

Co-Evolution of Pose and Mesh for 3D Human Body Estimation from Video

Yingxuan You¹ Hong Liu¹✉ Ti Wang¹ Wenhao Li¹ Runwei Ding¹ Xia Li²

¹Key Laboratory of Machine Perception, Shenzhen Graduate School, Peking University

²Department of Computer Science, ETH Zürich

{youyx, tiwang}@stu.pku.edu.cn, {hongliu, wenhaoli, dingrunwei}@pku.edu.cn, xia.li@inf.ethz.ch

Abstract

Despite significant progress in single image-based 3D human mesh recovery, accurately and smoothly recovering 3D human motion from a video remains challenging. Existing video-based methods generally recover human mesh by estimating the complex pose and shape parameters from coupled image features, whose high complexity and low representation ability often result in inconsistent pose motion and limited shape patterns. To alleviate this issue, we introduce 3D pose as the intermediary and propose a *Pose and Mesh Co-Evolution network (PMCE)* that decouples this task into two parts: 1) video-based 3D human pose estimation and 2) mesh vertices regression from the estimated 3D pose and temporal image feature. Specifically, we propose a two-stream encoder that estimates mid-frame 3D pose and extracts a temporal image feature from the input image sequence. In addition, we design a co-evolution decoder that performs pose and mesh interactions with the image-guided Adaptive Layer Normalization (AdaLN) to make pose and mesh fit the human body shape. Extensive experiments demonstrate that the proposed PMCE outperforms previous state-of-the-art methods in terms of both per-frame accuracy and temporal consistency on three benchmark datasets: 3DPW, Human3.6M, and MPI-INF-3DHP. Our code is available at <https://github.com/kasvii/PMCE>.

1. Introduction

Recovering 3D human mesh from an image or a video is an essential yet challenging task for many applications, such as human-robot interaction, virtual reality, and motion analysis. The challenges of this task arise from the 2D-to-3D ambiguity, cluttered background, and occlusions. Recently, many studies [8, 13, 16, 19, 22, 33] have been proposed to recover the 3D human mesh from a single image, which can generally be categorized into RGB-based methods and pose-based methods. RGB-based methods predict human mesh

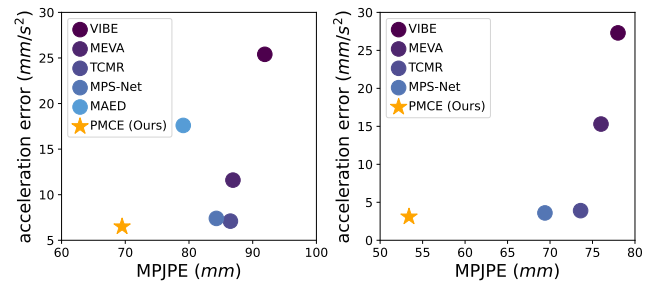


Figure 1: Comparison with video-based 3D human mesh recovery methods upon accuracy (MPJPE ↓) and temporal consistency (acceleration error ↓) on 3DPW [40] (left) and Human3.6M [12] (right) datasets.

end-to-end from image pixels, typically predicting the pose and shape parameters of the parametric human model (e.g., SMPL [27]) to generate the 3D human mesh. However, the representation ability of the parametric model is constrained by the limited pose and shape space [18, 19]. To overcome this limitation, non-parametric approaches have been proposed to predict the 3D coordinates of mesh vertices directly, which generally use Graph Convolutional Networks (GCNs) [8, 42] or Transformers [5, 24, 51] to capture the relations among vertices. In contrast, pose-based methods leverage 2D pose detectors [4, 36] as the front-end to recover human mesh from the detected 2D poses. With the significant advancements in 2D pose detection, pose-based methods have become increasingly robust and lightweight, making them popular for real-world applications [51].

Despite the significant progress in single-image 3D human mesh recovery, these methods still struggle to capture temporally consistent human motion from videos. To solve this problem, several works [7, 15, 28, 43] extend single image-based methods to video cases. They use a pre-trained Convolutional Neural Network (CNN) [17] to extract static features for each frame, then train a temporal network to predict SMPL parameters. Although these video-based methods significantly improve the temporal consistency of 3D human motion, there exists a trade-off between per-frame accuracy

✉ Corresponding author.

and motion smoothness for the following two main reasons: 1) The highly coupled image feature. The extracted image features in deep CNN layers are low-resolution and tightly coupled [37], which inevitably discard the spatial information in the image [41]. 2) The limited representation ability of the parametric human model. The SMPL model represents pose using 3D rotation, which might face periodicity and discontinuity issues [18], making pose prediction in videos more challenging. Besides, the local and swift mesh deformations described by the shape parameters are difficult to learn.

In recent years, video-based 3D human pose estimation from the detected 2D poses has achieved high pose accuracy and motion smoothness [21, 50], which inspires us that the skeleton sequence contains sufficient spatial and temporal pose information of human motion. But for the mesh recovery task, detailed shape information is needed, which can be acquired from the image features. Based on the above observations, we propose the **Pose and Mesh Co-Evolution network (PMCE)** to recover 3D human mesh from videos in a non-parametric way. We decouple the 3D human mesh recovery task into two consecutive parts: 1) video-based 3D pose estimation and 2) mesh vertices regression from 3D pose and image feature, where the latter is the focus. Specifically, in the first part, we propose a two-stream encoder. One stream takes a 2D pose sequence detected from input images to estimate the mid-frame 3D pose, and the other stream takes static image features extracted from images and aggregates them for a temporal image feature. In the second part, we design a co-evolution decoder that performs pose and mesh interactions with an image-guided Adaptive Layer Normalization (AdaLN). AdaLN adjusts the statistical characteristics of joint and vertex features to make the pose and mesh better fit the human body shape. As shown in Figure 1, compared to previous video-based methods, PMCE achieves better performance in terms of per-frame accuracy and temporal consistency on 3DPW [40] and Human3.6M [12] datasets.

In summary, our contributions are as follows:

- We propose a Pose and Mesh Co-Evolution network (PMCE) for recovering 3D human mesh from video. It decouples the task into two parts: video-based 3D pose estimation, and mesh vertices regression by image-guided pose and mesh co-evolution, achieving accurate and temporally consistent results.
- We design the co-evolution decoder that performs pose and mesh interactions guided by our proposed AdaLN. AdaLN adjusts the statistical characteristics of joint and vertex features based on the image feature to make them conform to the human body shape.
- Our method achieves state-of-the-art performance on challenging datasets like 3DPW, reducing MPJPE by 12.1%, PVE by 8.4% and acceleration error by 8.5%.

2. Related Work

Estimating 3D human pose from monocular video. Benefiting from the advance of 2D pose detection, video-based 3D human pose estimation often takes a 2D pose sequence detected from images by an off-the-shelf 2D pose detector to estimate the 3D pose. Recent years have witnessed the development of video-based 3D human pose estimation. Several methods have been proposed to explore spatial and temporal pose information through Temporal Convolution Networks (TCNs) [3, 26, 34, 48], and Transformer-based networks [20, 21, 50, 52], resulting in an impressive performance in terms of pose accuracy and motion smoothness. The achievements in video-based 3D human pose estimation demonstrate that the pose sequence contains sufficient spatial and temporal information, which can be leveraged to produce accurate and smooth 3D human pose motion.

Recovering 3D human mesh from a single image. 3D human mesh recovery aims to reconstruct 3D human pose and shape from monocular images. Previous studies can be divided into parametric methods and non-parametric methods. Parametric methods adopt the pre-trained parametric human model (*e.g.*, SMPL [27]) and estimate the pose and shape parameters to recover the human mesh. Early works [13] directly predict the parameter from the input image, but it is challenging to obtain accurate results. Recent works further introduce human body information extracted from the input image, such as 2D pose [33, 47], human part segmentation [16, 32], silhouette [33, 47], and depth map [10, 45, 49], then predict the SMPL parameters from them. However, parametric methods are constrained by the limited pose and shape representation space. To this end, non-parametric methods are proposed that directly regress the mesh vertices. They model the relations between mesh vertex by GCNs [8, 18, 42] or Transformers [5, 23, 24], which produce flexible human meshes. Although the single image-based methods have achieved significant performance in accuracy, it is hard for them to produce temporally consistent and smooth 3D human motion when applied to video cases.

Recovering 3D human mesh from monocular video. Different from image-based methods, video-based methods aim to produce accurate and temporally consistent human mesh from video frames. Most methods [7, 14, 15, 28, 43] extract the static image feature vectors using a pre-trained CNN [13] with global pooling operation and train the temporal networks, such as 1D CNNs [14], Gated Recurrent Units (GRUs) [7, 15, 28], and Transformers [43], to predict the pose and shape parameters for human mesh representation. For larger receptive fields, some methods [37, 41] take the image feature maps before the global pooling as input. Sun *et al.* [37] disentangle the skeletons from the image feature maps, and then aggregates per-frame skeleton and image feature vector by a bilinear transformation. MAED [41] takes the image patch sequence as input and

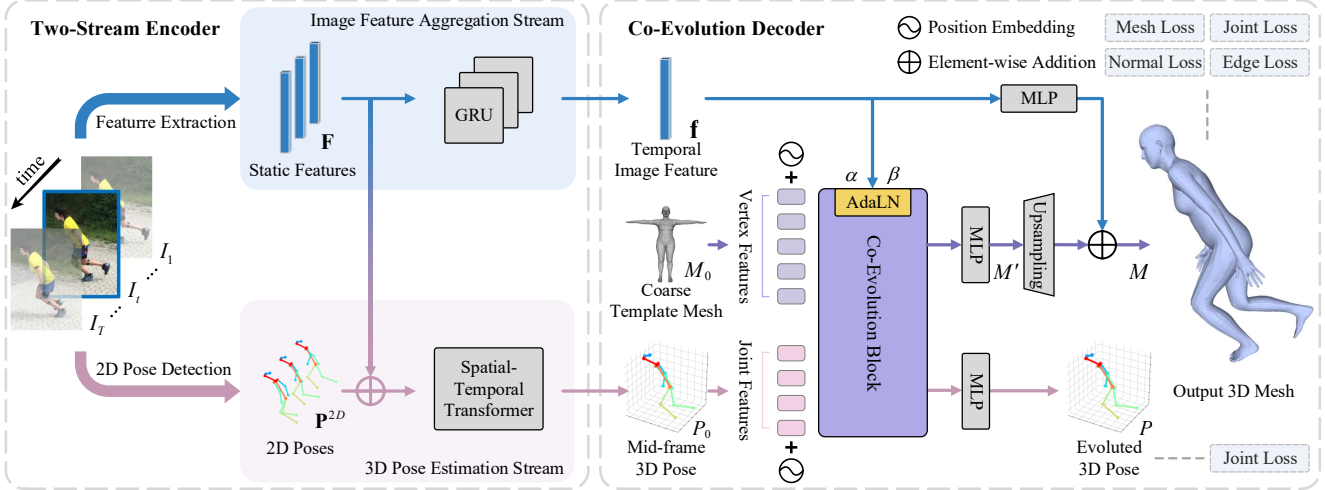


Figure 2: Architecture of the proposed PMCE. Given a video sequence, static image features are extracted by a pre-trained CNN [17], and 2D poses are detected by an off-the-shelf 2D pose detector [4, 46]. The two-stream encoder leverages dual parallel modules to generate a temporal feature and estimate the mid-frame 3D pose, respectively. Then, the co-evolution decoder regresses the mesh vertices through the pose and mesh interactions guided by our proposed AdaLN, which makes the pose and mesh fit the body shape.

uses a Transformer-based encoder-decoder framework to predict SMPL parameters. Despite the promising results of video-based methods, the low-resolution image features in the deep CNN layers inevitably discard detailed spatial information. Besides, their representation abilities are constrained by the limited parametric space, making it difficult to learn the slight and swift human motion in pose and shape domains. Therefore, existing video-based methods still struggle to recover accurate and temporally consistent 3D human mesh. In contrast, our method decouples this task into 3D pose estimation and non-parametric mesh regression by pose and mesh co-evolution, achieving better per-frame accuracy and temporal consistency.

3. Method

The overall architecture of our proposed Pose and Mesh Co-Evolution network (PMCE) is depicted in Figure 2, which comprises two sequential steps: 1) video-based 3D pose estimation and 2) mesh vertices regression from the 3D pose and temporal image feature. The former focuses on the human skeleton and predicts accurate and smooth human motion in terms of the pose. The latter leverages visual cues to supplement the information about the human body shape that recovers accurate human mesh and refines the predicted 3D pose, achieving pose and mesh co-evolution. Specifically, given an input video sequence $\mathbf{V} = \{\mathbf{I}_t\}_{t=1}^T$ with T frames, static image features $\mathbf{F} \in \mathbb{R}^{T \times 2048}$ are extracted by a pre-trained ResNet-50 [17], while 2D poses $\mathbf{P}^{2D} \in \mathbb{R}^{T \times J \times 2}$ are detected using an off-the-shelf 2D pose detector [4, 46], where J denotes the number of body joints. The two-stream

encoder applies dual parallel modules to generate a temporal image feature and estimate the mid-frame 3D pose, respectively. Then, the co-evolution decoder regresses the coordinates of mesh vertices from the 3D pose and temporal feature through pose and mesh interactions. Our proposed Active Layer Normalization (AdaLN) guides the interactions by adjusting the statistical characteristics of joint and vertex features based on the temporal image feature to make pose and mesh fit the human body shape. We elaborate on each part in the following subsections.

3.1. Preliminary

We design our network based on Transformer to model spatial-temporal pose relations and regress 3D human meshes. Thus, we first provide an overview of the fundamental components within the Transformer [39], including Multi-head Self-Attention (MSA), Multi-head Cross-Attention (MCA), and Layer Normalization (LN).

MSA. The Transformer’s input tokens $X \in \mathbb{R}^{n \times d}$ are linearly projected to queries $Q \in \mathbb{R}^{n \times d}$, keys $K \in \mathbb{R}^{n \times d}$ and values $V \in \mathbb{R}^{n \times d}$. The attention is calculated by the scaled dot product, which can be formulated as follows:

$$\text{Attention}(Q, K, V) = \text{Softmax}\left(\frac{QK^T}{\sqrt{d}}\right)V, \quad (1)$$

where n is the sequence length, and d is the dimension. Q, K, V are then split into h heads, each performing scaled dot-product attention in parallel. The final output is the concatenation of h heads mapped by a linear projection

matrix $W \in \mathbb{R}^{d \times d}$, which can be expressed as:

$$\begin{aligned} \text{MSA}(Q, K, V) &= \text{Concat}(H_1, H_2, \dots, H_h) W, \\ H_i &= \text{Attention}(Q_i, K_i, V_i), i \in [1, \dots, h], \end{aligned} \quad (2)$$

MCA. MCA has a similar structure to MSA. The difference is that, in MCA, the query is projected from X , while key and value are projected from another representation Y .

LN. LN [2] normalizes the activities of the neurons in a layer. Given an input token feature $x \in \mathbb{R}^d$, LN can be defined as a linear function adjusted by two trainable parameters, *i.e.*, scaling $\alpha \in \mathbb{R}^d$ and shifting $\beta \in \mathbb{R}^d$:

$$\text{LN}(x; \alpha, \beta) = \alpha \odot \left(\frac{x - \mu(x)}{\sigma(x)} \right) + \beta, \quad (3)$$

$$\mu(x) = \frac{1}{d} \sum_{i=1}^d x_i, \quad \sigma(x) = \sqrt{\frac{1}{d} \sum_{i=1}^d (x_i - \mu(x))^2}, \quad (4)$$

where \odot is element-wise multiplication, μ and σ denotes mean and standard deviation taken across the elements of x , respectively. x is first normalized by μ and σ then scaled and shifted by α and β .

3.2. Two-Stream Encoder

2D pose normalization by the full image. Previous methods [8, 13, 15, 16, 43] of 3D human mesh recovery typically follow the top-down manner where the region of humans is detected and cropped before being processed individually. This manner is effective in reducing background noise and simplifying feature extraction. However, the cropping operation discards the location information in the full image, which is essential to predict the global rotation in the original camera coordinate system [22]. To compensate for the loss of location information without introducing additional computation, we normalize the 2D pose $\mathbf{P}^{2D} \in \mathbb{R}^{T \times J \times 2}$ with respect to the full image instead of the cropped region [8, 51]. The 2D pose normalization is reformulated as:

$$\bar{\mathbf{P}}^{2D} = \frac{2 \times \mathbf{P}^{2D}}{w} - \left[1, \frac{h}{w}\right], \quad (5)$$

where w and h denote the width and height of the image, respectively. $\bar{\mathbf{P}}^{2D} \in \mathbb{R}^{T \times J \times 2}$ is the normalized 2D pose sequence within the range of -1 to 1.

3D pose estimation stream. Our 3D pose estimation stream is built upon the spatial-temporal Transformer (ST-Transformer) [41, 50], which is designed to estimate the mid-frame 3D pose from the 2D pose sequence (the illustration is provided in Sup. Mat.). Given the normalized 2D pose sequence $\bar{\mathbf{P}}^{2D}$, we first project it to the high-dimensional joint feature $\mathbf{X} \in \mathbb{R}^{T \times J \times C_1}$ with the feature dimension C_1 .

Secondly, we project the static image features $\mathbf{F} \in \mathbb{R}^{T \times 2048}$ to $\mathbf{F}' \in \mathbb{R}^{T \times C_1}$, and expand it to $T \times 1 \times C_1$, then add to the joint feature \mathbf{X} . To retain the positional information of spatial and temporal domains, we add the spatial embedding and the temporal embedding of the joint sequence to \mathbf{X} . The ST-Transformer consists of L_1 layers cascaded spatial Transformer and the temporal Transformer. The spatial Transformer aims to explore the spatial information among joints, which calculates the similarities between joint tokens in the same frame. Moreover, to capture the temporal relation among frames, the temporal Transformer reshapes the joint feature \mathbf{X} from $(T \times J \times C_1)$ to $(J \times T \times C_1)$. Thus, the attention matrix is calculated by the similarities between frame tokens of the same joint. Finally, we use a multi-layer perceptron (MLP) to transform the dimension from C_1 to 3 and fuse T frames to one to get the mid-frame 3D pose $P_0 \in \mathbb{R}^{J \times 3}$.

Image feature aggregation stream. The image feature aggregation stream aims to aggregate the static image features of T frames to get a mid-frame temporal feature. Given the static image feature sequence $\mathbf{F} \in \mathbb{R}^{T \times 2048}$, we use a bi-directional GRU [6] to capture the temporal information by recurrently updating the features of neighboring frames. After that, GRU yields an updated temporal image feature $\mathbf{f} \in \mathbb{R}^{2048}$ for mid-frame that aggregates the temporal information of all T frames.

3.3. Co-Evolution Decoder

The co-evolution decoder is proposed to recover the plausible human mesh in a non-parametric manner using the 3D pose P_0 , temporal image feature \mathbf{f} , and coarse template mesh M_0 (provided by SMPL [27]) as input, as shown in Figure 2. The mid-frame 3D pose and temporal image feature generated by the two-stream encoder provide complementary information on pose and shape, respectively. On the one hand, the 3D pose focuses on skeletal motion, which provides more precise and robust pose information than the image feature. On the other hand, the image feature contains visual cues, such as body shape and surface deformation, which are not available in the sparse 3D pose. The complementary information on human pose and shape is crucial for accurate 3D human mesh recovery. Furthermore, the body shape information, serving as a prior, can also refine the estimated 3D pose. Motivated by the above observation, we propose a co-evolution block that performs pose and mesh interactions with an Adaptive Layer Normalization (AdaLN). AdaLN adaptively adjusts the statistical characteristics of joint and vertex features based on the temporal image feature that guides the pose and mesh to fit the human body shape.

Adaptive layer normalization. Inspired by the style transfer task [11] that swaps style from a style image to a content image in the feature space by transferring feature statistics, we propose AdaLN to adaptively adjust the features of joints

and vertices towards the image feature \mathbf{f} . Different from LN [2], the scaling α and shifting β of AdaLN are generated from the image feature. The formulation can be written as:

$$\text{AdaLN}(x, \mathbf{f}) = \alpha(\mathbf{f}) \odot \left(\frac{x - \mu(x)}{\sigma(x)} \right) + \beta(\mathbf{f}), \quad (6)$$

$$\alpha(\mathbf{f}) = \text{MLP}_\alpha(\mathbf{f}), \quad \beta(\mathbf{f}) = \text{MLP}_\beta(\mathbf{f}), \quad (7)$$

where α and β are linearly transformed from the image feature \mathbf{f} by two MLPs, respectively. AdaLN takes the token feature x and image feature \mathbf{f} as inputs and adjusts the mean and standard deviation of x based on \mathbf{f} , allowing x adaptively matches to different \mathbf{f} . In this way, the shape information contained in the image feature can be injected into the joint and vertex features, while preserving their spatial structure.

Co-evolution block. As shown in Figure 3, we design the co-evolution block in a symmetric attention mechanism to model the interactions of pose and mesh. The inputs are the estimated pose $P_0 \in \mathbb{R}^{J \times 3}$, temporal image feature $\mathbf{f} \in \mathbb{R}^{2048}$, and coarse template mesh $M_0 \in \mathbb{R}^{V' \times 3}$ [27], where $V' = 431$ denotes the vertex number of the coarse mesh. The coordinate of each mesh vertex is reinitialized to that of its nearest 3D joint in P_0 , where the distances between joints and vertices are calculated from the template 3D pose and mesh [27]. We first linearly project the pose joints and mesh vertices to high-dimensional features $X_P \in \mathbb{R}^{J \times C_2}$ and $X_M \in \mathbb{R}^{V' \times C_2}$ respectively, then add position embedding to each other. Each feature is normalized by AdaLN with the image feature \mathbf{f} . After that, the vertex feature (joint feature for another branch) serves as the query Q , while joint feature (vertex feature for another branch) is regarded as key K and value V . They are fed to MCA with a residual connection to assist gradient propagation, which can be formulated as:

$$\begin{aligned} X_{M \rightarrow P} &= \text{MCA}(Q_P, K_M, V_M) + X_P, \\ X_{P \rightarrow M} &= \text{MCA}(Q_M, K_P, V_P) + X_M, \end{aligned} \quad (8)$$

where $X_{M \rightarrow P}$ and $X_{P \rightarrow M}$ are the cross-attention features interacting between joint and vertex representations. Then, the cross-attention features are merged by an MLP:

$$\begin{aligned} X'_P &= \text{MLP}(\text{AdaLN}(X_{M \rightarrow P}, \mathbf{f})) + X_{M \rightarrow P}, \\ X'_M &= \text{MLP}(\text{AdaLN}(X_{P \rightarrow M}, \mathbf{f})) + X_{P \rightarrow M}. \end{aligned} \quad (9)$$

Afterward, we perform self-interaction by MSA and MLP, which can be expressed as:

$$\begin{aligned} X''_{P/M} &= \text{MSA}(\text{AdaLN}(X'_{P/M}, \mathbf{f})) + X'_{P/M}, \\ X^{out}_{P/M} &= \text{MLP}(\text{AdaLN}(X''_{P/M}, \mathbf{f})) + X''_{P/M}, \end{aligned} \quad (10)$$

where $X^*_{P/M}$ denotes X^*_P or X^*_M . The joint features $X^{out}_P \in \mathbb{R}^{J \times C_2}$ and vertex feature $X^{out}_M \in \mathbb{R}^{V' \times C_2}$ are

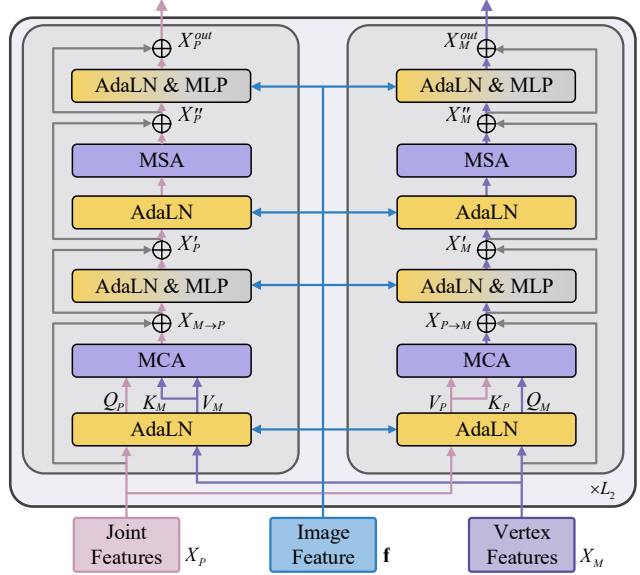


Figure 3: The co-evolution block is designed with a symmetric attention architecture to perform pose and mesh interactions. The joint and vertex features are adjusted by the AdaLN that guides pose and mesh to fit the human body shape based on the image feature.

regressed to the output pose $P \in \mathbb{R}^{J \times 3}$ and coarse mesh $M' \in \mathbb{R}^{V' \times 3}$ by MLPs, respectively. Then, we upsample the coarse mesh M' to the original mesh $M \in \mathbb{R}^{V \times 3}$ by a linear layer, where $V = 6890$. Finally, we add the vertex residuals projected from the temporal image feature \mathbf{f} by MLPs to compensate for the mesh details that are lost during upsampling operation, which can be written as:

$$M = \text{Upsampling}(M') + \text{MLP}_2(\text{MLP}_1(\mathbf{f})^T). \quad (11)$$

3.4. Loss Functions

Following [8, 51], the 3D pose estimation stream is trained using the 3D joint loss \mathcal{L}_{joint} to supervise the intermediate 3D pose P_0 . Then the whole network is supervised by four losses: mesh vertex loss \mathcal{L}_{mesh} , 3D joint loss \mathcal{L}_{joint} , surface normal loss \mathcal{L}_{normal} and surface edge loss \mathcal{L}_{edge} . The final loss is calculated as their weighted sum:

$$\mathcal{L} = \lambda_m \mathcal{L}_{mesh} + \lambda_j \mathcal{L}_{joint} + \lambda_n \mathcal{L}_{normal} + \lambda_e \mathcal{L}_{edge}, \quad (12)$$

where $\lambda_m=1$, $\lambda_j=1$, $\lambda_n=0.1$, and $\lambda_e=20$. More details about the loss functions are in the Sup. Mat.

4. Experiments

4.1. Datasets and Evaluation Metrics

Datasets. Following previous works [7, 41, 43], we adopt mixed 2D and 3D datasets for training. For 3D datasets,

Table 1: Evaluation of state-of-the-art methods on 3DPW, MPI-INF-3DHP, and Human3.6M datasets. All methods use pre-trained ResNet-50 [17] as the backbone to extract static features except MAED [41]. ‘*’ represents using ViT [9] as the backbone. **Bold**: best; Underline: second best.

Method	3DPW				MPI-INF-3DHP			Human3.6M		
	MPJPE ↓	PA-MPJPE ↓	PVE ↓	ACCEL ↓	MPJPE ↓	PA-MPJPE ↓	ACCEL ↓	MPJPE ↓	PA-MPJPE ↓	ACCEL ↓
HMMR (CVPR’19) [14]	116.5	72.6	139.3	15.2	-	-	-	-	56.9	-
VIBE (CVPR’20) [15]	91.9	57.6	99.1	25.4	103.9	68.9	27.3	65.9	41.5	18.3
MEVA (ACCV’20) [28]	86.9	54.7	-	11.6	96.4	65.4	11.1	76.0	53.2	15.3
TCMR (CVPR’21) [7]	86.5	52.7	102.9	<u>7.1</u>	97.6	63.5	<u>8.5</u>	62.3	41.1	5.3
MAED* (ICCV’21) [41]	<u>79.1</u>	45.7	<u>92.6</u>	17.6	<u>83.6</u>	<u>56.2</u>	-	<u>56.4</u>	<u>38.7</u>	-
MPS-NET (CVPR’22) [43]	84.3	52.1	99.7	7.4	96.7	62.8	9.6	69.4	47.4	<u>3.6</u>
PMCE (Ours)	69.5	<u>46.7</u>	84.8	6.5	79.7	54.5	7.1	53.5	37.7	3.1

3DPW [40], Human3.6M [12], and MPI-INF-3DHP [29] contain the annotations of 3D joints and SMPL parameters. For 2D datasets, COCO [25] and MPII [1] contain 2D joint annotation with pseudo SMPL parameters from NeuralAnnot [30]. For evaluation, we report the results on Human3.6M, 3DPW, and MPI-INF-3DHP to quantitatively compare with previous methods.

Metrics. We report results in the metrics of per-frame accuracy and temporal consistency. For accuracy evaluation, we adopt the Mean Per Joint Position Error (MPJPE), Procrustes-Aligned Mean Per Joint Position Error (PA-MPJPE), and Per Vertex Error (PVE), which measure the errors between the estimated results and the ground truth in millimeter (*mm*). For temporal evaluation, acceleration error (ACCEL) in mm/s^2 is used to report the smoothness of human motion.

4.2. Implementation Details

Following previous video-based methods [7, 15, 43], we set the sequence length T to 16 and use ResNet-50 pre-trained in SPIN [17] as the static image feature extractor. For 2D pose detectors, we adopt CPN [4] for Human3.6M and ViTPose [46] for 3DPW and MPI-INF-3DHP. The whole training is divided into two stages, and both are optimized by Adam. In the first stage, we train the 3D pose estimation stream with a batch size of 64 and a learning rate of 5×10^{-5} for 30 epochs. In the second stage, we load the weights of the 3D pose estimation stream and train the whole model end-to-end for 20 epochs with a batch size of 32 and a learning rate of 5×10^{-5} . In the 3D pose estimation stream, we use layer number $L_1 = 3$ and feature dimension $C_1 = 256$. In the co-evolution block, we set layer number $L_2 = 3$ and feature dimension $C_2 = 64$. The network is implemented by PyTorch on one NVIDIA RTX 3090 GPU.

4.3. Comparison with State-of-the-Art Methods

Comparison with video-based methods. Table 1 compares our method with the state-of-the-art video-based methods on the 3DPW, MPI-INF-3DHP, and Human3.6M datasets. The results show that our method outperforms previous video-

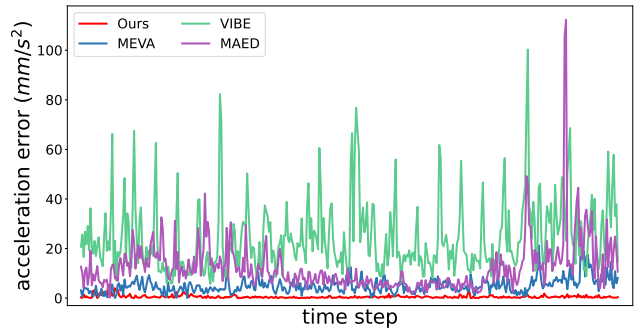


Figure 4: Comparison of the acceleration errors for VIBE, MEVA, MAED, and our PMCE.

based methods, bringing improvements by 12.1% (from 79.1 *mm* to 69.5 *mm*), 4.7% (from 83.6 *mm* to 79.7 *mm*), and 5.1% (from 56.4 *mm* to 53.5 *mm*) on the three datasets in the metric of MPJPE. Although MAED also makes great progress, it trains ViT [9] as the backbone and takes more fine-grained feature maps, which is training time-costing and memory-intensive [7]. In contrast, other methods [7, 15, 43] and our PMCE use ResNet-50 pre-trained by SPIN [17] to extract static feature vectors. Moreover, MAED has a trade-off between per-frame accuracy (PA-MPJPE) and temporal consistency (ACCEL). Specifically, when MAED reduces PA-MPJPE by 1 *mm*, it increases ACCEL by 11.1 mm/s^2 compared to our PMCE on 3DPW. To further evaluate the temporal consistency, Figure 4 compares the acceleration errors on the sequence ‘courtyard_drinking_00’ of 3DPW. Previous methods [15, 28, 41] reveal large acceleration errors, representing unsmooth and unstable motion estimations. In contrast, our method has relatively low acceleration errors, indicating the temporal consistency of our predictions. Overall, our method can estimate accurate and smooth 3D human motion from a video. The results demonstrate that first using the skeleton sequence to explore spatial-temporal pose information and then regressing the mesh vertices by interacting pose and shape information can effectively achieve per-frame accuracy and temporal consistency.

Table 2: Comparison with single pose-based methods. All methods are not trained on 3DPW.

Method	3DPW			
	MPJPE ↓	PA-MPJPE ↓	PVE ↓	ACCEL ↓
PQ-GCN (TCSVT'22) [42]	89.2	58.3	106.4	-
Pose2Mesh (ECCV'20) [8]	88.9	58.3	106.3	22.6
GTRS (ACM MM'22) [51]	88.5	58.9	106.2	25.0
PMCE (Ours)	81.6	52.3	99.5	6.8

Table 3: Ablation study for the designs of 3D pose estimation stream on Human3.6M. ‘Seq.’ means sequence input.

Method	MPJPE ↓	PA-MPJPE ↓	PVE ↓	ACCEL ↓
Cropped pose Seq.	58.2	40.8	66.3	3.3
Cropped pose Seq. + bbox info	56.4	38.6	65.0	3.4
Full image pose Seq.	53.5	37.7	61.3	3.1
Mid-frame full image pose	55.8	39.1	64.3	17.9

Comparison with single pose-based methods. We further compare our PMCE to single pose-based methods [8, 42, 51] since they also take a detected 2D pose from the image to estimate the intermediate 3D pose and then regress the mesh vertices from the estimated 3D pose, which is relevant to our method. The difference is that our PMCE expands the single pose-based framework to multiple frames and complements it with image cues for 3D human motion estimation. Table 2 compares their performance on the 3DPW dataset. Notably, the pose-based methods did not use the 3DPW training set during training, so all methods are not trained on 3DPW. Our method outperforms pose-based methods in per-frame accuracy and temporal consistency by a large margin, improving by 7.8% in MPJPE and 69.9% in ACCEL.

4.4. Ablation Study

Effectiveness of 3D pose estimation stream. There are two designs in the 3D pose estimation stream, *i.e.* 2D pose normalization by the full image and sequence input, which are respectively evaluated in Table 3. When normalizing the 2D pose in the cropped bounding box, the accuracy is low since the keypoints are located in regions with unfixed locations and scales. Then, adding bounding box information improves the performance, but the relations between keypoints and bounding boxes are still implicit and hard to explore. Furthermore, normalizing the 2D pose in full image achieves the best performance since the poses in different frames share the unified coordinate system, which provides clear location information. Besides, it can be observed that single-frame input (last line) yields low temporal consistency, while sequence inputs (top three lines) improve it significantly. This result suggests that the 3D pose estimation stream can effectively capture the temporal information of sequence and improves the temporal consistency.

Table 4: Ablation study of image features adding in different modules on 3DPW. ‘Enc.’ denotes the two-stream encoder. ‘Dec.’ denotes the co-evolution decoder.

Method	Intermediate Pose P_0	Output Mesh M	
	MPJPE ↓	MPJPE ↓	PVE ↓
Only pose	72.6	76.9	92.1
F in Enc.	70.5	72.2	89.4
F in Dec.	72.6	71.6	86.5
F in Enc. & Dec.	70.5	69.5	84.8

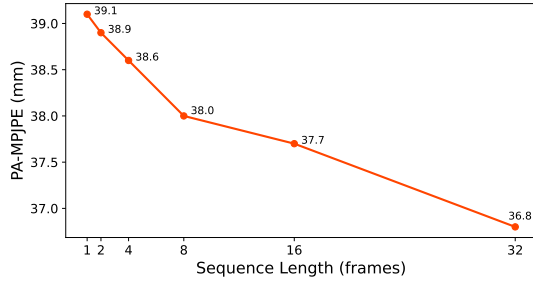
Table 5: Ablation study for the designs of the co-evolution decoder on 3DPW. We report the results of output 3D pose P and output mesh M . MPJPE of the input 3D pose is 70.5 *mm*.

Method	Output Pose P	Output Mesh M	
	MPJPE ↓	MPJPE ↓	PVE ↓
w/o interactions	70.6	72.3	87.7
Pose→Mesh	69.7	70.8	85.7
Mesh→Pose	68.8	71.6	87.6
Mesh↔Pose (Ours)	67.3	69.5	84.8
ALADIN [35]	71.3	72.8	88.6
LN	71.2	72.0	87.9
Bilinear [37]	70.5	71.3	87.1
AdaLN [11]	70.6	71.3	86.5
AdaNorm [44]	69.4	70.7	86.9
SEAN [53]	68.7	70.5	85.9
AdaLN (Ours)	67.3	69.5	84.8

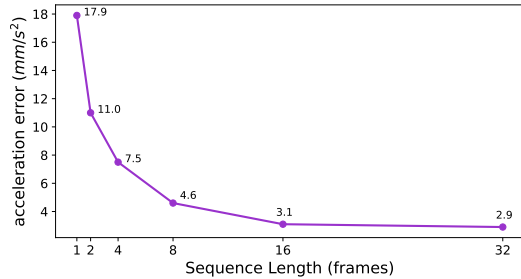
Effectiveness of image feature. Table 4 evaluates the impact of image features **F** on the intermediate 3D pose P_0 and the output mesh M . When only taking the 2D pose sequence as input, the mesh performance is low since the pose lacks shape information. Adding the image feature in the encoder or decoder can improve their individual output results, while using the image features in both the encoder and decoder achieves the best results, which indicates that the shape information contained in image features contributes to better results of intermediate 3D pose and output mesh.

Effectiveness of pose and mesh interactions. The co-evolution decoder is proposed to regress the mesh vertices from the estimated 3D pose and temporal image feature, containing pose and mesh interactions and AdaLN. In the first part of Table 5, it can be seen that the pose→mesh interaction improves the performance of output mesh M since the mesh can learn more accurate position information guided by the 3D pose. Meanwhile, the mesh→pose interaction enhances the performance of output 3D pose P , which indicates that the priors (*e.g.*, body shape, limb proportion) in mesh contributes to 3D pose refinement. Moreover, using bi-directional interactions further improves the performance of pose and mesh, which achieves their co-evolution.

Effectiveness of AdaLN. In the second part of Table 5, we evaluate the effect of AdaLN by degenerating it to



(a) Different sequence lengths under PA-MPJPE.



(b) Different sequence lengths under acceleration error.

Figure 5: Ablation studies on different sequence lengths in terms of PA-MPJPE and acceleration error on Human3.6M.

LN or replacing it with other feature fusion ways, including bilinear fusion [37] and some other normalization blocks [11, 35, 44, 53]. Specifically, bilinear fusion directly combines the image feature with joint and vertex features, which increases the complexity to learn the pose and shape information. The other normalization blocks focus on the fields of 2D image [11, 35, 53] or language [44], and design normalization approaches in different feature dimensions and statistical parameters towards their specific tasks, but not suitable for the task of 3D human body recovery. In contrast, AdaLN lies in its tailored adaptation for fitting pose and mesh to individual body shapes by adjusting the statistics of joint and vertex features towards the image features, which enhances the ability of normalization block in 3D human body recovery.

Impact of sequence lengths. For video-based methods, sequence length has a direct impact on the results. Figure 5 illustrates the influences of sequence length in (a) accuracy and (b) temporal consistency on the Human3.6M dataset. Increasing the sequence length can improve the performance in terms of both accuracy and temporal consistency. These results indicate that our method can effectively exploit spatial and temporal relations in sequence frames to estimate accurate and smooth 3D human motion. Note that, for a fair comparison with previous video-based methods [7, 15, 41], we choose the sequence length as 16 in our experiments.

Impact of 2D pose detections. 2D poses are used in the 3D pose estimation stream. To evaluate their impact, we

Table 6: Ablation study of different 2D pose detections on Human3.6M.

2D pose detection	MPJPE ↓	PA-MPJPE ↓	PVE ↓	ACCEL ↓
Detected by SH [31]	56.4	39.0	64.5	3.2
Detected by Detectron [34]	55.9	39.0	64.1	3.2
Detected by CPN [4]	53.5	37.7	61.3	3.1
Ground truth 2D pose	36.3	26.8	46.2	2.2

conduct experiments in different 2D poses, including the ground truth 2D poses and some detected 2D poses [4, 31, 34]. Table 6 illustrates the results on Human3.6M. Our method is robust to different 2D pose detections, and adopting the 2D pose detectors with higher precision can improve our reconstructed accuracy. The performance with ground truth 2D poses indicates the lower bound of the proposed PMCE, which shows the potential of our method to stay effective with the development of 2D pose detectors.

4.5. Qualitative evaluation

Visual comparison with MPS-Net. Figure 6 shows the qualitative comparison between the previous state-of-the-art method MPS-Net [43] and our PMCE on the in-the-wild 3DPW dataset. Our method can generate more plausible human meshes, especially on the arms and legs. When an occlusion occurs (the first sample in Row 2), our PMCE can infer the accurate result based on spatial and temporal relations in a video sequence.

Visualization of human body shape. Previous video-based methods [7, 15, 43] take 10-dimensional shape parameters of SMPL to control human body shape and supervise it in the range of mean shape, which limits their representation ability for various body shapes. On the contrary, our method relieves the limitation by directly regressing mesh vertices and fitting the mesh to the body shape through the proposed AdaLN. To verify the representation ability, we train and test MPS-Net [43] and our PMCE on SURREAL [38] dataset, which contains various body shapes. Figure 7 shows that MPS-Net tends to produce mean shapes, which reflect unreal body patterns. In contrast, the meshes generated by PMCE are more flexible and align better with the input images. It demonstrates that our PMCE has the ability to learn different human body shapes and recover plausible 3D human meshes.

5. Conclusion

This paper proposes the Pose and Mesh Co-Evolution network (PMCE), a new two-stage pose-to-mesh framework for recovering 3D human mesh from a monocular video. PMCE first estimates 3D human pose motion in terms of spatial and temporal domains, then performs image-guided pose and mesh interactions by our proposed AdaLN that injects body shape information while preserving their spatial structure. Extensive experiments on popular datasets show



Figure 6: Visual comparison between MPS-Net [43] (pink meshes) and our PMCE (blue meshes) on the challenging 3DPW dataset, which contains hard poses, fast motions, and occlusions. Our method can generate more plausible mesh results than MPS-Net.

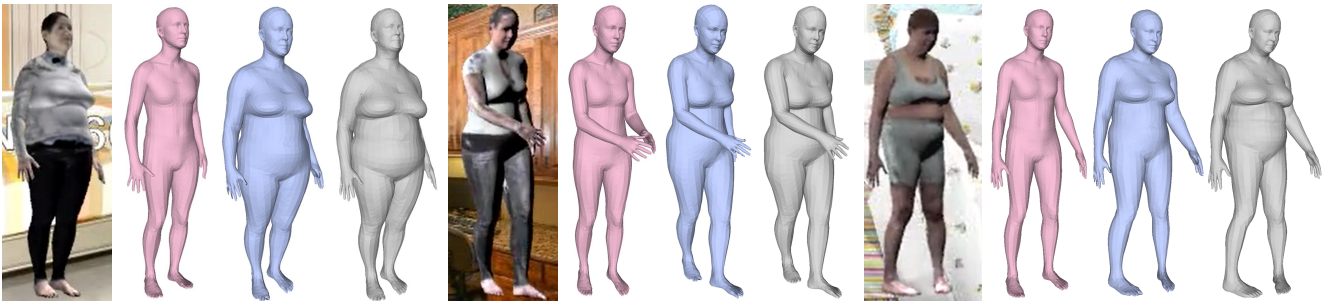


Figure 7: Mesh results in extreme shapes. Pink meshes are from MPS-Net [43], blue meshes are from PMCE, and grey meshes are ground truth. Our PMCE can fit the body shapes better.



Figure 8: Failure cases in self-contact and challenging pose.

that PMCE outperforms state-of-the-art methods in both per-frame accuracy and temporal consistency. We hope that our approach will spark further research in 3D human motion estimation considering both pose and shape consistency.

Limitation. Main limitation of our method comes from self-contacts and challenging poses, as shown in Figure 8. Due to the low contact perception and limited human body priors, the reconstructed meshes may be implausible. Our future work will explore the physical constraints of human body to alleviate this issue.

Acknowledgment. This work is supported by the National Natural Science Foundation of China (No. 62073004), Basic and Applied Basic Research Foundation of Guangdong (No. 2020A1515110370), Shenzhen Fundamental Research Program (GXWD20201231165807007-20200807164903001, No. JCYJ20200109140410340), the interdisciplinary doctoral grants (iDoc 2021-360) from the Personalized Health and Related Technologies (PHRT) of the ETH domain.

References

- [1] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2D human pose estimation: New benchmark and state of the art analysis. In *CVPR*, pages 3686–3693, 2014.
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [3] Tianlang Chen, Chen Fang, Xiaohui Shen, Yiheng Zhu, Zhili Chen, and Jiebo Luo. Anatomy-aware 3D human pose estimation with bone-based pose decomposition. *IEEE TCSVT*, 32(1):198–209, 2021.
- [4] Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, and Jian Sun. Cascaded pyramid network for multi-person pose estimation. In *CVPR*, pages 7103–7112, 2018.
- [5] Junhyeong Cho, Kim Youwang, and Tae-Hyun Oh. Cross-attention of disentangled modalities for 3D human mesh recovery with transformers. In *ECCV*, pages 342–359, 2022.
- [6] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *EMNLP*, 2014.
- [7] Hongsuk Choi, Gyeongsik Moon, Ju Yong Chang, and Kyoung Mu Lee. Beyond static features for temporally consistent 3D human pose and shape from a video. In *CVPR*, pages 1964–1973, 2021.
- [8] Hongsuk Choi, Gyeongsik Moon, and Kyoung Mu Lee. Pose2Mesh: Graph convolutional network for 3D human pose and mesh recovery from a 2D human pose. In *ECCV*, pages 769–787, 2020.
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: transformers for image recognition at scale. In *ICLR*, 2021.
- [10] Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. DensePose: Dense human pose estimation in the wild. In *CVPR*, pages 7297–7306, 2018.
- [11] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, pages 1501–1510, 2017.
- [12] 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 TPAMI*, 36(7):1325–1339, 2013.
- [13] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *CVPR*, pages 7122–7131, 2018.
- [14] Angjoo Kanazawa, Jason Y Zhang, Panna Felsen, and Jitendra Malik. Learning 3D human dynamics from video. In *CVPR*, pages 5614–5623, 2019.
- [15] Muhammed Kocabas, Nikos Athanasiou, and Michael J Black. VIBE: Video inference for human body pose and shape estimation. In *CVPR*, pages 5253–5263, 2020.
- [16] Muhammed Kocabas, Chun-Hao P Huang, Otmar Hilliges, and Michael J Black. PARE: Part attention regressor for 3D human body estimation. In *ICCV*, pages 11127–11137, 2021.
- [17] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3D human pose and shape via model-fitting in the loop. In *ICCV*, pages 2252–2261, 2019.
- [18] Nikos Kolotouros, Georgios Pavlakos, and Kostas Daniilidis. Convolutional mesh regression for single-image human shape reconstruction. In *CVPR*, pages 4501–4510, 2019.
- [19] Jiefeng Li, Chao Xu, Zhicun Chen, Siyuan Bian, Lixin Yang, and Cewu Lu. HybrIK: A hybrid analytical-neural inverse kinematics solution for 3D human pose and shape estimation. In *CVPR*, pages 3383–3393, 2021.
- [20] Wenhao Li, Hong Liu, Runwei Ding, Mengyuan Liu, Pichao Wang, and Wenming Yang. Exploiting temporal contexts with strided transformer for 3D human pose estimation. *IEEE TMM*, 25:1282–1293, 2023.
- [21] Wenhao Li, Hong Liu, Hao Tang, Pichao Wang, and Luc Van Gool. MHFormer: Multi-hypothesis transformer for 3D human pose estimation. In *CVPR*, pages 13147–13156, 2022.
- [22] Zhihao Li, Jianzhuang Liu, Zhensong Zhang, Songcen Xu, and Youliang Yan. CLIFF: Carrying location information in full frames into human pose and shape estimation. In *ECCV*, pages 590–606, 2022.
- [23] Kevin Lin, Lijuan Wang, and Zicheng Liu. End-to-end human pose and mesh reconstruction with transformers. In *CVPR*, pages 1954–1963, 2021.
- [24] Kevin Lin, Lijuan Wang, and Zicheng Liu. Mesh graphormer. In *ICCV*, pages 12939–12948, 2021.
- [25] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, pages 740–755, 2014.
- [26] Ruixu Liu, Ju Shen, He Wang, Chen Chen, Sen-ching Cheung, and Vijayan Asari. Attention mechanism exploits temporal contexts: Real-time 3D human pose reconstruction. In *CVPR*, pages 5064–5073, 2020.
- [27] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. SMPL: A skinned multi-person linear model. *ACM TOG*, 34(6):1–16, 2015.
- [28] Zhengyi Luo, S Alireza Golestaneh, and Kris M Kitani. 3D human motion estimation via motion compression and refinement. In *ACCV*, 2020.
- [29] 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 *3DV*, pages 506–516, 2017.
- [30] Gyeongsik Moon, Hongsuk Choi, and Kyoung Mu Lee. Neuralannot: Neural annotator for 3D human mesh training sets. In *CVPR*, pages 2299–2307, 2022.
- [31] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hour-glass networks for human pose estimation. In *ECCV*, pages 483–499, 2016.
- [32] Mohamed Omran, Christoph Lassner, Gerard Pons-Moll, Peter Gehler, and Bernt Schiele. Neural body fitting: Unifying deep learning and model based human pose and shape estimation. In *3DV*, pages 484–494, 2018.

- [33] Georgios Pavlakos, Luyang Zhu, Xiaowei Zhou, and Kostas Daniilidis. Learning to estimate 3D human pose and shape from a single color image. In *CVPR*, pages 459–468, 2018.
- [34] Dario Pavlo, Christoph Feichtenhofer, David Grangier, and Michael Auli. 3D human pose estimation in video with temporal convolutions and semi-supervised training. In *CVPR*, pages 7753–7762, 2019.
- [35] Dan Ruta, Saeid Motiian, Baldo Faieta, Zhe Lin, Hailin Jin, Alex Filipkowski, Andrew Gilbert, and John Collomosse. ALADIN: All layer adaptive instance normalization for fine-grained style similarity. In *ICCV*, pages 11926–11935, 2021.
- [36] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *CVPR*, pages 5693–5703, 2019.
- [37] Yu Sun, Yun Ye, Wu Liu, Wenpeng Gao, Yili Fu, and Tao Mei. Human mesh recovery from monocular images via a skeleton-disentangled representation. In *ICCV*, pages 5349–5358, 2019.
- [38] Gul Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J Black, Ivan Laptev, and Cordelia Schmid. Learning from synthetic humans. In *CVPR*, pages 109–117, 2017.
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 30, 2017.
- [40] Timo von Marcard, Roberto Henschel, Michael J Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3D human pose in the wild using imus and a moving camera. In *ECCV*, pages 601–617, 2018.
- [41] Ziniu Wan, Zhengjia Li, Maoqing Tian, Jianbo Liu, Shuai Yi, and Hongsheng Li. Encoder-decoder with multi-level attention for 3D human shape and pose estimation. In *ICCV*, pages 13033–13042, 2021.
- [42] Lei Wang, Xunyu Liu, Xiaoliang Ma, Jiayi Wu, Jun Cheng, and Mengchu Zhou. A progressive quadric graph convolutional network for 3D human mesh recovery. *IEEE TCSVT*, 2022.
- [43] Wen-Li Wei, Jen-Chun Lin, Tyng-Luh Liu, and Hong-Yuan Mark Liao. Capturing humans in motion: Temporal-attentive 3D human pose and shape estimation from monocular video. In *CVPR*, pages 13211–13220, 2022.
- [44] Jingjing Xu, Xu Sun, Zhiyuan Zhang, Guangxiang Zhao, and Junyang Lin. Understanding and improving layer normalization. *NeurIPS*, 32, 2019.
- [45] Peng Xu, Na Jiang, Jun Li, and Zhiping Shi. Learning attention map for 3D human recovery from a single RGB image. In *BMVC*, 2021.
- [46] Yufei Xu, Jing Zhang, Qiming Zhang, and Dacheng Tao. Vitpose: Simple vision transformer baselines for human pose estimation. *NeurIPS*, 35:38571–38584, 2022.
- [47] Zhenbo Yu, Junjie Wang, Jingwei Xu, Bingbing Ni, Chenglong Zhao, Minsi Wang, and Wenjun Zhang. Skeleton2Mesh: Kinematics prior injected unsupervised human mesh recovery. In *ICCV*, pages 8619–8629, 2021.
- [48] Ailing Zeng, Xiao Sun, Fuyang Huang, Minhao Liu, Qiang Xu, and Stephen Lin. SRNet: Improving generalization in 3D human pose estimation with a split-and-recombine approach. In *ECCV*, pages 507–523, 2020.
- [49] Hongwen Zhang, Jie Cao, Guo Lu, Wanli Ouyang, and Zhenan Sun. Learning 3D human shape and pose from dense body parts. *IEEE TPAMI*, 44(5):2610–2627, 2020.
- [50] 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 *CVPR*, pages 13232–13242, 2022.
- [51] Ce Zheng, Matias Mendieta, Pu Wang, Aidong Lu, and Chen Chen. A lightweight graph transformer network for human mesh reconstruction from 2D human pose. In *ACM MM*, 2022.
- [52] Ce Zheng, Sijie Zhu, Matias Mendieta, Taojiannan Yang, Chen Chen, and Zhengming Ding. 3D human pose estimation with spatial and temporal transformers. In *ICCV*, pages 11656–11665, 2021.
- [53] Peihao Zhu, Rameen Abdal, Yipeng Qin, and Peter Wonka. SEAN: Image synthesis with semantic region-adaptive normalization. In *CVPR*, pages 5104–5113, 2020.