

Root Pose Decomposition Towards Generic Non-rigid 3D Reconstruction with Monocular Videos

Yikai Wang¹ Yinpeng Dong^{1,2} Fuchun Sun¹✉ Xiao Yang¹

¹Beijing National Research Center for Information Science and Technology (BNRist),
 State Key Lab on Intelligent Technology and Systems,
 Department of Computer Science and Technology, Tsinghua University ²RealAI
 {yikaiw, dongyinpeng, fcsun}@tsinghua.edu.cn, yangxiao19@mails.tsinghua.edu.cn

Abstract

This work focuses on the 3D reconstruction of non-rigid objects based on monocular RGB video sequences. Concretely, we aim at building high-fidelity models for generic object categories and casually captured scenes. To this end, we do not assume known root poses of objects, and do not utilize category-specific templates or dense pose priors. The key idea of our method, Root Pose Decomposition (RPD), is to maintain a per-frame root pose transformation, meanwhile building a dense field with local transformations to rectify the root pose. The optimization of local transformations is performed by point registration to the canonical space. We also adapt RPD to multi-object scenarios with object occlusions and individual differences. As a result, RPD allows non-rigid 3D reconstruction for complicated scenarios containing objects with large deformations, complex motion patterns, occlusions, and scale diversities of different individuals. Such a pipeline potentially scales to diverse sets of objects in the wild. We experimentally show that RPD surpasses state-of-the-art methods on the challenging DAVIS, OVIS, and AMA datasets. We provide video results in <https://rpd-share.github.io>.

1. Introduction

The reconstruction of non-rigid (or deformable) 3D objects with monocular RGB videos is a long-standing and challenging task in computer vision and graphics [34]. It is needed for a variety of applications ranging from XR to robotics. Traditionally, a typical prior pipeline leverages template-based models such as human skeleton models SMPL [19], SMPL-X [25], GHUM(L) [42], and reconstruct models for specific categories like the human body or human face. These methods do not scale to diverse categories.

With the success of Neural Radiance Field (NeRF) [21], several representative works [5, 23] unify frames to a canonical space and do not rely on pre-defined skeleton models. Whilst a variety of methods have been proposed to improve the reconstruction fidelity, the performance degenerates when encountering large object deformations or movements. Meanwhile, they are not suitable for casual videos when background Structure from Motion (SfM) does not provide root poses for the object. As a typical work to address the issues, BANMo [41] initializes approximate camera poses by leveraging continuous surface embeddings (CSE) [22] or pre-trained DensePose models [12]. Nevertheless, CSE is acquired by annotations and only applies to specific categories in the training set, *e.g.*, quadruped animals. DensePose is learned with the aid of manual UV fields from SMPL and is thus also highly category-limited. Such a pipeline that relies on off-the-shelf pose or surface models of a certain categories does not generalize to reconstruct generic object categories.

As a result, it is an open problem when considering the non-rigid construction in the wild that might contain complicated factors, *e.g.*, multiple categories, complex motion patterns, individual diversities, or object occlusions. Based on casually captured monocular videos, our effort is devoted to building articulated models for generic categories, without explicitly incorporating priors that might limit the generalization of categories. Towards this goal, we propose **Root Pose Decomposition (RPD)**, a method for non-rigid 3D reconstruction based on monocular RGB videos. RPD does not rely on known camera poses or poses compensated by background-SfM, category-specific skeletons, or pre-trained dense pose models (*e.g.*, DensePose, CSE), while could achieve articulated reconstruction for objects with rapid object deformations, complex motion patterns, and large pose changes.

Concretely, RPD follows the common approach for non-rigid reconstruction that builds a canonical space for dif-

✉ Corresponding author: Fuchun Sun.

Method	Nerfies [23]	ViSER [40]	BANMo [41]	Ours
Dependency (\times is preferred)				
Known camera poses	\times	\times	\times	\times
Background-SfM*	\checkmark	\times	\times	\times
Pre-trained CSE/DensePose \ddagger	\times	\times	\checkmark	\times
Algorithm				
Canonical space sharing* $\dagger\ddagger$	Multi-frame	Multi-frame	Multi-frame	Multi-frame & multi-object
Warping function	\checkmark	\checkmark	\checkmark	\checkmark
Dense transformation field* $\dagger\ddagger$	SE(3)	\times	\times	Sim(3)
Linear skinning weights*	\times	\checkmark	\checkmark	\checkmark
Registration* $\dagger\ddagger$	Photometric	Self-supervised feature	CSE feature	Non-rigid point registration
Shape representation \dagger	Implicit	Mesh	Implicit	Implicit
Performance (\checkmark is preferred)				
Handling large deformation* \dagger	\times	\times	\checkmark	\checkmark
Reconstruction for generic categories \ddagger	\checkmark	\checkmark	\times	\checkmark
Multi-object occlusions & differences* $\dagger\ddagger$	\times	\times	\times	\checkmark

Table 1. Comparison with related methods. We highlight the differences with Nerfies, ViSER, and BANMo by *, \dagger , and \ddagger , respectively.

ferent frames with learned warping functions. Here, the warping function consists of a mapping function and also a shared root pose transformation per frame. The core idea of RPD lies in its decomposition of the root pose into local transformations consisting of rotations, translations, and scaling factors at observation points by building a dense Sim(3) field. On the one hand, unlike the rigid root pose, incorporating the local transformation provides more room for each point to fit the canonical model but simultaneously keeps the local transformation continuous with position, direction, and time. On the other hand, maintaining a shared root pose per frame stabilizes the learning process of warping functions.

Besides, towards the generic reconstruction in the wild, we propose techniques that are naturally compatible with RPD. As a result, RPD could be applied to multi-object scenarios with more challenging occluded objects. Apart from that, RPD also adapts to the complex scenes with different individuals that vary in shape, height, or scale. By category-level generalization and the ability to adapt to individual differences, we could reconstruct multiple objects simultaneously and no longer need to collect multiple video sequences that contain the exactly same object individual.

Experiments for evaluating RPD are conducted on three monocular video datasets, including two challenging multi-object segmentation datasets DAVIS [4] and OVIS [26], a reconstruction dataset AMA [36] with ground truth meshes for evaluation, and a casually collected dataset Casual [41]. Experimental results demonstrate the superiority of our design and the potential towards reconstructing large-scale generic object categories with monocular RGB videos.

A comparison with related methods is added in Table 1. To summarize, the contributions of our work are:

- We reveal that estimating per-frame root poses is a core factor for the non-rigid 3D reconstruction of generic categories, with few dependencies on category-specific templates/pose priors, known camera poses, or poses compensated by background-SfM.

- With this motivation, we propose an effective method (RPD) to estimate the per-frame root pose by decomposing the root pose into local transformations, and for the first time, we propose to leverage the success of the non-rigid point registration field for the non-rigid monocular 3D reconstruction.
- Apart from the promising results achieved by RPD, the method is compatible with handling multiple objects with occlusions and individual differences.

2. Related Work

We introduce recent methods of monocular non-rigid 3D reconstruction and non-rigid point cloud registration, since they are both related to our paper.

Monocular non-rigid reconstruction. A large number of works focus on this topic. For example, shape from template (SfT) [1, 10, 13] assumes a static template is given as a prior. Non-rigid structure from motion (NRSfM) [3, 15, 27] constructs objects without using 3D template priors, but might heavily rely on observed point trajectories [28, 31]. We refer readers to the survey [34] for a more detailed overview. Neural rendering methods are more related to ours. To build a higher-fidelity canonical space based on NeRF [21], Nerfies [23] formulates warping functions into a dense SE(3) field that encodes rigid motions with elastic regularization to encourage rigid local deformations. NDR [5] designs bijective mappings with a strictly invertible representation optimized given mainly RGB-D inputs. Nevertheless, many typical approaches [5, 18, 23, 24] assume that object movements are small and camera transformations are given, since learning warping functions is especially hard when encountering large pose changes and object movements, such as running. BANMo [41] models warping functions with linear skinning weights and root poses obtained by off-the-shelf templates (*e.g.*, SMPL [19]). However, the current pipeline with root poses highly relies on category-specific priors like CSE surface features [22]

(*e.g.*, for cats). Different from existing methods, our method handles large object deformations and movements without needing known camera poses, background-SfM, or pre-trained dense pose/CSE models.

Root pose estimation. For objects in monocular videos, “root pose” presents the global 3D orientation of each target object. It remains tough to estimate root poses given the ambiguity in reflective symmetry, textureless region, etc. ShSMesh [44] adopts view synthesis adversarial learning to optimize camera poses. U-CMR [11] proposes camera-multiplex that represents the distribution over cameras, which is yet category-specific. Both methods are designed for rigid objects. For non-rigid objects, NRSfM-related works [3, 8, 33] estimate poses from 2D point trajectories in a class-agnostic manner, but is weak at capturing long-range correspondences or estimating root poses in the wild. DOVE [38] adopts view space shape reflection. BANMo [41] relies on pre-trained CSE [22] and DensePose [12] to initialize root poses, and the category is limited to human or quadruped animals. Hence, BANMo does not apply to generic object categories and is prone to be affected by the transferred quality of pre-built surface/pose embeddings from large-scale datasets. A contemporary work L2G-NeRF [7] predicts accurate camera pose for bundle-adjusting NeRFs. ViSER [40] adopts optical flow that learns approximate pixel-surface embeddings for pose initialization. We have a similar pose initialization scheme as ViSER. Differently, we rectify root poses during training by non-rigid point registration.

Non-rigid point registration. This task aims to build a deformation field or estimate the point-to-point alignment from one point cloud to another [9]. For example, deformation graph [30] builds a sparsely sub-sampled graph from the surface and propagates deformation from node to surface. Leopard [16] learns partial point cloud mapping for guiding the global non-rigid registration [35]. Deformation-Pyramid [17] creates a multi-level deformation motion field for registration in general scenes. This work leverages the idea of non-rigid point registration to learn the decomposed transformation at each observation point.

3. Methodology

In this section, we first introduce the basic framework for non-rigid 3D reconstruction with monocular RGB videos in Sec. 3.1. We describe our proposed RPD, a method for handling large movements and diverse poses of non-rigid objects in Sec. 3.2, followed by its optimization and registration strategies in Sec. 3.3.

3.1. Preliminary: Non-rigid Reconstruction

Prior to going further, we first provide the notations and necessary preliminaries used in our paper. Given sequences of monocular RGB videos with totally T frames that consist

of one or multiple object identities for reconstruction, different object identities of the same category (*e.g.*, human) in different frames are supposed to share a common canonical space. To achieve this, each 3D point $\mathbf{x}_* \in \mathbb{R}^3$ in the canonical space corresponds to a 3D point $\mathbf{x}_t \in \mathbb{R}^3$ in the camera space from the t -th frame (termed as time t) within the T frames. Here, \mathbf{x}_t locates on the ray which emanates from a 2D pixel in the frame image. Inspired by the recent non-rigid reconstruction methods [5, 23, 41], we learn time-variant warping functions $\mathcal{W}_{t \rightarrow *}$ and $\mathcal{W}_{* \rightarrow t}$, satisfying

$$\mathbf{x}_* = \mathcal{W}_{t \rightarrow *}(x_t), \quad \mathbf{x}_t = \mathcal{W}_{* \rightarrow t}(\mathbf{x}_*). \quad (1)$$

Predicting articulated 3D models for non-rigid objects is now translated to reconstructing a static canonical space, where NeRF [21] could be applied to predict two properties of \mathbf{x}_* with Multi-Layer Perception (MLP), including its color $\mathbf{c}_t \in \mathbb{R}^3$ and density $\sigma \in [0, 1]$,

$$\mathbf{c}_t = \text{MLP}_c(\mathbf{x}_*, \mathbf{d}_t, \boldsymbol{\omega}_t), \quad \sigma = \phi_s(\text{MLP}_{\text{SDF}}(\mathbf{x}_*)), \quad (2)$$

where $\mathbf{d}_t \in \mathbb{R}^2$ and $\boldsymbol{\omega}_t$ denote the viewing direction and the latent appearance code [20] at time t , respectively. The object surface is extracted as the zero level-set of signed distance function (SDF), with $\phi_s(\cdot)$ being a unimodal density distribution to convert SDF to the density [37].

When given sampled \mathbf{x}_t from the camera space, the main optimization is performed based on $\mathcal{W}_{t \rightarrow *}$ and Eq. (2) by minimizing the color reconstruction loss and the silhouette reconstruction loss, following existing volume rendering pipelines [21, 43]. Besides, with the help of $\mathcal{W}_{* \rightarrow t}$, a cycle loss could be introduced to maintain the cycle consistency between deformed frames [5, 18].

In summary, the reconstruction of non-rigid objects combines building the canonical space and NeRF. Yet the way to determine warping functions remains challenging and attracts much research effort. As described in Sec. 2, typical approaches include formulating $\mathcal{W}_{t \rightarrow *}$ into a dense SE(3) field [23], or using a hyper-embedding if the surface undergoes topology changes [5, 24], usually under the assumption of small pose changes, small object movements, and known (or by background-SfM) camera parameters [5, 18, 23, 24]. Meanwhile, the reconstruction pipeline with pre-estimated root poses [41] requires reliable category-specific priors like DensePose CSE features [22].

3.2. Root Pose Decomposition

Supposing the warping function $\mathcal{W}_{t \rightarrow *}$ is composed by a mapping function $\mathcal{M}_{t \rightarrow *}$ and a root transformation matrix $\mathbf{G}_t \in \text{SE}(3)$ that uniformly rotates and translates the points at time t to the canonical space. A similar decomposition is also applied to $\mathcal{W}_{* \rightarrow t}$. Specifically, by Eq. (1), there is,

$$\mathbf{x}_* = \mathcal{M}_{t \rightarrow *}(\mathbf{G}_t \mathbf{x}_t), \quad \mathbf{x}_t = \mathbf{G}_t^{-1} \mathcal{M}_{* \rightarrow t}(\mathbf{x}_*). \quad (3)$$

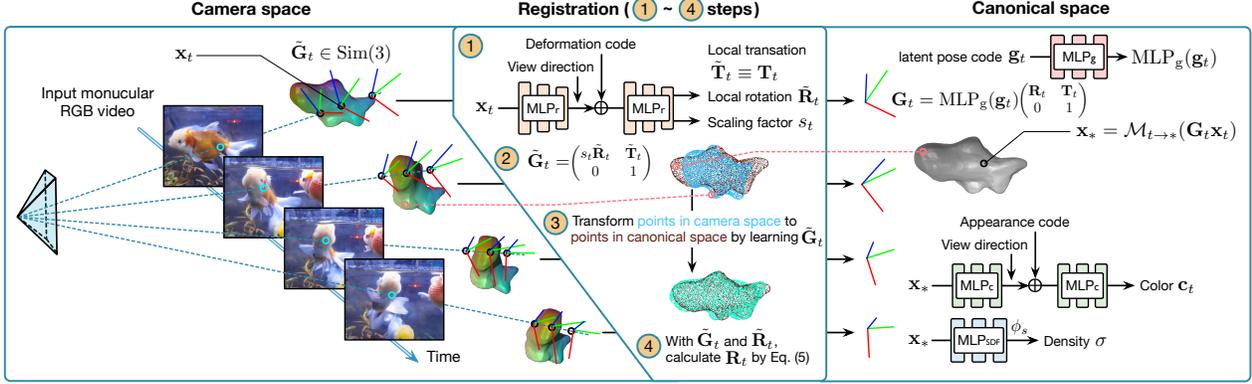


Figure 1. Overall architecture of our RPD for non-rigid monocular reconstruction. In RPD, the root pose transformation \mathbf{G}_t is computed by local transformations $\tilde{\mathbf{G}}_t$ at each point \mathbf{x}_t . Decomposing \mathbf{G}_t into a dense Sim(3) field allows more flexible registration to the canonical space. While maintaining a per-frame global \mathbf{G}_t instead of directly using $\tilde{\mathbf{G}}_t$ helps the optimization of $\mathcal{M}_{t \rightarrow *}$ more stable.

Optimizing a separate root transformation \mathbf{G}_t is to handle large deformations and movements of non-rigid objects in monocular videos. Intuitively, at time t , we expect that \mathbf{G}_t rotates and translates all observed points by fitting the corresponding points in the canonical space to a large extent, making it more focused when learning $\mathcal{M}_{t \rightarrow *}$. To estimate \mathbf{G}_t , existing methods such as PnP solutions might suffer from catastrophic failures given the non-rigidity of objects, and CSE-based pose estimation methods [41] limit the category-level generalization of the reconstruction. In view of this, we propose Root Pose Decomposition (RPD) that learns the root transformation matrix without needing pre-built continuous surface features or camera transformations. An overall architecture of RPD is depicted in Fig. 1.

We let the root transformation $\mathbf{G}_t = \text{MLP}_g(\mathbf{g}_t) \begin{pmatrix} \mathbf{R}_t & \mathbf{T}_t \\ 0 & 1 \end{pmatrix}$ consist of a rotation $\mathbf{R}_t \in \text{SO}(3)$ and a translation $\mathbf{T}_t \in \mathbb{R}^3$, where \mathbf{g}_t denotes the latent pose code. Specifically, \mathbf{g}_t is time-variant and represented by a Fourier embedding [21]. MLP_g takes as input the Fourier embedding and predicts poses with a rotation-translation head. At time t , we parametrize $\tilde{\mathbf{G}}_t$ with the composition of individual local transformations $\tilde{\mathbf{G}}_t = \begin{pmatrix} s_t \tilde{\mathbf{R}}_t & \tilde{\mathbf{T}}_t \\ 0 & 1 \end{pmatrix} \in \text{Sim}(3)$ at 3D points \mathbf{x}_t in the camera space. Here, $\tilde{\mathbf{R}}_t$ and $\tilde{\mathbf{T}}_t$ represent the local rotation and translation at \mathbf{x}_t , respectively. $s_t \in \mathbb{R}^+$ denotes a scaling factor that deals with scale variances of different individuals, as described in Sec. 3.4. We experimentally let $\tilde{\mathbf{T}}_t \equiv \mathbf{T}_t$ for all points but learn the per-point rotation matrix $\tilde{\mathbf{R}}_t \in \text{SO}(3)$ and the scaling factor s_t .

Suppose \mathbf{o}_* is the object center in the canonical space, which could be approximately obtained by averaging sampled points on the object surface given SDF. The corresponding object center at time t is then computed by $\mathbf{o}_t = \mathcal{W}_{* \rightarrow t}(\mathbf{o}_*)$. We experimentally find that setting \mathbf{o}_t to a fixed point already achieves promising performance. We employ a NeRF-style architecture that builds a dense Sim(3) field to calculate the rotation matrix $\tilde{\mathbf{R}}_t$ of each point \mathbf{x}_t around

the center \mathbf{o}_t and also estimate the scaling factor s_t , namely,

$$\tilde{\mathbf{R}}_t, s_t = \text{MLP}_r(\mathbf{x}_t - \mathbf{o}_t, \mathbf{d}'_t, \varphi_t), \quad (4)$$

where $\mathbf{d}'_t \in \mathbb{R}^2$ is the viewing direction of the vector $\mathbf{x}_t - \mathbf{o}_t$ at time t , and φ_t is the latent deformation code at time t .

Inspired by the hierarchical architecture for point registration [17], we parameterize MLP_r in Eq. (4) by a multi-level network where each level outputs the motion increments from its previous level. More details of its optimization will be provided in Sec. 3.3.

Denote \mathbf{x}_t^n as the n -th 3D point sampled along the camera ray that emanates from the 2D frame pixel, and denote i as the index of a ray. We approximate \mathbf{R}_t by considering the rotation matrices at sampled points of all rays. Since multiple works [2, 6, 23] utilize the Frobenius norm to penalize the deviation from the closest rotation, inspired by this, we attain \mathbf{R}_t by seeking a rotation matrix that satisfies,

$$\mathbf{R}_t = \arg \min_{\mathbf{R} \in \text{SO}(3)} \sum_{i,n} \tau^n \|\tilde{\mathbf{R}}_t^n - \mathbf{R}\|_F^2, \quad (5)$$

where $\|\cdot\|_F$ denotes the Frobenius norm, and τ^n represents the probability of \mathbf{x}_t^n being visible to the camera, calculated by $\tau^n = \prod_{m=1}^{n-1} \exp(-\sigma^m \delta^m) (1 - \exp(-\sigma^n \delta^n))$ with the density $\sigma^n = \sigma(\mathcal{W}_{t \rightarrow *})(\mathbf{x}_t^n)$ by Eq. (2) and the interval δ^n between the n -th sampled point and the next. Introducing τ^n to Eq. (5) encourages a point near the surface to acquire greater importance. Here, the parameters (e.g., τ^n , $\tilde{\mathbf{R}}_t^n$) are also related with i , and we omit the script for simplicity.

3.3. Optimization and Registration

Inspired by BANMo [41] that reduces the reconstruction complexity, we apply linear skinning weights with control points to represent $\mathcal{M}_{t \rightarrow *}$ and $\mathcal{M}_{* \rightarrow t}$ in Eq. (3). Here, we mainly focus on optimizing the root transformation and decomposed local transformations.

Let $\mathbf{M}_t \in \mathbb{R}^{3 \times 3}$ be the minimum point of the summation part in Eq. (5) regardless of the SO(3) restriction. We first calculate \mathbf{M}_t and then apply the singular-value decomposition on \mathbf{M}_t . To consider SO(3), we compute \mathbf{R}_t as the closest rotation matrix to \mathbf{M}_t based on its decompositions,

$$\mathbf{M}_t = \frac{1}{\sum_{i,n} \tau^n} \sum_{i,n} \tau^n \tilde{\mathbf{R}}_t^n, \quad \mathbf{M}_t = \mathbf{U}_t \boldsymbol{\Sigma}_t \mathbf{V}_t^T, \quad (6)$$

$$\mathbf{R}_t = \arg \min_{\mathbf{R} \in \text{SO}(3)} \|\mathbf{M}_t - \mathbf{R}\|_F^2 = \mathbf{U}_t \mathbf{V}_t^T. \quad (7)$$

To learn the per-point rotation matrix $\tilde{\mathbf{R}}_t$ and the scaling factor s_t , we sample a set (denoted by \mathbb{S}) of points in the camera space, where the points with high τ^n values acquire large sampling probabilities. Inspired by point registration methods (see Sec. 2 for more details), we encourage \mathbb{S} to be close to a set (denoted by \mathbb{T}) of sampled surface points in the canonical space by leveraging the chamfer distance,

$$\begin{aligned} \mathcal{L}_{\text{cd}} = & \frac{1}{|\mathbb{S}|} \sum_{\mathbf{x}_t^n \in \mathbb{S}} \min_{\mathbf{x}_* \in \mathbb{T}} \|\mathcal{M}_{t \rightarrow *}(\tilde{\mathbf{G}}_t^n \mathbf{x}_t^n) - \mathbf{x}_*\| \\ & + \frac{1}{|\mathbb{T}|} \sum_{\mathbf{x}_* \in \mathbb{T}} \min_{\mathbf{x}_t^n \in \mathbb{S}} \|\mathcal{M}_{t \rightarrow *}(\tilde{\mathbf{G}}_t^n \mathbf{x}_t^n) - \mathbf{x}_*\|, \end{aligned} \quad (8)$$

where the l_1 norm is adopted for the better partial-to-partial registration. By Eq. (8), we optimize $\tilde{\mathbf{R}}_t^n$ and s_t while leaving $\tilde{\mathbf{T}}_t^n$ and $\mathcal{M}_{t \rightarrow *}$ temporally frozen.

Following DeformationPyramid [17], we use a deformability regularization that encourages as-rigid-as-possible movement. Concretely, we define \mathcal{L}_{ela} which penalizes the deviation of the log singular values of \mathbf{M}_t from zero, *i.e.*,

$$\mathcal{L}_{\text{ela}} = \|\log \boldsymbol{\Sigma}_t\|_F^2. \quad (9)$$

Multi-level registration. For ease of optimization, the point registration process during pose decomposition is empirically achieved in a hierarchical manner. Specifically, MLP_r in Eq. (4) is a multi-level pyramid network [17]. We denote $\mathbf{x}_t^{(k)}$ as the transformed coordinate of the output of the k -th pyramid level at time t , where $k = 1, \dots, K$. Then $\mathbf{x}_t^{(k)}$ is obtained by a hierarchical deformation over the initial coordinate $\mathbf{x}_t^{(0)}$,

$$\mathbf{x}_t^{(k)} = \prod_{j=1}^k s_t^{(j)} \tilde{\mathbf{R}}_t^{(j)} \mathbf{x}_t^{(0)} + \tilde{\mathbf{T}}_t, \quad (10)$$

where $\tilde{\mathbf{T}}_t \equiv \mathbf{T}_t$ denotes the translation. $s_t^{(k)}$ and $\tilde{\mathbf{R}}_t^{(k)}$ are obtained by the k -th level output of MLP_r .

Let $\tilde{\mathbf{G}}_t^{(k)} = \begin{pmatrix} \prod_{j=1}^k s_t^{(j)} \tilde{\mathbf{R}}_t^{(j)} \mathbf{x}_t^{(0)} & \tilde{\mathbf{T}}_t \\ 0 & 1 \end{pmatrix}$, by substituting $\tilde{\mathbf{G}}_t^n$ in Eq. (8) with $\tilde{\mathbf{G}}_t^{n,(k)}$ which transforms from the point $\mathbf{x}_t^{n,(0)}$, we could formulate the k -th chamfer distance loss function, termed $\mathcal{L}_{\text{cd}}^{(k)}$.

We reformulate $\tilde{\mathbf{R}}_t$ by $\tilde{\mathbf{R}}_t = \prod_{k=1}^K \tilde{\mathbf{R}}_t^{(k)}$, and compute \mathbf{R}_t based on $\tilde{\mathbf{R}}_t$ following Eq. (7). In summary, Eq. (4) estimates multi-level $\tilde{\mathbf{R}}_t^{(k)}$ instead of directly estimating $\tilde{\mathbf{R}}_t$.

The total loss function \mathcal{L} is summarized as,

$$\mathcal{L} = \sum_{k=1}^K \mathcal{L}_{\text{cd}}^{(k)} + \mathcal{L}_{\text{ela}} + \mathcal{L}_{\text{cyc}}, \quad (11)$$

where the consistency loss \mathcal{L}_{cyc} is composed by a 2D part and a 3D part, which will be detailed in our appendix.

3.4. Discussions for Multi-object Scenarios

Monocular reconstruction methods usually assume a single target object in the scene [5] or utilize multiple video sequences that share the same individual [41]. Yet real scenarios are likely to contain multiple target objects, where object occlusions and individual differences are both common and challenging. We demonstrate that RPD could adapt to multi-object scenarios.

Object occlusions. We find our framework could handle object occlusions under certain circumstances. Given multiple target objects, we design an anti-occlusion silhouette reconstruction loss to deal with the object occlusion issue, as will be detailed in the appendix.

Individual differences. In complex scenarios with multiple non-rigid objects, different individuals of the same category could vary in height, volume, or scale. Instead of constructing several independent canonical models for different objects, we allow these individuals to share the canonical model and improve the data efficiency. **1)** This is partially realized by Sim(3) with the scaling factor s_t (incorporated in Sec. 3.2) which adjusts the size at local regions when performing the point matching in Eq. (8). **2)** When combined with the linear skinning weights, the ability to handle scale differences of our reconstruction method is further enhanced by shrinking or stretching the control points/bones. Qualitative results in Sec. 4.2 evaluate the effectiveness of handling individual differences in multi-object scenarios.

4. Experiments

Our experiments are conducted with monocular RGB videos from challenging DAVIS [4], OVIS [26], AMA [36], and Casual [41] datasets. We detail dataset settings and implementation details in Sec. 4.1. We compare the proposed RPD with state-of-the-art methods including Nerfies [23], ViSER [40], and BANMo [41] qualitatively in Sec. 4.2 and quantitatively in Sec. 4.3. In addition, we perform analytical experiments in Sec. 4.4 to verify the advantage of each component in RPD. More details of network architectures and visualizations are further provided in the appendix.



Figure 2. Reconstruction of humans trained with 15 video sequences (without needing SMPL [19] or pre-trained CSE [22]/DensePose [12]). Every two groups in the same row are collected from the same video sequence. Exhibited results cover complicated scenes containing multiple persons with differences in body size, large pose changes, fast movements like parkour, etc. All these examples share a common canonical space. A topdown novel view is provided for each group.

4.1. Experimental Settings

Datasets. DAVIS [4] and OVIS [26] are video datasets for object segmentation and they both provide dense 2D instance-level object annotations. Concretely, DAVIS contains 150 videos and 376 annotated objects. OVIS covers 25 categories and collects 901 video sequences with an average length of 12.77 seconds. In addition, OVIS contains more challenging videos that record non-rigid objects with complex motion patterns, rapid deformation, and heavy occlusion. AMA [36] dataset contains articulated videos with ground-truth meshes and could be utilized for computing quantitative results. To make a fair comparison with existing methods, ground-truth object silhouettes are utilized for training, and we adopt the video group named “swing”. Time synchronization and camera extrinsics are not used for training. The Casual dataset is collected by BANMo [41] and contains casually captured animals and humans.

Implementation details. We use the similar architecture with 8 layers for volume rendering as in NeRF [21], but deform observed objects into the canonical space. We initialize MLP_{SDF} as an approximate unit sphere [43] centered at \mathbf{o}_t (Eq. (4)). The dimension of each latent code embedding is chosen as $\omega_t \in \mathbb{R}^{64}$ in Eq. (2), $\mathbf{g}_t \in \mathbb{R}^{128}$ in Sec. 3.2, and $\varphi_t \in \mathbb{R}^{128}$ in Eq. (4). The time-variant latent code is represented by a Fourier embedding [21]. MLP_g in Sec. 3.2 is a 8-layer network that takes as input the Fourier embedded \mathbf{g}_t . In addition, MLP_r in Eq. (4) is a multi-level network

consisting of 9 pyramid levels [17] (*i.e.*, $K = 9$). Inspired by BANMo [41], we reduce the reconstruction complexity with linear skinning weights and parameterize $\mathcal{M}_{t \rightarrow *}$ and $\mathcal{M}_{* \rightarrow t}$ in Eq. (3) based on 25 control points. We follow ViSER [40] that learns approximate pixel-surface embeddings (which are not category-specific) to obtain reasonable initial poses based on the optical flow.

Optimization details. For DAVIS and OVIS datasets, we adopt the provided segmentation masks when calculating the silhouette reconstruction loss. These annotated masks are unfilled at occluded pixels, hence the discussed anti-occlusion strategy is still necessary especially given that both datasets contain heavy object occlusion. For the Casual dataset, differently, we predict segmentation masks by an off-the-shelf network PointRend [14]. For all datasets, we adopt VCN-robust [39] to compute the optical flow as required for optimization in Sec. 3.3. During training, we adopt OneCycle [29] as the learning rate scheduler with the initial value, maximum value, and final value being 2×10^{-5} , 5×10^{-4} , and 1×10^{-4} , respectively. For each experiment, we observe that the model could achieve high-fidelity performance by training 20k iterations, which takes approximately 9 hours on a single V100 GPU. To kick-start with a reasonable initial pose, we follow ViSER [40] that adopts optical flow for learning approximate pixel-surface embeddings (which are not category-specific). But due to our proposed designs, our performance significantly outperforms ViSER as will be compared in Sec. 4.2.



Figure 3. Qualitative comparison of our method (without needing pre-trained CSE [22]/DensePose [12]) vs BANMo [41] (with pre-trained CSE/DensePose). The experiment reconstructs three highly-occluded cats that walk in a circle. Segmentation masks are also illustrated. Red dashed circles highlight the regions where RPD outperforms BANMo. Best view in color and zoom in.

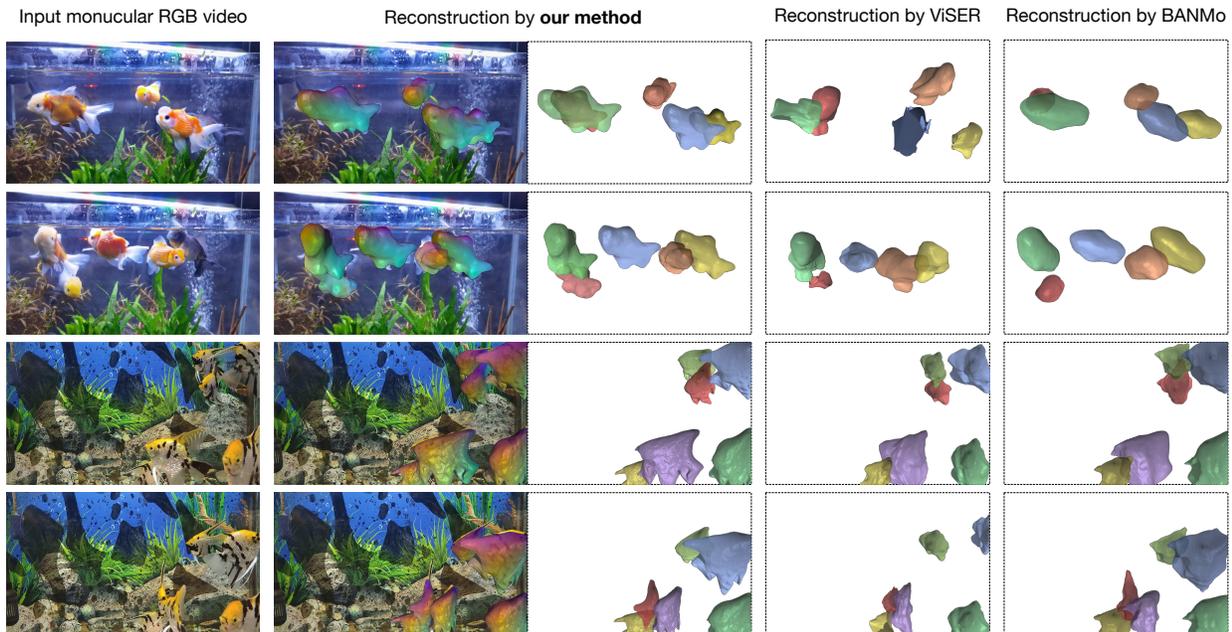


Figure 4. Qualitative comparison of our method with prior art ViSER [40] and BANMo [41]. We provide two groups of experiments that reconstruct fast-moving fish with complex motion patterns and self-occlusion. Each group is trained with **only 1 video sequence** with 11 seconds, 114 frames (top 2 rows) and 12 seconds, 120 frames (bottom 2 rows), respectively. Best view in color and zoom in.

4.2. Qualitative Results

To evaluate the effectiveness of our method, we choose hard video samples from the datasets to cover complex object deformations/movements, large pose changes, multi-object scenes with object occlusions, and multiple individuals of the same category yet with scale differences.

Qualitative results of humans. We first consider the non-rigid reconstruction of humans without using off-the-shelf templates (*e.g.*, SMPL [19]). Fig. 2 provides our illustration results of the human reconstruction, trained with 5 video sequences from the DAVIS dataset, and 10 video

sequences from the Casual dataset. All reconstructed individuals share a joint canonical space. Different from Nerfies [23] that assumes the root pose of an object is compensated by background-SfM, or BANMo [41] that relies on pre-trained human-specific pose priors (*e.g.*, DensePose), our RPD optimizes the per time object pose. From the first two rows of demonstrated results, we observe that RPD adapts to individual differences such as different heights, fat/skinny, and local shape scales. The reconstruction illustration in the following two rows indicates that the pose estimation by registration is relatively accurate even in scenes with large movements such as parkour and playing tennis.

Qualitative results of occluded cats. We then illustrate the reconstruction of cats in Fig. 3, which is jointly trained by 1 video sequence from the OVIS dataset and 3 video sequences from the Casual dataset. In addition, we compare our RPD with the state-of-the-art approach BANMo [41]. Differently, BANMo relies on pre-trained category-specific pose embeddings for the pose estimation, however, ours does not leverage such priors. We observe from the illustration that our reconstructed meshes are realistic and preserve geometric details, also showing fewer wrinkles on cat bodies. Besides, our superiority in articulated reconstruction is even more highlighted when considering object occlusions. As shown in the illustration, three cats are heavily occluded when walking around, while segmentation masks usually do not capture the occluded regions. Under the circumstance, BANMo fails to predict or reason the portions behind other objects, while our method presents accurate and realistic reconstruction in occluded regions.

Qualitative results of single-video fish. To demonstrate the advantage of our RPD on generic object categories such as invertebrates and fish, where off-the-shelf pre-trained pose networks are unavailable. In Fig. 4, we exhibit the reconstruction result of several fish, and each experiment is trained with only 1 video sequence (11 or 12 seconds) from the OVIS dataset. The reconstruction is performed in a multi-object scenario with more challenging dynamics, *e.g.*, multiple objects with heavy occlusion and complex motion patterns. We compare our reconstruction results with prior art including BANMo [41] and ViSER [40] by downloading their open-sourced codebases. Similar to our method, both baseline methods utilize 2D supervision of videos including segmentation masks and optical flows. We find that removing the pre-trained CSE embeddings in BANMo simply leads to catastrophic failures. As a result, for BANMo, we still preserve its original implementation of dense pose priors. For a fair comparison, we perform mesh upsampling and smoothing for ViSER to improve its resolution during rendering. It is observed that BANMo might degenerate into a smooth shape with scarce geometric details, which is probably due to the lack of pose priors of the proper object category. Our method clearly outperforms both baseline methods by the holistic articulated shape reconstruction and the local preservation of details, showing the ability to reconstruct generic object categories even under limited data.

4.3. Quantitive Results

We consider chamfer distance and F-score as 3D metrics to quantify the performance on AMA [36] and Hands [41]. Chamfer distance calculates the average distance between ground-truth and points at the reconstructed surface by finding the nearest neighbour matches. We report the F-score at distance thresholds 2% of the longest edge of the axis-aligned object bounding box [32]. Experimental results are

Method	Training Data	AMA-swing data		Hands data	
		CSE Pose	CD (↓) F@2% (↑)	Real Pose CD (↓) F@2% (↑)	Real Pose CD (↓) F@2% (↑)
Nerfies [23]	Single-video	✓	22.6 13.2	-	-
BANMo [41]	Single-video	✓	9.4 56.8	✓	10.3 39.5
RPD	Single-video	✓	9.0 57.4	-	-
BANMo	Single-video	✗	28.3 10.2	✗	25.1 14.4
RPD	Single-video	✗	11.4 52.7	✗	13.0 30.9
ViSER [40]	Multi-video	✓	15.7 52.2	✓	16.8 21.3
BANMo	Multi-video	✓	9.1 57.0	✓	7.5 49.6
RPD	Multi-video	✓	8.5 58.7	-	-
ViSER	Multi-video	✗	27.7 10.3	✗	27.3 10.1
BANMo	Multi-video	✗	24.2 12.5	✗	18.7 19.4
RPD	Multi-video	✗	10.1 53.9	✗	8.2 47.0

Table 2. Quantitive results on single/multi-video AMA-swing [36] and Hands [41] datasets. Evaluation metrics: Chamfer distance (CD) and F-score (F@2%) averaged over all frames.

Method	Training Data	CSE Pose	Dataset	mIoU (% ↑)
BANMo [41]	Multi-video	✓	DAVIS (Human)	72.7
BANMo	Multi-video	✗	DAVIS (Human)	35.1
RPD	Multi-video	✗	DAVIS (Human)	75.9
BANMo	Multi-video	✓	OVIS (Cats)	76.2
BANMo	Multi-video	✗	OVIS (Cats)	53.6
RPD	Multi-video	✗	OVIS (Cats)	87.7
BANMo	Single-video	✗	OVIS (Fish)	49.1
ViSER [40]	Single-video	✗	OVIS (Fish)	40.4
RPD	Single-video	✗	OVIS (Fish)	70.2

Table 3. Quantitative results of mIoU (%), computed with estimated 2D foreground masks and ground truth masks.

provided in Table 2. By comparison, our RPD achieves better quantitative performance in all settings, and shows significant improvement when the pre-trained category-specific CSE poses [22] and the ground truth poses are not available, *e.g.*, 53.9 vs 12.5 (AMA-swing data) and 47.0 vs 19.4 (Hands data) for the multi-video setting. Besides, based on the results of our RPD with/without CSE, we find that RPD also benefits from a proper pose prior for initialization.

Besides 3D metrics, we also report 2D mIoU (averaged over all frames) between the annotated segmentation masks and estimated masks. To avoid mistakenly computing mIoU at occlusion pixels, we only consider two classes (foreground and background) for all datasets. Although mIoU is not a standard metric for 3D scenes, it reflects the reconstruction fidelity to some extent, especially for multi-object scenarios where different individuals share the canonical space. Results in Table 3 indicate that our RPD demonstrates better performance in all three dataset settings. In particular, RPD outperforms both baselines by a large margin when CSE pose initialization for the given object category (*e.g.*, fish) is not available, such as 70.2 vs 49.1 when compared with BANMo. These results show the superiority of our method in capturing high-fidelity geometric details while not relying on category-specific dense pose priors.

4.4. Ablation Studies

Applying RPD separately to ViSER/BANMo: We further demonstrate that RPD can also serve as a standalone component to rectify root poses for existing methods, *e.g.*, ViSER [40] and BANMo [41]. Results are provided in Table 4. We observe that combining RPD with baseline meth-

Method	AMA-swing data			Hands data		
	CSE pose	CD (\downarrow)	F@2% (\uparrow)	Real pose	CD (\downarrow)	F@2% (\uparrow)
ViSER [40]	✓	15.7	52.2	✓	16.8	21.3
BANMo [41]	✓	9.1	57.0	✓	7.5	49.6
ViSER + RPD	✓	12.2 (-3.5)	53.6 (+1.4)	-	-	-
BANMo + RPD	✓	8.8 (-0.3)	58.2 (+1.2)	-	-	-
ViSER	✗	27.7	10.3	✗	27.3	10.1
BANMo	✗	24.2	12.5	✗	18.7	19.4
ViSER + RPD	✗	13.9 (-13.8)	51.5 (+41.2)	✗	19.0 (-8.3)	17.3 (+7.2)
BANMo + RPD	✗	11.0 (-14.2)	53.2 (+40.7)	✗	8.8 (-9.9)	46.1 (+26.7)

Table 4. Quantitive results of applying RPD to ViSER/BANMo, on multi-video AMA-swing [36] and Hands [41] datasets. Evaluation metrics follow Table 2.

Pose Error	Chamfer	F-score	Pose Error	Chamfer	F-score
	Distance (cm, \downarrow)	@2% (% , \uparrow)		Distance (cm, \downarrow)	@2% (% , \uparrow)
0°	9.6	55.9	45°	9.6	55.7
90°	10.3	54.7	135°	16.4	50.2

Table 5. Sensitivity to inaccurate initial poses on AMA-swing. Evaluation metrics: Chamfer distance (cm) and F-score (%).

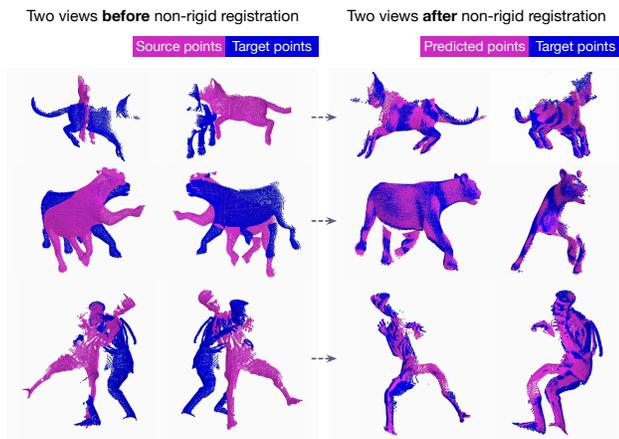


Figure 5. Evaluation of the multi-level pyramid structure for non-rigid point registration, performed on 4DMatch-F [16].

ods achieves huge gains when there are **no** CSE/real poses. For example, BANMo+RPD obtains 40.7/26.7 F-score improvements over BANMo itself when CSE is not available.

Sensitivity to pose initialization. In Sec. 4.1, we describe the scheme of root pose initialization. In Table 5, we inject different levels of Gaussian noise into initial poses, resulting in averaged rotation errors of 45°, 90°, or 135°. RPD appears to be relatively stable up to 90° pose error.

Non-rigid point registration. We show the effectiveness of multi-level pyramid structure and perform non-rigid point registration on 4DMatch-F [16], as shown in Fig. 5. Even the initial pose error is large (about 90°), the registration process still achieves a proper matching.

Reconstruction w.r.t. training steps. Fig. 6 illustrates meshes that are extracted from the canonical space during training, and compares the performance with/without RPD. The reconstruction with RPD converges quickly within only 5k training steps, while the reconstruction without RPD fails to handle the front-back human pose ambiguity.

Reconstruction w.r.t. pyramid levels. Fig. 7 compares the performance under various pyramid levels. It demonstrates that a larger capability of pose estimation network

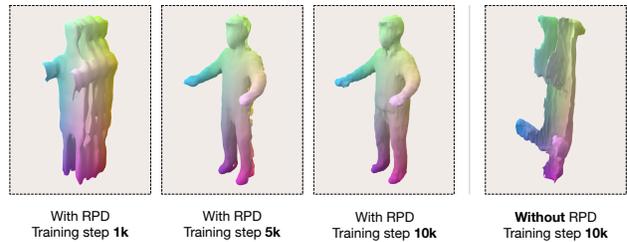


Figure 6. Meshes of the canonical space with/without RPD at different training steps.

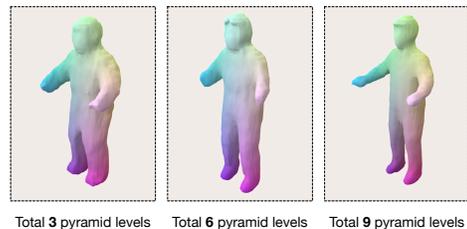


Figure 7. Meshes of the canonical space under different pyramid level settings. All meshes are extracted at the same training step.



Figure 8. (a) Different #control points. (b) Rendered color images.

(with the increase of levels) leads to a more elaborate 3D reconstruction, which evaluates the effectiveness of our RPD.

Reconstruction w.r.t. numbers of control points. The choice of the number of control points is ablated in Fig. 8(a), where using more control points achieves to a more elaborate result. We show rendered color images in Fig. 8(b).

5. Conclusion

We propose Root Pose Decomposition (RPD) to reconstruct generic non-rigid objects from monocular videos. RPD decomposes per-frame root pose by learning a dense neural field with local transformations, which are rectified through the point registration to a shared canonical space. RPD has few dependencies on category-specific templates, background-SfM, or pre-trained dense poses. Meanwhile, RPD demonstrates promising performance for objects with complex deformations/movements, and in multi-object scenarios containing occlusions and individual differences. Limitation: The method is limited when only a narrow range of poses is covered throughout input videos.

Acknowledgement

This work is funded by the Major Project of the New Generation of Artificial Intelligence (No. 2018AAA0102900) and the Sino-German Collaborative Research Project Crossmodal Learning (NSFC 62061136001/DFG TRR169). Y. Wang and Y. Dong are supported by the Shuimu Tsinghua Scholar Program.

References

- [1] Adrien Bartoli, Yan Gérard, François Chadebecq, Toby Collins, and Daniel Pizarro. Shape-from-template. *TPAMI*, 2015. 2
- [2] Sofien Bouaziz, Sebastian Martin, Tiantian Liu, Ladislav Kavan, and Mark Pauly. Projective dynamics: fusing constraint projections for fast simulation. *ACM Trans. Graph.*, 2014. 4
- [3] Christoph Bregler, Aaron Hertzmann, and Henning Biermann. Recovering non-rigid 3d shape from image streams. In *CVPR*, 2000. 2, 3
- [4] Sergi Caelles, Jordi Pont-Tuset, Federico Perazzi, Alberto Montes, Kevis-Kokitsi Maninis, and Luc Van Gool. The 2019 DAVIS challenge on VOS: unsupervised multi-object segmentation. *CoRR*, abs/1905.00737, 2019. 2, 5, 6
- [5] Hongrui Cai, Wanquan Feng, Xuetao Feng, Yan Wang, and Juyong Zhang. Neural surface reconstruction of dynamic scenes with monocular RGB-D camera. In *NeurIPS*, 2022. 1, 2, 3, 5
- [6] Isaac Chao, Ulrich Pinkall, Patrick Sanan, and Peter Schröder. A simple geometric model for elastic deformations. *ACM Trans. Graph.*, 2010. 4
- [7] Yue Chen, Xingyu Chen, Xuan Wang, Qi Zhang, Yu Guo, Ying Shan, and Fei Wang. Local-to-global registration for bundle-adjusting neural radiance fields. In *CVPR*, 2023. 3
- [8] Yuchao Dai, Hongdong Li, and Mingyi He. A simple prior-free method for non-rigid structure-from-motion factorization. In *CVPR*, 2012. 3
- [9] Bailin Deng, Yuxin Yao, Roberto M. Dyke, and Juyong Zhang. A survey of non-rigid 3d registration. *Comput. Graph. Forum*, 2022. 3
- [10] Mathias Gallardo, Daniel Pizarro, Toby Collins, and Adrien Bartoli. Shape-from-template with curves. *IJCV*, 2020. 2
- [11] Shubham Goel, Angjoo Kanazawa, and Jitendra Malik. Shape and viewpoint without keypoints. In *ECCV*, 2020. 3
- [12] Riza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *CVPR*, 2018. 1, 3, 6, 7
- [13] Navami Kairanda, Edith Tretschk, Mohamed Elgharib, Christian Theobalt, and Vladislav Golyanik. Phi-sft: Shape-from-template with a physics-based deformation model. In *CVPR*, 2022. 2
- [14] Alexander Kirillov, Yuxin Wu, Kaiming He, and Ross B. Girshick. Pointrend: Image segmentation as rendering. In *CVPR*, 2020. 6
- [15] Chen Kong and Simon Lucey. Deep non-rigid structure from motion. In *ICCV*, 2019. 2
- [16] Yang Li and Tatsuya Harada. Leopard: Learning partial point cloud matching in rigid and deformable scenes. In *CVPR*, 2022. 3, 9
- [17] Yang Li and Tatsuya Harada. Non-rigid point cloud registration with neural deformation pyramid. In *NeurIPS*, 2022. 3, 4, 5, 6
- [18] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *CVPR*, 2021. 2, 3
- [19] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: a skinned multi-person linear model. *ACM Trans. Graph.*, 2015. 1, 2, 6, 7
- [20] Ricardo Martin-Brualla, Noha Radwan, Mehdi S. M. Sajjadi, Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *CVPR*, 2021. 3
- [21] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 1, 2, 3, 4, 6
- [22] Natalia Neverova, David Novotný, Marc Szafraniec, Vasil Khalidov, Patrick Labatut, and Andrea Vedaldi. Continuous surface embeddings. In *NeurIPS*, 2020. 1, 2, 3, 6, 7, 8
- [23] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B. Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *ICCV*, 2021. 1, 2, 3, 4, 5, 7, 8
- [24] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T. Barron, Sofien Bouaziz, Dan B. Goldman, Ricardo Martin-Brualla, and Steven M. Seitz. Hypernerf: a higher-dimensional representation for topologically varying neural radiance fields. *ACM Trans. Graph.*, 2021. 2, 3
- [25] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3d hands, face, and body from a single image. In *CVPR*, 2019. 1
- [26] Jiyang Qi, Yan Gao, Yao Hu, Xinggang Wang, Xiaoyu Liu, Xiang Bai, Serge J. Belongie, Alan L. Yuille, Philip H. S. Torr, and Song Bai. Occluded video instance segmentation: A benchmark. *IJCV*, 2022. 2, 5, 6
- [27] Chris Russell, João Fayad, and Lourdes Agapito. Dense non-rigid structure from motion. In *3DIMPVT*, 2012. 2
- [28] Peter Sand and Seth J. Teller. Particle video: Long-range motion estimation using point trajectories. *IJCV*, 2008. 2
- [29] Leslie N Smith and Nicholay Topin. Super-convergence: Very fast training of neural networks using large learning rates. In *Artificial intelligence and machine learning for multi-domain operations applications*, volume 11006, pages 369–386. SPIE, 2019. 6
- [30] Robert W. Sumner, Johannes Schmid, and Mark Pauly. Embedded deformation for shape manipulation. *ACM Trans. Graph.*, 2007. 3
- [31] Narayanan Sundaram, Thomas Brox, and Kurt Keutzer. Dense point trajectories by gpu-accelerated large displacement optical flow. In *ECCV*, 2010. 2
- [32] Maxim Tatarchenko, Stephan R. Richter, René Ranftl, Zhuwen Li, Vladlen Koltun, and Thomas Brox. What do single-view 3d reconstruction networks learn? In *CVPR*. Computer Vision Foundation / IEEE, 2019. 8
- [33] Carlo Tomasi and Takeo Kanade. Shape and motion from image streams under orthography: a factorization method. *IJCV*, 1992. 3
- [34] Edith Tretschk, Navami Kairanda, Mallikarjun B. R., Rishabh Dabral, Adam Kortylewski, Bernhard Egger, Marc Habermann, Pascal Fua, Christian Theobalt, and Vladislav

- Golyanik. State of the art in dense monocular non-rigid 3d reconstruction. *CoRR*, abs/2210.15664, 2022. 1, 2
- [35] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 3
- [36] Daniel Vlasic, Ilya Baran, Wojciech Matusik, and Jovan Popovic. Articulated mesh animation from multi-view silhouettes. *ACM Trans. Graph.*, 2008. 2, 5, 6, 8, 9
- [37] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In *NeurIPS*, 2021. 3
- [38] Shangzhe Wu, Tomas Jakab, Christian Rupprecht, and Andrea Vedaldi. DOVE: learning deformable 3d objects by watching videos. *CoRR*, abs/2107.10844, 2021. 3
- [39] Gengshan Yang and Deva Ramanan. Volumetric correspondence networks for optical flow. In *NeurIPS*, 2019. 6
- [40] Gengshan Yang, Deqing Sun, Varun Jampani, Daniel Vlasic, Forrester Cole, Ce Liu, and Deva Ramanan. Viser: Video-specific surface embeddings for articulated 3d shape reconstruction. In *NeurIPS*, 2021. 2, 3, 5, 6, 7, 8, 9
- [41] Gengshan Yang, Minh Vo, Natalia Neverova, Deva Ramanan, Andrea Vedaldi, and Hanbyul Joo. Banmo: Building animatable 3d neural models from many casual videos. In *CVPR*, 2022. 1, 2, 3, 4, 5, 6, 7, 8, 9
- [42] Lingchen Yang, Zefeng Shi, Youyi Zheng, and Kun Zhou. Dynamic hair modeling from monocular videos using deep neural networks. *ACM Trans. Graph.*, 2019. 1
- [43] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Ronen Basri, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. In *NeurIPS*, 2020. 3, 6
- [44] Yufei Ye, Shubham Tulsiani, and Abhinav Gupta. Shelf-supervised mesh prediction in the wild. In *CVPR*, 2021. 3