Multi-Person 3D Pose and Shape Estimation via Inverse Kinematics and Refinement

Junuk Cha¹[0000-0003-2321-2797]</sup>, Muhammad Saqlain^{1,2†}[0000-0001-5877-6432], GeonU Kim^{1‡}, Mingyu Shin^{1,3‡}, and Seungryul Baek¹[0000-0002-0856-6880]</sup>

¹UNIST, South Korea ²eSmart Systems, Norway ³Yeongnam Univ., South Korea

Abstract. Estimating 3D poses and shapes in the form of meshes from monocular RGB images is challenging. Obviously, it is more difficult than estimating 3D poses only in the form of skeletons or heatmaps. When interacting persons are involved, the 3D mesh reconstruction becomes more challenging due to the ambiguity introduced by person-to-person occlusions. To tackle the challenges, we propose a coarse-to-fine pipeline that benefits from 1) inverse kinematics from the occlusion-robust 3D skeleton estimation and 2) Transformer-based relation-aware refinement techniques. In our pipeline, we first obtain occlusion-robust 3D skeletons for multiple persons from an RGB image. Then, we apply inverse kinematics to convert the estimated skeletons to deformable 3D mesh parameters. Finally, we apply the Transformer-based mesh refinement that refines the obtained mesh parameters considering intra- and inter-person relations of 3D meshes. Via extensive experiments, we demonstrate the effectiveness of our method, outperforming state-of-the-arts on 3DPW, MuPoTS and AGORA datasets.

Keywords: Multi-person, 3D mesh reconstruction, Transformer

1 Introduction

Recovering 3D human body meshes for a single person or multi-person from a monocular RGB image has made great progress in recent years [3, 10, 12, 17, 23, 27, 28, 30–33, 38, 39, 62, 65, 69, 73, 71]. The technique is essential to understand people's behaviors, intentions and person-to-person interactions. It has a wide range of real-world applications such as human motion imitation [41], virtual try on [48], motion capture [46], action recognition [5, 58, 67], etc.

Recently, deep convolutional neural network-based mesh reconstruction methods [6, 10, 12, 17, 23, 27, 28, 30–33, 38, 39, 62, 65, 69, 73, 71] have shown the practical performance on in-the-wild scenes [21, 25, 44, 45]. Most of the existing 3D human body pose and shape estimation approaches [6, 10, 17, 27, 28, 30–33, 38, 39, 69] achieved promising results for single-person cases. Generally, firstly they

This research was conducted when Dr. Saqlain was the post-doctoral researcher at UNIST[†], and when Mr. Kim and Mr. Shin were undergraduate interns at UNIST[‡].



Fig. 1: Example outputs from our pipeline: (a) input RGB image, (b) initial skeleton estimation results obtained from the input image, (c) initial meshes obtained from the inverse kinematics process, (d) refined meshes obtained from the refinement Transformer, (e, f) top- and side-views for the refined meshes.

crop the area with a person in an input image using bounding-box and then extract features for each detected person, which are further used for 3D human mesh regression.

Some of the recent studies [26–28, 30–33, 36, 39, 65, 71] reconstruct each person 3D mesh individually for multi-person 3D mesh reconstruction using the same bounding-boxes detector [4, 18, 56]. Multiple persons can create severe person-to-person or person-to-environmental occlusions, erroneous monocular depth and diverse human body appearance which results in performance ambiguity in crowded scenes, while in these methods, proper modules that tackle the interacting persons have not been established yet. A few recent methods [23, 73] applied direct regression for multiple persons which do not require individual person detection. Sun et al. [62] used body center heatmaps as the target representation to identify mesh parameter map. However, without applying the human detection, the human pose estimation result is frequently affected by unimportant pixels and it frequently fails to capture scale variations, which result in the inferior performance.

In parallel, there have been efforts to reduce the ambiguity of estimating 3D meshes from an RGB image. However in the aspect of the pose recovery, 3D body mesh recovery methods [27, 30, 31, 33] still fall behind the 3D skeleton or heatmap estimation methods [8, 9, 22, 61]. One drawback of 3D skeleton estimation method is that it cannot reconstruct the full 3D body mesh. Recently, Li et al. [36] proposed an inverse kinematics method for single-person mesh reconstruction to recover 3D meshes from 3D skeletons. This approach is promising since it is able to deliver good poses obtained from 3D skeleton estimator to the 3D mesh reconstruction pipeline.

To tackle the multi-person 3D body mesh reconstruction task, we propose a coarse-to-fine pipeline that first estimates 3D skeletons, reconstruct 3D meshes

from 3D skeletons via inverse kinematics and refine the initial 3D mesh parameters via relation-aware refinement. Inspired by [60], our 3D skeleton estimator involves metric-scale heatmaps and is trained by both relative and absolute positional 3D poses to be robust to occlusions. By extending the IK process [36] towards the multi-person scenario, we are able to obtain the initial 3D meshes for multiple persons from 3D skeletons; while the accuracy is limited especially for interacting person cases. To compensate for the limitation, we propose the relation-aware Transformer to refine the initial mesh parameters considering intra- and inter-person 3D mesh relationships. The Fig. 1 shows example outputs for intermediate steps. To summarize, our contributions are as follows:

- We propose a coarse-to-fine multi-person 3D body mesh reconstruction pipeline that first estimates 3D skeletons and then delivers it toward 3D meshes via inverse kinematics. To make our pipeline robust to interacting persons, we borrowed the occlusion-robust techniques for 3D skeleton estimation.
- To further boost the performance, we propose the Transformer-based architecture for relation-aware mesh refinement to refine the initial mesh parameters considering intra- and inter-person relationships.
- Extensive comparisons are conducted involving three challenging multi-person 3D body pose benchmarks (i.e. 3DPW, MuPoTS and AGORA) and we have demonstrated the state-of-the-art performance on each benchmark. Via ablation studies, we prove that each component works in the meaningful way.

2 Related Works

Single-person 3D mesh regression. There is a long history of methods for predicting 3D human body meshes from monocular RGB images or video frames [16]. Recently, there has been quick advancement in this field thanks to SMPL [42] which provides a low dimensional parameterization of the 3D human body mesh. Here we focus on a 3D body mesh regressing by adopting a parametric model like SMPL from a monocular RGB image. Bogo et al. [3] represented an optimization-based method called SMPLify by fitting SMPL on the detected 2D body joints iteratively. However, this optimization-based approach is comparatively time-consuming and struggle with the higher inference time per input frame.

Some recent studies [34, 51, 54] use deep neural networks for SMPL parameters regression from images in a two-stage manner, which have been effective and can generate more accurate mesh reconstruction outputs in the presence of large-scale 3D datasets. They first determine intermediate renderings such as silhouettes and 2D keypoints from input images and then map them to the SMPL parameters. Impressive results have been achieved for in-the-wild images by applying diverse weak supervision signals such as semantic segmentation [69], texture consistency [53], efficient temporal features [30, 64, 66], 2D pose [11, 27, 35], motion dynamics [28], etc.

More recently, Li et al. [36] proposed a 3D human body pose and shape estimation method by collaborating the 3D keypoints and body meshes. Authors

introduced an inverse kinematics process to find the relative rotations using twist-and-swing decomposition which estimates targeted body joint locations.

Multi-person 3D skeleton regression. There have been variety of methods [13, 47, 49, 57] that tackle the 3D body pose estimation for multi-person: Zanfir et al. [57] proposed LCR-Net that consists of localization, classification, and regression modules. The localization module detects multi-persons from a single image. The classification module classifies the detected human into several anchor-poses. Finally, the regression module refines the anchor-poses. Mehta et al. [47] proposed a single-shot method for multi-person 3D pose estimation from a single image. In addition, they introduced the MuCo-3DHP dataset which has multi-person interactions and occlusions images. Moon et al. [49] proposed top-down method for 3D multi-person pose estimation from a monocular RGB image. This method consists of human detection, absolute 3D human root localization, and root-relative 3D single-person pose estimation modules. Dong et al. [13] used the multi-view images for estimating the multi-person 3D pose. They proposed a coarse-to-fine method lifting the 2D joints to the 3D joints. They obtained the 2D joints candidates from [4]. The initial 3D joints are triangulated from 2D joints candidates of different camera views of the same image. In addition, the initial 3D joints are updated using the prior information using the SMPL [42] model.

Recent multi-person 3D pose regression works [7, 55, 60, 72] tackled a variety of issues such as developing attention-based mechanism dedicated to the 3D pose estimation problem which considers 3D-to-2D projection process [72], combining the top-down and bottom-up networks [7], developing the tracking-based for multi-person [55] and so on. Sárándi et al. [60] recently proposed a metric-scale 3D pose estimation method that is robust to truncations. It is able to reason about the out-of-image joints well. Also, this method is robust to occlusion and bounding-box noise.

Multi-person 3D mesh regression. There have been few works [12, 23, 62, 63, 70, 73] that concern the multi-person 3D body mesh regression: The approaches could be categorized into two: bottom-up and top-down methods.

Bottom-up methods [23, 62, 63, 73] perform multi-person detection and 3D mesh reconstruction simultaneously. Zhang et al. [73] proposed a Body Mesh as Point (BMP) using a multi-scale 2D center map grid-level representation, which locates selective persons at the grid cell's center. Sun et al. [62] proposed a ROMP, which creates parameter maps (i.e. body center heatmap, camera map and mesh parameter map) for 2D human body detection, body positioning and 3D body mesh parameter regression, respectively. Jiang et al. [23] represented a coherent reconstruction of multiple humans (CRMH) model, which utilizes the Faster R-CNN based RoI-aligned feature of all persons to estimate SMPL parameters. They further defined the position relevance between multiple persons through a depth ordering-aware loss and an interpenetration. Sun et al. [63] further introduced Bird's-Eye-View (BEV) representation for reasoning the multiperson body centers and depth simultaneously and combining them to estimate 3D body positions.

Top-down methods [12, 70] first detect each individual person in the frame using bounding-boxes and then estimate the 3D mesh parameters of each detected person. They are basically similar to the single-person 3D mesh reconstruction pipeline; however different in that they provide dedicated modules or loss functions for the multi-person scenario. For example, Zanfir et al. [70] proposed a 3D mesh reconstruction method to firstly infer 3D skeletons of each person and group estimated skeletons to infer the final 3D meshes for multi-person. Choi et al. [12] proposed a method for combining early-stage image features and estimated 2D pose heatmaps which are robust to occlusions, to reconstruct 3D meshes for multiple persons.

Bottom-up methods are frequently affected by unimportant image pixels and suffer from scale variations. They further fail to detect small persons since the person detection is not powerful enough compared to that of top-down methods. On the contrary, in top-down methods, proper modules that tackle the interacting persons have not been established yet. In this paper, we take the top-down approach to secure the robustness and propose to use the Transformer architecture to consider the interacting scenario.

3 Method

Our aim is to reconstruct 3D meshes $\{\mathbf{M}^i\}_{i=1}^M$ of the multiple persons in an RGB image **I**, where M denotes the number of persons in **I**. To achieve this goal, we propose the coarse-to-fine reconstruction pipeline as in Fig. 2 that 1) first estimates the 3D skeletons $\{\mathbf{P}^i\}_{i=1}^M$, 2) obtains the deformable 3D mesh parameters from 3D skeletons via the inverse kinematics (IK) process and 3) refines the initial 3D meshes $\{\mathbf{M}^i\}_{i=1}^M$ using the Transformer architecture that considers intra-person and inter-person relationships. In the remainder of this section, we will elaborate each step in detail.

SMPL body model. For the 3D mesh representation, we use the SMPL deformable 3D mesh model [42] for its compact representation. Variations of the SMPL model [42] are controlled by pose $\boldsymbol{\theta} \in \mathbb{R}^{24\times 6}$ and shape parameters $\boldsymbol{\beta} \in \mathbb{R}^{1\times 10}$. The pose and shape parameters contain 3D rotational information of 24 human body joints in 6D representation and the top-10 principal component analysis coefficients of the 3D shape space, respectively. Using the differentiable mapping between SMPL parameters (i.e. $\boldsymbol{\theta}$ and $\boldsymbol{\beta}$) and the 3D body mesh $\mathbf{M} = \{\mathbf{v}, \mathbf{f}\}$ defined in [42], we can differentiably obtain the 3D body mesh \mathbf{M} from $\boldsymbol{\theta}$ and $\boldsymbol{\beta}$, where $\mathbf{v} \in \mathbb{R}^{6,890\times 3}$, $\mathbf{f} \in \mathbb{R}^{13,776\times 3}$ denote vertices having 6,890 vertices, 13,776 triangular faces that are defined by 3 vertices.

3.1 Initial 3D Skeleton Estimation

We take the top-down approach for 3D skeleton estimation that first detect bounding boxes of the humans and estimate 3D skeletons within each bounding box. Following [43, 50, 59, 60], we constituted the person detector using the YOLOv4 [2] to obtain the cropped image $\mathbf{X} \in \mathbb{R}^{256 \times 256 \times 3}$ from an image I. In



Fig. 2: The schematic diagram of our framework: We first detect persons from an image I and crop it to X and the image encoder extracts image features \mathbf{F}_{img} from X. Then, initial 3D skeletons P are estimated via the initial 3D skeleton estimator f^{P} and SMPL parameters $\boldsymbol{\Theta}_{init}$ are reconstructed via the inverse kinematic process, involving the twist angle and shape estimator f^{TS} (GAP denotes global average pooling layer). Finally, we refine the initial SMPL parameters by inputting the image features \mathbf{F}_{img} and $\boldsymbol{\Theta}_{init}$ to the relation-aware refiner f^{Ref} to produce the refined mesh parameters $\boldsymbol{\Theta}_{\text{ref}}$. The final 3D mesh M is obtained from the refined SMPL parameters $\boldsymbol{\Theta}_{\text{ref}}$. The blue boxes denote involved loss functions (L_{P} , L_{TS} , L_{mesh} , L_{adv} and L_{pose}).

order to develop the initial 3D skeleton estimation network $f^{\mathrm{P}} : \mathbf{X} \to \mathbf{P} \in \mathbb{R}^{K \times 3}$ aiming to use it for the inverse kinematics (IK) process, it is necessary to align the output dimension K with the SMPL model [42]: We are required to set K as 24 to align it with the SMPL model which uses 24-dimensional pose parameters. The reason behind this is that the IK process we use (see Sec. 3.2) requires to calculate the SMPL pose parameters $\boldsymbol{\theta}$ by comparing the 3D skeletons to the SMPL template skeletons having 24 joints.

To further obtain occlusion-robust 3D skeletons, we follow the architecture and loss functions of the recent 3D skeletal estimation approach [60] which utilizes the metric-scale heatmaps for 3D skeleton estimation and use both absolutescale 3D skeletons and image aligned skeletons as the supervision. Within $f^{\rm P}$, we applied ResNet [19] as a feature extractor that predicts image features $\mathbf{F}_{\rm img} \in \mathbb{R}^{8\times8\times2,048}$. They are fed to a 1x1 convolutional layer to extract 3D heatmaps that can produce the root-relative 3D skeletons $\mathbf{P}_{\rm rel}$. In parallel, the image features $\mathbf{F}_{\rm img}$ are fed to a 1x1 convolutional layer to obtain image-scale 2D heatmaps which further produces the 2D skeletons $\mathbf{P}_{\rm img}$. Finally, the absolute 3D skeletons \mathbf{P} are differentiably calculated by combining \mathbf{P}_{img} and \mathbf{P}_{rel} with camera intrinsics as in [60].

3.2 Initial 3D Mesh Reconstruction via Inverse Kinematics.

We define the inverse kinematics (IK) as the process that reveals angle $\boldsymbol{\theta}$ and shape $\boldsymbol{\beta}$ parameters of the SMPL [42] model from estimated 3D skeletons **P**. The angle parameter $\boldsymbol{\theta}$ could be obtained by finding the relative rotation matrix **R** that rotates the template skeletons $\mathbf{T} = \{\mathbf{t}_k\}_{k=1}^K$ to locate it on estimated initial 3D skeletons $\mathbf{P} = \{\mathbf{p}_k\}_{k=1}^K$. To reconstruct this, we use the same formula as [36] that decompose the relative rotation matrix with twist and swing angles. **Reconstructing swing angles** $\boldsymbol{\alpha}$. The axis of swing rotation \mathbf{n}_k which is perpendicular to \mathbf{t}_k and \mathbf{p}_k and the swing angle $\boldsymbol{\alpha} = \{\alpha_k\}_{k=1}^K$ are expressed as:

$$\mathbf{n}_{k} = \frac{\mathbf{t}_{k} \times \mathbf{p}_{k}}{\|\mathbf{t}_{k} \times \mathbf{p}_{k}\|}, \quad \cos \alpha_{k} = \frac{\mathbf{t}_{k} \cdot \mathbf{p}_{k}}{\|\mathbf{t}_{k}\| \|\mathbf{p}_{k}\|}, \quad \sin \alpha_{k} = \frac{\mathbf{t}_{k} \times \mathbf{p}_{k}}{\|\mathbf{t}_{k}\| \|\mathbf{p}_{k}\|}$$
(1)

By the Rodrigues formula, swing rotation matrix \mathbf{R}_k^{sw} can be derived from the axis \mathbf{n}_k and the angle α_k as follows:

$$\mathbf{R}_{k}^{\mathrm{sw}} = \mathbf{I} + \sin \alpha_{k} [\mathbf{n}_{k}]_{\times} + (1 - \cos \alpha_{k}) [\mathbf{n}_{k}]_{\times}^{2}$$
⁽²⁾

where \mathbf{I} is 3×3 identity matrix and $[\mathbf{n}_k]_{\times}$ is the skew-symmetric matrix of \mathbf{n}_k . **Reconstructing twist angles** ϕ and shape parameter β . While swing angles α could be obtained from \mathbf{t}_k and \mathbf{p}_k using Eq. 1; it is hard to find closedform equations for twist angles ϕ . Furthermore, estimating shape parameter β is non-trivial. To bypass the challenges, similarly to [36], we involve the network called as twist angle and shape estimator $f^{\text{TS}} : \mathbf{F}_{\text{img}} \to [\phi, \beta_{\text{init}}]$ that estimates twist angle $\phi = \{\phi_k\}_{k=1}^K$ and shape parameter β_{init} from image features \mathbf{F}_{img} . To resolve the discontinuity issue, it directly estimates cosine value c_{ϕ_k} and sine value s_{ϕ_k} of twist angle instead of estimating ϕ_k . The axis of twist rotation is \mathbf{t}_k and thus, twist rotation matrix \mathbf{R}_k^{tw} can be derived by the axis \mathbf{t}_k and the angle ϕ as follows:

$$\mathbf{R}_{k}^{\mathrm{tw}} = \mathbf{I} + \frac{\sin \phi_{k}}{\|\mathbf{t}_{k}\|} [\mathbf{t}_{k}]_{\times} + \frac{(1 - \cos \phi_{k})}{\|\mathbf{t}_{k}\|^{2}} [\mathbf{t}_{k}]_{\times}^{2},$$
(3)

where $[\mathbf{t}_k]_{\times}$ is the skew-symmetric matrix of \mathbf{t}_k .

Finally, the relative rotation matrix \mathbf{R} can be determined as follows:

$$\mathbf{R} = \mathbf{R}^{\mathrm{sw}} \mathbf{R}^{\mathrm{tw}}.\tag{4}$$

After obtaining the rotation matrix **R**, we convert it to 6D rotation representation and obtain the pose parameters θ_{init} . We initialize the camera parameter **C**_{init} as [0.9, 0, 0] and use the constant values during the inverse kinematics step.

3.3 3D Mesh Refinement via Relation-Aware Transformer

The relation-aware refiner $f^{\text{Ref}} : [\mathbf{F}_{\text{img}}, \boldsymbol{\Theta}_{\text{init}}] \to \boldsymbol{\Theta}_{\text{ref}}$ is proposed to refine the initial SMPL parameters based on the vision Transformer architecture [14]. Its input is the concatenation of image features \mathbf{F}_{img} and SMPL parameters $\boldsymbol{\Theta}_{\text{init}} = [\boldsymbol{\theta}_{\text{init}}; \boldsymbol{\beta}_{\text{init}}; \mathbf{C}_{\text{init}}]$ which are obtained from the IK process. We use $N \times K$ as the sequence length of the Transformer where N is the maximum number of people for the input and K is the number of joints for one person. By rearranging and concatenating image features \mathbf{F}_{img} with $\boldsymbol{\Theta}_{\text{init}}$, we generate the $(N \times K) \times 2,067$ array as the input to the Transformer (see supplemental for details). We obtain the $\Delta \boldsymbol{\Theta}_{\text{ref}}$ from the Transformer and the final SMPL parameter is obtained as follows:

$$\boldsymbol{\Theta}_{\rm ref} = \boldsymbol{\Theta}_{\rm init} + \Delta \boldsymbol{\Theta}_{\rm ref}.$$
 (5)

From the refined parameter $\boldsymbol{\Theta}_{ref} = [\boldsymbol{\theta}_{ref}; \boldsymbol{\beta}_{ref}; \mathbf{C}_{ref}]$, 3D meshes **M** are obtained, and corresponding 3D skeletons \mathbf{P}_{ref} are further obtained by applying the mesh-to-joint regressor [42] to mesh vertices.

When constituting the Transformer, we use the masking input patch as METRO [38]: randomly 0 to 30% of input patches are masked and this makes the Transformer learn non-local interactions. We select not to use the positional embedding while using the masking scheme from results in Table 4.

Sampling interacting persons. The number of persons M varies depending on the image I; while the relation-aware refiner f^{Ref} requires to fix N which is the maximum number of persons in the input. We set N as 3 according to the ablation study shown in Table 4. For images having less than N persons (M < N), we apply Transformer once by simply zero-padding unoccupied inputs, while for images having more than N persons (M > N), we need to apply Transformers multiple times by sampling the interacting persons. The sampling scheme during training and testing stages are proposed as follows: At training, we randomly sample multiple data consisting of N persons so that Transformer can see various combinations as epochs go. At testing, we run f^{Ref} exactly Mtimes, getting results once for each person. At each run, we set each person as the target to refine, inputting N - 1 closest persons as contexts.

3.4 Training Method

We use PyTorch to implement our pipeline. A single NVIDIA TITAN GPU is used for each experiment with a batch size of 64. The Adam optimizer [29] is used for the optimization with a learning rate of 5×10^{-5} for relation-aware Transformer and 1×10^{-4} for all other networks, respectively. We decrease the learning rate exponentially by a factor of 0.9 per each epoch. Total 100 epochs are executed for completely training our network.

To train the proposed initial 3D skeleton estimation network $f^{\rm P}$, twist angle and shape estimation network $f^{\rm TS}$ and relation-aware refiner $f^{\rm Ref}$, we used the loss L defined as follows:

$$L(f^{\rm P}, f^{\rm TS}, f^{\rm Ref}) = L_{\rm P}(f^{\rm P}) + L_{\rm TS}(f^{\rm TS}) + L_{\rm Ref}(f^{\rm Ref}).$$
 (6)

Each loss term is detailed in the remainder of this subsection.

Skeleton loss $L_{\mathbf{P}}$. We use multiple L1 losses using 2D and 3D skeletons in absolute and relative coordinate spaces to train the initial skeleton estimation network $f^{\mathbf{P}}$ using the loss $L_{\mathbf{P}}$ as follows:

$$L_{\rm P}(f^{P}) = \|\mathbf{P} - \hat{\mathbf{P}}_{\rm abs}^{\rm 3D}\|_{1} + \|\mathbf{P}_{\rm rel} - \hat{\mathbf{P}}_{\rm rel}^{\rm 3D}\|_{1} + \|\mathbf{P}_{\rm img} - \hat{\mathbf{P}}^{\rm 2D}\|_{1} + \|\Pi(\mathbf{P}_{\rm rel}) - \hat{\mathbf{P}}^{\rm 2D}\|_{1}$$
(7)

where $\hat{\mathbf{P}}_{abs}^{3D}$, $\hat{\mathbf{P}}_{rel}^{3D}$ and $\hat{\mathbf{P}}^{2D}$ are ground-truth absolute 3D skeletons, relative 3D skeletons and 2D skeletons, respectively. Π is an orthographic projection. **Twist angle and shape loss** L_{TS} . We use the loss L_{TS} to train the twist angle and shape estimator f^{TS} as follows:

$$L_{\rm TS}(f^{\rm TS}) = L_{\rm angle}(f^{\rm TS}) + L_{\rm shape}(f^{\rm TS})$$
(8)

where

$$L_{\text{angle}}(f^{\text{TS}}) = \frac{1}{K} \sum_{k=1}^{K} \| (c_{\phi_k}, s_{\phi_k}) - (\cos \hat{\phi}_k, \sin \hat{\phi}_k) \|_2, \tag{9}$$

$$L_{\text{shape}}(f^{\text{TS}}) = \|\boldsymbol{\beta}_{\text{init}} - \hat{\boldsymbol{\beta}}\|_2, \tag{10}$$

 $\hat{\phi}_k$ denotes the ground-truth twist angle and $\hat{\beta}$ denotes the ground-truth SMPL shape parameters.

Refinement loss L_{Ref} . We use the loss L_{Ref} combining several losses to train our relation-aware refinement network f^{Ref} and additional discriminators $D = \{D_{\theta}, D_{\beta}\}$ as follows:

$$L_{\text{Ref}}(f^{\text{Ref}}, D) = L_{\text{mesh}}(f^{\text{Ref}}) + L_{\text{pose}}(f^{\text{Ref}}) + L_{\text{adv}}(f^{\text{Ref}}) + L_{\text{adv}}(D)$$
(11)

where

$$L_{\text{mesh}}(f^{\text{Ref}}) = \|\boldsymbol{\theta}_{\text{ref}} - \hat{\boldsymbol{\theta}}\|_2 + \|\boldsymbol{\beta}_{\text{ref}} - \hat{\boldsymbol{\beta}}\|_2$$
(12)

enforces the estimated pose θ_{ref} and shape β_{ref} parameters close to the groundtruth pose $\hat{\theta}$ and shape $\hat{\beta}$ parameters of the meshes,

$$L_{\text{pose}}(f^{\text{Ref}}) = \|\mathbf{P}_{\text{ref}} - \hat{\mathbf{P}}_{\text{ref}}^{3\text{D}}\|_{2}^{2} + \|\Pi(\mathbf{P}_{\text{ref}}) - \hat{\mathbf{P}}^{2\text{D}}\|_{2}^{2}$$
(13)

enforces estimated 3D skeletons \mathbf{P}_{ref} and its orthographic projection $\Pi(\mathbf{P}_{ref})$ close to ground-truth 3D and 2D skeletons $(\hat{\mathbf{P}}_{ref}^{3D}, \hat{\mathbf{P}}^{2D})$, respectively,

$$L_{\rm adv}(D) = \|D_{\theta}(\theta_{\rm ref}) - 0\|_2 + \|D_{\theta}(\theta_{\rm real}) - 1\|_2 + \|D_{\beta}(\beta_{\rm ref}) - 0\|_2 + \|D_{\beta}(\beta_{\rm real}) - 1\|_2$$
(14)

trains discriminators D_{θ} , D_{β} to classify real SMPL parameter θ_{real} and β_{real} as real (i.e. 1) and estimated SMPL parameter θ_{ref} and β_{ref} as fake (i.e. 0) and

$$L_{\rm adv}(f^{\rm Ref}) = \|D_{\theta}(\theta_{\rm ref}) - 1\|_2 + \|D_{\beta}(\beta_{\rm ref}) - 1\|_2$$
(15)

enforces the estimated θ_{ref} and β_{ref} become realistic to deceive two discriminators D_{θ} and D_{β} to say that it is the real sample (i.e. 1).

Table 1: SOTA	comparis	ons on 3DP	Table 2: SOTA Comparisons on AGORA.							
Method	$MPJPE(\downarrow) PA-MPJPE(\downarrow) PVE(\downarrow)$			А	11	Matched				
HMR [27]	130.0	76.7	-	\mathbf{Method}						
temporal HMR [28]	116.5	72.6	139.3		$NMVE(\downarrow)$	$NMJE(\downarrow)$	$MVE(\downarrow)$	$MPJPE(\downarrow)$		
BMP [73]	104.1	63.8	119.3	ROMP [62]	227.3	236.6	161.4	168.0		
SPIN [33]	96.6	59.2	116.4							
VIBE [30]	93.5	56.5	116.4	HMR [27]	217.0	226.0	173.6	180.5		
ROMP(Resnet-50) [62]	89.3	53.5	105.6	SPIN [33]	216.3	223.1	168.7	175.1		
ROMP(HRNet-32) [62]	85.5	53.3	103.1	D MAD [71]	000.0	007 4	100.0	174.0		
PARE(Resnet-50) [31]	82.9	52.3	99.7	Pymaf [71]	200.2	207.4	168.2	174.2		
PARE(HRNet-50) [31]	82.0	50.9	97.9	EFT [26]	196.3	203.6	159.0	165.4		
3DCrowdNet [12]	82.8	52.2	100.2	U		100 5		156 1		
HybrIK [36]	80.0	48.8	94.5	Hydrik [50]	-	100.0	-	150.1		
METRO [38]	77.1	47.9	88.2	PARE [31]	167.7	174.0	140.9	146.2		
MeshLeTemp [65]	74.8	46.8	86.5	SPEC [32]	126.8	133 7	106.5	119.3		
Mesh Graphormer [39]	74.7	45.6	87.7	51 10 [52]	120.0	100.1	100.5	112.0		
Ours	66.0	39.0	76.3	Ours	104.5	110.4	86.7	91.6		

Table 3: 3DPCK relevant on MuPoTS-3D dataset for all sequences. The above table shows accuracy only for all groundtruths. The below table shows accuracy only for matched groundtruths.

Method-3DPCK(\uparrow)	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10	S11	S12	S13	S14	S15	S16	S17	S18	S19	S20	Avg
Accuracy for all groundtruths																					
Jiang et al. [23]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	69.1
ROMP [62]	89.8	73.1	67.2	68.4	78.9	41.0	68.7	68.2	70.1	85.4	69.2	63.2	66.5	60.9	78.1	77.4	75.1	80.7	74.0	61.1	71.9
SPEC [32]	87.2	69.4	69.0	71.5	78.5	63.8	69.1	66.2	71.5	85.7	69.2	63.2	66.5	60.9	78.1	77.4	75.1	80.7	74.0	61.1	71.9
BMP [73]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	73.8
PARE [31]	87.7	65.4	66.4	67.7	81.9	62.5	64.9	69.9	73.8	88.5	80.1	79.2	74.5	62.9	81.6	84.5	89.6	83.7	73.7	66.5	75.3
Moon et al. [49]	94.4	77.5	79.0	81.9	85.3	72.8	81.9	75.7	90.2	90.4	79.2	79.9	75.1	72.7	81.1	89.9	89.6	81.8	81.7	76.2	81.8
Metrabs [60]	94.0	82.6	88.4	86.5	87.3	76.2	85.9	66.9	85.8	92.9	81.8	89.9	77.6	68.5	85.6	92.3	89.3	85.1	78.2	71.6	83.3
Cheng et al. [7]	93.4	91.3	84.7	83.3	89.1	85.2	95.4	92.1	89.5	93.1	85.4	85.7	89.9	90.1	88.8	93.7	92.2	87.9	89.7	91.9	89.6
Ours	97.3	84.7	91.1	89.9	92.9	89.8	92.2	87.1	89.1	94.0	88.6	92.9	84.6	80.4	94.3	96.7	98.8	91.5	86.1	76.7	89.9
Accuracy only for matched groundtruths.																					
Jiang et al. [23]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	72.2
ROMP [62]	92.1	81.9	69.8	69.1	85.9	43.2	69.3	70.7	70.1	85.4	69.2	63.2	68.0	63.6	78.1	77.6	75.4	80.7	74.5	72.9	74.0
SPEC [32]	88.1	78.5	69.6	71.6	81.0	63.8	69.1	77.4	71.5	85.7	69.2	63.2	68.0	63.6	78.1	77.6	75.4	80.7	74.5	72.9	74.0
BMP [73]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	75.3
PARE [31]	88.8	77.1	66.7	67.7	83.2	62.5	64.9	77.9	73.8	88.5	80.1	79.2	76.4	65.8	81.6	85.4	89.6	83.7	74.2	82.5	77.5
Moon et al. [49]	94.4	78.6	79.0	82.1	86.6	72.8	81.9	75.8	90.2	90.4	79.4	79.9	75.3	81.0	81.0	90.7	89.6	83.1	81.7	77.3	82.5
Metrabs [60]	94.0	86.5	89.0	87.1	91.1	77.4	90.2	75.7	85.8	92.9	86.0	90.7	83.8	82.0	85.6	94.3	89.8	89.6	86.5	91.7	87.5
Cheng et al. [7]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	89.6
Ours	97.6	94.1	90.7	89.6	95.2	88.7	94.2	88.4	89.2	93.7	89.1	93.2	86.6	90.5	94.4	97.4	98.5	91.9	86.1	84.8	91.7

4 Experiments

Setup. We involved multiple datasets to train our model. We used Human3.6M [21], MPI-INF-3DHP [45], LPS [24], MSCOCO [40], MPII [1] datasets as the training data, which are the same setting as Kanazawa et al. [27]. Additionally, MuCo-3DHP [47], CMU-Panoptic [25], SAIL-VOS [20], SURREAL [68], AIST++ [37] are used to calculate the $L_P(f^P)$. For evaluation, we used 3DPW [44], MuPoTs [47], and AGORA [52]: The 3DPW dataset is an outdoor 3D human pose benchmark involving real sequences. It contains diverse subjects, various backgrounds and occlusion scenario. The MuPoTS dataset contains both real indoor and outdoor sequences having multiple persons occluding each other. AGORA is the synthetic benchmark having multi-person within it. Image contains many persons with various clothes, ages and ethnicities.

Measures. For 3DPW dataset, we measure the accuracy of ours with widely used evaluation metrics to compare with others: MPJPE, PA-MPJPE and MVE: The MPJPE is the mean per joint position error which is calculated based on the Euclidean distance between ground-truth and estimated joint positions. For this,

11

Table 4: Ablation study of the effectiveness of IK, refiner, positional embedding, masking input patch and comparison between the different number of Transformer's input persons on 3DPW.

Method	$MPJPE(\downarrow)$	$PA-MPJPE(\downarrow)$	$MVE(\downarrow)$	Method	$MPJPE(\downarrow)$	$PA-MPJPE(\downarrow)$	$MVE(\downarrow)$
Ours w/o IK, w/o Ref	71.8	42.5	-	Ours (N=1)	66.9	39.4	77.0
Ours w/o Ref	67.3	39.3	77.5	Ours (N=2)	66.7	39.3	76.7
Ours w/ positional embedding	68.3	39.4	78.3	Ours (N=3)	66.0	39.0	76.3
Ours w/o masking input patch	67.1	39.0	76.6	Ours (N=4)	66.6	39.4	77.2



Fig. 3: Attention visualization: (Row 1) input image, (Row 2-5) part-based attentions obtained within an intra-person in the column 1, those among two persons in the column 2 and those among three persons in the column 3-4, (Row 6) initial mesh obtained from inverse kinematics, (Row 7) refined mesh after mesh refinement module. We visualize the self-attentions between a specified joint and all other joints, where brighter color and thicker line indicate stronger attentions.

pelvis joint is aligned by moving the estimated pelvis joint to the ground-truth pelvis joint. The PA-MPJPE is Procrustes Aligned MPJPE which is calculated

similarly to MPJPE; however it is measured after rigidly aligning the estimated joints to ground-truth joints via Procrustes Analysis [15]. The MVE is the mean vertex error which is calculated as the Euclidean distance between ground-truth and estimated mesh vertices. For MuPoTS dataset, we measure the performance of our methods using 3DPCK. The 3DPCK is the 3D percentage of correct keypoints. It counts the keypoints as correct when the Euclidean distance between the estimated joint position and its ground-truth is within a threshold. We used 150mm as the threshold following [49]. For AGORA dataset, we measure the performance of our methods on AGORA using MPJPE, MVE, NMJE and NMVE. The MPJPE and MVE are measured on matched detections. The NMJE and NMVE are normalized MPJPE and MVE by F1 score to punish misses and false alarms in the detection.

Baselines. In our experiments, we have involved several state-of-the-art 3D body mesh reconstruction [26–28, 30–33, 36, 38, 39, 65, 71] pipelines for single persons to compare with ours: HMR [27] and SPIN [33] are the pioneering works that first tried to infer SMPL pose and shape parameters using the CNN network for a single person. Temporal HMR [28] and VIBE [30] developed to further utilize the temporal information from a video. METRO [38] used transformer architecture for non-parametric mesh reconstruction. Mesh Graphormer [39] combines self-attention and graph convolution network for mesh reconstruction. Mesh-LeTemp [65] proposed the learnable template which reflects not only vertexvertex interactions but also the human pose and body shape. PyMAF [71] uses a pyramidal mesh alignment feedback loop to refine the mesh based on the mesh-image alignment mechanism. EFT [26] trains HMR architecture with a large-scale dataset having pseudo ground-truths. SPEC [32] estimates the perspective camera to accurately infer the 3D mesh coordinates. PARE [31] learns the body-part-guided attention masks to be robust to occlusions. HybrIK [36] estimates SMPL pose and shape parameters from estimated 3D skeletons via inverse kinematics. The single-person mesh reconstruction methods take the topdown approach using bounding boxes obtained from YOLOv4 [2] for AGORA dataset, and using those obtained from ground-truth 3D skeletons for MuPoTS and 3DPW datasets, respectively.

Several multi-person frameworks are also involved for the comparisons [12, 23, 62, 73]: Jiang et al. [23], ROMP [62] and BMP [73] are bottom-up methods that estimate the multi-person SMPL pose and shape parameters at once and simultaneously localizes multi-person instances and predicts 3D body meshes in a single stage, respectively. Jiang et al. [23] proposed inter-penetration loss to avoid collision and depth ordering loss for the rendering. 3DCrowdNet [12] is a top-down method that proposed to concatenate image features and 2D pose heatmaps to exploit the 2D pose-guided features for better accuracy. We also involved 3D skeleton estimation approaches [7, 49, 60]: Moon et al. [49]'s work that estimates the absolute root position and root-relative 3D skeletons focusing on camera distance. Cheng et al. [7]'s work that integrates top-down method and bottom-up methods for estimating better 3D skeletons. Metrabs [60] that

is robust to truncation/occlusion variations thanks to the metric-scale heatmap representation.

Results. We compared ours to state-of-the-art algorithms on three challenging datasets (i.e. 3DPW, MuPoTS and AGORA). The results are summarized in Tables 1, 2 and 3: From Tables 1 and 2, we could observe that ours obtained the superior performance compared to previous mesh reconstruction works. We obtained even better performance than works exploiting temporal information [28, 30] and several multi-person 3D mesh reconstruction methods [12, 23, 62, 73]. Also, in Table 1, we also compared ours to HybrIK [36] that applied the inverse kinematics process on the pre-estimated 3D skeleton results. We outperforms it by successfully extending it towards the multi-person scenario. As shown in Table 3, we have achieved the state-of-the-art accuracy on MuPoTS. Note that the 3D skeleton estimation methods [7, 49, 60] are also included in the comparison and they produced superior performance than 3D mesh estimation approaches [31, 32, 62, 73]. However, our method outperforms them by delivering good pose accuracy from the initial 3D skeleton estimator, while reconstructing both poses and shapes in the form of 3D meshes.

Fig. 3 shows the visualization for the attention learned in the relation-aware refinement network f^{Ref} . It learns the attentions for intra-person parts as in the column 1 and learns the inter-person attentions among at most N = 3persons as in columns 2 through 4. From the visualization, we could see that the refinement network f^{Ref} refines the initial meshes (in the 6th row) a lot towards refined meshes (in the 7th row). Fig. 4 shows the qualitative comparisons to competitive state-of-the-arts [31, 32, 62]. Ours faithfully reconstructs 3D human bodies with diverse artifacts while others frequently fails to capture the details. Ablation study. We conduct an ablation study for several design choices. The Table 4 shows the ablation results on 3DPW dataset: 'Ours w.o IK, w/o Ref', 'Ours w/o Ref', 'Ours w/ positional embedding', 'Ours w/o masking input patch' denote our results obtained without inverse kinematics process and refinement module which are 3D skeletons, our results without the refinement module. our results with positional embedding and our results without masking input patches, respectively. From the results, we can see that inverse kinematics and relation-aware refinements consistently increase the accuracy of our pipeline. Furthermore, we decided not to use positional embedding while using the masking input patch scheme. 'Ours (N=1)' through 'Ours (N=4)' denote experiments conducted by varying the maximum number of persons (i.e. N) of the Transformer input. We observe that N = 3 works best.

5 Conclusion

In this paper, we proposed a coarse-to-fine pipeline for the multi-person 3D mesh reconstruction task, which first estimates occlusion-robust 3D skeletons, then reconstructs initial 3D meshes via the inverse kinematic process and finally refines them based on the relation-aware refiner considering intra- and inter-person relationships. By extensive experiments, we find that our idea of



Fig. 4: Qualitative comparisons on 3DPW (Row 1), AGORA (Rows 2-3) and MuPoTS (Rows 4-5) datasets. Red circles highlight wrongly estimated parts.

delivering the accurate occlusion-robust 3D poses to 3D meshes, and refining initial mesh parameters of interacting persons indeed works: Our pipeline consistently outperforms multiple 3D skeleton-based, 3D mesh-based baselines and each component proposed works meaningfully for the intended scenario.

Acknowledgements. This work was supported by IITP grants (No. 2021-0-01778 Development of human image synthesis and discrimination technology below the perceptual threshold; No. 2020-0-01336 Artificial intelligence graduate school program(UNIST); No. 2021-0-02068 Artificial intelligence innovation hub; No. 2022-0-00264 Comprehensive video understanding and generation with knowledge-based deep logic neural network) and the NRF grant (No. 2022R1F1A1074828), all funded by the Korean government (MSIT).

References

- 1. Andriluka, M., Pishchulin, L., Gehler, P., Schiele, B.: 2d human pose estimation: New benchmark and state of the art analysis. In: CVPR (2014)
- Bochkovskiy, A., Wang, C.Y., Liao, H.Y.M.: Yolov4: Optimal speed and accuracy of object detection. arXiv:2004.10934 (2020)
- Bogo, F., Kanazawa, A., Lassner, C., Gehler, P., Romero, J., Black, M.J.: Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In: ECCV (2016)
- Cao, Z., Hidalgo, G., Simon, T., Wei, S.E., Sheikh, Y.: Openpose: realtime multiperson 2d pose estimation using part affinity fields. TPAMI (2019)
- 5. Cha, J., Saqlain, M., Kim, D., Lee, S., Lee, S., Baek, S.: Learning 3d skeletal representation from transformer for action recognition. IEEE Access (2022)
- Cha, J., Saqlain, M., Lee, C., Lee, S., Lee, S., Kim, D., Park, W.H., Baek, S.: Towards single 2d image-level self-supervision for 3d human pose and shape estimation. Applied Sciences (2021)
- 7. Cheng, Y., Wang, B., Yang, B., Tan, R.T.: Monocular 3d multi-person pose estimation by integrating top-down and bottom-up networks. In: CVPR (2021)
- Cheng, Y., Yang, B., Wang, B., Tan, R.T.: 3d human pose estimation using spatiotemporal networks with explicit occlusion training. In: AAAI (2020)
- Cheng, Y., Yang, B., Wang, B., Yan, W., Tan, R.T.: Occlusion-aware networks for 3d human pose estimation in video. In: ICCV (2019)
- 10. Choi, H., Moon, G., Chang, J.Y., Lee, K.M.: Beyond static features for temporally consistent 3d human pose and shape from a video. In: CVPR (2021)
- 11. Choi, H., Moon, G., Lee, K.M.: Pose2mesh: Graph convolutional network for 3d human pose and mesh recovery from a 2d human pose. In: ECCV (2020)
- 12. Choi, H., Moon, G., Park, J., Lee, K.M.: 3dcrowdnet: 2d human pose-guided3d crowd human pose and shape estimation in the wild. arXiv:2104.07300 (2021)
- 13. Dong, Z., Song, J., Chen, X., Guo, C., Hilliges, O.: Shape-aware multi-person pose estimation from multi-view images. In: ICCV (2021)
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. ICLR (2021)
- 15. Gower, J.C.: Generalized procrustes analysis. Psychometrika (1975)
- Guan, P., Weiss, A., Balan, A.O., Black, M.J.: Estimating human shape and pose from a single image. In: ICCV (2009)
- Guler, R.A., Kokkinos, I.: Holopose: Holistic 3d human reconstruction in-the-wild. In: CVPR (2019)
- 18. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: ICCV (2017)
- He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: ECCV (2016)
- Hu, Y.T., Chen, H.S., Hui, K., Huang, J.B., Schwing, A.G.: Sail-vos: Semantic amodal instance level video object segmentation-a synthetic dataset and baselines. In: CVPR (2019)
- Ionescu, C., Papava, D., Olaru, V., Sminchisescu, C.: Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. TPAMI (2013)
- Iskakov, K., Burkov, E., Lempitsky, V., Malkov, Y.: Learnable triangulation of human pose. In: ICCV (2019)

- 16 Cha et al.
- Jiang, W., Kolotouros, N., Pavlakos, G., Zhou, X., Daniilidis, K.: Coherent reconstruction of multiple humans from a single image. In: CVPR (2020)
- 24. Johnson, S., Everingham, M.: Clustered pose and nonlinear appearance models for human pose estimation. In: BMVC (2010)
- Joo, H., Liu, H., Tan, L., Gui, L., Nabbe, B., Matthews, I., Kanade, T., Nobuhara, S., Sheikh, Y.: Panoptic studio: A massively multiview system for social motion capture. In: ICCV (2015)
- Joo, H., Neverova, N., Vedaldi, A.: Exemplar fine-tuning for 3d human model fitting towards in-the-wild 3d human pose estimation. In: 3DV (2021)
- 27. Kanazawa, A., Black, M.J., Jacobs, D.W., Malik, J.: End-to-end recovery of human shape and pose. In: CVPR (2018)
- Kanazawa, A., Zhang, J.Y., Felsen, P., Malik, J.: Learning 3d human dynamics from video. In: CVPR (2019)
- 29. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. ICLR (2015)
- Kocabas, M., Athanasiou, N., Black, M.J.: Vibe: Video inference for human body pose and shape estimation. In: CVPR (2020)
- Kocabas, M., Huang, C.H.P., Hilliges, O., Black, M.J.: Pare: Part attention regressor for 3d human body estimation. In: ICCV (2021)
- 32. Kocabas, M., Huang, C.H.P., Tesch, J., Müller, L., Hilliges, O., Black, M.J.: Spec: Seeing people in the wild with an estimated camera. In: ICCV (2021)
- Kolotouros, N., Pavlakos, G., Black, M.J., Daniilidis, K.: Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In: ICCV (2019)
- Kolotouros, N., Pavlakos, G., Daniilidis, K.: Convolutional mesh regression for single-image human shape reconstruction. In: CVPR (2019)
- Kundu, J.N., Rakesh, M., Jampani, V., Venkatesh, R.M., Babu, R.V.: Appearance consensus driven self-supervised human mesh recovery. In: ECCV (2020)
- Li, J., Xu, C., Chen, Z., Bian, S., Yang, L., Lu, C.: Hybrik: A hybrid analyticalneural inverse kinematics solution for 3d human pose and shape estimation. In: CVPR (2021)
- 37. Li, R., Yang, S., Ross, D.A., Kanazawa, A.: Learn to dance with aist++: Music conditioned 3d dance generation. arXiv:2101.08779 (2021)
- 38. Lin, K., Wang, L., Liu, Z.: End-to-end human pose and mesh reconstruction with transformers. In: CVPR (2021)
- 39. Lin, K., Wang, L., Liu, Z.: Mesh graphormer. In: ICCV (2021)
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV (2014)
- 41. Liu, W., Piao, Z., Min, J., Luo, W., Ma, L., Gao, S.: Liquid warping gan: A unified framework for human motion imitation, appearance transfer and novel view synthesis. In: ICCV (2019)
- Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: Smpl: A skinned multi-person linear model. TOG (2015)
- Ludl, D., Gulde, T., Curio, C.: Enhancing data-driven algorithms for human pose estimation and action recognition through simulation. IEEE transactions on intelligent transportation systems (2020)
- von Marcard, T., Henschel, R., Black, M.J., Rosenhahn, B., Pons-Moll, G.: Recovering accurate 3d human pose in the wild using imus and a moving camera. In: ECCV (2018)
- Mehta, D., Rhodin, H., Casas, D., Fua, P., Sotnychenko, O., Xu, W., Theobalt, C.: Monocular 3d human pose estimation in the wild using improved cnn supervision. In: 3DV (2017)

- 46. Mehta, D., Sotnychenko, O., Mueller, F., Xu, W., Elgharib, M., Fua, P., Seidel, H.P., Rhodin, H., Pons-Moll, G., Theobalt, C.: Xnect: Real-time multi-person 3d motion capture with a single rgb camera. TOG (2020)
- 47. Mehta, D., Sotnychenko, O., Mueller, F., Xu, W., Sridhar, S., Pons-Moll, G., Theobalt, C.: Single-shot multi-person 3d pose estimation from monocular rgb. In: 3DV (2018)
- Mir, A., Alldieck, T., Pons-Moll, G.: Learning to transfer texture from clothing images to 3d humans. In: CVPR (2020)
- 49. Moon, G., Chang, J.Y., Lee, K.M.: Camera distance-aware top-down approach for 3d multi-person pose estimation from a single rgb image. In: ICCV (2019)
- 50. Ning, G., Pei, J., Huang, H.: Lighttrack: A generic framework for online top-down human pose tracking. In: CVPR workshop (2020)
- Omran, M., Lassner, C., Pons-Moll, G., Gehler, P., Schiele, B.: Neural body fitting: Unifying deep learning and model based human pose and shape estimation. In: 3DV (2018)
- 52. Patel, P., Huang, C.H.P., Tesch, J., Hoffmann, D.T., Tripathi, S., Black, M.J.: Agora: Avatars in geography optimized for regression analysis. In: CVPR (2021)
- 53. Pavlakos, G., Kolotouros, N., Daniilidis, K.: Texturepose: Supervising human mesh estimation with texture consistency. In: ICCV (2019)
- 54. Pavlakos, G., Zhu, L., Zhou, X., Daniilidis, K.: Learning to estimate 3d human pose and shape from a single color image. In: CVPR (2018)
- Reddy, N.D., Guigues, L., Pishchulin, L., Eledath, J., Narasimhan, S.G.: Tessetrack: End-to-end learnable multi-person articulated 3d pose tracking. In: CVPR (2021)
- Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: CVPR (2016)
- 57. Rogez, G., Weinzaepfel, P., Schmid, C.: Lcr-net: Localization-classificationregression for human pose. In: CVPR (2017)
- 58. Saqlain, M., Kim, D., Cha, J., Lee, C., Lee, S., Baek, S.: 3dmesh-gar: 3d human body mesh-based method for group activity recognition. Sensors (2022)
- Sárándi, I., Linder, T., Arras, K.O., Leibe, B.: Synthetic occlusion augmentation with volumetric heatmaps for the 2018 eccv posetrack challenge on 3d human pose estimation. arXiv:1809.04987 (2018)
- 60. Sárándi, I., Linder, T., Arras, K.O., Leibe, B.: Metrabs: Metric-scale truncationrobust heatmaps for absolute 3d human pose estimation. IEEE Transactions on Biometrics, Behavior, and Identity Science (2020)
- Sun, X., Xiao, B., Wei, F., Liang, S., Wei, Y.: Integral human pose regression. In: ECCV (2018)
- Sun, Y., Bao, Q., Liu, W., Fu, Y., Black, M.J., Mei, T.: Monocular, one-stage, regression of multiple 3d people. In: ICCV (2021)
- Sun, Y., Liu, W., Bao, Q., Fu, Y., Mei, T., Black, M.J.: Putting people in their place: Monocular regression of 3d people in depth. arXiv:2112.08274 (2021)
- 64. Sun, Y., Ye, Y., Liu, W., Gao, W., Fu, Y., Mei, T.: Human mesh recovery from monocular images via a skeleton-disentangled representation. In: ICCV (2019)
- Tran, T.Q., Than, C.C., Nguyen, H.T.: Meshletemp: Leveraging the learnable vertex-vertex relationship to generalize human pose and mesh reconstruction for in-the-wild scenes. arXiv:2202.07228 (2022)
- 66. Tung, H.Y.F., Tung, H.W., Yumer, E., Fragkiadaki, K.: Self-supervised learning of motion capture. NeurIPS (2017)
- 67. Varol, G., Laptev, I., Schmid, C., Zisserman, A.: Synthetic humans for action recognition from unseen viewpoints. IJCV (2021)

- 18 Cha et al.
- Varol, G., Romero, J., Martin, X., Mahmood, N., Black, M.J., Laptev, I., Schmid, C.: Learning from synthetic humans. In: CVPR (2017)
- 69. Xu, Y., Zhu, S.C., Tung, T.: Denserac: Joint 3d pose and shape estimation by dense render-and-compare. In: ICCV (2019)
- 70. Zanfir, A., Marinoiu, E., Zanfir, M., Popa, A.I., Sminchisescu, C.: Deep network for the integrated 3d sensing of multiple people in natural images. NeurIPS (2018)
- Zhang, H., Tian, Y., Zhou, X., Ouyang, W., Liu, Y., Wang, L., Sun, Z.: Pymaf: 3d human pose and shape regression with pyramidal mesh alignment feedback loop. In: ICCV (2021)
- 72. Zhang, J., Cai, Y., Yan, S., Feng, J., et al.: Direct multi-view multi-person 3d pose estimation. NeurIPS (2021)
- 73. Zhang, J., Yu, D., Liew, J.H., Nie, X., Feng, J.: Body meshes as points. In: CVPR (2021)