ACM Trans. Graph. p. 80 (2007)
- Tang, J., Xu, D., Jia, K., Zhang, L.: Learning parallel dense correspondence from spatio-temporal descriptors for efficient and robust 4d reconstruction. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 6022–6031 (2021)
- Tretschk, E., Tewari, A., Golyanik, V., Zollhöfer, M., Lassner, C., Theobalt, C.: Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In: Int. Conf. Comput. Vis. pp. 12959–12970 (2021)
- Vlasic, D., Baran, I., Matusik, W., Popović, J.: Articulated mesh animation from multi-view silhouettes. ACM Trans. Graph. 27(3), 97:1–97:9 (Aug 2008)
- Wu, Z., Li, K., Lai, Y.K., Yang, J.: Global as-conformal-as-possible non-rigid registration of multi-view scans. In: Int. Conf. Multimedia and Expo. pp. 308–313. IEEE (2019)
- Yang, J., Guo, D., Li, K., Wu, Z., Lai, Y.K.: Global 3d non-rigid registration of deformable objects using a single rgb-d camera. IEEE Trans. Image Process. 28(10), 4746–4761 (2019)
- Yao, Y., Deng, B., Xu, W., Zhang, J.: Quasi-newton solver for robust non-rigid registration. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 7600–7609 (2020)
- Yao, Y., Deng, B., Xu, W., Zhang, J.: Fast and robust non-rigid registration using accelerated majorization-minimization. IEEE Trans. Pattern Anal. Mach. Intell. (2023)

- 18 Y. Yao, S. Ren et al.
- Yenamandra, T., Tewari, A., Bernard, F., Seidel, H.P., Elgharib, M., Cremers, D., Theobalt, C.: i3dmm: Deep implicit 3d morphable model of human heads. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 12803–12813 (2021)
- 57. Zeng, A., Song, S., Nießner, M., Fisher, M., Xiao, J., Funkhouser, T.: 3dmatch: Learning local geometric descriptors from rgb-d reconstructions. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 1802–1811 (2017)
- Zeng, Y., Qian, Y., Zhu, Z., Hou, J., Yuan, H., He, Y.: Corrnet3d: Unsupervised end-to-end learning of dense correspondence for 3d point clouds. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 6052–6061 (2021)
- Zhang, B., Li, J., Deng, X., Zhang, Y., Ma, C., Wang, H.: Self-supervised learning of implicit shape representation with dense correspondence for deformable objects. In: Int. Conf. Comput. Vis. pp. 14268–14278 (2023)
- Zuffi, S., Kanazawa, A., Jacobs, D.W., Black, M.J.: 3d menagerie: Modeling the 3d shape and pose of animals. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 6365–6373 (2017)