MHR-Net: Multiple-Hypothesis Reconstruction of Non-Rigid Shapes from 2D Views

Haitian Zeng^{1,2}, Xin Yu¹, Jiaxu Miao³, and Yi Yang³

¹ University of Technology Sydney, ² Baidu Research, ³ Zhejiang University, haitian.zeng@student.uts.edu.au,xin.yu@uts.edu.au jiaxu.miao@yahoo.com,yangyics@zju.edu.cn

Abstract. We propose MHR-Net, a novel method for recovering Non-Rigid Shapes from Motion (NRSfM). MHR-Net aims to find a set of reasonable reconstructions for a 2D view, and it also selects the most likely reconstruction from the set. To deal with the challenging unsupervised generation of non-rigid shapes, we develop a new Deterministic Basis and Stochastic Deformation scheme in MHR-Net. The non-rigid shape is first expressed as the sum of a coarse shape basis and a flexible shape deformation, then multiple hypotheses are generated with uncertainty modeling of the deformation part. MHR-Net is optimized with reprojection loss on the basis and the best hypothesis. Furthermore, we design a new Procrustean Residual Loss, which reduces the rigid rotations between similar shapes and further improves the performance. Experiments show that MHR-Net achieves state-of-the-art reconstruction accuracy on Human3.6M, SURREAL and 300-VW datasets.

Keywords: NRSfM, multiple-hypothesis

1 Introduction

Recovering 3D structures from multiple 2D views is a classic and important task in computer vision. Non-Rigid Structure-from-Motion (NRSfM), which aims at reconstructing deformable shapes, is a challenging task and has been studied for decades.

The major difficulty of NRSfM is the ambiguity of solutions due to arbitrary deformation of shapes. Most of the NRSfM methods are based on the assumption of Bregler et al. [9] where the deformable shape is a linear combination of a small number of atom shapes. This assumption greatly reduces the degree of freedom in NRSfM, yet it is still not enough for researchers to reach a deterministic and closed-form solution. Prior work of Akhter et al. [25] reveals that the local minimas grows exponentially with the basis number, and the reconstructed shapes from most of local minimas deviate significantly from ground truth. And Dai et al. [18, 16] demonstrate that the rank minimization method also leads to multiple minimas in perspective cases. These reveal that there are usually multiple solutions that all minimize the cost function of a "prior-free" NRSfM, but it

is generally intractable to find all those solutions due to the inherent complexity of NRSfM ambiguity.

Nevertheless, in most cases we only need to obtain several most reasonable hypotheses since they are valuable in practical scenarios, while searching for all ambiguous solutions exhaustively is not necessary. This leads us to focus on finding multiple high-quality hypotheses for the NRSfM problem.

Multiple hypotheses are usually modeled with uncertainty or generative models like CVAE [59], MDN [36] or CGAN [37]. However, these conventional modeling methods are supervised by 3D ground-truth, which is not available for a NRSfM problem. Moreover, in NRSfM, a naive ensemble of independent models is prone to decomposition ambiguity [18], and a variational autoencoder is also found hard to train [66].

To overcome the above challenges, we propose a novel MHR-Net for Multiple-Hypothesis Reconstruction of non-rigid shapes. Different from a standard model which outputs one reconstruction for a single input, MHR-Net is capable to produce multiple reasonable solutions and one best solution. We develop several critical designs for the successful generation of multiple hypotheses. Firstly, one non-rigid shape is expressed as the sum of a *basis* and a *deformation*. The basis is the coarse and shared structure among all shapes, while the deformation accounts for the diverse and flexible parts of shapes. This shape expression enhances the representation capability of MHR-Net when trained with an intermediate reprojection loss on the basis. Based on this expression, we further propose a novel Deterministic Basis and Stochastic Deformation (DBSD) scheme for multiple hypotheses generation. Specifically, MHR-Net estimates one basis in a standard deterministic manner and multiple deformations in a stochastic way. Then the multiple reconstructions are obtained by adding the basis and deformations. To optimize MHR-Net, we adopt a pseudo "hindsight" loss which is to select a hypothesis with the minimal reprojection error and calculates the standard loss function on the selected hypothesis. In inference, the model produces the best hypothesis in the same way. The DBSD scheme not only enables MHR-Net to produce multiple high-quality solutions of NRSfM, but also further enhances the accuracy of the reconstruction.

Moreover, we develop a new Procrustean Residual Loss to regularize the reconstruction and reduce undesirable rigid rotations in a differentiable and efficient way. Experiments on Human3.6M, 300-VW and SURREAL datasets demonstrate state-of-the-art reconstruction accuracy of MHR-Net. Finally, we show that MHR-Net is capable to produce multiple possible solutions of 3D human poses and largely-deforming regions of dense human body meshes.

We summarize our contributions as follows:

- We propose the novel MHR-Net for NRSfM. To the best of our knowledge, it is the first method that produces multiple high-quality hypotheses for non-rigid shape reconstruction in one model.
- We introduce a deterministic basis and stochastic deformation scheme together with a intermediate loss and a pseudo hindsight loss. These designs are



Fig. 1. (a) A standard function-based NRSfM model maps the 2D input to a single 3D output. (b) The proposed MHR-Net is aware of the uncertainty of 2D-to-3D mapping. Given an extra noise vector, MHR-Net is capable to output multiple reasonable reconstructions of the 2D input.

effective for the challenging unsupervised uncertainty modeling of multiple 3D shapes.

 We develop a novel Procrustean Residual Loss for NRSfM regularization, and it further improves the shape recovery accuracy of MHR-Net.

2 Related Works

NRSfM. Recovering the deforming 3D shape from multiple 2D views is known as the Non-Rigid Structure-from-Motion problem (NRSfM), which is firstly put forward by Bregler et al. [9]. NRSfM is a highly ambiguous problem, and therefore enforcing correct constraints is key to solving this problem. Previous works propose various effective approaches to the non-rigid shape reconstruction, including rank minimization [17, 41], smooth trajectories [7, 21, 20, 10], manifold [22, 31], metric projection [42], sparsity-based methods [28], energy minimization [52, 57], inextensibility [64, 15], isometry [43], deep models [28, 40, 47, 69], procrustean normal distribution [34], consensus [13], hierarchical priors [61], force-based and physic-based methods [3, 6], union of subspaces [70, 33, 5, 1, 4], piecewise methods [19, 60], and many other breaking-through methods [32, 11, 50, 2, 53, 49, 44, 24].

Much attention has been paid to the uniqueness and determinacy of NRSfM. Xiao et al. [68] show that selecting a set of frames as the basis can lead to a unique closed-form solution. Akhter et al. [25] argue that the orthogonal constraints of rotations is indeed sufficient for a unique solution except for a rigid rotation and the major difficulty lies in the optimization. Dai et al. [17] propose a block-matrix rank-minimizing method and analyze whether their method leads to a unique solution or multiple solutions. Park et al. [45] provide a geometric analysis showing that the quality of sequential reconstruction is affected by the relative motion of a point and the camera, and propose a novel measure reconstructability to indicate the reconstruction quality. Valmadre et al. [62] propose a deterministic approach to 3D human pose recovery by using the rigid sub-structure of human body.

Multiple-Hypothesis 3D Pose Estimation. The ambiguity of monocular 3D human pose estimation has been noticed early [58]. Li et al. [36] use a mixture density network or a Conditional GAN [37] to output a set of plausible 3D poses from a 2D input. Sharma et al. [55] propose to solve the ill-posed 2D-to-3D lifting problem with CVAE [59] and Ordinal Ranking. Wehrbein et al. [67] use Normalizing Flows to model the deterministic 3D-2D projection and solve the ambiguous inverse 2D-3D lifting problem. The major difference between our work and multiple-hypotheses 3D pose estimation is that our model is trained without 3D ground truth.

3 Preliminary

In the classic non-rigid structure-from-motion problem, given N_f frames 2D observations $\{W_i\}_{i=1}^{N_f}$ of a deformable object as input, we are interested in factorizing $W_i \in \mathbb{R}^{2 \times N_p}$ into a camera matrix $M_i \in \mathbb{R}^{2 \times 3}$ and a shape matrix $S_i \in \mathbb{R}^{3 \times N_p}$ such that:

$$\mathbf{W}_i = \mathbf{M}_i \mathbf{S}_i. \tag{1}$$

Here, we suppose that S_i is centered at zero such that the translation term is cancelled, and N_p stands for the number of points. M_i is the composition of a projection matrix $\Pi \in \mathbb{R}^{2\times 3}$ and a rotation matrix $R_i \in SO(3)$ so that $M_i = \Pi R_i$. For orthographic projection, Π is simply $[I_2 \mathbf{0}]$. In this work, we suppose that the camera projection Π is known, allowing us to focus on the estimation of rotation R_i .

In the recent progress of NRSfM [28, 40, 47], M_i and S_i are modeled as functions of the input W_i . One typical paradigm [40, 14] is to first extract features from W_i using a backbone network $\mathcal{H}(W_i)$ like [39], and then to estimate the M_i and S_i with different network branches \mathcal{F}_0 and \mathcal{G} subsequent to \mathcal{H} :

$$S_i = \mathcal{F}_0(\mathcal{H}(W_i)), \quad M_i = \mathcal{G}(\mathcal{H}(W_i)).$$
 (2)

Modeling the factorization as a function enables NRSfM methods to be optimized on large-scale datasets, and allows models to directly perform reconstruction on unseen data. To train such models, the cost function usually contains a data term and a regularization term, represented as:

$$\mathcal{L} = \sum_{i=1}^{N_f} \mathcal{L}_{\text{data}}(\mathbf{W}_i, \mathbf{M}_i, \mathbf{S}_i) + \mathcal{L}_{\text{reg}}(\mathbf{M}_i, \mathbf{S}_i),$$
(3)

where the data term \mathcal{L}_{data} is usually the reprojection error $||W_i - M_i S_i||$ and the regularization term is versatile.



Fig. 2. An overview of the proposed MHR-Net. MHR-Net uses a backbone network \mathcal{H} to extract features from the 2D input W_i . The camera rotation R_i is estimated by the rotation layer. Next, the basis shape S_i^B and its coefficients α_i are estimated by Ψ_B and Ψ_{α} respectively. To generate multiple hypotheses, beta layer Ψ_{β} takes α_i , $\mathcal{H}(W_i)$ and random noise $\{z^1, \dots, z^{N_m}\}$ as inputs, and outputs a set of deformation coefficients $\{\beta_i^1, \dots, \beta_i^{N_m}\}$. Then the deformation coefficients are passed to the deformation layer Ψ_D to produce deformations $\{S_i^{D,1}, \dots, S_i^{D,N_m}\}$. By adding each deformation to S_i^B , multiple hypotheses $\{S_i^1, \dots, S_i^{N_m}\}$ are obtained. Then, the hypothesis with the smallest reprojection error is selected as the best hypothesis S_i^* . Loss functions are calculated on S_i^B and S_i^* .

4 Proposed Method

4.1 Multiple Hypothesis Reconstruction Network - Overview

We aim to develop a prediction function \mathcal{F} that can output N_m reconstructions for a single input W_i :

$$\mathcal{F}(\mathbf{W}_i) = \{\mathbf{S}_i^1, \mathbf{S}_i^2, \cdots, \mathbf{S}_i^{N_m}\},\tag{4}$$

and each of these reconstructions is supposed to minimize the cost function in Eq. 3. As the inherent ambiguity of NRSfM is complex and there exists a large number of poor ambiguous solutions [25], we also expect the hypotheses to be as accurate as possible among all solutions.

However, generating multi-hypothesis reconstruction for NRSfM is challenging for several reasons: (1) Without 3D ground-truth as supervision, the ambiguous 2D-to-3D mappings cannot be learned using standard generative models like CVAE [59], Conditional GAN [37] or Normalizing Flow [67]. (2) Multiple hypotheses easily suffer from the decomposition ambiguity of NRSfM [18], i.e. multiple solutions are trivial if they are related by a certain rotation G inserted in the decomposition $W_i = M_i G G^{-1} S_i$.

We introduce a novel Multiple Hypothesis Reconstruction Network (MHR-Net), which takes a step towards a multiple-hypothesis NRSfM model. MHR- Net overcomes the above difficulties in a simple and effective way, and is capable to produce multiple accurate reconstructions and one best hypothesis. Next, we describe the hypothesis generation scheme in Sec. 4.2, the optimization in Sec. 4.3 and the regularization in Sec. 4.4.

4.2 Deterministic Basis and Stochastic Deformation for Hypothesis Generation

Traditionally, the 3D shape S_i is represented as a linear combination of K_b atom shapes $B_k \in \mathbb{R}^{3 \times N_p}$:

$$\mathbf{S}_{i} = \sum_{k=1}^{K_{b}} (\alpha_{i,k} \otimes \mathbf{I}_{3}) \mathbf{B}_{k}, \tag{5}$$

where $\alpha_{i,k}$ is the weight of k-th atom shapes in S_i , and \otimes is the Kronecker product. This widely-accepted representation implicitly assumes the low-rankness of all estimated 3D shapes, and it is one of keys to successful recovery of non-rigid shapes. However, the capacity of this shape representation is limited when applied to modern large-scale datasets since they usually contain millions of frames or thousands of keypoints. A naive way is to enlarge the capacity by increasing the dimension K_b , but it also bring undesirable degrees of freedom to the full shape, leading to a regularization-flexibility dillema.

Inspired by previous works [1, 8, 57], we develop a new shape representation for MHR-Net. We posit that the deformable shape is comprised of a *basis* S_i^B and a *deformation* S_i^D , written as:

$$\mathbf{S}_{i} = \underbrace{\sum_{k=1}^{K_{b}} (\alpha_{i,k} \otimes \mathbf{I}_{3}) \mathbf{B}_{k}}_{\mathbf{S}_{i}^{B}} + \underbrace{\sum_{l=1}^{K_{d}} (\beta_{i,l} \otimes \mathbf{I}_{3}) \mathbf{D}_{l}}_{\mathbf{S}_{i}^{D}}, \tag{6}$$

where $D_l \in \mathbb{R}^{3 \times N_p}$ denotes the *l*-th atom deformation, $\beta_{i,l}$ is the weight of D_l in S_i , and $K_b < K_d$. Note that the mean shape for short sequences in [57] is a special case of Eq. 6 with $K_b = 1, \alpha_i = 1$. This basis-deformation modeling enhances the shape representation capability of MHR-Net with hierarchical flexibility. Intuitively, S_i^B is used to capture the low-rank common part of the 3D shapes, while S_i^D fits the diverse small deformations with higher flexibility.

Based on the basis-deformation expression of shapes, we now introduce a novel Deterministic Basis and Stochastic Deformation (DBSD), which is the core scheme for multiple hypothesis generation in MHR-Net. We assume that the variation of N_m accurate hypotheses appears only in S_i^D . While Wang et al. [66] find that introducing randomness to the estimation of full 3D shapes with VAE [27] is not viable, the proposed partial uncertainty in S_i^D makes MHR-Net overcome the training difficulty. More importantly, we find that DBSD can lead to an even better recovery accuracy with the training strategy in Sec. 4.3. Specifically, MHR-Net generates multiple reconstructions by estimating one basis in the standard deterministic way and multiple deformations with stochastic variations. For the deterministic part, coefficients $\alpha_i = [\alpha_{i,1} \cdots \alpha_{i,K_b}]^{\mathrm{T}}$ are estimated with a convolutional layer Ψ_{α} :

$$\alpha_i = \Psi_\alpha(\mathcal{H}(\mathbf{W}_i)). \tag{7}$$

For the stochastic deformation part, the coefficients $\beta_i = [\beta_{i,1} \cdots \beta_{i,K_d}]^{\mathrm{T}}$ are calculated by a convolutional layer Ψ_{β} which takes features $\mathcal{H}(\mathrm{W}_i)$, coefficients α_i and a noise vector $z \sim \mathrm{N}(0,\mathrm{I})$ as inputs:

$$\beta_i = \Psi_\beta(\mathcal{H}(\mathbf{W}_i), \alpha_i, z). \tag{8}$$

By sampling N_m noise vectors $\{z^1, \dots, z^{N_m}\}$ from the isotropic Gaussian with dimension \dim_z and passing them to Ψ_β , we have a collection $\{\beta_i^1, \dots, \beta_i^{N_m}\}$. Next, the basis and multiple deformations are produced by two following convolutional layers Ψ_B and Ψ_D :

$$\mathbf{S}_i^B = \Psi_B(\alpha_i), \quad \mathbf{S}_i^{D,m} = \Psi_D(\beta_i^m). \tag{9}$$

Note that the atoms B_k and D_l are learned as the parameters of Ψ_B and Ψ_D . Finally, the multiple hypothesis reconstructions $\{S_i^1, \dots, S_i^{N_m}\}$ are generated by:

$$\mathbf{S}_i^m = \mathbf{S}_i^B + \mathbf{S}_i^{D,m},\tag{10}$$

where $m = 1 \cdots N_m$.

For camera R_i estimation, we avoid the decomposition ambiguity by simply estimating one rotation matrix for all hypothesis. We follow [40] to predict the rotation matrix using Rodrigues' Rotation Formula, which is parameterized by a 3 dimensional output of a convolutional layer built upon the extracted features. Given Π , we obtain $M_i = \Pi R_i$.

4.3 Optimizing with Intermediate Loss and Best Hypothesis Selection

We introduce two effective methods for optimizing MHR-Net on the data term \mathcal{L}_{data} .

Intermediate Loss. First, we consider the optimization of a single hypothesis S_i^m produced by MHR-Net from W_i . The proposed basis-deformation expression of shapes in Eq. 6 is the summation of two linear subspaces. In practice, MHR-Net is prone to using only one flexible subspace S_i^D for shape expression if we adopt the standard reprojection loss $||W_i - M_i S_i^m||$. To make the basis-deformation expression work as expected, we propose to add an extra reprojection loss to the intermediate results of reconstruction, i.e. the basis S_i^B . This enforces MHR-Net to produce a low-rank and coarse prediction S_i^B of the 3D shape, thus letting S_i^D focus on the small and diverse residuals. The extra intermediate loss is written as $||W_i - M_i S_i^B||$.

Best Hypothesis Selection. Now we consider the optimization of all hypotheses. There are several approaches to train a model with multiple predictions, including minimizing losses of all predictions, selecting one hypothesis with mixture density [36], *etc.* In MHR-Net, we choose to use a pseudo "hindsight" loss [35, 51]. The original "hindsight" loss choose a prediction that is the closest to the ground truth and then calculate the standard single-prediction loss. As in NRSfM we do not have any 3D ground truth, we heuristically use the reprojection error as the criterion to select the best hypothesis among our predictions. Together with the intermediate loss, \mathcal{L}_{data} is written as:

$$\mathcal{L}_{\text{data}} = \lambda_B \left\| \mathbf{W}_i - \mathbf{M}_i \mathbf{S}_i^B \right\| + \lambda_F \min_m \left\| \mathbf{W}_i - \mathbf{M}_i \mathbf{S}_i^m \right\|,\tag{11}$$

where λ_B, λ_F are balancing factors and $\lambda_B + \lambda_F = 1$. Compared to other multiprediction training approaches, we find that the proposed strategy brings a better reconstruction accuracy to all hypotheses. Moreover, when inferring a single best reconstruction S_i^* of the input, S_i^* is naturally obtained with the same best hypothesis strategy.

Discussion. (1) The insight of successful hypothesis generation in MHR-Net is to constrain the norm of the flexible deformation subspace. Although the proposed shape representation (Eq. 6) has a large degree of freedom (S_i^D could cause a maximum of 2^{9K_d} local minimas [25]), Eq. 11 implicitly constrains the norm of S_i^D to be relatively small compared to S_i^B when the balancing factors are chosen as $\lambda_B = 0.8$, $\lambda_F = 0.2$. On the other hand, one can choose a smaller λ_B and adding a diversity loss as in [37] for generating more diverse hypotheses. (2) Moreover, we found that the model with DBSD has a slightly higher \mathcal{L}_{data} and a lower \mathcal{L}_{reg} compared to the deterministic model, which indicates that DBSD leads to a better regularized model. The diverse hypotheses reduce overfitting of MHR-Net and enhance the generalization capability.

4.4 Procrustean Residual Loss for Regularization

In this section, we introduce a novel Procrustean Residual Loss for regularizing the non-rigid shape reconstruction.

Motivation. Reducing the rigid motions between reconstructed shapes is one of the keys to successful NRSfM. In the previous work of Novotny et al. [40], the Transversal property is proposed to characterize a space where a shape is enforced to appears in a canonical view, such that the effects of rigid motions between same shapes are removed. Implemented with an auxiliary neural network, the Transversal property is effective in performing non-rigid reconstructions.

However, the Transversal property is still restricted to aligning only identical shapes. That means shapes with small differences are not guaranteed to be aligned in a Transversal set. In other words, the effect of a rigid motion is not removed for similar (but not identical) shapes. As regularization on similar shapes is shown to be useful in a recent work [69], we are motivated to reduce the rigid motion between similar shapes.

To achieve this objective, we first define two distance measures:

Definition 1. Given two non-degenerated shapes $S_i, S_j \in \mathbb{R}^{3 \times N_p}$ and the optimal rotation \mathbb{R}^* aligning S_i to S_j , the Original distance δ_{ori} is $\|S_i - S_j\|_F$, and the Procrustean distance δ_{pro} is $\|\mathbb{R}^*S_i - S_j\|_F$.

Here, a shape is non-degenerated if $\operatorname{rank}(S_i) = 3$, $\|\cdot\|_F$ denotes the Frobenius norm, and the optimal rotation \mathbb{R}^* is obtained with the orthogonal Procrustes [54].

Next, we use the two defined distances to: (1) test whether S_i and S_j are similar or not; (2) if S_i, S_j are similar, measure the effect of rigid motions. In step (1), S_i, S_j are considered to be similar if $\delta_{pro} < \epsilon$, where ϵ is a hyperparameter of similarity threshold. We use the Procrustean distance in this step since it is agnostic of rigid rotations. In step (2), we propose to measure the effect of rigid motions with the *Procrustean Residual* δ_{res} , calculated as:

$$\delta_{res} = \delta_{ori} - \delta_{pro}.$$
 (12)

The Procrustean Residual indicates how much the Original distance can be reduced with rigid motions, and δ_{res} reaches zero if and only if S_i and S_j are already optimally aligned (i.e. $R^* = I$). Therefore, the undesirable rigid motion between S_i, S_j is reduced when we minimize δ_{res} .

Loss design. We now introduce the Procrustean Residual Loss for NRSfM regularization. This loss function realizes the minimization⁴ of δ_{res} and is developed in a differentiable way.

Given two shapes S_i , S_j randomly sampled from the network prediction batch, the optimal rotation that aligns S_i to S_j is calculated with the orthogonal Procrustes [54]:

$$\mathbf{R}_{i,j}^* = \mathbf{V}\mathbf{U}^{\mathrm{T}}, \quad \mathbf{S}_i\mathbf{S}_j^{\mathrm{T}} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^{\mathrm{T}},\tag{13}$$

where $U\Sigma V^{T}$ is the singular value decomposition of $S_{i}S_{i}^{T}$.

Next, we calculate the (normalized) Procrustean distance and the Procrustean Residual as follows:

$$\bar{\delta}_{\text{pro}} = \frac{\left\|\mathbf{R}_{i,j}^*\mathbf{S}_i - \mathbf{S}_j\right\|_{\mathbf{F}}}{\left\|\mathbf{S}_j\right\|_{\mathbf{F}}}, \quad \bar{\delta}_{\text{res}} = \frac{\left\|\mathbf{S}_i - \mathbf{S}_j\right\|_{\mathbf{F}}}{\left\|\mathbf{S}_j\right\|_{\mathbf{F}}} - \bar{\delta}_{\text{pro}}.$$
(14)

We normalize the differences with $\|S_j\|_F$ to make the loss numerically stable. The Procrustean Residual Loss is:

$$\mathcal{L}_{\rm res}(\mathbf{S}_i, \mathbf{S}_j) = \rho(\bar{\delta}_{\rm pro}, \epsilon) \cdot \bar{\delta}_{\rm res}, \tag{15}$$

where $\rho(\bar{\delta}_{\rm pro}, \epsilon) = 1$ if $\bar{\delta}_{\rm pro} < \epsilon$, else $\rho = 0$.

The practical problem of the proposed loss function is that \mathcal{L}_{res} contains a non-differentiable operation SVD in Eq. 13. To make \mathcal{L}_{res} differentiable, one effective way is to use the Lagrange multiplier method on Lie Group [46, 47] for a closed-form partial derivative, and another way is to leverage modern auto-grad

⁴ Note that in general δ_{res} cannot be reduce to exactly zero for all pairs simultaneously.

libraries where the numeric computation of SVD is differentiable [28, 66, 69]. In this paper, we choose to use a simple alternative approach by detaching $R_{i,j}^*$ from the computation graph, namely $\mathbf{R}^*_{i,j}$ is viewed as a constant matrix. In such way, \mathcal{L}_{res} only involves standard differentiable operations of inputs S_i, S_j , which frees us from the calculation of SVD gradient and keeps the model computationally efficient.

For the regularization term $\mathcal{L}_{\mathrm{reg}}$ of MHR-Net, we apply a canonicalization loss $\mathcal{L}_{cano}{}^{5}$ [40] to the deterministic basis S_{i}^{B} and \mathcal{L}_{res} to S_{i}^{*} , leading to:

$$\mathcal{L}_{\rm reg} = \mathcal{L}_{\rm cano} + \lambda_{\rm res} \mathcal{L}_{\rm res}, \tag{16}$$

where λ_{res} is the weight of Procrustean Residual Loss. Although only including \mathcal{L}_{res} in \mathcal{L}_{reg} is possible and produces good results, we empirically find that using two losses jointly leads to a better performance.

$\mathbf{5}$ Experiments

We evaluate the proposed MHR-Net in two aspects: (1) The reconstruction accuracy of the best hypothesis. (2) The multiple hypothesis reconstructions. We also make an in-depth analysis of proposed components.

5.1**Datasets and Experimental Setups**

Human3.6M [26]. It is the largest 3D human pose dataset with a total of 3.6 million frames. It contains 15 different activities performed by 7 professional actors and captured by four cameras. We follow the common protocols to use five subjects (S1, S5-8) as the training set and two subjects (S9 and S11) as the testing set. We adopt the widely-used pre-processing from Pavllo et al. [48].

300VW [56]. The 300VW is a large-scale facial landmark dataset. It has a total of 114 sequences with 2D annotations of 68 landmarks. Following [47], we use the subset of 64 sequences from 300VW, and divide them into a training set and a testing set of 32 sequences respectively. As 3D ground-truth is not provided. we follow [47] to adopt the results from [12] as 3D ground-truths.

SURREAL [63]. The SURREAL dataset contains 6 million synthetic human images with large variations in shape, view-point and pose. The 6,890 dense 3D points are obtained by fitting SMPL [38] to CMU MOCAP dataset. Following [47, 66], the training and testing sets include 5,000 and 2,401 frames selected from the full dataset, respectively.

Metrics. We adopt the following two metrics:

(1) MPJPE: It stands for the mean per joint error, which is calculated as $\frac{1}{N_{\pi}} \|\mathbf{S}_{i} - \mathbf{S}_{i}^{\text{gt}}\|_{1}$. To address the reflection ambiguity, we follow [40, 47, 66] to report the minimal error with ground-truth between original and flipped shapes. (2) Normalized Error (NE): It reflects the relative estimation error and is computed by: $\frac{\left\|\mathbf{S}_{i}-\mathbf{S}_{i}^{\text{gt}}\right\|_{\text{F}}}{\left\|\mathbf{S}_{\text{gt}}\right\|_{\text{F}}}.$

⁵ Please refer to the supplementary material or [40] for details.

Methods (Ortho.)	Direct.	Discuss	Eating	Greet	Phone	Pose	Purch.	Sitting	SittingD.	Smoke	Photo	Wait	Walk	WalkD.	WalkT.	Mean
CSF2 [23]	87.2	90.1	96.1	95.9	102.9	92.1	99.3	129.8	136.7	99.5	120.1	95.2	90.8	102.4	89.2	101.6
SPM [18]	65.3	68.7	82.0	70.1	95.3	65.1	71.9	117.0	136.0	84.3	88.9	71.2	59.5	73.3	68.3	82.3
C3DPO [40]	56.1	55.6	62.2	66.4	63.2	62.0	62.9	76.3	85.8	59.9	88.7	63.3	71.1	70.7	72.3	67.8
PRN [47]	65.3	58.2	60.5	73.8	60.7	71.5	64.6	79.8	90.2	60.3	81.2	67.1	54.4	61.2	65.6	66.7
MHR-Net (Ours)	60.3	54.3	55.5	67.9	67.7	69.5	61.3	69.7	83.2	67.6	85.3	61.7	61.9	63.4	68.2	65.8
Methods (Persp.)	Direct.	Discuss	Eating	Greet	Phone	Pose	Purch.	Sitting	SittingD.	Smoke	Photo	Wait	Walk	WalkD.	WalkT.	Mean
PoseGAN [30]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	130.9
SFC [29]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	218.0
Consensus [13]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	120.1
DNRSFM [28]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	101.6
Wang et al. [65]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	86.2
C3DPO [40]	96.8	85.7	85.8	107.1	86.0	96.8	93.9	94.9	96.7	86.0	124.3	90.7	95.2	93.4	101.3	95.6
PRN [47]	93.1	83.3	76.2	98.6	78.8	91.7	81.4	87.4	91.6	78.2	104.3	89.6	83.0	80.5	95.3	86.4
PAUL [66]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	88.3
PoseDict [14]	74.6	82.9	77.0	86.7	80.0	79.2	94.2	88.4	124.0	77.1	103.8	80.8	78.8	94.2	78.3	85.6
ITES [14]	77.6	77.3	77.1	77.3	77.3	77.4	77.3	77.2	77.3	77.1	77.1	77.5	77.3	77.2	77.5	77.2
MHR-Net (Ours)	62.8	68.3	62.2	73.9	73.7	67.0	70.2	76.7	100.0	71.5	90.0	72.3	68.8	80.2	71.0	72.6

Table 1. Quantitative results on Human3.6M Dataset.

5.2 Main Results

In this subsection and Sec. 5.3, we treating MHR-Net as a single-prediction model by using the best hypothesis S_i^* . We report the standard NRSfM evaluation results of MHR-Net on three datasets.

For Human3.6M, we test the performance of MHR-Net under two settings: orthographic camera and perspective camera. The major competitors of MHR-Net are state-of-the-art deep NRSfM models, including C3DPO [40], DNRSFM [28], PRN [47], PAUL [66], ITES [14]. In Table. 1, we report the MPJPE of all frames and 15 individual activities on the test set. We also includes classic methods like Consensus [13], SFC [29] for comparison. As shown in Tab. 1, MHR-Net outperforms all competing methods overall in both orthographic and perspective experiments. These results on the challenging Human3.6M dataset verify the effectiveness of MHR-Net on reconstructing highly-flexible human poses.

For 300VW dataset and SURREAL dataset, we compare MHR-NET with C3DPO [40], PRN [47], PR [46] and PAUL [66]. The Normalized Error results are shown in Tab. 3 and Tab. 2 respectively. These outcomes validate that MHR-Net is capable to perform accurate reconstruction of both facial landmarks and dense meshes. It is noting that MHR-Net recovers the dense point clouds of SURREAL dataset without splitting it into several subsets, unlike [47]. This is achieved by avoiding the burdensome SVD of matrices whose scales are related to N_p . With the differentiable design, MHR-Net shares the same level of scalability as SVD-free methods while achieving better performance.

5.3 Ablation Study

We show the effectiveness of the important designs in MHR-Net. We conduct the experiments on perspective Human 3.6M dataset. BD and IL in Tab. 4 are short for Basis-Deformation (Eq. 6) and Intermediate Loss (Sec. 4.3).

Basis-Deformation and Intermediate Loss. We setup several ablated models: (1) Baseline: We use a modified PoseDict [14] as baseline. We replace the

H. Zeng et al.

Madal	NE	Model NE
Model	INE	Model NE
C3DPO [40]	0.3509	CSF2 [23] 0.2751
PRN [47]	0.1377	PR [46] 0.2730
PAUL [66]	0.1236	C3DPO [40] 0.1715
MHR-Net (w/o 4	$C_{\rm res}$) 0.1388	PRN [47] 0.1512
MHR-Net	0.1192	MHR-Net 0.1007
Table 2 Regults of	n SUBBEAL	Table 3 Results on 300

	IV.	ш	1-IN6	эt
	1	0	D	1.

Table 3. Results on 300VW.

No.	BD	IL	DBSD	Optim.	$\mathcal{L}_{\mathrm{res}}$	MPJPE
1	X	X	X	X	X	83.5
2	1					198.3
3	1	1				75.5
4	1	1	1	Best		73.7
5	1	1	1	Worst		78.4
6	1	1	1	MD		77.9
7	1	1	1	Worst	1	76.2
8	1	1	1	MD	1	77.3
9	✓	1	1	Best	1	72.6

 Table 4. Ablation study results.

invariance loss of PoseDict with the canonicalization loss [40], and it works slightly better than PoseDict. In Baseline, only the basis and camera are estimated. (2) Baseline with Basis-Deformation. The deformation here is implemented deterministically. (3) Baseline with Basis-Deformation and Intermediate Loss. Comparing the results of (1) and (2) in Tab. 4, we observe the degradation of performance. This implies that a naive extension of the deformation subspace will harm the regularization of low-rankness and lead to the failure of non-rigid reconstruction. By adding an intermediate loss in (3), the MPJPE is greatly reduced from 83.5 to 75.5, which is already better than the MPJPE (77.2) of the most competitive method ITES.

Stochastic Deformation and Hypothesis Optimization Strategy. We now use the deterministic basis and stochastic deformation with best hypothesis selection strategy, indicated by (4) in Tab. 4. We compare the proposed design with two alternatives: (5) Worst Hypothesis. In this strategy, we choose to optimize the hypothesis with the largest re-projection error, which is the opposite of the best hypothesis strategy. The intuition of Worst Hypothesis is that it tries to minimize the upper bound of errors. (6) Mixture Density (MD) [36], where a hypothesis is selected by sampling from a learned mixture density. As we do not have ground-truth labels for training, we use the uniform categorical distribution instead of a learned distribution. The results in Tab. 4 demonstrates that the combination of DBSD and the Best Hypothesis strategy produce the best performance.

12



Fig. 3. Point-wise variation of deformations.

Procrustean Residual Loss. We show the effectiveness of the Procrustean Residual Loss \mathcal{L}_{res} . As reported in Tab. 4, the full model (9) has a better MPJPE of 72.6. Despite the fact that the previous state-of-the-art method ITES uses two networks and the prior knowledge of human pose, MHR-Net outperforms ITES by 4.6 MPJPE. By adding \mathcal{L}_{res} to models with alternative MD and Worst Hypothesis, these two models (7) and (8) also outperform the corresponding models (5) and (6) without \mathcal{L}_{res} . Moreover, the improvement by using \mathcal{L}_{res} is also significant on the dense mesh dataset SURREAL, as indicated in Tab. 2.

5.4 Analysis of Stochastic Deformation

Point-wise variation of deformations. We measure the variation of stochastic deformations of point j as $\frac{1}{N_f} \sum_{i=1}^{N_f} \max_{m,n} \left\| \mathbf{S}_{i,j}^{D,m} - \mathbf{S}_{i,j}^{D,n} \right\|_{\mathbf{F}}$, where the subscript j denotes the j-th point. As shown in Fig. 3, the variation of deformations decreases as the training processes. One tentative interpretation is that MHR-Net searches solutions with a more diverse set of hypotheses in the early stage of training, and produces less diverse (or more confident) hypotheses after convergence. Moreover, the variation also differs between points. We observe the largest variations at Left/Right Hand/Elbow/Foot, which is consistent with the common sense that these are most flexible parts of body.

Accuracy of other hypotheses. To verify the accuracy of non-best hypotheses, we evaluate the MPJPE of the worst (largest reprojection error) hypothesis on Human3.6M. Compared to the best hypothesis, the results of worst hypothesis show a decline of -0.5 and -0.1 MPJPE at epoch 10 and epoch 50 with $N_m = 50$. This verifies that other hypotheses from MHR-Net are also accurate. Visualizing largely-deformed regions of meshes. Another advantage of MHR-Net is that we can use $||S_i^D||$ as an indicator of the degree of deformations. We visualize the largely-deformed regions of reconstructed dense point clouds from SURREAL dataset, as shown in Fig. 5. The visualization clearly illustrates the largely deformed parts of the body, which is helpful for better understanding of non-rigid reconstruction.



Fig. 4. Visualization of hypotheses $(N_m = 3)$ on Human3.6M test set.



Fig. 5. Visualization of reconstructions on SURREAL. Points with large deformations are marked in red. Best viewed in color.

6 Conclusion

We propose MHR-Net, a novel approach for reconstructing non-rigid shapes from 2D observations. To the best of our knowledge, MHR-Net is the first method produces multiple high-quality hypotheses for NRSfM. With the proposed DBSD scheme and optimization strategy, MHR-Net is capable to generate multiple reconstructions and achieves state-of-the-art shape recovery accuracy. Furthermore, we introduce a Procrustean Residual Loss for enhancing performance.

Acknowledgement

This work was partially done when Haitian Zeng interned at Baidu Research. This work was partially supported by ARC DP200100938. We thank Dr. Sungheon Park for sharing the SURREAL dataset. We thank all reviewers and area chairs for their valuable feedback.

References

- 1. Agudo, A.: Unsupervised 3d reconstruction and grouping of rigid and non-rigid categories. IEEE TPAMI (2020)
- Agudo, A., Agapito, L., Calvo, B., Montiel, J.M.M.: Good vibrations: A modal analysis approach for sequential non-rigid structure from motion. In: CVPR (2014)
- 3. Agudo, A., Moreno-Noguer, F.: Learning shape, motion and elastic models in force space. In: ICCV (2015)
- 4. Agudo, A., Moreno-Noguer, F.: Recovering pose and 3d deformable shape from multi-instance image ensembles. In: ACCV (2016)
- 5. Agudo, A., Moreno-Noguer, F.: DUST: dual union of spatio-temporal subspaces for monocular multiple object 3d reconstruction. In: CVPR. pp. 1513–1521 (2017)
- Agudo, A., Moreno-Noguer, F., Calvo, B., Montiel, J.M.M.: Sequential non-rigid structure from motion using physical priors. IEEE TPAMI (2016)
- Akhter, I., Sheikh, Y., Khan, S., Kanade, T.: Trajectory space: A dual representation for nonrigid structure from motion. IEEE TPAMI 33(7), 1442–1456 (2011)
- Bartoli, A., Gay-Bellile, V., Castellani, U., Peyras, J., Olsen, S.I., Sayd, P.: Coarseto-fine low-rank structure-from-motion. In: CVPR (2008)
- Bregler, C., Hertzmann, A., Biermann, H.: Recovering non-rigid 3d shape from image streams. In: CVPR. pp. 2690–2696 (2000)
- Bue, A.D.: Adaptive non-rigid registration and structure from motion from image trajectories. IJCV 103(2), 226–239 (2013)
- Bue, A.D., Smeraldi, F., Agapito, L.: Non-rigid structure from motion using ranklet-based tracking and non-linear optimization. Image Vis. Comput. 25(3), 297–310 (2007)
- 12. Bulat, A., Tzimiropoulos, G.: How far are we from solving the 2d & 3d face alignment problem? (and a dataset of 230, 000 3d facial landmarks). In: ICCV (2017)
- 13. Cha, G., Lee, M., Oh, S.: Reconstruct as far as you can: Consensus of non-rigid reconstruction from feasible regions. IEEE TPAMI pp. 1–1 (2019)
- Chenxin, X., Siheng, C., Maosen, L., Ya, Z.: Invariant teacher and equivariant student for unsupervised 3d human pose estimation. In: AAAI. pp. 3013–3021 (2021)
- Chhatkuli, A., Pizarro, D., Collins, T., Bartoli, A.: Inextensible non-rigid shapefrom-motion by second-order cone programming. In: CVPR. pp. 1719–1727 (2016)
- Dai, Y., Li, H.: Rank minimization or nuclear-norm minimization: Are we solving the right problem? In: International Conference on Digital Image Computing: Techniques and Applications (2014)
- 17. Dai, Y., Li, H., He, M.: A simple prior-free method for non-rigid structure-frommotion factorization. In: CVPR (2012)
- Dai, Y., Li, H., He, M.: A simple prior-free method for non-rigid structure-frommotion factorization. IJCV 107(2), 101–122 (2014)
- Fayad, J., Agapito, L., Bue, A.D.: Piecewise quadratic reconstruction of non-rigid surfaces from monocular sequences. In: ECCV. pp. 297–310 (2010)
- Fragkiadaki, K., Salas, M., Arbeláez, P.A., Malik, J.: Grouping-based low-rank trajectory completion and 3d reconstruction. In: NeurIPS. pp. 55–63 (2014)
- Gotardo, P.F.U., Martinez, A.M.: Computing smooth time trajectories for camera and deformable shape in structure from motion with occlusion. IEEE TPAMI 33(10), 2051–2065 (2011)
- Gotardo, P.F.U., Martinez, A.M.: Kernel non-rigid structure from motion. In: ICCV. pp. 802–809 (2011)

- 16 H. Zeng et al.
- Gotardo, P.F.U., Martínez, A.M.: Non-rigid structure from motion with complementary rank-3 spaces. In: CVPR (2011)
- 24. Iglesias, J.P., Olsson, C., Ornhag, M.V.: Accurate optimization of weighted nuclear norm for non-rigid structure from motion. In: CVPR (2020)
- Ijaz, A., Yaser, S., Sohaib, K.: In defense of orthonormality constraints for nonrigid structure from motion. In: CVPR (2009)
- Ionescu, C., Papava, D., Olaru, V., Sminchisescu, C.: Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. IEEE TPAMI 36(7), 1325–1339 (2014)
- 27. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: ICLR (2014)
- Kong, C., Lucey, S.: Deep non-rigid structure from motion with missing data. IEEE TPAMI pp. 1–1 (2020)
- Kong, C., Zhu, R., Kiani, H., Lucey, S.: Structure from category: A generic and prior-less approach. In: 3DV. pp. 296–304 (2016)
- Kudo, Y., Ogaki, K., Matsui, Y., Odagiri, Y.: Unsupervised adversarial learning of 3d human pose from 2d joint locations. CoRR abs/1803.08244 (2018)
- Kumar, S.: Jumping manifolds: Geometry aware dense non-rigid structure from motion. In: CVPR. pp. 5346–5355 (2019)
- Kumar, S., Cherian, A., Dai, Y., Li, H.: Scalable dense non-rigid structure-frommotion: A grassmannian perspective. In: CVPR. pp. 254–263 (2018)
- Kumar, S., Dai, Y., Li, H.: Spatio-temporal union of subspaces for multi-body non-rigid structure-from-motion. Pattern Recognition 71, 428–443 (2017)
- Lee, M., Cho, J., Choi, C.H., Oh, S.: Procrustean normal distribution for non-rigid structure from motion. IEEE TPAMI 39(7), 1388–1400 (2017)
- Lee, S., Purushwalkam, S., Cogswell, M., Ranjan, V., Crandall, D.J., Batra, D.: Stochastic multiple choice learning for training diverse deep ensembles. In: NeurIPS (2016)
- 36. Li, C., Lee, G.H.: Generating multiple hypotheses for 3d human pose estimation with mixture density network. In: CVPR (2019)
- 37. Li, C., Lee, G.H.: Weakly supervised generative network for multiple 3d human pose hypotheses. In: BMVC (2020)
- Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: a skinned multi-person linear model. ACM TOG 34(6), 248:1–248:16 (2015)
- Martinez, J., Hossain, R., Romero, J., Little, J.J.: A simple yet effective baseline for 3d human pose estimation. In: ICCV (2017)
- 40. Novotny, D., Ravi, N., Graham, B., Neverova, N., Vedaldi, A.: C3DPO: canonical 3d pose networks for non-rigid structure from motion. In: ICCV (2019)
- 41. Paladini, M., Bartoli, A., Agapito, L.: Sequential non-rigid structure-from-motion with the 3d-implicit low-rank shape model. In: ECCV (2010)
- Paladini, M., Bue, A.D., Stosic, M., Dodig, M., Xavier, J., Agapito, L.: Factorization for non-rigid and articulated structure using metric projections. In: CVPR (2009)
- 43. Parashar, S., Pizarro, D., Bartoli, A.: Isometric non-rigid shape-from-motion with riemannian geometry solved in linear time. IEEE TPAMI **40**(10), 2442–2454 (2018)
- 44. Parashar, S., Salzmann, M., Fua, P.: Local non-rigid structure-from-motion from diffeomorphic mappings. In: CVPR (2020)
- 45. Park, H.S., Shiratori, T., Matthews, I.A., Sheikh, Y.: 3d reconstruction of a moving point from a series of 2d projections. In: ECCV (2010)
- 46. Park, S., Lee, M., Kwak, N.: Procrustean regression: A flexible alignment-based framework for nonrigid structure estimation. IEEE TIP 27(1), 249–264 (2018)

- 47. Park, S., Lee, M., Kwak, N.: Procrustean regression networks: Learning 3d structure of non-rigid objects from 2d annotations. In: ECCV. pp. 1–18 (2020)
- 48. Pavllo, D., Feichtenhofer, C., Grangier, D., Auli, M.: 3d human pose estimation in video with temporal convolutions and semi-supervised training. In: CVPR (2019)
- 49. Probst, T., Paudel, D.P., Chhatkuli, A., Gool, L.V.: Incremental non-rigid structure-from-motion with unknown focal length. In: ECCV (2018)
- 50. Rabaud, V.C., Belongie, S.J.: Re-thinking non-rigid structure from motion. In: CVPR (2008)
- 51. Rupprecht, C., Laina, I., DiPietro, R.S., Baust, M.: Learning in an uncertain world: Representing ambiguity through multiple hypotheses. In: ICCV (2017)
- 52. Russell, C., Fayad, J., Agapito, L.: Energy based multiple model fitting for nonrigid structure from motion. In: CVPR (2011)
- Salzmann, M., Pilet, J., Ilic, S., Fua, P.: Surface deformation models for nonrigid 3d shape recovery. IEEE TPAMI 29(8), 1481–1487 (2007)
- 54. Schonemann, P.H.: A generalized solution of the orthogonal procrustes problem. Psychometrika (1966)
- 55. Sharma, S., Varigonda, P.T., Bindal, P., Sharma, A., Jain, A.: Monocular 3d human pose estimation by generation and ordinal ranking. In: ICCV (2019)
- Shen, J., Zafeiriou, S., Chrysos, G.G., Kossaifi, J., Tzimiropoulos, G., Pantic, M.: The first facial landmark tracking in-the-wild challenge: Benchmark and results. In: ICCVW (2015)
- Sidhu, V., Tretschk, E., Golyanik, V., Agudo, A., Theobalt, C.: Neural dense nonrigid structure from motion with latent space constraints. In: ECCV. pp. 204–222 (2020)
- Sminchisescu, C., Triggs, B.: Covariance scaled sampling for monocular 3d body tracking. In: CVPR (2001)
- 59. Sohn, K., Lee, H., Yan, X.: Learning structured output representation using deep conditional generative models. In: NeurIPS (2015)
- Taylor, J., Jepson, A.D., Kutulakos, K.N.: Non-rigid structure from locally-rigid motion. In: CVPR. pp. 2761–2768 (2010)
- Torresani, L., Hertzmann, A., Bregler, C.: Nonrigid structure-from-motion: Estimating shape and motion with hierarchical priors. IEEE TPAMI 30(5), 878–892 (2008)
- Valmadre, J., Lucey, S.: Deterministic 3d human pose estimation using rigid structure. In: ECCV (2010)
- Varol, G., Romero, J., Martin, X., Mahmood, N., Black, M.J., Laptev, I., Schmid, C.: Learning from synthetic humans. In: CVPR (2017)
- Vicente, S., Agapito, L.: Soft inextensibility constraints for template-free non-rigid reconstruction. In: ECCV. pp. 426–440 (2012)
- Wang, C., Kong, C., Lucey, S.: Distill knowledge from nrsfm for weakly supervised 3d pose learning. In: ICCV (2019)
- Wang, C., Lucey, S.: PAUL: procrustean autoencoder for unsupervised lifting. In: CVPR (2021)
- 67. Wehrbein, T., Rudolph, M., Rosenhahn, B., Wandt, B.: Probabilistic monocular 3d human pose estimation with normalizing flows. In: ICCV (2021)
- Xiao, J., Chai, J., Kanade, T.: A closed-form solution to non-rigid shape and motion recovery. IJCV 67(2), 233–246 (2006)
- 69. Zeng, H., Dai, Y., Yu, X., Wang, X., Yang, Y.: PR-RRN: pairwise-regularized residual-recursive networks for non-rigid structure-from-motion. In: ICCV (2021)
- Zhu, Y., Huang, D., Torre, F.D.L., Lucey, S.: Complex non-rigid motion 3d reconstruction by union of subspaces. In: CVPR. pp. 1542–1549 (2014)