

Deep Non-Rigid Structure-from-Motion Revisited: Canonicalization and Sequence Modeling

Hui Deng¹, Jiawei Shi¹, Zhen Qin², Yiran Zhong³, Yuchao Dai¹ *

¹ School of Electronics and Information, Northwestern Polytechnical University

² TapTap,

³ OpenNLPLab

{denghui986, sjw2018}@mail.nwpu.edu.cn, {zhenqin950102, zhongyiran}@gmail.com, daiyuchao@nwpu.edu.cn

Abstract

Non-Rigid Structure-from-Motion (NRSfM) is a classic 3D vision problem, where a 2D sequence is taken as input to estimate the corresponding 3D sequence. Recently, deep neural networks have greatly advanced the task of NRSfM. However, existing deep NRSfM methods still have limitations in handling the inherent sequence property and motion ambiguity associated with the NRSfM problem. In this paper, we revisit deep NRSfM from two perspectives to address the limitations of current deep NRSfM methods: (1) canonicalization and (2) sequence modeling. We propose an easy-to-implement per-sequence canonicalization method as opposed to the previous per-dataset canonicalization approaches. With this in mind, we propose a sequence modeling method that combines temporal information and subspace constraints. As a result, we have achieved a more optimal NRSfM reconstruction pipeline compared to previous efforts. The effectiveness of our method is verified by testing the sequence-to-sequence deep NRSfM pipeline with corresponding regularization modules on several commonly used datasets.

Introduction

Non-Rigid Structure-from-Motion (NRSfM) is one of the classical tasks in 3D computer vision, which aims at recovering the 3D deforming shape sequences from 2D observation sequences. There has been a lot of work trying to construct a reasonable model for this task. Various traditional mathematical methods (Akhter et al. 2008; Akhter, Sheikh, and Khan 2009; Dai, Li, and He 2014; Zhu et al. 2014) based on the factorization framework (Tomasi and Kanade 1992), represented by Bregler *et al.* (Bregler, Hertzmann, and Biermann 2000), have produced remarkable results in modeling the deformation sequences.

Along with the rise of machine learning techniques and the advantages shown by neural network models in terms of inference speed and computational accuracy, more people have started to combine deep learning with non-rigid 3D reconstruction (Novotny et al. 2019; Kong and Lucey 2021; Wang, Lin, and Lucey 2020; Zeng et al. 2022; Wang and Lucey 2021; Zeng et al. 2021; Park, Lee, and Kwak 2020), and some progress has been made. These methods have cer-

tain advantages over traditional methods in terms of computational accuracy, inference speed, and handling of shape ambiguity. However, they also encountered the same problem as traditional methods, namely motion ambiguity, where ambiguity interferes with the direction of convergence of the network, which in turn leads to poor estimation results. To cope with this problem, C3dpo (Novotny et al. 2019) proposes the concept of transversal property and designs the canonicalization network to suppress the ambiguity throughout the data set. On the other hand, PRN (Park, Lee, and Kwak 2020) represents a sequential regularization method based on Procrustean Analysis (PA) and an elegant training method that eliminates the need for the network to perform complex PA during inference. However, the implementation and computational complexity of this network in training are relatively high. At the same time, the backbone network of these methods (Novotny et al. 2019; Park, Lee, and Kwak 2020; Sidhu et al. 2020; Zeng et al. 2021) only performs lifting on 2D observations and lacks modeling of continuous 2D observation sequences.

Recently, Deng *et al.* (Deng et al. 2022) (denoted as Seq2Seq) designs sequence a reconstructing module by referring to the traditional mathematical modeling of sequences, and has achieved promising results. However, instead of tackling canonicalization in a sequential manner, this method follows (Novotny et al. 2019) and deals with the problem throughout the dataset. In addition, this method models the sequence with little consideration of the positional information in the sequence structure, and only an absolute positional encoding bias is used to inject temporal information into the sequence reconstruction process.

In this paper, we propose to revisit the deep NRSfM problem. We argue that dealing with shape ambiguity in a *sequence-by-sequence* manner is more effective than over the entire dataset, reducing the effect of ambiguity on the estimation results. We propose a sequential canonicalization based on General Procrustean Analysis (GPA). This method focuses on a parameter-free GPA layer, which avoids high training overhead and eliminates the need for complex GPA operations during inference. In conjunction with the canonicalization for sequences, we design a sequence modeling method that can utilize sequence temporal more effectively. We argue that using temporal bias alone to introduce temporal information for sequence reconstruction is insufficient,

*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

because there is no guarantee that the network will understand what the bias represents.

Our main contributions are summarized as follows:

- A sequence-by-sequence canonicalization method with an easy-to-implement GPA Layer, effectively reducing the impact of the inherent ambiguity of NRSfM on estimation accuracy;
- A method that can efficiently incorporate temporal information to reconstruct non-rigid sequences for better shape sequence estimation results;
- We conducted experiments on a substantial amount of data, and the results confirm the validity of our method in most scenarios.

Related Work

Classical NRSfM Methods

Bregler *et al.* (Bregler, Hertzmann, and Biermann 2000) first introduced factorization framework (Tomasi and Kanade 1992) for Non-rigid Structure-from-Motion. Xiao *et al.* (Xiao, Chai, and Kanade 2004) pointed out that the recovery of shape basis has inherent ambiguity problems. Akhter *et al.* (Akhter, Sheikh, and Khan 2009) showed that the non-rigid 3D shapes can be recovered uniquely by suitable optimization method under the orthogonal constraint. Subsequent researchers have proposed many effective NRSfM methods, including using DCT basis to recover motion trajectory (Akhter *et al.* 2008; Gotardo and Martinez 2011), using low-rank constraint to regularize 3D structure recovery (Paladini, Bartoli, and Agapito 2010; Dai, Li, and He 2014; Kumar 2020; Kumar and Van Gool 2022; Shi, Deng, and Dai 2024), using hierarchical prior (Torresani, Hertzmann, and Bregler 2008), using union-of-subspace constraint (Zhu *et al.* 2014; Kumar, Dai, and Li 2016, 2017), using procrustean alignment to refine camera motion (Lee *et al.* 2013; Lee, Choi, and Oh 2014; Park, Lee, and Kwak 2017), using consensus prior (Cha *et al.* 2019), using manifold to model non-rigid 3D surface (Kumar *et al.* 2018; Parashar, Pizarro, and Bartoli 2018; Parashar, Salzmann, and Fua 2020), and many other quintessential methods (Marques and Costeira 2008; Paladini *et al.* 2009; Garg, Roussos, and Agapito 2013a)

Deep NRSfM Methods

Novotny *et al.* (Novotny *et al.* 2019) proposed the first deep learning framework for NRSfM. They design a decomposition network to predict shape and motion separately, and use canonical loss to fix the estimated 3D shapes in the canonical coordinate. Kong *et al.* (Kong and Lucey 2021) builds a hierarchical network for non-rigid 3D reconstruction. Wang *et al.* (Wang, Lin, and Lucey 2020) extended this method to the perspective projection model. Park *et al.* (Park, Lee, and Kwak 2020) extended the Procrustean Regression (Park, Lee, and Kwak 2017) to the deep learning method, avoiding the complex camera motion estimation problem. Zeng *et al.* (Zeng *et al.* 2021) further strengthened the inter-frame structure connection by introducing a regularization constraint for the pairwise alignment of 3D shapes. Wang *et al.* (Wang and Lucey 2021) propose to use the orthographic-N-point

algorithm to regularize auto-encoder to obtain better reconstruction. Sidhu *et al.* (Sidhu *et al.* 2020) designs the trajectory energy function for imposing subspace constraints, accompanied by a regular term for temporal smoothing. However, the method does not combine subspace constraints with sequential temporal constraints. Recently, Zeng *et al.* (Zeng *et al.* 2022) focused on the depth ambiguity problem and modeled the generation of multi-hypothesis through the introduction of Gaussian noise.

Deep Sparse NRSfM Methods often use the network to regress the coordinates of 3D keypoints from 2D annotations, while dense NRSfM tasks often use more complex representation to represent 3D structures. Yang *et al.* (Yang *et al.* 2021a,b, 2022) proposed a framework that establishes the correspondence between long-sequence observation points by optical flow. In this paper, we focus on the reconstruction of sparse non-rigid 3D structures.

Sequence Modeling

Neural networks in sequence processing problems can improve processing efficiency and have been widely used in many sequence-to-sequence tasks (Hinton *et al.* 2012; Krizhevsky, Sutskever, and Hinton 2017; He *et al.* 2017). Recently, the proposal of Transformer (Vaswani *et al.* 2017) has greatly promoted the development of the sequence modeling network. Transformers have been extended to many fields, such as human pose estimation (Kocabas, Athanasiou, and Black 2020; Zhang *et al.* 2022), action recognition (Choi *et al.* 2019), NRSfM (Deng *et al.* 2022), *etc.*

Some methods do not adopt the Transformer structure, but try to deal with sequence problems from other perspectives. Wu *et al.* (Wu *et al.* 2022) used the time series analysis method to calculate the spectrum information of the sequence. Qin *et al.* (Qin *et al.* 2023) proposed to use the relative position encoded Toeplitz matrix to replace the attention to realize the function of the token mixer. In deep non-rigid 3D reconstruction tasks, some methods strengthen the network’s understanding of sequences. Park *et al.* (Park, Lee, and Kwak 2020) used the Procrustean alignment to design an optimization equation and extended it to a loss function. Zeng *et al.* (Zeng *et al.* 2021) constructed the internal constraints of the 3D structures through pairwise contrastive and consistent regularization.

Preliminary

Method

We first introduce the definition of this problem and review some classical solutions to provide possible thinking paths. The following equation gives the general form of NRSfM:

$$\begin{aligned} \mathbf{W} &= \begin{bmatrix} \mathbf{\Pi}_1 & & \\ & \ddots & \\ & & \mathbf{\Pi}_F \end{bmatrix} \begin{bmatrix} \mathbf{R}_1 & & \\ & \ddots & \\ & & \mathbf{R}_F \end{bmatrix} \begin{bmatrix} \mathbf{S}_1 \\ \vdots \\ \mathbf{S}_F \end{bmatrix} \\ &= \mathbf{\Pi} \mathbf{R} \mathbf{S}, \mathbf{W} \in \mathbb{R}^{2F \times P}, \mathbf{R} \in \mathbb{R}^{3F \times 3F}, \mathbf{S} \in \mathbb{R}^{3F \times P}, \end{aligned} \quad (1)$$

where $\mathbf{\Pi}_i = [1 \ 0 \ 0; 0 \ 1 \ 0]$ is the orthogonal projection matrix, \mathbf{W} is the matrix composed of 2D observations of continuously deforming objects, \mathbf{S} is the 3D shape matrix corresponding to the 2D observations, and \mathbf{R} is the motion

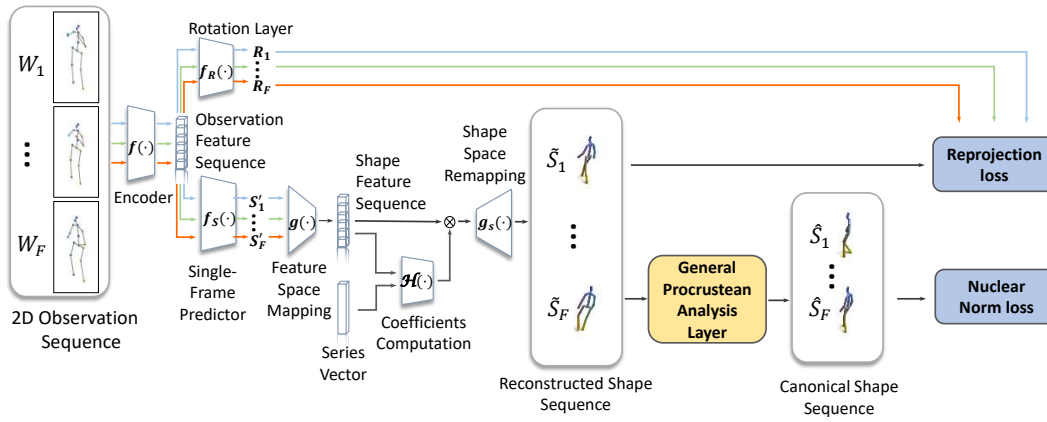


Figure 1: An overview of deep NRSfM pipeline with proposed shape sequence reconstruction and GPA layer. The whole pipeline consists of three parts: the single-frame shape/rotation predictor, the shape sequence reconstruction stage and general procrustean analysis layer. The General Procrustean Analysis layer is a parameter-free layer, which is not needed for inference when training is finished.

matrix. It should be noted that in the above modeling, no distinction is made between the ego-motion of the camera and the motion of the object:

$$\mathbf{W} = \mathbf{IIR}_{motion} \mathbf{S} = \mathbf{IIR}_{camera} \mathbf{R}_{shape} \mathbf{S}. \quad (2)$$

The underdetermination defect caused by this ambiguity of coupled motion has resulted in the loss of reconstruction accuracy in traditional methods and deep methods, and how to effectively deal with this problem has become one of the focuses of discussion in various works.

The solution model defined by Eq. (1) is usually transformed into an optimization model with additional constraints, and one of the most representative work (Dai, Li, and He 2014) gives the optimization condition by analyzing the rank of the sequence matrix to suppress ambiguity:

$$\min_{\mathbf{S}} \|\mathbf{S}^\sharp\|_*, s.t. \mathbf{W} = \mathbf{IIRS}, \mathbf{S} \in \mathbb{R}^{3F \times P}, \mathbf{S}^\sharp \in \mathbb{R}^{F \times 3P}, \quad (3)$$

where \mathbf{S}^\sharp is a reshuffled form of \mathbf{S} as described in Dai *et al.* (Dai, Li, and He 2014). The most important feature of this model is that no additional priori assumptions are introduced for the solution, but rather a low-rank constraint. Later, Zhu *et al.* (Zhu *et al.* 2014) noted on the basis of Dai *et al.* (Dai, Li, and He 2014) that this low-rank constraint on deforming sequences can be extended to a union-of-subspace representation. This work assumes that deforming shapes in the same sequence should come from one or more intersecting shape spaces, each of which is spanned by a set of shape basis:

$$\min_{\mathbf{D}} \|\mathbf{D}\|_*, s.t. \mathbf{W} = \mathbf{IIRS}, \mathbf{S} = \mathbf{DC}, \quad (4)$$

where \mathbf{D} is a dictionary matrix containing multiple bases that can be tensed into different subspace shapes, and \mathbf{C} is a coefficient matrix. This sequence modeling method referred to LRR (Liu *et al.* 2012) brings new effective constraints to the solution process, refining the regularity of the low-rank constraint on the estimation results.

However, these modeling of sequences are rarely discussed in subsequent deep methods, and the mainstream deep methods (Novotny *et al.* 2019; Kong and Lucey 2021; Wang and Lucey 2021) reconstruct the non-rigid shape sequences in a *single-frame lifting* way as follows:

$$\mathbf{S}'_i = f_S(f(\mathbf{W}_i), \Theta_S), \mathbf{S}'_i \in \mathbb{R}^{3 \times P}, \quad (5)$$

$$\mathbf{R}_i = f_R(f(\mathbf{W}_i), \Theta_R), \mathbf{R}_i \in \mathbb{R}^{3 \times 3}, \quad (6)$$

where Θ is learnable parameters, and $f: \mathbb{R}^{3 \times P} \rightarrow \mathbb{R}^D$ is an encoder module. Whilst such methods heavily rely on a network's ability to recall the relevant 3D structure from a sole frame of data, it is evidently apparent that they are incapable of utilizing constraints from *sequence context, temporal information, and the sequence structure itself* which can be provided to facilitate the reconstruction process.

In this paper, we revisit the canonicalization and the sequence structure self-expressive regularity under sequence modeling, and design modules that can effectively impose both constraints to obtain a deeply non-rigid reconstruction pipeline that can effectively exploit the sequence structure information. The entire pipeline is shown in Fig. 1, and we will then present first how to perform canonicalization in Sec. , followed by sequence modeling in Sec. .

General Procrustean Analysis Layer

To reduce the ambiguity described as Eq. (2) for reduce the loss of accuracy on estimating shapes, we design a novel module $\mathcal{G}: \mathbb{R}^{F \times 3P} \rightarrow \mathbb{R}^{F \times 3P}$ that constrains the sequence estimated by the network to the canonical coordinate system originally proposed by C3DPO (Novotny *et al.* 2019).

Definition 1 A deforming shape set $\mathcal{S} = \{\mathbf{S}_1, \dots, \mathbf{S}_n\}$ is in canonical coordinate, $\forall \mathbf{S}, \mathbf{S}' \in \mathcal{S}$ and $\mathbf{R} \in SO(3)$ $\mathbf{S} = \mathbf{RS}'$ i.f.f. $\mathbf{S} = \mathbf{S}', \mathbf{R} = \mathbf{I}_3$.

Unlike C3dpo, we define the input sequence rather than the entire dataset as a deforming shape set \mathcal{S} .

Thus, the previously coupled motion decomposition problem is treated as solving only one motion \mathbf{R}_{motion} . It

should be noted that the motion \mathbf{R}_{motion} here is actually no longer the camera motion but a combination of the origin camera motion and shape rigid motion, $\mathbf{R}_{motion} = \mathbf{R}_{camera}\mathbf{R}_{shape}$, since NRSfM focuses more on the recovery of the nonrigid shapes.

In contrast to another Deep NRSfM classical work C3dpo(Novotny et al. 2019) Seq2Seq (Deng et al. 2022) differs in that it does not need random sampling in the rotation space $SO(3)$ for each frame of 3D shape and training the network to be invariant to arbitrary rotational transformations. The context layer in Seq2Seq grants greater focus on estimating the sequence structure. The specific principle is shown in Fig. 2.

Inspired by PND (Lee et al. 2013), PR (Park, Lee, and Kwak 2017) and PRN (Park, Lee, and Kwak 2020), we design a parameter-free GPA Layer, which uses Generalized Procrustean Analysis (Gower 1975) to model the canonical coordinate and employs an iterative approach to obtain aligned shapes. The goal of the GPA layer is to eliminate rigid motion between any two frames as much as possible in the least squares sense. The GPA layer takes 3D sequence $\tilde{\mathbf{S}}$ as input and outputs the aligned 3D shapes $\hat{\mathbf{S}}$, which is defined as follows:

$$\hat{\mathbf{S}} = \hat{\mathbf{R}}\tilde{\mathbf{S}}, \text{ and } \left\{ \hat{\mathbf{R}}_i \right\}_{i=1}^F = \arg \min_{\mathbf{R}_i} \sum_{i=1}^F \left\| \mathbf{R}_i \tilde{\mathbf{S}}_i - \bar{\mathbf{S}} \right\|_F^2, \quad (7)$$

where $\bar{\mathbf{S}}$ is the mean shape of aligned shapes, which is updated during the iteration, and $\hat{\mathbf{S}}$ is the reconstructed 3D shape which will be introduced in Sec. Sequence Reconstruction. The aligned shapes $\hat{\mathbf{S}}$ are under a fixed coordinate and can naturally be added with some regularization, such as low-rank *etc.*

PRN (Park, Lee, and Kwak 2020) gives an elegant derivation of how to introduce constraints using GPA in gradient backpropagation, but is accompanied by a high computational overhead. Here we have an approach that is simpler in terms of implementation, but has a similar level of effectiveness. We refer to the iterative gradient backpropagation procedure of (Geng et al. 2021) and use the one-step gradient strategy to approximate the gradient of the GPA layer during the iteration.

Next, we show how the gradient of this computational procedure should be propagated. Suppose that the regularization term constraint acting on $\hat{\mathbf{S}}$ is $\mathcal{G}(\cdot)$, then the gradient can be calculated according to the chain rule similar to (Park, Lee, and Kwak 2020):

$$\frac{\partial \mathcal{G}}{\partial \tilde{\mathbf{S}}} = \left\langle \frac{\partial \mathcal{G}}{\partial \hat{\mathbf{S}}}, \frac{\partial \hat{\mathbf{S}}}{\partial \tilde{\mathbf{S}}} \right\rangle, \quad (8)$$

where $\langle \cdot, \cdot \rangle$ represents the inner product. Once \mathcal{G} is defined, $\frac{\partial \mathcal{G}}{\partial \tilde{\mathbf{S}}}$ can be directly derived. Since computing $\mathbf{S} \rightarrow \hat{\mathbf{S}}$ is an iterative process with an unknown number of iterations, we represent $\frac{\partial \hat{\mathbf{S}}}{\partial \tilde{\mathbf{S}}}$ into iterative form according to the chain rule:

$$\frac{\partial \hat{\mathbf{S}}}{\partial \tilde{\mathbf{S}}} = \left\langle \frac{\partial \hat{\mathbf{S}}}{\partial \mathbf{S}^k}, \left\langle \frac{\partial \mathbf{S}^{k-1}}{\partial \mathbf{S}^{k-2}}, \left\langle \dots, \left\langle \frac{\partial \mathbf{S}^2}{\partial \mathbf{S}^1}, \frac{\partial \mathbf{S}^1}{\partial \mathbf{S}} \right\rangle \dots \right\rangle \right\rangle \quad (9)$$

where \mathbf{S}^k is the intermediate state in the iterative process. This back-propagation process allows the shape sequence $\tilde{\mathbf{S}}$ estimated by our backbone network to be in the canonical coordinate as much as possible.

Sequence Reconstruction

In the previous section, we introduce the GPA layer, which aims to reduce motion ambiguity by aligning the shape sequences to the canonical coordinate, and in this section, we describe what properties should exist for the sequence. Inspired by Eq. (4), sequence can be represented as follows:

$$\mathbf{S} = \mathbf{C}\mathbf{S}, \mathbf{S} = [\mathbf{S}'_1, \dots, \mathbf{S}'_F]^\top, \mathbf{S}'_i \in \mathbb{R}^{3P \times 1}. \quad (10)$$

This indicates that any shape can be derived through a linear combination of other shapes. In this form, it is found that the identity matrix \mathbf{I} is a trivial case for coefficient matrix \mathbf{C} , *i.e.* shapes in the sequence are not related to each other. However, the study conducted by (Zhu et al. 2014) indicates that the shapes with close distances in the deformation sequence are present in the same subspace. This guarantees the absence of trivial cases, meaning that the deformation shape sequence obtains the **self-expressive** property.

To guarantee that the estimated shape sequence has self-expressive properties, we could rebuild the sequence by applying a coefficient matrix connected to it:

$$\mathbf{S} = \mathbf{C}\mathbf{S}', \mathbf{C} = \mathcal{H}(\mathbf{S}'), \quad (11)$$

$$s.t. \mathbf{S}' = [\mathbf{S}'_1, \dots, \mathbf{S}'_F]^\top, \mathbf{S}'_i \in \mathbb{R}^{3P \times 1}. \quad (12)$$

Where the mapping $\mathcal{H}(\cdot) : \mathbb{R}^{F \times 3P} \rightarrow \mathbb{R}^{F \times F}$ is referenced from the (Qin et al. 2023). This mapping is designed to encode shape temporal position and similarity in the coefficient matrix at the same time. To make the network more flexible, we here first map the shapes to the feature space before imposing this regularity:

$$\tilde{\mathbf{S}} = g_s(\mathcal{H}(g(\mathbf{S}'), \mathbf{L}, \Theta)g(\mathbf{S}')), \quad (13)$$

where \mathbf{S}' is the deforming sequence output from the predictor $f_S(\cdot, \Theta_S)$ in Eq. (5), Θ is the learnable parameter, \mathbf{L} is a series vector representing the sequence order, *i.e.* $\mathbf{L} = [1, 2, \dots, F]^\top$. To make the network more flexible, we impose this constraint in the feature space. $g : \mathbb{R}^{F \times 3P} \rightarrow \mathbb{R}^{F \times D}$ maps the shape to the feature space, and $g_s : \mathbb{R}^{F \times D} \rightarrow \mathbb{R}^{F \times 3P}$ remaps the deformation sequence from the feature space back to the shape space.

From the above self-expressive properties of the deformation sequence, it is clear that the shapes in the sequence are linearly correlated with each other. Therefore, the deformation sequence matrix \mathbf{S} should be in canonical coordinate described by Definition 1 and rank defect. Therefore, we perform the canonicalization with GPA Layer and the nuclear norm loss on the shape sequence, $\|\hat{\mathbf{S}}\|_*$. The nuclear norm alone cannot point to the correct learning direction, so we use the re-projection error as a data term, and finally obtain the total loss function that oversees the entire pipeline training process as follows:

$$\mathcal{L} = \alpha \|\mathbf{W} - \mathbf{I}\mathbf{R}\tilde{\mathbf{S}}\|_2 + \beta \|\hat{\mathbf{S}}\|_*, \quad (14)$$

where α and β are hyperparameters used to the balance of loss functions. After analysing the results of our experiments, we hereby choose the loss weight as $\alpha = 9, \beta = 0.1$.




Figure 2: The left side of the figure shows the canonicalization method in (Novotny et al. 2019), which performs random rotations for each frame, *i.e.*, training the network over the *entire dataset* to determine the canonical coordinate. The right side shows our method to align for *each sequence*, explicitly helping the network to determine the canonical coordinate.

Implementation Details

The network has three parts. The first part is a single-frame predictor that can obtain an intermediate sequence (predictor structure is the same as C3dpo(Novotny et al. 2019)). In the second part, the intermediate sequence obtained is mapped to the feature space, and then regularized by the context layer to ensure that the sequence can have the self-expressive property, and then mapped back to the shape space. Finally, the output is regularized by the GPA layer for ambiguity reduction, and the reprojection error and nuclear norm are calculated using the final output for training supervision.

The module g is a one-layer linear layer, the input dimension is $3 \times P$, the output dimension is 128. g_s contains a gated linear unit(Shazeer 2020) and a four-layer MLP. \mathcal{H} contains a gated Toeplitz unit (Gtu) as (Qin et al. 2023). In addition, the GPA layer does not contain any learnable parameters, so it can be left out of the computation during testing. To clearly present the structure of our network module, we have provided more detailed information in our **supplementary material**. The experimental settings are, GPA convergence threshold is $1e^{-8}$, and the maximum number of iterations is 100. The weight of the loss function is $\alpha = 9, \beta = 0.1$, the sequence length is 32, the batch is 256, and 4 Nvidia RTX 3080Ti are used for training.

Experiments

We commence this section with a thorough depiction of the dataset employed as well as the metrics. Following this, the quantitative outcomes will be presented, and ultimately, the ablation experiments will be conducted, accompanied by a performance discussion.

Dataset

Human3.6M. This classic dataset contains a large number of human motion sequences annotated with 3D ground truth which are extracted by motion capture systems (Ionescu et al. 2014). Following the setup of C3dpo(Novotny et al. 2019), two variants of the dataset are used: 1. Input 2D keypoints by the projection of ground truth 3D keypoints for training and testing (marked as *GT-H36M*); 2. Input 2D

keypoints detected by HRNet (Sun et al. 2019) (marked as *HR-H36M*). We follow the protocol of (Kudo et al. 2018) and evaluate the mean per-joint position error(MPJPE) and Stress(Novotny et al. 2019) over 17 joints.

InterHand2.6M. This dataset contains a large number of highly deformed hand poses (Moon et al. 2020). The hand pose sequence has different deformation characteristics from the human pose sequence, and in order to verify the generality of the NRSfM method, this dataset is introduced for experimental comparison. We evaluate the mean per-joint position error (MPJPE) and Stress over 21 joints.

3DPW. 3DPW(von Marcard et al. 2018) is a dataset commonly used for *human pose estimation*, which is less commonly used in NRSfM, and less frequently mentioned in previous work, so we follow the evaluation protocol of (Kudo et al. 2018). We also report the mean per-joint position error and Stress over 24 joints.

CMU MOCAP. For a fair comparison with existing methods, we follow the setup of (Zeng et al. 2021) to preprocess the CMU MOCAP dataset¹. We follow the splitting protocol of (Kong and Lucey 2021), where 1/5 action sequences are used for testing, and the rest are used for training. For evaluation, we use the same criterion as (Zeng et al. 2021) and report the reconstruction error e_{3D} on shapes.

Short Sequence. We choose some classical short sequence datasets used in the traditional methods, the Kinect dataset(Varol et al. 2012) has 1503 points per frame and in total 191 frames; the Rug dataset(Garg, Roussos, and Agapito 2013b) has 3912 points per frame and in total 159 frames. These datasets have a very small number of frames but have much denser 2D observations.

Evaluation Metric. Two major evaluation metrics were mentioned, according to the object we want to compare, we follow C3dpo(Novotny et al. 2019) and PRRRN(Zeng et al. 2021) to choose different metrics on different data, they are: $MPJPE(\mathbf{S}_i, \mathbf{S}_i^*) = \frac{1}{P} \sum_{j=1}^P \|\mathbf{S}_{ij} - \mathbf{S}_{ij}^*\|_2$ and $e_{3D}(\mathbf{S}_i, \mathbf{S}_i^*) = \frac{\|\mathbf{S}_i - \mathbf{S}_i^*\|_2}{\|\mathbf{S}_i^*\|_2}$ where the \mathbf{S}_{ij}^* denote the position of j -th key point in the 3D shape of the ground truth of i

¹CMU Dataset. available at <http://mocap.cs.cmu.edu/>

Methods	GT-H36M		HR-H36M		I26M		3DPW	
	MPJPE	Stress	MPJPE	Stress	MPJPE	Stress	MPJPE	Stress
PRN(Park, Lee, and Kwak 2020)	86.4	-	-	-	-	-	-	-
PAUL(Wang and Lucey 2021)	88.3	-	-	-	-	-	-	-
ITES(Xu et al. 2021)	77.2	-	-	-	-	-	-	-
PoseDict(Xu et al. 2021)	85.7	-	-	-	-	-	-	-
C3dpo(Novotny et al. 2019)	95.6	41.5	110.8	56.3	9.8	6.2	77.5	35.4
DNRSfM(Kong and Lucey 2021)	109.9	35.9	121.4	72.4	13.8	8.5	184.5	288.1
Seq2Seq(Deng et al. 2022)	79.8	33.8	98.5	49.6	8.9	6.1	109.7	49.2
MHR(Zeng et al. 2022)	72.1	36.4	93.4	49.9	29.1	12.2	79.6	34.7
Ours	66.1	25.9	84.5	45.0	8.6	5.7	74.4	31.8

Table 1: Experimental results on the Human3.6M datasets with ground truth 2D keypoint(Marked as GT-H36M) and HRNet detected 2D keypoints(Marked as HR-H36M), InterHand2.6M(Marked as I26M) and 3DPW dataset. We report the mean per joint position error (MPJPE) over the set of test actions. Our method achieves state-of-the-art performance on these datasets.

	Methods	S07	S20	S23	S33	S34	S38	S39	S43	S93
		All	CSF(Gotardo and Martinez 2011)	1.231	1.164	1.238	1.156	1.165	1.188	1.172
	URN(Cha, Lee, and Oh 2019)	1.504	1.770	1.329	1.205	1.305	1.303	1.550	1.434	1.601
	CNS(Cha et al. 2019)	0.310	0.217	0.184	0.177	0.249	0.223	0.312	0.266	0.245
	C3DPO(Novotny et al. 2019)	0.226	0.235	0.342	0.357	0.354	0.391	0.189	0.351	0.246
	Seq2Seq(Deng et al. 2022)	<u>0.072</u>	0.122	0.137	0.158	0.142	0.093	<u>0.090</u>	0.108	<u>0.129</u>
	MHR(Zeng et al. 2022)	0.384	0.461	0.495	0.474	0.479	0.434	0.430	0.495	0.478
	Ours	0.056	<u>0.146</u>	<u>0.168</u>	<u>0.169</u>	<u>0.201</u>	<u>0.098</u>	0.082	<u>0.116</u>	0.122
Unseen	DNRSFM	0.097	0.219	0.264	0.219	0.209	0.137	0.127	0.223	0.164
	PR-RRN	0.061	<u>0.167</u>	0.249	0.254	0.265	0.108	0.028	0.080	0.242
	C3DPO	0.286	0.361	0.413	0.421	0.401	0.263	0.330	0.491	0.325
	Seq2Seq	0.081	0.139	0.196	0.191	0.195	<u>0.097</u>	0.089	0.139	<u>0.151</u>
	MHR	0.398	0.467	0.511	0.474	0.471	0.435	0.428	0.506	0.476
	Ours	<u>0.069</u>	0.188	<u>0.228</u>	<u>0.251</u>	<u>0.213</u>	0.096	<u>0.084</u>	<u>0.121</u>	0.120

Table 2: Results on the long sequences of the CMU motion capture dataset. We follow the comparison in (Zeng et al. 2021). Our result achieve the comparable result on this dataset with the state-of-the art methods. We additionally discovered that on certain specific datasets, the benefits displayed by our created modules were not apparent. Thus, we conducted further experiments and analysed them in the Discussion section in Appendix.

-th frame and also we follow C3dpo (Novotny et al. 2019) to report $\text{Stress}(\mathbf{S}_i, \mathbf{S}_i^*) = \sum_{j < k} \frac{\|\|\mathbf{S}_{ij} - \mathbf{S}_{ik}\|_2 - \|\mathbf{S}_{ij}^* - \mathbf{S}_{ik}^*\|_2\|}{P(P-1)}$.

Quantitative Result

Table 1 reports results on Human3.6M with ground truth keypoint (**GT-H36M**) and results on Human3.6M with detected keypoints (**HR-H36M**). We compare our result with several deep non-rigid reconstruction methods including PRN(Park, Lee, and Kwak 2020), PAUL(Wang and Lucey 2021), ITES(Xu et al. 2021), PoseDict(Xu and Dunn 2021), C3DPO(Novotny et al. 2019), DNRSfM(Kong and Lucey 2021), MHR(Zeng et al. 2022), Pre(Deng et al. 2022), where we have directly quoted some experimental results from MHR. The result validates the effectiveness of our method on highly flexible shape sequences. The visualization of the result is shown in Fig. 3. We also test this method on Inter-Hand 2.6M as Table 1, the results validate the effectiveness of our method on dataset with category other than human.

Our method is also compared with several methods on the CMU Mocap dataset in Table 2. Since we use the

same preprocessing method and testing strategy as PR-RRN(Zeng et al. 2021), we directly cite the experiment results of several methods including CSF(Gotardo and Martinez 2011), URN(Cha, Lee, and Oh 2019), CNS(Cha et al. 2019) from PR-RRN, and test C3DPO(Novotny et al. 2019), MHR(Zeng et al. 2022), DNRSFM(Kong and Lucey 2021), Seq2Seq(Deng et al. 2022) which achieves state-of-the-art (SOTA) reconstruction accuracy on CMU MOCAP, and our approach achieves comparable precision to the SOTA.

We also test method on the classical nonrigid dataset 3DPW. The results are reported in Table 1, showing the advantages of our method. Also, we test method on the Kinect and Rug dataset for dense reconstruction. The result is partially visualized in Fig. 4, and quantitative results are in Appendix. Although the size of these dataset are small that it is not suited for deep methods, our method still guarantees results and is more stable compared to (Deng et al. 2022).

We additionally discovered that on certain specific datasets and we conducted further experiments and analyzed them in the Discussion section in Appendix.

Methods	Settings							Result	
	Series	GPA	Context	Nuclear	Mean	Canonical	PRN	MPJPE	Stress
all-one Series Vector		✓	✓	✓				90.4	36.5
w/o GPA Layer			✓	✓				121.1	61.5
w/o Context Layer	✓	✓		✓				166.7	76.8
w/o Nuclear Norm	✓	✓	✓					111.2	45.8
Ours Canonical	✓		✓	✓		✓		79.8	32.6
Ours Mean	✓		✓	✓	✓			83.4	33.2
Ours PRN	✓		✓	✓			✓	118.7	50.8
Ours	✓	✓	✓	✓				66.1	25.9

Table 3: Ablation studies performed on the Human3.6M dataset, **all-one Series Vector** representation does not take into account the temporal order at all when calculating the coefficient matrix, and its results are poor. On the other hand, the GPA Layer, Context Layer, and low rank constraint, loss of any of the constraints affect the results. **Ours Canonical** delegates do not use GPA Layer, but use **canonical loss** from C3dpo(Novotny et al. 2019), which can be found to not work well as a replacement. In addition, if we do not use GPA Layer but design a mean shape loss to replace this layer, the result as shown in **Ours Mean** and this setting does not achieve the same performance as GPA Layer. **Ours PRN** replace GPA Layer with **PRN gradient**(Park, Lee, and Kwak 2020).The proposed modules should be used together for best outcomes.




Figure 3: The visualization on different methods and ablation on the Human3.6M(first row) and InterHand2.6M(second row), more visualization are included in supplementary.

Ablation Study

To verify the effectiveness of our regularization module and the related regularization loss, the ablation result is reported in Table 3. The results (**Ours w/o GPA Layer**, **Ours w/o Context Layer**) show that the self-expressive constraint without the alignment constraint fails to work, and vice versa. We compare the results of using the all-one Series Vector (**Our all-one Series Vector**, setting the Series Vector $\mathbf{L} = 1, \dots, 1$) and the normal vector, the results show




Figure 4: This dense reconstruction visualization demonstrates that our method still works on dense settings.

that the correct temporal series vector in context layer leads to better final shape. To verify the effectiveness of the GPA layer, we replace it with the GPA loss function. The results (**Ours w/o Nuclear Norm**) prove that the rank constraint for the shape sequence is effective and it can provide a positive effect for the network to find the correct learning direction. We test the Canonical loss of C3dpo (Marked as **Ours Canonical**) and the mean shape loss designed with reference to GPA respectively (**Ours Mean**) as follows:

$$\mathcal{L}_{mean} = \sum_{i=1}^F \left\| \frac{\text{trace}(\hat{\mathbf{R}}_i) - 1}{2} \right\|_1, \quad (15)$$

$$\hat{\mathbf{R}}_i = \arg \min_{\mathbf{R}_i} \|\mathbf{R}_i \mathbf{S}_i - \bar{\mathbf{S}}\|_2, \quad \bar{\mathbf{S}} = \frac{1}{F} \sum_{i=1}^F \mathbf{S}_i, \quad (16)$$

where $\hat{\mathbf{R}}_i$ is obtained by the SVD-based least square algorithm. To compare the difference between PRN and GPA Layer, we conducted **Ours PRN** by removing the GPA Layer and performing canonicalization as PRN. The results show that the GPA Layer performs better. Additionally, during the experiments, it was discovered that PRN has a more intricate implementation in the computation of the gradient compared to the GPA layer.

In addition, we investigate where to supervise the reconstructed sequences, the details of which we put in the appendix due to space constraints.

Conclusion

In this paper, we revisited deep NRSfM from the perspective of canonicalization and sequence modeling. Based on our analysis, we designed two modules that can be utilized for deep NRSfM which adhere to the canonical rule. To verify the efficacy of the modules, several representative datasets are used to run the reconstruction pipeline. Experimental results confirm the validity of our approach. The findings also highlight a limitation of our method on smaller datasets. We will address this issue in our future research.

Acknowledgments

This work was partly supported by the National Natural Science Foundation of China (62271410) and the Fundamental Research Funds for the Central Universities. School of Electronics and Information, Northwestern Polytechnical University and Shaanxi Key Laboratory of Information Acquisition and Processing, Xi'an, China.

References

- Akhter, I.; Sheikh, Y.; and Khan, S. 2009. In defense of orthogonality constraints for nonrigid structure from motion. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 1534–1541.
- Akhter, I.; Sheikh, Y.; Khan, S.; and Kanade, T. 2008. Non-rigid structure from motion in trajectory space. In *Adv. Neural Inform. Process. Syst. (NIPS)*, 41–48.
- Bregler, C.; Hertzmann, A.; and Biermann, H. 2000. Recovering non-rigid 3D shape from image streams. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 690–696.
- Cha, G.; Lee, M.; Cho, J.; and Oh, S. 2019. Reconstruct as far as you can: Consensus of non-rigid reconstruction from feasible regions. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI)*, 43(2): 623–637.
- Cha, G.; Lee, M.; and Oh, S. 2019. Unsupervised 3d reconstruction networks. In *Int. Conf. Comput. Vis. (ICCV)*, 3849–3858.
- Choi, J.; Gao, C.; Messou, J. C.; and Huang, J.-B. 2019. Why can't i dance in the mall? learning to mitigate scene bias in action recognition. *Adv. Neural Inform. Process. Syst. (NIPS)*, 32.
- Dai, Y.; Li, H.; and He, M. 2014. A simple prior-free method for non-rigid structure-from-motion factorization. *Int. J. Comput. Vis. (IJCV)*, 107(2): 101–122.
- Deng, H.; Zhang, T.; Dai, Y.; Shi, J.; Zhong, Y.; and Li, H. 2022. Deep Non-rigid Structure-from-Motion: A Sequence-to-Sequence Translation Perspective. *arXiv preprint arXiv:2204.04730*.
- Garg, R.; Roussos, A.; and Agapito, L. 2013a. Dense variational reconstruction of non-rigid surfaces from monocular video. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 1272–1279.
- Garg, R.; Roussos, A.; and Agapito, L. 2013b. A variational approach to video registration with subspace constraints. *Int. J. Comput. Vis. (IJCV)*, 104: 286–314.
- Geng, Z.; Guo, M.-H.; Chen, H.; Li, X.; Wei, K.; and Lin, Z. 2021. Is Attention Better Than Matrix Decomposition? In *Int. Conf. Learn. Represent. (ICLR)*.
- Gotardo, P. F.; and Martinez, A. M. 2011. Computing smooth time trajectories for camera and deformable shape in structure from motion with occlusion. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI)*, 33(10): 2051–2065.
- Gower, J. C. 1975. Generalized procrustes analysis. *Psychometrika*, 40: 33–51.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *Int. Conf. Comput. Vis. (ICCV)*, 2961–2969.
- Hinton, G.; Deng, L.; Yu, D.; Dahl, G. E.; Mohamed, A.-r.; Jaitly, N.; Senior, A.; Vanhoucke, V.; Nguyen, P.; Sainath, T. N.; et al. 2012. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal processing magazine*, 29(6): 82–97.
- Ionescu, C.; Papava, D.; Olaru, V.; and Sminchisescu, C. 2014. Human3.6M: Large Scale Datasets and Predictive Methods for 3D Human Sensing in Natural Environments. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI)*, 36(7): 1325–1339.
- Kocabas, M.; Athanasiou, N.; and Black, M. J. 2020. Vibe: Video inference for human body pose and shape estimation. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 5253–5263.
- Kong, C.; and Lucey, S. 2021. Deep Non-Rigid Structure from Motion with Missing Data. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI)*, 43(12): 4365–4377.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2017. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6): 84–90.
- Kudo, Y.; Ogaki, K.; Matsui, Y.; and Odagiri, Y. 2018. Unsupervised adversarial learning of 3D human pose from 2D joint locations. In *Eur. Conf. Comput. Vis. (ECCV)*.
- Kumar, S. 2020. Non-rigid structure from motion: Prior-free factorization method revisited. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 51–60.
- Kumar, S.; Cherian, A.; Dai, Y.; and Li, H. 2018. Scalable dense non-rigid structure-from-motion: A grassmannian perspective. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 254–263.
- Kumar, S.; Dai, Y.; and Li, H. 2016. Multi-body non-rigid structure-from-motion. In *In Proc. of the International Conf. on 3D Vision (3DV)*, 148–156.
- Kumar, S.; Dai, Y.; and Li, H. 2017. Spatio-temporal union of subspaces for multi-body non-rigid structure-from-motion. *Pattern Recognition(PR)*, 71: 428–443.
- Kumar, S.; and Van Gool, L. 2022. Organic Priors in Non-rigid Structure from Motion. In *Eur. Conf. Comput. Vis. (ECCV)*, 71–88.
- Lee, M.; Cho, J.; Choi, C.-H.; and Oh, S. 2013. Procrustean normal distribution for non-rigid structure from motion. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 1280–1287.
- Lee, M.; Choi, C.-H.; and Oh, S. 2014. A Procrustean Markov Process for Non-Rigid Structure Recovery. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*.
- Liu, G.; Lin, Z.; Yan, S.; Sun, J.; Yu, Y.; and Ma, Y. 2012. Robust recovery of subspace structures by low-rank representation. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI)*, 35(1): 171–184.
- Marques, M.; and Costeira, J. 2008. Optimal shape from motion estimation with missing and degenerate data. In *2008 IEEE Workshop on Motion and Video Computing*, 1–6.
- Moon, G.; Yu, S.-I.; Wen, H.; Shiratori, T.; and Lee, K. M. 2020. Interhand2.6m: A dataset and baseline for 3d interacting hand pose estimation from a single rgb image. In *Eur. Conf. Comput. Vis. (ECCV)*, 548–564.

- Novotny, D.; Ravi, N.; Graham, B.; Neverova, N.; and Vedaldi, A. 2019. C3dpo: Canonical 3d pose networks for non-rigid structure from motion. In *Int. Conf. Comput. Vis. (ICCV)*, 7688–7697.
- Paladini, M.; Bartoli, A.; and Agapito, L. 2010. Sequential non-rigid structure-from-motion with the 3d-implicit low-rank shape model. In *Eur. Conf. Comput. Vis. (ECCV)*, 15–28.
- Paladini, M.; Del Bue, A.; Stosic, M.; Dodig, M.; Xavier, J.; and Agapito, L. 2009. Factorization for non-rigid and articulated structure using metric projections. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2898–2905.
- Parashar, S.; Pizarro, D.; and Bartoli, A. 2018. Isometric Non-Rigid Shape-from-Motion with Riemannian Geometry Solved in Linear Time. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI)*, 40(10): 2442–2454.
- Parashar, S.; Salzmann, M.; and Fua, P. 2020. Local non-rigid structure-from-motion from diffeomorphic mappings. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2059–2067.
- Park, S.; Lee, M.; and Kwak, N. 2017. Procrustean regression: A flexible alignment-based framework for nonrigid structure estimation. *IEEE Trans. Image Process. (TIP)*, 27(1): 249–264.
- Park, S.; Lee, M.; and Kwak, N. 2020. Procrustean regression networks: Learning 3d structure of non-rigid objects from 2d annotations. In *Eur. Conf. Comput. Vis. (ECCV)*, 1–18.
- Qin, Z.; Han, X.; Sun, W.; He, B.; Li, D.; Li, D.; Dai, Y.; Kong, L.; and Zhong, Y. 2023. Toeplitz Neural Network for Sequence Modeling. In *Int. Conf. Learn. Represent. (ICLR)*.
- Shazeer, N. 2020. Glu variants improve transformer. *arXiv preprint arXiv:2002.05202*.
- Shi, J.; Deng, H.; and Dai, Y. 2024. Non-rigid Structure-from-Motion: Temporally-smooth Procrustean Alignment and Spatially-variant Deformation Modeling. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 21446–21455.
- Sidhu, V.; Tretschk, E.; Golyanik, V.; Agudo, A.; and Theobalt, C. 2020. Neural dense non-rigid structure from motion with latent space constraints. In *Eur. Conf. Comput. Vis. (ECCV)*, 204–222.
- Sun, K.; Xiao, B.; Liu, D.; and Wang, J. 2019. Deep High-Resolution Representation Learning for Human Pose Estimation. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 5693–5703.
- Tomasi, C.; and Kanade, T. 1992. Shape and motion from image streams under orthography: a factorization method. *Int. J. Comput. Vis. (IJCV)*, 9(2): 137–154.
- Torresani, L.; Hertzmann, A.; and Bregler, C. 2008. Non-rigid structure-from-motion: Estimating shape and motion with hierarchical priors. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI)*, 30(5): 878–892.
- Varol, A.; Salzmann, M.; Fua, P.; and Urtasun, R. 2012. A constrained latent variable model. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2248–2255.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Adv. Neural Inform. Process. Syst. (NIPS)*, 5998–6008.
- von Marcard, T.; Henschel, R.; Black, M.; Rosenhahn, B.; and Pons-Moll, G. 2018. Recovering Accurate 3D Human Pose in The Wild Using IMUs and a Moving Camera. In *European Conference on Computer Vision (ECCV)*.
- Wang, C.; Lin, C.-H.; and Lucey, S. 2020. Deep nrsfm++: Towards 3d reconstruction in the wild. In *In Proc. of the International Conf. on 3D Vision (3DV)*, 12–22.
- Wang, C.; and Lucey, S. 2021. PAUL: Procrustean Autoencoder for Unsupervised Lifting. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 434–443.
- Wu, H.; Hu, T.; Liu, Y.; Zhou, H.; Wang, J.; and Long, M. 2022. TimesNet: Temporal 2D-Variation Modeling for General Time Series Analysis. *arXiv preprint arXiv:2210.02186*.
- Xiao, J.; Chai, J.-x.; and Kanade, T. 2004. A closed-form solution to non-rigid shape and motion recovery. In *Eur. Conf. Comput. Vis. (ECCV)*, 573–587.
- Xu, C.; Chen, S.; Li, M.; and Zhang, Y. 2021. Invariant teacher and equivariant student for unsupervised 3d human pose estimation. In *AAAI*, volume 35, 3013–3021.
- Xu, X.; and Dunn, E. 2021. GTT-Net: Learned Generalized Trajectory Triangulation. In *Int. Conf. Comput. Vis. (ICCV)*, 5795–5804.
- Yang, G.; Sun, D.; Jampani, V.; Vlastic, D.; Cole, F.; Chang, H.; Ramanan, D.; Freeman, W. T.; and Liu, C. 2021a. Lasr: Learning articulated shape reconstruction from a monocular video. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 15980–15989.
- Yang, G.; Sun, D.; Jampani, V.; Vlastic, D.; Cole, F.; Liu, C.; and Ramanan, D. 2021b. Viser: Video-specific surface embeddings for articulated 3d shape reconstruction. *Adv. Neural Inform. Process. Syst. (NIPS)*, 34: 19326–19338.
- Yang, G.; Vo, M.; Neverova, N.; Ramanan, D.; Vedaldi, A.; and Joo, H. 2022. BANMo: Building Animatable 3D Neural Models From Many Casual Videos. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2863–2873.
- Zeng, H.; Dai, Y.; Yu, X.; Wang, X.; and Yang, Y. 2021. PR-RRN: Pairwise-Regularized Residual-Recursive Networks for Non-rigid Structure-from-Motion. In *Int. Conf. Comput. Vis. (ICCV)*, 5600–5609.
- Zeng, H.; Yu, X.; Miao, J.; and Yang, Y. 2022. MHR-Net: Multiple-Hypothesis Reconstruction of Non-Rigid Shapes from 2D Views. In *Eur. Conf. Comput. Vis. (ECCV)*, 1–17.
- Zhang, J.; Tu, Z.; Yang, J.; Chen, Y.; and Yuan, J. 2022. Mixste: Seq2seq mixed spatio-temporal encoder for 3d human pose estimation in video. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 13232–13242.
- Zhu, Y.; Huang, D.; De La Torre, F.; and Lucey, S. 2014. Complex non-rigid motion 3d reconstruction by union of subspaces. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 1542–1549.