

Unsupervised 3D Pose Estimation with Non-Rigid Structure-from-Motion Modeling

Haorui Ji¹ Hui Deng² Yuchao Dai² Hongdong Li¹

¹The Australian National University

²School of Electronics and Information, Northwestern Polytechnical University

{haorui.ji, hongdong.li}@anu.edu.au, denghui986@foxmail.com, daiyuchao@gmail.com

Abstract

Most existing 3D human pose estimation work rely heavily on the powerful memory capability of networks to obtain suitable 2D-3D mappings from the training data. Few works have studied the modeling of human posture deformation in motion. In this paper, we propose a new modeling method for human pose deformations and design an accompanying diffusion-based motion prior. Inspired by the field of non-rigid structure-from-motion, we divide the task of reconstructing 3D human skeletons in motion into the estimation of a 3D reference skeleton, and a frame-by-frame skeleton deformation. A mixed spatial-temporal **NRSfMformer** is used to simultaneously estimate the 3D reference skeleton and the skeleton deformation of each frame from 2D observations sequence, and then sum them up to obtain the pose of each frame. Subsequently, a loss term based on the diffusion model is used to ensure that the pipeline learns the correct prior motion knowledge. Finally, we have evaluated our proposed method on mainstream datasets and obtained superior results outperforming the state-of-the-art.

1. Introduction

Multi-frame 3D Human pose estimation (HPE) aims to recover human 3D poses from 2D image sequences. It has a wide range of applications such as augmented reality, robotics, sports analysis and film industry, *etc.* [10, 14, 15, 35]. 3D HPE can be viewed as a special instance of NRSfM (non-rigid structure-from-motion) and thus its solution may be benefited from the advances in the field of structure-from-motion. Fig. 1 demonstrates a typical pipeline for solving the 3D HPE problem with NRSfM. In this paper, we will use the term "pose", "skeleton" or human "shape" interchangeably without loss of generality.

Non-rigid structure-from-motion is significantly more challenging than its rigid counterpart [23] because human bodies can go through arbitrary rigid movements and non-

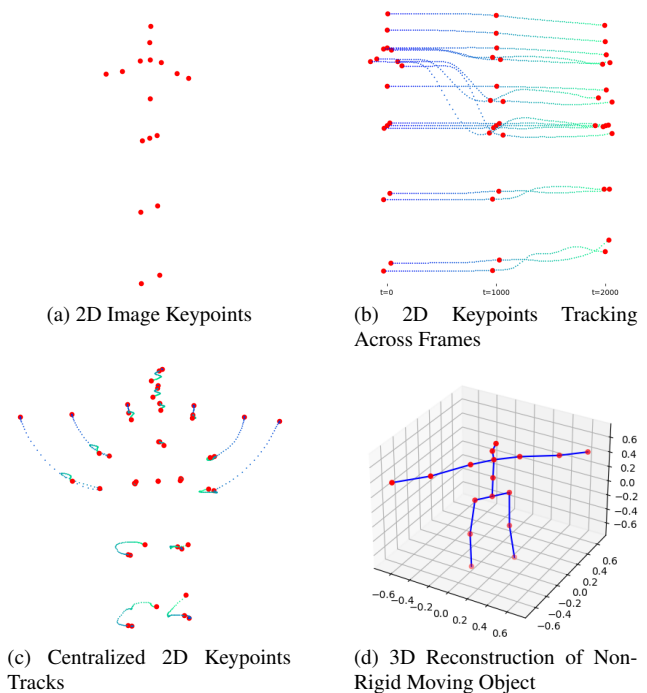


Figure 1. A visual illustration of basic pipeline setup for solving 3D HPE problem with NRSfM factorization. Given a sequence of 2D keypoints to form a track that describes the motion of a non-rigidly deforming subject, NRSfM aims to reconstruct the structure and pose of this subject in 3D space at each timestamp.

rigid articulations across different viewpoints, therefore introducing multiple ambiguities: (i) depth ambiguity: single 2D measurement can correspond to multiple 3D poses and the variance across viewpoints cannot be modeled by mere multi-view geometry. (ii) motion ambiguity: relative rigid motion between the human body and the camera is entangled so that absolute human pose and motion decomposition cannot be achieved. These ambiguities make NRSfM essentially an underconstrained problem and additional constraints are required to make it solvable and computationally tractable.

Accordingly, NRSfM methods focus on imposing priors in shapes and trajectories to ensure the uniqueness and correctness of the solution. On one hand, traditional optimization-based methods use hand-crafted constraints to reveal geometric properties of the deforming shapes to better regularize the optimization process. On the other hand, in recent years learning-based NRSfM methods gradually boost the performance. They focus more on using the expressiveness of neural networks to directly learn a reasonable shape and motion prior from massive amounts of data.

In this paper, we aim to address the 3D HPE problem in an NRSfM manner. While most existing 3D HPE methods rely heavily on the powerful memory capability of the neural networks to obtain suitable 2D-3D lifting mappings from the training data, few have studied the modeling of human posture deformation in motion. Inspired by the NRSfM problem, we design a geometric-aware 3D HPE pipeline, called RMNRD (standing for: Rigid-Motion-Non-Rigid-Deformation), that divides the human pose sequence lifting problem into the estimation of a 3D sequential reference skeleton, and a frame-by-frame skeleton deformation. In addition, to alleviate the inherent depth ambiguity, we propose to utilize a diffusion-based pose generator and leverage the prior knowledge encoded within to enforce our network to predict reasonable 3D poses conditioned on 2D observations but without any 3D supervision. In summary, our contributions are as follows:

- We propose RMNRD, an NRSfM-style strategy to model the 3D human pose estimation by dividing the sequence lifting problem into the estimation of sequential reference skeleton, skeleton deformation of each frame, and corresponding rigid transformation.
- We propose NRSfMFormer, a mixed spatial-temporal architecture that implements the RMNRD strategy, estimating the reference skeleton, skeleton deformation, and rigid transformation simultaneously.
- We utilize diffusion models trained on large human pose generation datasets to embed knowledge prior. It enforces our network to predict 3D poses that are consistent with the prior but without any ground truth 3D annotations used in training.
- Our method achieves comparably good performance across multiple benchmarks compared with State-Of-The-Art(SOTA) deep NRSfM methods, indicating the effectiveness of our method in dealing with sequential data.

2. Related Work

Unsupervised 3D HPE In contrast to supervised 3D HPE, unsupervised setting does not allow the usage of ground

truth 3D annotations in training, which makes the problem much more challenging due to the lack of data to learn the one-on-one 2D-3D mapping. Recent methods utilize multi-view information as a supervision signal and try to alleviate the depth ambiguity problem. Rhodin *et al.* [26] firstly propose to learn a geometry-aware skeleton structure constructed via multi-view images. Other methods [4, 34] bypass the reliance on multi-view inputs and construct self-consistency in 3D space in a loop-closure manner with random rotation perturbation. Other priors have also been considered such as adversarial training [4], kinematic human skeleton structure [17], or fixed scale bone-length ratios [39] to enforce the predicted human pose to be realistic and accurate.

Non-Rigid Structure-from-Motion Recovering 3D shape and motion from 2D keypoint sequences has a long research history. Many works have constructed different sequence modeling to solve this problem. For example, deforming shapes should span a low-rank space [1, 2, 7], deform smoothly [11], lie in a union of linear subspaces [44], or satisfy certain probability distribution [18, 19], etc. Other priors such as modeling object trajectories with DCT bases [1] and Dai *et al.*'s block-matrix rank-minimization [7] have also been explored. These optimization-based methods are sensitive to initialization, and the calculation of optimization-based methods is time-consuming. Recently, researchers resort to neural networks to learn shape and motion priors directly from large-scale data. Novotny *et al.* [22] and follow-up approaches [8, 36, 40, 41] introduced the transversal property to constrain shapes to a canonical view. In contrast, Park *et al.* [24] targeted NRSfM through shape alignment, constraining motion parameters so that sequential 3D structures will align to a reference frame. Furthermore, [16, 35] successfully incorporate perspective projection in learning-based NRSfM framework and is able to handle missing data scenarios. These two sets of methods inspire our Rigid-Motion-Non-Rigid-Deformation (RMNRD) modeling strategy, which takes entire sequences as input for more informed canonical view generation.

Diffusion-based generative models Diffusion models are a family of generative models that are inspired by non-equilibrium thermodynamics [12, 32]. They define a Markov chain of diffusion steps to slowly add random noise to data and then learn to reverse the diffusion process to construct desired data samples from the noise. For their simple designs and flexibility in describing the underlying structure in arbitrary data distribution, diffusion models have sparked remarkable progress in numerous areas, including image generation [21, 25, 27, 28], object detection [5], semantic segmentation [3, 38], etc. In addition to the success in the 2D image domain, more recent works have extended diffusion models to 3D content generation

with multi-modality inputs. For instance, [6, 31] apply diffusion models to 3D human pose estimation to resolve depth ambiguity, and [29, 33, 43] use them as priors to guide 3D motion generation. In this paper, we use the diffusion model to define a plausible distribution in 3D and 2D joint locations, resulting in a novel method for lifting.

3. Method

The overall pipeline is illustrated in Fig. 2. In the following sections, we'll get down to each individual building block. Specifically, we'll introduce the novel RMNRD strategy for modeling the human pose deformation in motion in Sec.3.2. Next, we'll clarify how to utilize diffusion priors to resolve shape ambiguities in Sec.3.3. In addition, we'll address our overall network design and geometric constraints used to facilitate the optimization process in Sec.3.4 and 3.5.

3.1. Problem Statement

We first review the NRSfM-styled 3D HPE problem formulation that will be used later to describe our approach. Consider P landmarks on a human skeleton tracked across F frames. Observation of these landmarks from each viewpoint is represented by a measurement matrix $\mathbf{W} \in \mathbb{R}^{2F \times P}$ where each element $\mathbf{W}_i \in \mathbb{R}^{2 \times P}, i = 1, \dots, F$ is generated through a camera projection Π process from the 3D point set $\mathbf{S}_i \in \mathbb{R}^{3 \times P}$, and corresponding rigid motion $(\mathbf{R}_i, \mathbf{T}_i) \in \text{SO}(3) \times \text{T}(3)$. Based on the linear model [2], each non-rigid shape \mathbf{S}_i can be described as a linear combination of a basis set $\{\mathbf{B}_{i1}, \mathbf{B}_{i2}, \dots, \mathbf{B}_{iK}\}$ with weight coefficients $\{c_{i1}, c_{i2}, \dots, c_{iK}\}$.

$$\mathbf{S}_i = \sum_{k=1}^K c_{ik} \mathbf{B}_{ik}, c_{ik} \in \mathbb{R}, \mathbf{B}_{ik} \in \mathbb{R}^{3 \times P} \quad (1)$$

Assuming the measurement matrix and shape structure are already centralized, rigid motion will reduce to pure rotation, eliminating the impact of translation term. If further consider an orthographic camera model, the camera projection will become a linear function, yielding the following formulation

$$\begin{aligned} \mathbf{W} &= \begin{bmatrix} x_{11} & \cdots & x_{1P} \\ \vdots & & \vdots \\ x_{F1} & \cdots & x_{FP} \end{bmatrix} = \begin{bmatrix} \mathbf{R}_1 \mathbf{S}_1 \\ \vdots \\ \mathbf{R}_F \mathbf{S}_F \end{bmatrix} \\ &= \begin{bmatrix} c_{11} \mathbf{R}_1 & \cdots & c_{1K} \mathbf{R}_1 \\ \vdots & & \vdots \\ c_{F1} \mathbf{R}_F & \cdots & c_{FK} \mathbf{R}_F \end{bmatrix} \begin{bmatrix} \mathbf{B}_1 \\ \vdots \\ \mathbf{B}_K \end{bmatrix} \\ &= \mathbf{R}(\mathbf{C} \otimes \mathbf{I}_3) \mathbf{B} = \mathbf{M} \mathbf{B} \end{aligned} \quad (2)$$

In Eqn.2, $\mathbf{R} = \text{blkdiag}(\mathbf{R}_1, \dots, \mathbf{R}_F) \in \mathbb{R}^{2F \times 3F}$ is the camera motion matrix, $\mathbf{B} \in \mathbb{R}^{3F \times P}$ and $\mathbf{C} \in \mathbb{R}^{F \times K}$ are

the matrix form of base shapes and corresponding weights coefficients. The goal of NRSfM task is to recover both motion and shape matrices \mathbf{M}, \mathbf{B} from the measurement matrix \mathbf{W} .

3.2. Disentangle Rigid Motion and Non-Rigid Deformation

One of the biggest challenges in NRSfM is the non-uniqueness in the factorization of the measurement matrix $\mathbf{W} = \mathbf{R} \mathbf{S}$ as in Eqn.2. Intuitively, we can interpret this ambiguity in the sense that when both the human body and camera are in motion, it's impossible to discern whether the observed rigid motion originates from the human or the camera solely through 2D measurements as $\mathbf{W} = \mathbf{R} \mathbf{S} = \mathbf{R}' \mathbf{R}_{\text{camera}} \mathbf{R}_{\text{human}} \mathbf{S}'$. This inherent ambiguity largely accounts for why non-rigid factorization is substantially more challenging than its rigid counterpart.

In this paper, we propose a way to alleviate this ambiguity by separating out the rigid motions. Specifically, we will keep the rigid pose fixed in a canonical form throughout the entire sequence. This means that we only consider variations of the human body due to non-rigid deformations, while rigid movements are not permitted. To accomplish this, it is essential that the network is capable of disregarding shape changes due to rigid motion or changes in viewpoint, and instead, concentrates on predicting shape variations resulting from non-rigid deformations within a consistent canonical view. This is facilitated by employing the concept of the transversal property introduced by [22]. The central idea is to construct a set $\mathcal{X}_0 \in \mathbb{R}^{P \times 3}$ that for any given pair $\mathbf{X}, \mathbf{X}' \in \mathcal{X}_0$, if they are related by a rotation $\mathbf{X}' = \mathbf{R} \mathbf{X}$, then $\mathbf{X} = \mathbf{X}'$.

In single-frame methods like [22] and its follow-up works, they explicitly learn a canonicalization mapping $\Psi : \mathbb{R}^{P \times 3} \rightarrow \mathbb{R}^{P \times 3}$ such that for any arbitrary perturbation $\mathbf{R} \in \text{SO}(3)$ and shape $\mathbf{S} \in \mathbb{R}^{P \times 3}$, it fulfills $\mathbf{S} = \Psi(\mathbf{R} \mathbf{S})$. In this way, the network is encouraged to maintain consistency in how reconstructions are executed. However, since their input is single-frame, they have no choice but to aggregate information of all shapes from the dataset. We argue that this way of constructing the transversal set is inappropriate because variations of shapes from separate sequences follow different patterns and should not be grouped together. The transversal set may not suitable for all poses in dataset, while it should only be maintained within the same input sequence. Therefore, to incorporate the transversal property correctly in sequential data, we don't directly apply Ψ to each individual frame. Instead, we turn to utilize Procrustes analysis, which enables alignment of a set of shapes as closely as possible to a common reference one through

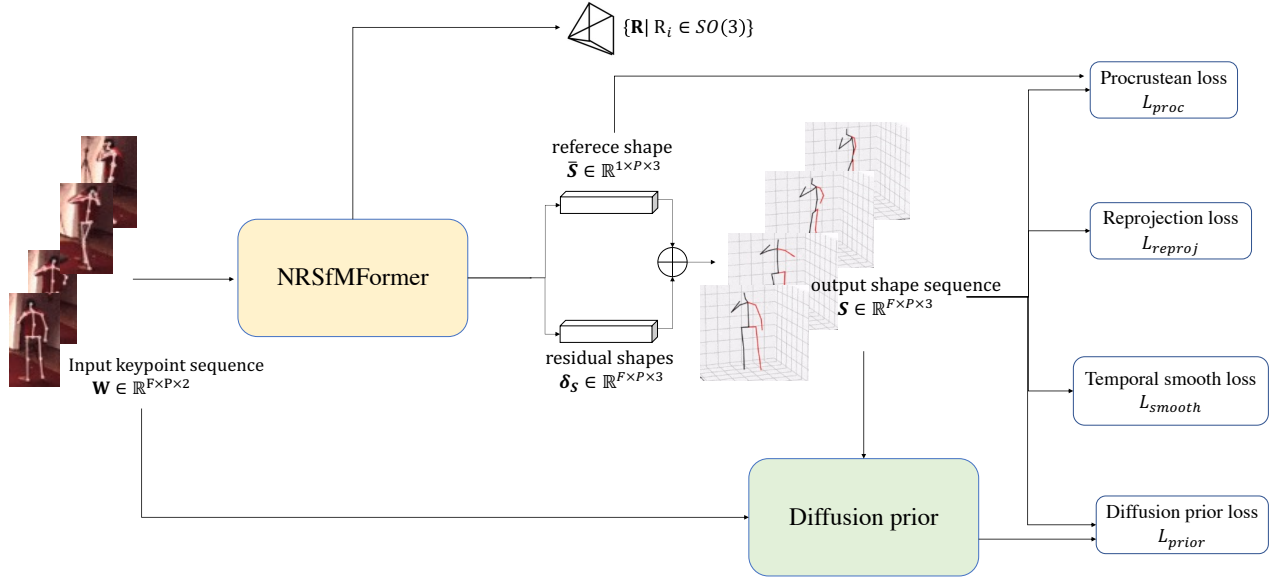


Figure 2. Overview of our framework and training pipeline. Two core components of the framework are the NRSfMFormer, a transformer-based context modeling network, and a pretrained 3D pose generation diffusion model. We apply the typical factorization paradigm but only to the sequence input and regularize it with additional constraints introduced in the loss functions in the figure. The diffusion model serves as a prior and is only used during training. In inference time, we only use the NRSfMFormer part.

rigid transformations.

$$\begin{aligned} \bar{\mathbf{S}}, \mathbf{R}_i &= \arg \min_{\bar{\mathbf{S}}, \mathbf{R}_i} \|\mathbf{R}_i \mathbf{S}_i - \bar{\mathbf{S}}\| \\ \text{s.t. } \mathbf{R}_i^T \mathbf{R}_i &= \mathbf{I} \end{aligned} \quad (3)$$

Here, $\bar{\mathbf{S}}$ represents the reference shape that all other shapes are to be aligned with, and \mathbf{R}_i denotes the corresponding aligning rotation matrices. [24] has demonstrated that the set of shapes aligned through Procrustes analysis constitutes a transversal set. Therefore when applying it to sequential data, we can maintain the transversal property within each sequence without being influenced by others, which satisfies our design.

Building on this concept, we introduce the Rigid-Motion-Non-Rigid-Deformation (RMNRD) modeling strategy. Given a sequence of shapes $\{\mathbf{S}_1, \mathbf{S}_2, \dots, \mathbf{S}_F\}$, we stabilize their rigid poses by superimposing all shapes onto a common reference shape using Procrustes analysis and construct each original shape \mathbf{S}_i as the sum of the common reference shape $\bar{\mathbf{S}}$ and its corresponding residual deforming shape $\delta_{\mathbf{S}_i}$:

$$\mathbf{S}_i = \bar{\mathbf{S}} + \delta_{\mathbf{S}_i}, i = 1, \dots, F \quad (4)$$

where $\bar{\mathbf{S}}$ can be computed from Eqn.3. This modeling strategy not only provides a way of resolving the factorization

ambiguity but also bolsters the shape representation capabilities by introducing hierarchical flexibility. Conceptually, the reference shape $\bar{\mathbf{S}}$ is able to capture the low-rank common structure among the sequential 3D shapes, while the residual deforming shape $\delta_{\mathbf{S}_i}$ accommodates diverse minor deformations with enhanced adaptability.

3.3. Measurement-Conditioned Pose Generation Priors

To tackle with the inherent ambiguity of the NRSfM pipeline, we propose to utilize a diffusion-based pose generator to serve as knowledge prior. Denoising Diffusion Probabilistic Models (DDPM) are a type of generative models that learn the underlying distribution over training data samples. We build our method upon the recent diffusion-based 3D human pose generation model [31] for its efficiency in generating high-quality human poses with 2D observation guidance. The whole framework consists of a forward diffusion process and a reverse sampling process.

Forward diffusion process perturbs the ground truth 3D pose y_0 drawn from an unknown data distribution $p_{data}(y)$ by gradually adding independent Gaussian noise $\epsilon \sim N(0, I)$ and finally corrupts to a noise pose y_t after t step that follows the predefined prior distribution, e.g. isotropic Gaussian distribution.

Reverse diffusion process samples from the predefined

prior distribution and reverse the forward diffusion process to get a plausible pose that satisfies the original data distribution $p_{data}(y)$. The whole diffusion process is trained by minimizing the difference between original pose and pose generated from reverse sampling process

$$\mathbb{E}_{y_0 \sim p_{data}(y), t, \epsilon \sim N(0, I)} [\|y_0 - g_\omega(y_t, c, t)\|_2^2] \quad (5)$$

where t is the diffusion timestep, y_0 is an arbitrary 3D pose drawn from the data distribution and y_t is the noised pose to time t with gaussian noise ϵ . g_ω is the denoising network with parameters ω . The diffusion model also takes a conditioning input c which serves as a guidance in the reverse process. In our problem setup the condition is the 2D measurements.

Once the diffusion model is well-trained, we'll fix the parameters ω and use it to instruct the NRSfM optimization process $\mathbf{M}, \mathbf{S} = f_\theta(\mathbf{W})$. We can optimize for the predicted shape \mathbf{S} to follow the pose distribution priors conditioned on \mathbf{W} and can also backpropagate gradients to the NRSfM-Former and thus get a stochastic gradient descent on θ .

3.4. NRSfMFormer

We implement the RMNRD framework with NRSfM-Former, a Transformer encoder-based architecture. Its proficiency in modeling dependencies between inputs, regardless of time intervals, makes it highly suitable for leveraging the structure of non-rigid scenes along the temporal dimension.

The NRSfMFormer is composed of a mixed spatio-temporal encoder for feature extraction and three subnetwork modules that parameterize the components as discussed in previous sections.

The overview of the NRSfMFormer is depicted in Fig.3. The network takes the 2D sequence measurement $\mathbf{W} \in \mathbb{R}^{F \times P \times 2}$ with F frames and P keypoints as input. We first project the input sequence to a high-dimensional feature $\mathbf{Z} \in \mathbb{R}^{F \times P \times D}$, which can be interpreted at two levels: it consists of F frame features $\mathbf{Z}_i \in \mathbb{R}^{P \times D}, i = 1, \dots, F$ and each frame feature consists of P keypoint features $\mathbf{Z}_{ij} \in \mathbb{R}^D, j = 1, \dots, P$. Then we apply a position embedding for both spatial and temporal domains to retain the positional information.

Different from previous work using the Transformer structure, NRSfMFormer adopts a novel way of generating a unified feature representation that incorporates information required by both human skeleton estimation and 6DoF rigid transformation estimation tasks. Specifically, NRSfM-Former considers reference shape as one of the properties of the pose sequence, and similarly, the rigid transformation between each camera position and the human skeleton is a property that is bound to the pose in each individual frame. Enlightened by ViT's design of *[class]* token [9], we attach two learnable embeddings to both frame-level and

keypoint-level features so that the final input to the encoder is:

$$\mathbf{Z} = [\mathbf{Z}_0, \mathbf{Z}_1, \mathbf{Z}_2, \dots, \mathbf{Z}_F] \in \mathbb{R}^{(F+1) \times (P+1) \times D} \quad (6)$$

The design intuition behind this is that the embedded token added on the frame feature serves as the semantic representation for the whole sequence ($\mathbf{Z}_0 = x_{seq}$). Similarly, the embedded token added on the keypoint feature serves as the representation for regressing the object's 6DoF rigid transformation ($\mathbf{Z}_{i0} = x_{pose}$).

Following the design scheme from [42], we stack L Mixed Spatio-Temporal Encoder block coherence feature extraction. It alternately uses Spatial Transformer Block (STB) and Temporal Transformer Block (TTB) to learn spatial and temporal correlation with multi-head self-attention. STB computes the attention between joints and aims to learn the body joint relations within each frame. On the other hand, TTB computes the attention between frames and focuses on learning the global temporal correlation of each joint across frames.

Passing through the encoder blocks, we now have feature $\mathbf{Z}^L = [Z_0^L, Z_1^L, Z_2^L, \dots, Z_F^L]$ where the superscript indicates which encoder layer we're extracting the feature from. On top of that, subsequent MLP-based modules **Rigid Motion Module**, **Deforming Shape Module** and **Reference Shape Module** are used to recover the geometric meaning of corresponding features from the latent space. Specifically, the rigid motion module and the reference shape module are applied to the extra learnable embedded tokens Z_{i0}^L and Z_0^L to regress 6DoF rigid transformation and reference shape. The deforming shape module is applied to the remaining individual representations from Transformer encoder output $\{Z_{i1}^L, Z_{i2}^L, \dots, Z_{iP}^L\}, i = 1, \dots, F$ to predict the residual shape corresponding to each 2D input measure with respect to the reference shape.

$$\begin{aligned} \bar{\mathbf{S}} &= \text{MLP}(Z_0^L) \\ \delta \mathbf{S}_i &= \text{MLP}(Z_i^L) \quad i = 1, \dots, F \\ \mathbf{R}_i &= \text{MLP}(Z_{i0}^L) \quad i = 1, \dots, F \end{aligned} \quad (7)$$

3.5. Loss Design

To solve the NRSfM problem defined in Sec.3.1, we propose to minimize a differentiable loss function with respect to motion parameters $\mathbf{R} = \{\mathbf{R}_1, \mathbf{R}_2, \dots, \mathbf{R}_F\}$ and human pose $\mathbf{S} = \{\mathbf{S}_1, \mathbf{S}_2, \dots, \mathbf{S}_F\}$. In the following section, we'll introduce each part separately.

Reprojection loss Reprojection loss is utilized to enforce that the network learns accurate motion and shape estimations in a self-supervised manner, ensuring alignment with 2D observations.

$$\mathcal{L}_{reproj} = \frac{1}{F} \sum_{i=1}^F \|\mathbf{IIR}_i \mathbf{S}_i - \mathbf{W}_i\|_F^2 \quad (8)$$

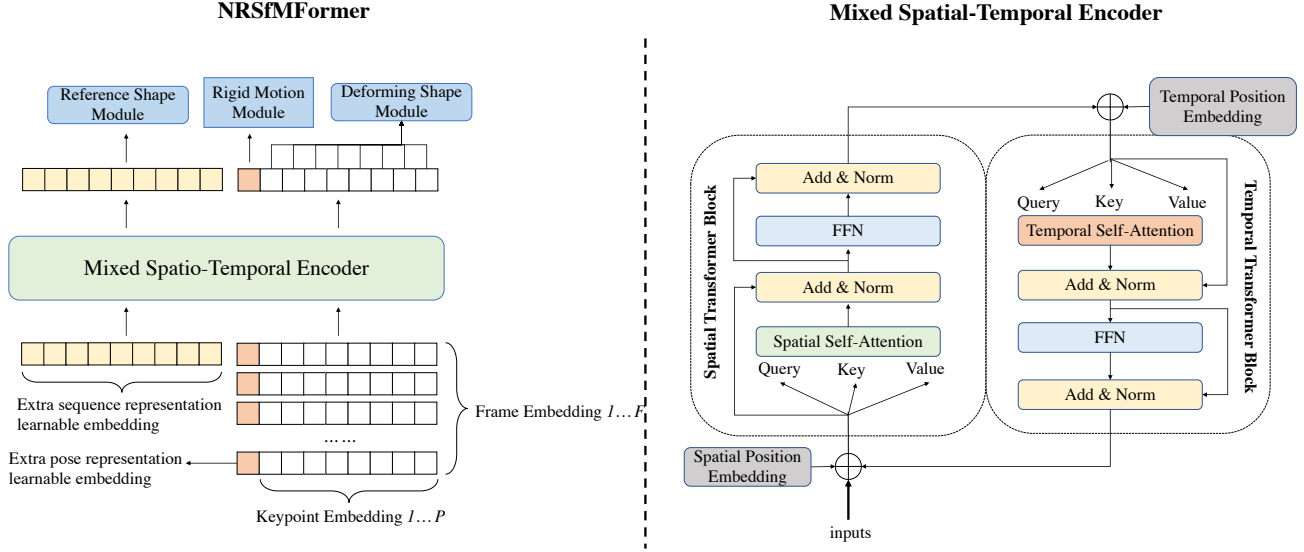


Figure 3. Overview of NRSfMFormer. Input is the vector consisting of features from each frame in a video sequence. After being linearly embedded and adding sequence position embeddings, we feed the resulting sequence of vectors to a mixed spatio-temporal transformer encoder. In order to perform reference shape regression, we’ll need a sequence representation independent from individual frame features. Therefore, we add an extra learnable sequence representation token to the input. Similarly, we add another learnable embedding to each keypoint feature, which serves as the representation for regressing the object’s 6DoF pose. The illustration of the Transformer encoder was inspired by [42]

Procrustean alignment loss According to the analysis in Sec.3.2, one crucial ingredient of the RMNRD modeling strategy is that all shapes in the sequence should align to a common reference shape through Procrustes analysis so that the transversal property can be retained.

Specifically, with our predicted shapes, reference shapes, and residual shapes \mathbf{S}_i^{pred} , $\bar{\mathbf{S}}^{pred}$, $\delta_{\mathbf{S}_i}^{pred}$, we can follow the procedure described in algorithm 1 for solving a Procrustean alignment problem to associate $\bar{\mathbf{S}}^{pred}$ and \mathbf{S}_i^{pred} as close as possible. This enables us to guarantee that by imposing constraints based on Procrustean alignment, for a sequence of shapes whose variations only exist in non-rigid deformations, we can always eliminate the ambiguity caused by rigid motions.

$$\mathcal{L}_{proc} = \frac{1}{F} \sum_{i=1}^F \|\mathbf{S}_i^{pred} - \mathbf{S}_i^*\| \quad (9)$$

where \mathbf{S}_i^* is the shape retrieved after Procrustean alignment.

Diffusion prior loss As discussed in Sec.3.3, the output of NRSfMFormer should satisfy certain prior distribution defined by pretrained diffusion model g_ω .

$$\begin{aligned} \mathcal{L}_{prior} &= \mathbb{E}_{t, \epsilon \sim N(0, I)} [\|y_0 - g_\omega(y_t, \mathbf{W}, t)\|_2^2] \\ y_0 &= f_\theta(\mathbf{W}) \end{aligned} \quad (10)$$

Algorithm 1 Procrustean Alignment

Input: Shape sequence $\{\mathbf{S}_i\}$, reference shape $\bar{\mathbf{S}}$

Output: Aligned shape sequence $\{\mathbf{S}_i^*\}$

Centralize each shape as $\mathbf{S}_i \leftarrow \mathbf{S}_i - \mathbf{S}_i \mathbf{1} \mathbf{1}^T$

Centralize reference shape as $\bar{\mathbf{S}} \leftarrow \bar{\mathbf{S}} - \bar{\mathbf{S}} \mathbf{1} \mathbf{1}^T$

Update the rotations that can align each shape to the reference shape $\mathbf{R}_i = \mathbf{U}_i \mathbf{V}_i^T$, where $\bar{\mathbf{S}}^T \mathbf{S}_i = \mathbf{U}_i \mathbf{A} \mathbf{V}_i^T$ is obtained through singular value decomposition.

Update aligned shapes as $\mathbf{S}_i^* \leftarrow \mathbf{R}_i \mathbf{S}_i$

This term is similar to Eqn.5, despite that y_0 is not sampled arbitrarily from the data distribution but is the output of our NRSfMFormer. In addition, we use the 2D measurement matrix \mathbf{W} as the condition during the diffusion process.

Temporal smoothness loss As we’re using sequence data as input, we propose to consider the temporal smoothness assumption as the object moves and deforms continuously in the physical world.

$$\mathcal{L}_{smooth} = \sum_{i=2}^F \lambda_{\mathbf{R}} \|\mathbf{R}_i - \mathbf{R}_{i-1}\|^2 + \lambda_{\mathbf{S}} \|\mathbf{S}_i - \mathbf{S}_{i-1}\|^2 \quad (11)$$

Combining all loss functions introduced above, our total

loss is as follows:

$$\mathcal{L} = \beta_1 \mathcal{L}_{reproj} + \beta_2 \mathcal{L}_{proc} + \beta_3 \mathcal{L}_{prior} + \beta_4 \mathcal{L}_{smooth} \quad (12)$$

where $\beta_1, \beta_2, \beta_3, \beta_4$ are trade-off weights hyperparameters.

4. Experiments

In this section, we’ll introduce our experimental results evaluated on different large-scale human pose estimation datasets. We first introduce the datasets and evaluation metrics we are using and then the implementation details. Next, qualitative experimental results and comparisons with other methods are reported. Finally, we analyze the influence of different components of the proposed method in detail through ablation studies.

4.1. Datasets and Evaluation Protocols

Human3.6M [13] is a large dataset containing various types of human motion sequences annotated with 3D ground truth extracted using motion capture systems. The 3D human skeleton in Human3.6M is modeled as a 17-joint skeleton, subjects S1, S5, S6, S7, and S8 are applied during training, while the S9 and S11 are used for testing.

MPI-INF-3DHP [20] is a recently popular dataset that features in more complex scenarios including indoor, green screen, and outdoor scenes. The training set contains 8 activities performed by 8 actors, and the test set covers 7 activities. Following [30], we use the valid frames provided by the official for testing.

Evaluation metrics We report the mean per joint position error (MPJPE) in millimeters. Since an unsupervised setting does not contain metric data we scale the reconstructed 3D pose to match ground truth, commonly known as N-MPJPE. Besides, Procrustes alignment (including scaling), known as PA-MPJPE is also reported. For 3DHP datasets, we also report the Percentage of Correct Keypoints (PCK) thresholded at 150mm and the Area Under the Curve (AUC).

4.2. Implementation Details

Our model is implemented in a two-stage training scheme. In the first stage, we train the DDPM model for 40 epochs with AdamW optimizer and a learning rate of $5e^{-5}$. In the second stage, we freeze the diffusion model and only train the NRSfMFormer part. We adopt a similar architecture as that of Zhang et al. [42], which consists of 4 mixed spatio-temporal encoder blocks. All the regression heads are single-layer MLP. All parameters are initialized randomly from the zero-mean Gaussian distribution with $\sigma = 0.001$. During training, the input sequence length is 27 and we use AdamW optimizer with a mini-batch size of 128 to update the parameters. We trained 50 epochs in total with $5e^{-5}$ initial learning rate, and the learning rate will

Unsupervised 3D HPE	N-MPJPE ↓	PA-MPJPE ↓
Chen <i>et al.</i> [4]	-	51
Yu <i>et al.</i> [39]	85.3	42.0
Wang <i>et al.</i> [34]	64.0	36.7
NRSfMFormer	<u>41.6</u>	<u>33.1</u>
NRSfM	N-MPJPE ↓	PA-MPJPE ↓
Park <i>et al.</i> [24]	86.4	-
Wang <i>et al.</i> [36]	88.3	-
Novotny <i>et al.</i> [22]	95.6	-
Xu <i>et al.</i> [37]	77.2	-
Deng <i>et al.</i> [8]	79.8	-
Zeng <i>et al.</i> [41]	72.6	-
NRSfMFormer	<u>41.6</u>	<u>33.1</u>

Table 1. 3D HPE quantitative results on Human3.6M dataset. All results are reported in millimeters (mm). The diffusion prior model is trained on Human3.6M and GT 2D keypoints inputs are used.

reduce by a factor of 0.99 after each epoch. Our experiments are performed on one NVIDIA 1080Ti GPU, and the training takes about 20 hours to complete.

4.3. Quantitative & Qualitative results

Results on Human3.6M We compare our model with unsupervised 3D HPE methods as well as NRSfM methods that are evaluated on the Human3.6M dataset. All experiments use the GT 2D keypoints as input and our diffusion prior is pretrained on the training set of Human3.6M. As shown in Table 1, we obtain 41.6.4mm in N-MPJPE and 33.1mm in PA-MPJPE, which outperforms any existing SOTA methods in both unsupervised 3D HPE and NRSfM by a large margin. Be advised that most of the comparing methods use single-frame data as input to predict the corresponding 3D pose. In contrast, our method explicitly models the underlying sequence structure with the temporal constraint, taking advantage of the information ignored by the other methods and thus achieving improved results. We also show some qualitative results in Fig.4 comparing with the SOTA NRSfM method [41]. From the figure we can observe that our model has more accurate predictions especially on edge joints of the human body, like hands, feet, etc. This contributes to the performance improvements of our model.

Results on 3DHP. Following previous works [4, 34, 39], we use ground truth 2D poses as input. As shown in Table2, we achieve 46.3mm in PA-MPJPE, which outperforms existing SOTA performance by 14%. Although there exists some gap in terms of PCK, our method prevails in AUC by a margin up to nearly 40%, which still demonstrates the effectiveness of our method.




Figure 4. Visualization of several methods on Human3.6M with ground truth 2D keypoint as input.

Unsupervised 3D HPE	PA-MPJPE ↓	PCK ↑	AUC ↑
Chen <i>et al.</i> [4]	-	71.7	36.3
Yu <i>et al.</i> [39]	-	86.2	51.7
Wang <i>et al.</i> [34]	54.0	86.0	50.1
NRSfMFormer	<u>46.3</u>	73.4	<u>68.1</u>

Table 2. Our results (14-joint) on MPI-INF-3DHP with metrics as in Wang *et al.* [34]. All results are reported in millimeters (mm). GT 2D keypoints inputs are used.

4.4. Ablation Study

In this section, we study the impact of different building blocks in our model. All results are evaluated on Human3.6M dataset and with the same configuration and training scheme.

We set up several ablated models to study the effect of RMNRD modeling strategy, temporal consistency, procrustean alignment, and diffusion prior. As shown in Table 3, we observe the degradation of performance by removing any of the components, indicating the effectiveness of our general framework design. Among all the compositions, we can see that the most crucial one is the diffusion model. Without the pretrained prior, performance will drop dramatically and errors will even double. This reveals the fact that resolving the depth ambiguity and predicting a plausible 3D pose should be the main concern when solving an NRSfM problem. Other constraints like motion ambiguities

Ablations	N-MPJPE ↓	PA-MPJPE ↓
w/o RMNRD strategy	46.7	36.5
w/o temporal consistency	42.3	33.8
w/o Procrustean alignment	44.6	36.7
w/o diffusion prior	77.3	54.4
NRSfMFormer	<u>41.6</u>	<u>33.1</u>

Table 3. Ablation studies performed on the Human3.6M dataset and temporal information will also have respective contributions but the impacts are not as significant.

5. Conclusion

In this paper, we aim to solve the sequential unsupervised 3D human pose estimation problem with Non-Rigid Structure-from-Motion modeling. Inspired by the field of non-rigid structure-from-motion, we design a geometric-aware pipeline that uses a new modeling scheme for human pose deformations, which divides the task of 3D HPE into the estimation of the 3D sequence reference skeleton and a frame-by-frame skeleton deformation. We also propose to use a diffusion-based shape prior to resolve the inherent depth ambiguity. One limitation of this work is that NRSfM is essentially a universal method that can be applied to diverse scenarios. However, in this work we’re only focusing on its application to the human body and not fully utilizing the potential of either NRSfM or diffusion prior. In future work and will expand our test subjects to more general cases.

References

- [1] Ijaz Akhter, Yaser Sheikh, Sohaib Khan, and Takeo Kanade. Nonrigid structure from motion in trajectory space. *Advances in neural information processing systems*, 21, 2008. [2](#)
- [2] Christoph Bregler, Aaron Hertzmann, and Henning Biermann. Recovering non-rigid 3d shape from image streams. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition. CVPR 2000 (Cat. No. PR00662)*, volume 2, pages 690–696. IEEE, 2000. [2](#), [3](#)
- [3] Emmanuel Asiedu Brempong, Simon Kornblith, Ting Chen, Niki Parmar, Matthias Minderer, and Mohammad Norouzi. Denoising pretraining for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4175–4186, 2022. [2](#)
- [4] Ching-Hang Chen, Amrith Tyagi, Amit Agrawal, Dylan Drover, Rohith Mv, Stefan Stojanov, and James M Rehg. Unsupervised 3d pose estimation with geometric self-supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5714–5724, 2019. [2](#), [7](#), [8](#)
- [5] Shoufa Chen, Peize Sun, Yibing Song, and Ping Luo. Diffusionet: Diffusion model for object detection. *arXiv preprint arXiv:2211.09788*, 2022. [2](#)
- [6] Hai Ci, Mingdong Wu, Wentao Zhu, Xiaoxuan Ma, Hao Dong, Fangwei Zhong, and Yizhou Wang. Gfpose: Learning 3d human pose prior with gradient fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4800–4810, 2023. [3](#)
- [7] Yuchao Dai, Hongdong Li, and Mingyi He. A simple prior-free method for non-rigid structure-from-motion factorization. *International Journal of Computer Vision*, 107:101–122, 2014. [2](#)
- [8] Hui Deng, Tong Zhang, Yuchao Dai, Jiawei Shi, Yiran Zhong, and Hongdong Li. Deep non-rigid structure-from-motion: A sequence-to-sequence translation perspective. *arXiv preprint arXiv:2204.04730*, 2022. [2](#), [7](#)
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [5](#)
- [10] Vladislav Golyanik, Aman S Mathur, and Didier Stricker. Nrsfm-flow: Recovering non-rigid scene flow from monocular image sequences. In *BMVC*, 2016. [1](#)
- [11] Paulo FU Gotardo and Aleix M Martinez. Computing smooth time trajectories for camera and deformable shape in structure from motion with occlusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(10):2051–2065, 2011. [2](#)
- [12] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. [2](#)
- [13] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence*, 36(7):1325–1339, 2013. [7](#)
- [14] Zhongyu Jiang, Haorui Ji, Samuel Menaker, and Jenq-Neng Hwang. Golfpose: Golf swing analyses with a monocular camera based human pose estimation. In *2022 IEEE International Conference on Multimedia and Expo Workshops (ICMEW)*, pages 1–6. IEEE, 2022. [1](#)
- [15] Muhammed Kocabas, Nikos Athanasiou, and Michael J Black. Vibe: Video inference for human body pose and shape estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5253–5263, 2020. [1](#)
- [16] Chen Kong and Simon Lucey. Deep non-rigid structure from motion with missing data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(12):4365–4377, 2020. [2](#)
- [17] Jogendra Nath Kundu, Siddharth Seth, MV Rahul, Mugalodi Rakesh, Venkatesh Babu Radhakrishnan, and Anirban Chakraborty. Kinematic-structure-preserved representation for unsupervised 3d human pose estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11312–11319, 2020. [2](#)
- [18] Minsik Lee, Jungchan Cho, Chong-Ho Choi, and Songhwai Oh. Procrustean normal distribution for non-rigid structure from motion. In *Proceedings of the IEEE Conference on computer vision and pattern recognition*, pages 1280–1287, 2013. [2](#)
- [19] Minsik Lee, Chong-Ho Choi, and Songhwai Oh. A procrustean markov process for non-rigid structure recovery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1550–1557, 2014. [2](#)
- [20] Dushyant Mehta, Helge Rhodin, Dan Casas, Pascal Fua, Oleksandr Sotnychenko, Weipeng Xu, and Christian Theobalt. Monocular 3d human pose estimation in the wild using improved cnn supervision. In *2017 international conference on 3D vision (3DV)*, pages 506–516. IEEE, 2017. [7](#)
- [21] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. [2](#)
- [22] David Novotny, Nikhila Ravi, Benjamin Graham, Natalia Neverova, and Andrea Vedaldi. C3dpo: Canonical 3d pose networks for non-rigid structure from motion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7688–7697, 2019. [2](#), [3](#), [7](#)
- [23] Onur Özyeşil, Vladislav Voroninski, Ronen Basri, and Amit Singer. A survey of structure from motion*. *Acta Numerica*, 26:305–364, 2017. [1](#)
- [24] Sunghoon Park, Minsik Lee, and Nojun Kwak. Procrustean regression networks: Learning 3d structure of non-rigid objects from 2d annotations. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIX 16*, pages 1–18. Springer, 2020. [2](#), [4](#), [7](#)
- [25] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image gen-

- eration with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 2
- [26] Helge Rhodin, Mathieu Salzmann, and Pascal Fua. Unsupervised geometry-aware representation for 3d human pose estimation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 750–767, 2018. 2
- [27] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 2
- [28] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022. 2
- [29] Yonatan Shafir, Guy Tevet, Roy Kapon, and Amit H Bermano. Human motion diffusion as a generative prior. *arXiv preprint arXiv:2303.01418*, 2023. 3
- [30] Wenkang Shan, Zhenhua Liu, Xinfeng Zhang, Shanshe Wang, Siwei Ma, and Wen Gao. P-stmo: Pre-trained spatial temporal many-to-one model for 3d human pose estimation. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part V*, pages 461–478. Springer, 2022. 7
- [31] Wenkang Shan, Zhenhua Liu, Xinfeng Zhang, Zhao Wang, Kai Han, Shanshe Wang, Siwei Ma, and Wen Gao. Diffusion-based 3d human pose estimation with multi-hypothesis aggregation. *arXiv preprint arXiv:2303.11579*, 2023. 3, 4
- [32] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 2
- [33] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H Bermano. Human motion diffusion model. *arXiv preprint arXiv:2209.14916*, 2022. 3
- [34] Bastian Wandt, James J Little, and Helge Rhodin. Elepose: Unsupervised 3d human pose estimation by predicting camera elevation and learning normalizing flows on 2d poses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6635–6645, 2022. 2, 7, 8
- [35] Chaoyang Wang, Chen-Hsuan Lin, and Simon Lucey. Deep nrsfm++: Towards unsupervised 2d-3d lifting in the wild. In *2020 International Conference on 3D Vision (3DV)*, pages 12–22. IEEE, 2020. 1, 2
- [36] Chaoyang Wang and Simon Lucey. Paul: Procrustean autoencoder for unsupervised lifting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 434–443, 2021. 2, 7
- [37] Chenxin Xu, Siheng Chen, Maosen Li, and Ya Zhang. Invariant teacher and equivariant student for unsupervised 3d human pose estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 3013–3021, 2021. 7
- [38] Jiarui Xu, Sifei Liu, Arash Vahdat, Wonmin Byeon, Xiaolong Wang, and Shalini De Mello. Open-vocabulary panoptic segmentation with text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2955–2966, 2023. 2
- [39] Zhenbo Yu, Bingbing Ni, Jingwei Xu, Junjie Wang, Chenglong Zhao, and Wenjun Zhang. Towards alleviating the modeling ambiguity of unsupervised monocular 3d human pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8651–8660, 2021. 2, 7, 8
- [40] Haitian Zeng, Yuchao Dai, Xin Yu, Xiaohan Wang, and Yi Yang. Pr-rnn: pairwise-regularized residual-recursive networks for non-rigid structure-from-motion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5600–5609, 2021. 2
- [41] Haitian Zeng, Xin Yu, Jiayu Miao, and Yi Yang. Mhr-net: Multiple-hypothesis reconstruction of non-rigid shapes from 2d views. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part II*, pages 1–17. Springer, 2022. 2, 7
- [42] Jinlu Zhang, Zhigang Tu, Jianyu Yang, Yujin Chen, and Jun-song Yuan. Mixste: Seq2seq mixed spatio-temporal encoder for 3d human pose estimation in video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13232–13242, 2022. 5, 6, 7
- [43] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiondiffuse: Text-driven human motion generation with diffusion model. *arXiv preprint arXiv:2208.15001*, 2022. 3
- [44] Yingying Zhu, Dong Huang, Fernando De La Torre, and Simon Lucey. Complex non-rigid motion 3d reconstruction by union of subspaces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1542–1549, 2014. 2