

# PSVT: End-to-End Multi-person 3D Pose and Shape Estimation with Progressive Video Transformers

Zhongwei Qiu<sup>1,3,4</sup>, Qiansheng Yang<sup>2</sup>, Jian Wang<sup>2</sup>, Haocheng Feng<sup>2</sup>, Junyu Han<sup>2</sup>,  
 Errui Ding<sup>2</sup>, Chang Xu<sup>3</sup>, Dongmei Fu<sup>1,4</sup>, Jingdong Wang<sup>2</sup>

<sup>1</sup>School of Automation and Electrical Engineering, University of Science and Technology Beijing

<sup>2</sup>Baidu, <sup>3</sup>University of Sydney, <sup>4</sup>Beijing Engineering Research Center of Industrial Spectrum Imaging

## Abstract

Existing methods of multi-person video 3D human Pose and Shape Estimation (PSE) typically adopt a two-stage strategy, which first detects human instances in each frame and then performs single-person PSE with temporal model. However, the global spatio-temporal context among spatial instances can not be captured. In this paper, we propose a new end-to-end multi-person 3D Pose and Shape estimation framework with progressive Video Transformer, termed PSVT. In PSVT, a spatio-temporal encoder (STE) captures the global feature dependencies among spatial objects. Then, spatio-temporal pose decoder (STPD) and shape decoder (STSD) capture the global dependencies between pose queries and feature tokens, shape queries and feature tokens, respectively. To handle the variances of objects as time proceeds, a novel scheme of progressive decoding is used to update pose and shape queries at each frame. Besides, we propose a novel pose-guided attention (PGA) for shape decoder to better predict shape parameters. The two components strengthen the decoder of PSVT to improve performance. Extensive experiments on the four datasets show that PSVT achieves state-of-the-art results.




Figure 1. Comparison of multi-stage and end-to-end framework. (a) Existing video-based methods [4, 16, 49, 50] perform single-person pose and shape estimation (SPSE) on the cropped areas by temporal modeling, such as Gated Recurrent Units (GRUs). (b) PSVT achieves end-to-end multi-person pose and shape estimation in video with spatial-temporal encoder (STE) and decoder (STD).

## 1. Introduction

Multi-person 3D human Pose and Shape Estimation (PSE) from monocular video requires localizing the 3D joint coordinates of all persons and reconstructing their human meshes (e.g. SMPL [29] model). As an essential task in computer vision, it has many applications including human-robot interaction detection [22], virtual reality [33], and human behavior understanding [8], etc. Although remarkable progress has been achieved in PSE from videos [4, 41, 50, 51] or images [5, 44, 45], capturing multi-person spatio-temporal relations of pose and shape simultaneously is still challenging since the difficulty in modeling long-range global interactions.

To tackle this challenge, as shown in Figure 1 (a), existing methods [4, 16, 50, 51] employ a detection-based strategy of firstly detecting each human instance, then cropping the instance area in each frame and feeding it into the temporal model, such as the recurrent neural network [4, 6, 16]. However, this framework can not capture the spatial relationship among human instances in an image and has the limitation of extracting long-range global context. Besides, the computational cost is expensive since it is proportional to the number of instances in image and it needs extra tracker [50] to identify each instance. Other temporal smoothing methods [15, 47] adopt a post-processing module to align the shape estimated by image-based PSE ap-

proaches [5, 15, 23, 44, 45]. However, they can not capture temporal information directly from visual image features and lack the ability of long-range global interactions. These multi-stage methods split space and time dimensions and can not be end-to-end optimized.

To strengthen the long-range modeling ability, recently developed Transformer models [7, 46] have been introduced in PSE. The Transformer-based mesh reconstruction approaches [24, 25, 36, 54] take each human joint as a token and capture the relationship of human joints by attention mechanism. However, the global context among different persons in spatio-temporal dimensions has not been explored. Other Transformer-based human pose estimation approaches [27, 57] explore the spatio-temporal context of human joints for single-person, but not on the multi-person mesh. Besides, these methods focus on capturing the relations among human joints, while ignoring the relations between human poses and shapes.

To tackle the above problems, we propose an end-to-end multi-person 3D Pose and Shape estimation framework with Video Transformer, termed PSVT, to capture long-range spatio-temporal global interactions in the video. As shown in Figure 1 (b), PSVT formulates the human instance localization and fine-grained pose and mesh estimation as a set prediction problem as [3, 42]. First, PSVT extracts a set of spatio-temporal tokens from the deep visual features and applies a spatio-temporal encoder (STE) on these visual tokens to learn the relations of feature tokens. Second, given a set of pose queries, a progressive spatio-temporal pose decoder (STPD) learns to capture the relations of human joints in both spatial and temporal dimensions. Third, with the guidance of pose tokens from STPD, a progressive spatio-temporal shape decoder (STSD) learns to reason the relations of human mesh and pose in both spatial and temporal dimensions and further estimates the sequence 3D human mesh based on the spatio-temporal global context. Compared with previous shape estimation works [4, 5, 44, 45, 50, 51], PSVT achieves end-to-end multi-person 3D pose and shape estimation in video.

In PSVT, different from previous methods, we propose a novel progressive decoding mechanism (PDM) for sequence decoding and pose-guided attention (PGA) for decoder. PDM takes the output tokens from the last frame as the initialized queries for next frame, which enables better sequence decoding for STPD and STSD. PGA aligns the pose tokens and shape queries and further computes the cross-attention with feature tokens from encoder. With the guidance of pose tokens, shape estimation can be more accurate. Our contributions can be summarized as follows:

- We propose a novel video Transformer framework, termed PSVT, which is the first end-to-end multi-person 3D human pose and shape estimation framework with video Transformer.

- We propose a novel progressive decoding mechanism (PDM) for the decoder of video Transformer, which updates the queries at each frame in the attention block to improve the pose and shape decoding.
- We propose a novel pose-guided attention (PGA), which can capture the spatio-temporal relations among pose tokens, shape tokens, and feature tokens to improve the performance of shape estimation.
- Extensive experiments on four benchmarks show that PSVT achieves new state-of-the-art results.

## 2. Related Work

### 2.1. Imaged-based 3D Human PSE

Image-based 3D human PSE methods [5, 14, 38, 39, 44, 49, 55] estimate the 3D pose, shape, and camera parameters from single RGB image, further to decode human mesh by SMPL model [29]. They can be divided into single-person and multi-person methods.

**Single-Person.** HMR [14] proposes an end-to-end single-person human mesh recovery framework by estimating the parameters of the SMPL model. Based on SMPL model, 2D heatmaps and silhouettes are used as the prior information to improve mesh estimation by [37]. Following this framework, more prior knowledge [55] or stronger backbone network [49] are adopted to improve the performance of shape estimation. To deal with the occlusion problem, 3DCrowdNet [5] estimates robust 3D human mesh from in-the-wild crowded scenes by using 2D keypoints heatmaps as the key cues. Although these methods achieve great performance on 3D human PSE, the performances of these methods rely on the accuracy of human detection.

**Multi-Person.** Most existing multi-person methods [5, 19, 35, 40, 55] adopt a multi-stage framework to tackle multi-person problem, which firstly conducts human detection and follows a single-person mesh estimation model. However, the multi-stage framework is low-efficiency and can not be optimized in an end-to-end fashion. To solve this problem, ROMP [44] achieves a novel one-stage pipeline of multi-person image 3D human pose and shape estimation, which directly estimates multiple 2D maps for 2D human detection, positioning, and mesh parameter regression. To capture the relative depth of multiple persons in image, BEV [45] uses an additional imaginary Bird’s-Eye-View representation to explicitly reason about depth. The single-stage methods [44, 45] are high-efficiency, but cannot handle the small object well.

Although the image-based 3D human pose and shape estimation methods achieve remarkable progress, they lack competitiveness compared with video-based methods because temporal information is important to improve the performance of 3D pose and shape estimation.

## 2.2. Video-based 3D Human PSE

Video-based 3D human pose and shape estimation methods [4, 16, 49–51] can extract more temporal context to keep the consistency of pose and shape on the time dimension. Usually, these video-based methods follow a pipeline of first detecting human instances, and then conducting temporal modeling for a single instance by temporal models, such as recurrent neural networks and 3D convolution. Typically, VIBE [16] builds a temporal encoder by bidirectional gated recurrent units (GRU) to encode the static feature from the backbone network into temporal features, further to regress SMPL parameters by a regressor. TCMR [4] uses a GRU-based temporal encoder to extract temporal information with three different encoding strategies. GLAMR [51] formulates a four-stage framework with global motion optimization to tackle the occlusion problem in video human mesh recovery. Besides, MPS-Net [50] captures human motion among different temporal frames and fuses these motion features to estimate pose and mesh parameters by attention-based feature extraction and integration modules.

These video-based methods are multi-stage and perform single-person temporal modeling in the pose and shape estimation stage. Despite the improvements that have been achieved by these temporal models, they are limited to the human detector and cannot be optimized by an end-to-end scheme. Moreover, they cannot capture the depth information between different human instances due to crop images.

## 2.3. Transformers in 3D Human PSE

Recently, Transformer-based models [17, 24, 25, 42, 49, 51] have been introduced in human pose and shape estimation since their strong long-term modeling capabilities of sequence. They can be divided into two categories: capturing joint relations [17, 24, 25] and capturing temporal relations [49, 51]. METRO [24] and Mesh Graphformer [25] design end-to-end single-person Transformer and Graph-based Transformer to capture the human joints relations, respectively. However, they cannot process the video problem. MEAD [49] uses a spatio-temporal Transformer to handle the spatio-temporal relations of human joints. Besides, GLAMR [51] uses the attention mechanism to deal with the occlusion problem in video. However, existing Transformer-based methods only focus on the single-person pose and estimation problem, and use attention to capture the relations of human joints in spatial or temporal dimensions. Moreover, they cannot capture the relation of multiple human instances in video and their multi-stage framework is expensive. In this paper, we propose a novel end-to-end Transformer-based multi-person human pose and shape estimation framework to capture the relation of human instances in global spatial and temporal dimensions. Besides, we propose new pose-guided attention to decode the shape parameters of human instances.

## 3. Method

In this section, we first introduce 3D human shape reconstruction models SMPL [29] and SMPL+A [34]. Then, we propose PSVT, 3D human Pose and Shape estimation with progressive Video Transformer. We build a video transformer baseline with vanilla attention. After that, we extend the baseline to a progressive video transformer with progressive decoding mechanism and pose-guided shape attention.

### 3.1. SMPL-based Pose and Shape Estimation

SMPL [29] is a widely-used 3D body shape estimation method, which parameterizes human mesh into low-dimensional parameters. SMPL has been extended to SMIL [9] and SMPL+A [34] to tackle infant and age problems. Following [34], we adopt the SMPL+A model, which can output human 3D mesh  $\mathcal{M}(\theta, \beta, \alpha) \in \mathbb{R}^{6890 \times 3}$  with 3D pose  $\theta$ , shape  $\beta$ , and age offset  $\alpha \in [0, 1]$ . The pose parameters  $\theta \in \mathbb{R}^{6 \times 22}$  include the 6D rotations of the 22 body joints in the SMPL+A model. The parameters  $\beta \in \mathbb{R}^{10}$  are the top-10 PCA coefficients of shape space.

Given the a video  $V = \{I^t | t \in [1, T]\}$  including  $T$  frames, where  $I^t \in \mathbb{R}^{H \times W \times 3}$  means  $t^{th}$  frame of height  $H$  and width  $W$ , multi-person video 3D human pose and shape estimation aims to output 3D joints  $J = \{J_i^t | t \in [1, T], i \in [1, N]\}$  and mesh  $\mathcal{M} = \{\mathcal{M}_i^t | t \in [1, T], i \in [1, N]\}$  for each instance in the video, where  $N$  represents the number of instances in an image.  $J_i^t \in \mathbb{R}^{K \times 3}$  represents the 3D joints of person  $i$  in  $t^{th}$  frame, which can be generated as  $J_i^t = \mathcal{W}\mathcal{M}_i^t$ , where  $\mathcal{W} \in \mathbb{R}^{K \times 6890}$  is weight matrix to map mesh into joints.  $K$  is the number of joints.

### 3.2. Video Transformer Baseline

Different from existing multi-stage video-based methods [4, 16, 49–51], we propose an end-to-end multi-person 3D pose and shape estimation framework, which consists of backbone network, spatio-temporal encoder, and decoder.

**Framework.** The framework of the video transformer baseline is shown in Figure 2 (a). Given the input video  $V$ , deep features  $F^t \in \mathbb{R}^{\frac{H}{s} \times \frac{W}{s} \times C}$  are extracted for each frame by backbone network HRNet [43], where  $s$  and  $C$  represent scale and feature dim, respectively. Then,  $F^t$  are extracted as tokens  $\tau_e^t \in \mathbb{R}^{L \times D}$  by patch embedding with a patch size of  $B \times B$ .  $L$  represents the number of tokens and  $D$  is embedding dim. For the video including  $T$  frames,  $T \times L$  tokens  $\tau_e = \{\tau_e^t | t \in [1, T]\}$  are generated. These tokens are sent into a spatio-temporal encoder. After that,  $L$  pose queries and  $L$  shape queries for each frame, which have the same spatial positions with tokens  $\tau_e^t$ , are sent into the pose decoder and shape decoder to reason joints heatmaps and mesh parameters, respectively.

**Encoder.** The feature maps  $F$  are extracted as tokens  $\tau_e \in \mathbb{R}^{T \times L \times D}$ , which serve as the input of the transformer

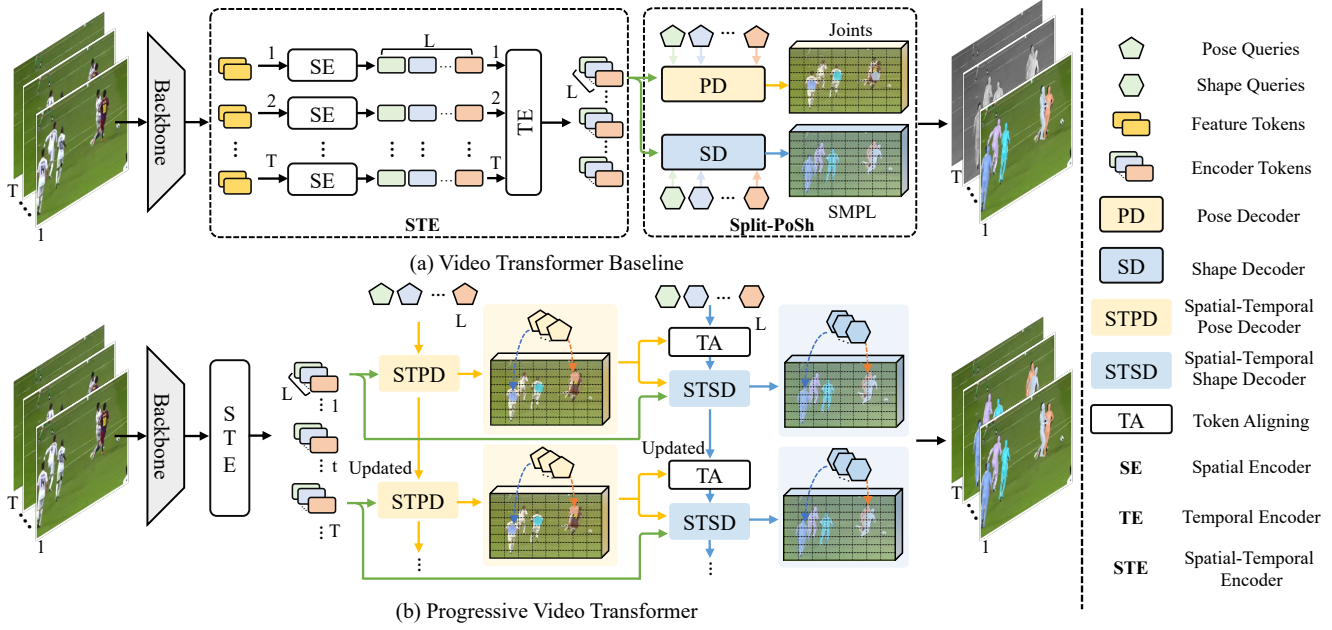


Figure 2. The overview of (a) Video Transformer Baseline (VTL) and (b) Progressive Video Transformer for Pose and Shape estimation (PSVT). A spatial-temporal encoder captures global feature interactions for the decoders of VTL and PSVT. VTL splits the pose decoder and shape decoder (Split-PoS) to localize human joints and regress the shape parameter of the SMPL model. PSVT adopts a progressive decoding mechanism and pose-guided shape decoder.

encoder. To capture the relations among tokens in both spatial and temporal dimensions and reduce the computational costs, we adopt divided attention for video encoder, which firstly computes the self-attention among the  $L$  spatial tokens and then follows the self-attention among  $T$  tokens in the same spatial place. This scheme is termed Spatio-Temporal Encoder (STE) since it captures the global context in both spatial and temporal dimensions.

For the self-attention of spatio-temporal encoder, we adopt the vanilla multi-head attention following [46]. Let  $A$  denote basic attention, and it is computed as:

$$A(Q, K, V) = \text{softmax}\left(\frac{Q \cdot K^T}{\sqrt{d}}\right) \cdot V, \quad (1)$$

where  $Q, K, V$ , and  $d$  are queries, keys, values, and feature dim of tokens, respectively. Dividing tokens into  $h$  groups, then the multi-head attention(MHA) can be formulated as:

$$\begin{aligned} \text{MHA}(Q, K, V) &= P(\text{Concat}(\text{head}_1, \dots, \text{head}_h)) \\ \text{s.t. } \text{head}_i &= A(Q_i, K_i, V_i), i \in [1, h], \end{aligned} \quad (2)$$

where  $P(\cdot)$  represents linear projection function.  $\text{Concat}$  is the operation of concatenating features along feature axis.

**Decoder.** For pose and shape estimation, we build pose decoder and shape decoder with attention as Equation 2, termed **Split-PoS** since the pose decoder and shape decoder are split. For a frame  $I^t$  in the video, given

$L$  pose queries  $Q_{pose}^t$ , pose decoder computes the cross-attention between pose queries and feature tokens  $\tau_e^t$  from encoder. The attention in pose decoder can be denoted as  $\text{MHA}_{pose}(Q_{pose}^t, \tau_e^t, \tau_e^t)$ . The output tokens from attention  $\text{MHA}_{pose}$  are used to regress joints 2D heatmaps  $M_{2D} \in \mathbb{R}^{K \times \frac{H}{s} \times \frac{W}{s}}$ , joints offsets  $M_o \in \mathbb{R}^{3 \times K \times \frac{H}{s} \times \frac{W}{s}}$ , and camera depth map  $M_d \in \mathbb{R}^{1 \times \frac{H}{s} \times \frac{W}{s}}$  by a multi-layer perceptron. Finally, the tokens at the positions with top  $N$  heat scores are used to predict the locations of  $N$  persons.

Similar to pose decoder, shape decoder computes the attention  $\text{MHA}_{shape}(Q_{shape}^t, \tau_e^t, \tau_e^t)$  between shape queries and feature tokens, then output shape parameters maps  $M_s \in \mathbb{R}^{143 \times \frac{H}{s} \times \frac{W}{s}}$ . Each token in  $M_s$  contains pose parameters of 6D rotations  $\theta$ , shape parameters  $\beta$ , and age offset  $\alpha$ , which is further used to generate SMPL+A mesh. The tokens at the same top  $N$  positions are used to regress  $\mathcal{M}^t = \{\mathcal{M}_i^t(\theta, \beta, \alpha) | i \in [1, N]\}$  with the help of their corresponding position and depth  $(x_i, y_i, d_i)$ . Then, the 3D joints  $J^t = \{J_i^t | i \in [1, N]\}$  can be obtained as the process introduced in Section 3.1.

### 3.3. Progressive Video Transformer

Compared with video Transformer baseline, progressive video Transformer has two differences. 1) The sequence decoding process is based on recurrent structure, in which the pose and shape queries in each frame are updated with the prior output tokens from the last step. This mechanism en-



ables decoder to capture the global relationship and reduce the computation costs of spatio-temporal attention. Meanwhile, the updated queries can improve the performance of the decoder since the appearance of human instances may change significantly as time proceeds. 2) We propose pose-guided shape attention (PGA) in shape decoder. PGA improves the performance of shape decoder since the depth and joint information in pose tokens can guide shape decoder to generate better mesh parameters. The framework of progressive video Transformer is shown in Figure 2 (b), which includes backbone network, encoder, and decoder.

**Progressive Decoding Mechanism.** In each frame of the video, human instances have different appearances and pose. Therefore, using the same queries to decode pose and shape tokens is not intuitive. To tackle this problem, we adopt the recurrent structure with an attention-based decoder to decode the sequence of pose and shape tokens.

As shown in Figure 2 (b), given the initialized pose queries  $Q_{pose}^t$ , and features tokens  $\tau_e^t$  from STE at  $t^{th}$  frame, the pose tokens  $\tau_{pose}^t$  are firstly decoded from spatio-temporal pose decoder (STPD) by computing the pose cross-attention between  $Q_{pose}^t$  and  $\tau_e^t$ , denoted as

$$\begin{aligned} \tau_{pose}^t &= \text{STPD}(\hat{Q}_{pose}^t, \tau_e^t), \\ \text{s.t. } \hat{Q}_{pose}^t &= \psi(Q_{pose}^t, \tau_{pose}^{t-1}), \end{aligned} \quad (3)$$

where  $\hat{Q}_{pose}^t$  represents the updated pose queries by fusing the decoded pose tokens  $\tau_{pose}^{t-1}$  at  $(t-1)^{th}$  frame and the initialized pose queries  $Q_{pose}^t$ ,  $t \in [1, T]$ .  $\psi(\cdot)$  is the linear projection layer to fuse tokens.  $\text{STPD}(\cdot)$  represents the operation of spatio-temporal pose decoder.

For shape decoder, the inputs include shape queries  $Q_{shape}^t$ , decoded pose tokens  $\tau_{pose}^t$ , and feature tokens  $\tau_e^t$  from STE.  $Q_{shape}^t$  and  $\tau_{pose}^t$  are firstly aligned and then sent into recurrent STSD to decode shape tokens  $\tau_{shape}^t$ . This process can be formulated as

$$\begin{aligned} \tau_{shape}^t &= \text{STSD}(\hat{Q}_{shape}^t, \tau_e^t), \\ \text{s.t. } \hat{Q}_{shape}^t &= \text{TA}(\psi(Q_{shape}^t, \tau_{shape}^{t-1}), \tau_{pose}^t), \end{aligned} \quad (4)$$

where  $\hat{Q}_{shape}^t$  represents the updated shape queries, which is computed by token aligning on shape tokens  $\psi(Q_{shape}^t, \tau_{shape}^{t-1})$  and pose tokens  $\tau_{pose}^t$ .  $\text{TA}(\cdot)$  represents the operation of token aligning (TA), which is a cross-attention with shape tokens as queries and pose tokens as keys and values.  $\psi(\cdot)$  is the linear projection layer to fuse shape queries and shape tokens at  $(t-1)^{th}$  frame,  $t \in [1, T]$ .  $\text{STSD}(\cdot)$  represents spatio-temporal shape decoder.

**Spatio-Temporal Pose Decoder.** Given pose queries  $Q_{pose}$  and feature tokens  $\tau_e$  from STE, Spatio-Temporal Pose Decoder (STPD) aims to decode the pose tokens  $\tau_{pose} = \{\tau_{pose}^t | t \in [1, T]\}$ . The structure of STPD is shown in Figure 3 (a). First,  $L$  pose queries are updated

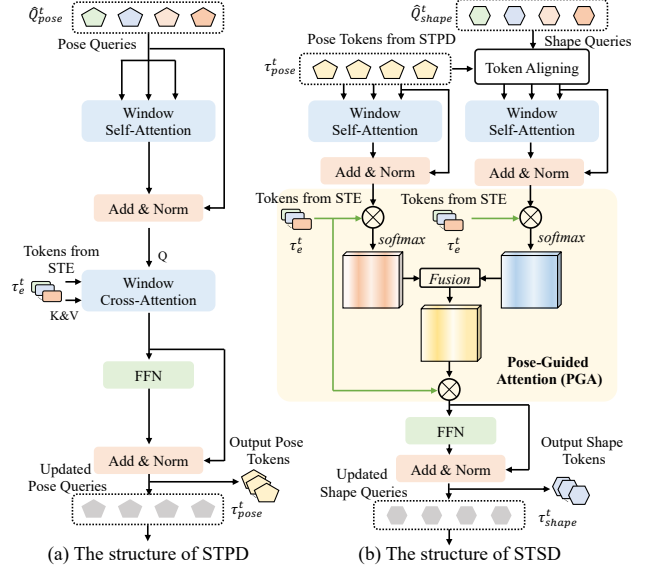


Figure 3. The structure of (a) STPD and (b) STSD. STPD adopts window self-attention [28] to capture the relations between spatial objects and window cross-attention to capture the interactions between pose queries and feature tokens. STPD uses PGA to capture the relations between shape queries and feature tokens with the guidance of pose tokens.  $\otimes$  represents matrix multiplication.

to  $\hat{Q}_{pose}^t$  as Equation 3, further are fed into a window-based pose self-attention module to capture the relations among  $N$  human instances, which can provide the spatial depth context among these objects. After the pose self-attention, the output tokens  $\tilde{Q}_{pose}^t$  at  $t^{th}$  frame can be generated as

$$\tilde{Q}_{pose}^t = \text{MHA}(\hat{Q}_{pose}^t, \hat{Q}_{pose}^t, \hat{Q}_{pose}^t) + \hat{Q}_{pose}^t, \quad (5)$$

where  $\text{MHA}(\cdot)$  is multi-head self-attention as Equation 2.

Then, output tokens  $\tilde{Q}_{pose}^t$  are fed into a cross-attention module as queries, which extract the relation between pose queries and feature tokens  $\tau_e^t$  from STE. After cross-attention, a feed-forward network (FFN) is used to regress the output tokens  $\tau_{pose}^t$ . This process can be formulated as

$$\begin{aligned} \tau_{pose}^t &= \text{FFN}(v_{pose}) + v_{pose}, \\ \text{s.t. } v_{pose} &= \text{MHA}(\tilde{Q}_{pose}^t, \tau_e^t, \tau_e^t), \end{aligned} \quad (6)$$

where we ignore the layer norm for simplification. FFN consists of two linear layers. For each frame in video, the output tokens  $\tau_{pose}^t \in \mathbb{R}^{L \times D}$  are used to regress 2D heatmaps  $M_{2D}$ , 3D offsets  $M_o$ , and camera depth map  $M_d$  as the pose decoder in video Transformer baseline.

**Spatio-Temporal Shape Decoder.** Given the shape queries  $Q_{shape}$ , pose tokens  $\tau_{pose}$ , and feature tokens  $\tau_e$  from STE, Spatio-Temporal Shape Decoder (STSD) aims to decode the shape tokens  $\tau_{shape} = \{\tau_{shape}^t | t \in [1, T]\}$ . Different from pose decoder, a novel Pose-Guided Attention (PGA) module is proposed to decode shape parameters.

The structure of STSD is shown in Figure 3 (b). First,  $L$  shape queries are updated to  $\hat{Q}_{shape}^t$  by token aligning with pose tokens  $\tau_{pose}^t$ . Then, a window-based pose self-attention module and a window-based shape self-attention module are applied on updated shape queries and pose tokens to generate shape tokens  $\tilde{Q}_{shape}^t$  and pose tokens  $\tilde{\tau}_{pose}^t$ , respectively. The process can be formulated as

$$\tilde{\eta} = \text{MHA}(\eta, \eta, \eta) + \eta, \eta \in \{\hat{Q}_{shape}^t, \tau_{pose}^t\}, \quad (7)$$

where  $\eta$  is wildcard character for pose tokens  $\tau_{pose}^t$  and shape tokens  $\hat{Q}_{shape}^t$ . After pose and shape self-attention,  $\tilde{Q}_{shape}^t$  and  $\tilde{\tau}_{pose}^t$  are fed into PGA module to compute cross-attention with feature tokens  $\tau_e^t$ , respectively. The pose and shape cross-attention matrices ( $\chi_{pose}^t, \chi_{shape}^t$ ) are fused by linear projection layer  $fc$ , further to recover the shape tokens. The process of PGA can be formulated as

$$\begin{aligned} \tau_{shape}^t &= \text{FFN}(v_{shape}) + v_{shape}, \\ v_{shape} &= fc(\chi_{pose}^t \oplus \chi_{shape}^t) \otimes \tau_e^t, \\ \chi_{pose}^t &= \text{softmax}(\tilde{\tau}_{pose}^t \otimes (\tau_e^t)^\top), \\ \chi_{shape}^t &= \text{softmax}(\tilde{Q}_{shape}^t \otimes (\tau_e^t)^\top), \end{aligned} \quad (8)$$

where  $\oplus$  indicates concatenating matrices.  $\otimes$  is matrix multiplication. For each frame in video, the output tokens  $\tau_{shape}^t \in \mathcal{R}^{L \times D}$  are used to regress mesh parameter maps  $M_s$  as the shape decoder in video Transformer baseline.

### 3.4. Loss Function

Each token in pose maps can localize one object. Following [44, 45], combined with pose maps ( $M_{2d}, M_o$ , and  $M_d$ ) and shape maps  $M_s$ , we can generate  $N$  SMPL mesh  $\mathcal{M}(\theta, \beta, \alpha)$  for  $N$  instances. For each person, the 3D joints  $J$  can be obtain from  $\mathcal{M}(\theta, \beta, \alpha)$ . Following previous works [14, 44, 45], we use pose losses and mesh losses to supervise outputs. Pose losses include  $L_2$  losses of heatmaps  $\mathcal{L}_{2D}$ , offsets  $\mathcal{L}_o$ , depths  $\mathcal{L}_d$ . mesh losses include  $L_2$  losses of pose parameters  $\mathcal{L}_\theta$ , shape parameters  $\mathcal{L}_\beta$  and age parameters  $\mathcal{L}_\alpha$ . Besides, the  $L_2$  loss of projected 3D joints  $\mathcal{L}_J$  and Mixture Gaussian prior loss of SMPL model in [19] are also used. The total loss  $\mathcal{L}$  is

$$\mathcal{L} = W_{pose} \mathcal{L}_{pose} + W_{mesh} \mathcal{L}_{mesh} + W_J \mathcal{L}_J + W_p \mathcal{L}_p, \quad (9)$$

where  $\mathcal{L}_{pose} = \mathcal{L}_{2D} + \mathcal{L}_o + \mathcal{L}_d$  and  $\mathcal{L}_{mesh} = \mathcal{L}_\theta + \mathcal{L}_\beta + \mathcal{L}_\alpha$ .  $W_{pose}$ ,  $W_{mesh}$ ,  $W_J$ , and  $W_p$  are the weights for pose losses, mesh losses, joints loss, and prior loss, respectively.

## 4. Experiments

### 4.1. Implementation Details

Following previous works [45, 50], we use HRNet as the backbone network. For training, we use the two-stage training strategy as [45]. For the first stage, PSVT is pre-trained

on two basic 3D datasets (Human3.6M [10] and Muco-3DHP [31]) and four 2D datasets (COCO [26], MPII [2], LSP [12], and CrowdPose [20]). Then, PSVT is fine-tuned on the target 3D datasets (RH [45], AGORA [34], 3DPW [48]) with loading the pre-trained parameters for better convergence. When training, PSVT is trained with 120 epochs for first-stage training and 60 epochs for second-stage training. Learning rate is  $5e-5$ . Batch size is 64 and the input size of image is  $512 \times 512$ . Following [44], loss weights  $W_{pose}$ ,  $W_{mesh}$ ,  $W_J$ ,  $W_p$  are set to 160, 1, 360, and 1.6, respectively. 8 Tesla V100 GPUs are used for training.

### 4.2. Datasets and Evaluation Metrics

**RH dataset.** RH dataset is a multi-person dataset [45] for evaluating depth reasoning, which includes about 7.6K images with weak annotations of over 24.8K people. Following [45], the percentage of correct depth relations (PCDR<sup>0.2</sup>) with a threshold of 0.2m is used as the metric to evaluate the accuracy of depth reasoning.

**AGORA dataset.** AGORA dataset is a multi-person dataset [34], which contains 14K training images and 3D testing images with accurate annotations of body meshes and 3D translations. It contains 4240 high-realism textured scans since the images are synthetic. For evaluation, mean per joint position error (MPJPE) and mean vertex error (MVE) are used to evaluate the accuracy of 3D pose and mesh estimation. Considering the missing detections, normalized mean joints error (NMJE) and normalized mean vertex error (NMVE) are also used for evaluation.

**CMU Panoptic dataset.** CMU Panoptic [13] is a large-scale multi-person dataset for 3D human pose and shape estimation. The images are captured by multiple cameras. Following previous works [16, 44, 45], we evaluate PSVT on the testing set of CMU Panoptic without using training set for fair comparison. The evaluation metric is MPJPE.

**3DPW dataset.** 3DPW [48] is a outdoor multi-person dataset, which contains 22K images and 35K images for training and testing, respectively. For comparison with previous works [4, 16, 45, 50], PA-MPJPE, MPJPE, and MPVE are used as evaluation metrics.

### 4.3. Comparison with the State-of-the-art Methods

We compare PSVT with existing methods on two image-based and video-based benchmarks.

**Evaluation on RH Dataset.** To evaluate the monocular depth reasoning ability of PSVT on the in-the-wild images, we compare PSVT with other SOTA methods on RH [45] dataset. The results are shown in Table 1. Compared with single-person method [32] and other multi-person methods [11, 44, 45, 56], PSVT achieves a PCDR<sup>0.2</sup> of 71.23% and outperforms them on each category (Baby, Kid, Teenager, and Adult). Although BEV [44] uses the 3D representation combined with bird’s eye view features

Method	PCDR <sup>0.2</sup> (%) $\uparrow$				
	Baby	Kid	Teen	Adult	All
3DMPPE [32]	39.33	51.42	60.91	57.95	57.47
CRMH [11]	34.74	48.37	59.11	55.47	54.83
SMAP [56]	31.58	40.29	47.35	41.65	41.55
ROMP [44]	30.08	48.41	51.12	55.34	54.81
BEV [45]	60.77	67.09	66.07	69.71	68.27
<b>PSVT (Ours)</b>	<b>64.00</b>	<b>71.29</b>	<b>70.45</b>	<b>71.95</b>	<b>71.23</b>

Table 1. The comparison of accuracy on RH [45] dataset.

Method	Matched $\downarrow$		All $\downarrow$	
	MVE	MPJPE	NMVE	NMJE
PARE [17]	140.9	146.2	167.7	174.0
SPIN [34]	148.9	153.4	193.4	199.2
SPEC [18]	106.5	112.3	126.8	133.7
ROMP [44]	103.4	108.1	113.6	118.8
BEV [45]	100.7	105.3	108.3	113.2
<b>PSVT (Ours)</b>	<b>94.1</b>	<b>97.7</b>	<b>101.2</b>	<b>105.1</b>

Table 2. The comparison of mean errors on AGORA [34] dataset.

to enhance the depth reasoning ability, PSVT has stronger depth reasoning ability than BEV since the pose-guided attention, which captures the global spatial context between different human instances.

**Evaluation on AGORA Dataset.** To evaluate the 3D pose and mesh estimation, we test PSVT on the AGORA dataset as previous works [17, 18, 34, 44, 45]. As shown in Table 2, for the matched objects, PSVT achieves 94.1 and 97.7 in MVE and MPJPE, respectively. For all objects, PSVT achieves 101.2 and 105.1 in NMVE and NMJE, respectively. Compared with SOTA method [45], PSVT outperforms BEV on all objects by 6.6% and 7.2% in NMVE and NMJE, respectively. These results demonstrate the effectiveness of PSVT on pose and mesh estimation.

**Evaluation on CMU Dataset.** To evaluate the effectiveness of PSVT on videos, we compare PSVT with other SOTA methods on CMU dataset [13]. As shown in Table 3, PSVT outperforms previous works and achieves 105.7mm in MPJPE, which has relative gain of 3.5% and shows the stronger generalization ability of PSVT since PSVT is not been trained on the CMU dataset.

**Evaluation on 3DPW Dataset.** We compare PSVT with other methods on the 3DPW dataset [48]. Without using temporal information, PSVT outperforms all image-based methods and achieves 45.7, 75.5, and 84.9 in PA-MPJPE, MPJPE, and MPVE, respectively. Moreover, PSVT achieves new SOTA results by using 9 frames. Compared with Transformer-based methods [17, 49–51], PSVT achieves better results in shape estimation (84.0 in MPVE) since PGA and progressive decoding mechanism.

Method	Haggl.	Mafia	Ultim.	Pizza	Mean $\downarrow$
MubyNet [53]	141.4	152.3	145.0	162.5	150.3
MSC [52]	140.0	165.9	150.7	156.0	153.4
CRMH [11]	129.6	133.5	153.0	156.7	143.2
ROMP [44]	110.8	122.8	141.6	137.6	128.2
3DCrowdNet [5]	109.6	135.9	129.8	135.6	127.3
BEV [45]	90.7	103.7	<b>113.1</b>	125.2	109.5
<b>PSVT (Ours)</b>	<b>88.7</b>	<b>97.9</b>	115.2	<b>121.1</b>	<b>105.7</b>

Table 3. The comparison of MPJPE on CMU [13] dataset.

Method	Frame	Type	PA-MPJPE	MPJPE	MPVE	
PyMAF [55]	1	SP	58.9	92.8	110.1	
HybrIK [21]			48.8	80.0	94.5	
METRO* [24]			47.9	77.1	88.2	
PARE* [17]			46.5	<b>74.5</b>	88.6	
ROMP [44]	1	MP	47.3	76.7	93.4	
BEV [45]			46.9	78.5	92.3	
<b>PSVT (Ours)</b>	1	MP	<b>45.7</b>	75.5	<b>84.9</b>	
HMMR [15]	20	SP	72.6	116.5	139.3	
MEVA [30]	90		54.7	86.9	-	
VIBE [16]	16		56.9	90.2	109.5	
TCMR [4]	16		52.7	86.5	102.9	
MPS-Net* [50]	16		52.1	84.3	99.7	
VIBE+ $D_m$ [16]	16		51.9	82.9	99.1	
GLAMR* [51]	75		51.1	-	-	
MAED* [49]	16		45.7	79.1	92.6	
<b>PSVT (Ours)</b>	9		MP	<b>43.5</b>	<b>73.1</b>	<b>84.0</b>

Table 4. The results on 3DPW [48] dataset. \* indicates Transformer-based method. SP means single-person method. MP means multi-person method, which input is the whole image including multiple persons. MP is more difficult and efficient.

Method	Type	Params (M)	FLOPs (G)	Model (MB)	PA-MPJPE
VIBE [16]	SP	72.43	42.17	776	56.9
MEVA [30]	SP	85.72	42.46	858.8	54.7
TCMR [4]	SP	108.89	42.99	1073	52.7
MPS-Net [50]	SP	39.63	42.45	331	52.1
BEV [45]	MP	35.86	48.89	144	46.9
<b>PSVT (Ours)</b>	MP	38.36	68.67	164	43.5

Table 5. Comparison of FLOPs, Parameters, and model size. FLOPs is computed with image size of  $512 \times 512$  and backbone network of HRNet-32. The FLOPs of SP methods will increase manyfold with the human number in the image.

**Params, FLOPs and Model Size.** The comparisons of parameters, FLOPs, and model size are shown in Table 5. Compared with the single-person video-based method MPS-Net [50], PSVT achieves better performance with a smaller model size and parameter. Besides, PSVT can output multi-person poses and shapes from video in one stage while other video-based methods need two stages. Compared with BEV [45], PSVT achieves better results with comparable parameters and computational costs.

#	Type	Encoder	Decoder	PA-MPJPE	MPJPE	MPVE
1	Image	✗	Conv	49.8	86.2	96.9
2		S	Conv	48.9	84.1	92.5
3		S	Split-PoSh	47.2	79.4	88.6
4		S	PGA	45.7	75.5	84.9
5	Video	ST	Conv	48.2	83.4	90.3
6		ST	Split-PoSh	46.9	78.3	87.5
7		ST	PGA	44.1	74.2	84.5
8		ST	Progressive+PGA	43.5	73.1	84.0

Table 6. The ablation study of PSVT on 3DPW [48] dataset. S: Spatial; ST: Spatio-Temporal. Conv, Split-PoSh, PGA and Progressive represent convolutional head, splitting pose and shape decoder in video Transformer baseline, pose-guided attention and progressive PGA in the decoder of PSVT, respectively.

#	TA	W-PSA	W-SSA	PA-MPJPE	MPJPE	MPVE
1	✗			46.2	76.6	85.1
2		✗		46.4	77.1	85.8
3			✗	46.0	76.1	85.0
4				45.7	75.5	84.9

Table 7. The ablation study of main components in STSD of PSVT on 3DPW [48]. ✗ means the component is removed.

#### 4.4. Ablation Study

**Different Encoders and Decoders** The ablation study of different encoders and decoders is shown in Table 6. With the same spatial encoder, PSVT with PGA achieves better results of 45.7 in PA-MPJPE, 75.5 in MPJPE, and 84.9 in MPVE. Compared with Split-PoSh in video Transformer baseline, the PGA in PSVT brings relative improvements of 3.2 %, 4.9%, and 4.2% in PA-MPJPE, MPJPE, and MPVE, respectively. With spatio-temporal encoder, the progressive decoding mechanism and PGA enable PSVT to learn more temporal information, which helps PSVT achieve 43.5 in PA-MPJPE, 73.1 in MPJPE, and 84.0 in MPVE. These results verify the effectiveness of the proposed PDM and PGA in PSVT.

**Component-wise Analysis in STSD** The STSD of PSVT includes four modules: Token Aligning (TA), Window Pose Self-Attention (W-PSA), Window Shape Self-Attention (W-SSA), and Pose-Guided Attention (PGA). The ablation studies of these components are shown in Table 7. The pose-guided attention with TA, W-PSA, and W-SSA achieves the best results.

**Visualization Analysis** To evaluate the generalization ability of PSVT on the in-the-wild videos, we test PSVT on the videos from PoseTrack [1] dataset. We compare PSVT with the image-based method and video-based method. Compared with BEV [45], as shown in Figure 4, PSVT shows better results in crowded scenarios due to the stronger




Figure 4. The comparison of BEV [45] and PSVT. PSVT shows better predictions since PGA captures global interactions between spatial objects. Red circles indicate the wrong predictions.




Figure 5. The comparison between MPS-Net [50] and our PSVT. PSVT shows better predictions since the captured global spatio-temporal context. Yellow circles and red circles indicate the missing objects and wrong predictions, respectively.

ability to model global spatial context. Compared with MPS-Net [50], as shown in Figure 5, PSVT performs better results on the video since PSVT has stronger ability of modeling global spatio-temporal context.

## 5. Conclusion

In this paper, we propose PSVT, the first end-to-end progressive video Transformer for multi-person 3D human pose and shape estimation, which captures the global context dependencies among different objects in both spatial and temporal dimensions. To handle the variances of objects as time proceeds and improve the performance of mesh estimation, we propose a progressive decoding mechanism and pose-guided attention for PSVT. Extensive experiments on four datasets show the effectiveness of proposed components, which leads PSVT to achieve state-of-the-art results.



## References

- [1] Mykhaylo Andriluka, Umar Iqbal, Eldar Insafutdinov, Leonid Pishchulin, Anton Milan, Juergen Gall, and Bernt Schiele. Posetrack: A benchmark for human pose estimation and tracking. In *CVPR*, pages 5167–5176, 2018. 8
- [2] 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. 6
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, pages 213–229. Springer, 2020. 2
- [4] 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. 1, 2, 3, 6, 7, 8
- [5] Hongsuk Choi, Gyeongsik Moon, JoonKyu Park, and Kyoung Mu Lee. Learning to estimate robust 3d human mesh from in-the-wild crowded scenes. In *CVPR*, pages 1475–1484, 2022. 1, 2, 7
- [6] Carl Doersch and Andrew Zisserman. Sim2real transfer learning for 3d human pose estimation: motion to the rescue. *NeurIPS*, 32, 2019. 1
- [7] 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. *ICLR*, 2021. 2
- [8] Thomas Gatt, Dylan Seychell, and Alexiei Dingli. Detecting human abnormal behaviour through a video generated model. In *ISPA*, pages 264–270. IEEE, 2019. 1
- [9] Nikolas Hesse, Sergi Pujades, Javier Romero, Michael J Black, Christoph Bodensteiner, Michael Arens, Ulrich G Hofmann, Uta Tacke, Mijna Hadders-Algra, Raphael Weinberger, et al. Learning an infant body model from rgb-d data for accurate full body motion analysis. In *MICCAI*, pages 792–800. Springer, 2018. 3
- [10] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *TPAMI*, 36(7):1325–1339, 2013. 6
- [11] Wen Jiang, Nikos Kolotouros, Georgios Pavlakos, Xiaowei Zhou, and Kostas Daniilidis. Coherent reconstruction of multiple humans from a single image. In *CVPR*, pages 5579–5588, 2020. 6, 7
- [12] Sam Johnson and Mark Everingham. Learning effective human pose estimation from inaccurate annotation. In *CVPR*, pages 1465–1472. IEEE, 2011. 6
- [13] Hanbyul Joo, Hao Liu, Lei Tan, Lin Gui, Bart Nabbe, Iain Matthews, Takeo Kanade, Shohei Nohuhara, and Yaser Sheikh. Panoptic studio: A massively multiview system for social motion capture. In *ICCV*, pages 3334–3342, 2015. 6, 7
- [14] 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. 2, 6
- [15] Angjoo Kanazawa, Jason Y Zhang, Panna Felsen, and Jitendra Malik. Learning 3d human dynamics from video. In *CVPR*, pages 5614–5623, 2019. 1, 7
- [16] Muhammed Kocabas, Nikos Athanasiou, and Michael J Black. Vibe: Video inference for human body pose and shape estimation. In *CVPR*, pages 5253–5263, 2020. 1, 3, 6, 7, 8
- [17] 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. 3, 7
- [18] Muhammed Kocabas, Chun-Hao P Huang, Joachim Tesch, Lea Müller, Otmar Hilliges, and Michael J Black. Spec: Seeing people in the wild with an estimated camera. In *ICCV*, pages 11035–11045, 2021. 7
- [19] 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. 2, 6
- [20] Jiefeng Li, Can Wang, Hao Zhu, Yihuan Mao, Hao-Shu Fang, and Cewu Lu. Crowdpose: Efficient crowded scenes pose estimation and a new benchmark. In *CVPR*, pages 10863–10872, 2019. 6
- [21] 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. 7
- [22] Yong-Lu Li, Xinpeng Liu, Han Lu, Shiyi Wang, Junqi Liu, Jiefeng Li, and Cewu Lu. Detailed 2d-3d joint representation for human-object interaction. In *CVPR*, pages 10166–10175, 2020. 1
- [23] 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*, 2022. 1
- [24] Kevin Lin, Lijuan Wang, and Zicheng Liu. End-to-end human pose and mesh reconstruction with transformers. In *CVPR*, pages 1954–1963, 2021. 2, 3, 7
- [25] Kevin Lin, Lijuan Wang, and Zicheng Liu. Mesh graphormer. In *ICCV*, pages 12939–12948, 2021. 2, 3
- [26] 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. Springer, 2014. 6
- [27] 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. 2
- [28] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pages 10012–10022, 2021. 5
- [29] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *TOG*, 34(6):1–16, 2015. 1, 2, 3

- [30] Zhengyi Luo, S Alireza Golestaneh, and Kris M Kitani. 3d human motion estimation via motion compression and refinement. In *ACCV*, 2020. 7, 8
- [31] Dushyant Mehta, Oleksandr Sotnychenko, Franziska Mueller, Weipeng Xu, Srinath Sridhar, Gerard Pons-Moll, and Christian Theobalt. Single-shot multi-person 3d pose estimation from monocular rgb. In *3DV*, pages 120–130. IEEE, 2018. 6
- [32] Gyeongsik Moon, Ju Yong Chang, and Kyoung Mu Lee. Camera distance-aware top-down approach for 3d multi-person pose estimation from a single rgb image. In *ICCV*, pages 10133–10142, 2019. 6, 7
- [33] Mathias Parger, Chengcheng Tang, Yuanlu Xu, Christopher David Twigg, Lingling Tao, Yijing Li, Robert Wang, and Markus Steinberger. Unoc: Understanding occlusion for embodied presence in virtual reality. *TVCG*, 2021. 1
- [34] Priyanka Patel, Chun-Hao P Huang, Joachim Tesch, David T Hoffmann, Shashank Tripathi, and Michael J Black. Agora: Avatars in geography optimized for regression analysis. In *CVPR*, pages 13468–13478, 2021. 3, 6, 7
- [35] Georgios Pavlakos, Nikos Kolotouros, and Kostas Daniilidis. Texturepose: Supervising human mesh estimation with texture consistency. In *ICCV*, pages 803–812, 2019. 2
- [36] Georgios Pavlakos, Jitendra Malik, and Angjoo Kanazawa. Human mesh recovery from multiple shots. In *CVPR*, pages 1485–1495, 2022. 2
- [37] 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. 2
- [38] Zhongwei Qiu, Kai Qiu, Jianlong Fu, and Dongmei Fu. Learning recurrent structure-guided attention network for multi-person pose estimation. In *ICME*, pages 418–423. IEEE, 2019. 2
- [39] Zhongwei Qiu, Kai Qiu, Jianlong Fu, and Dongmei Fu. Weakly-supervised pre-training for 3d human pose estimation via perspective knowledge. *PR*, page 109497, 2023. 2
- [40] Zhongwei Qiu, Qiansheng Yang, Jian Wang, and Dongmei Fu. Dynamic graph reasoning for multi-person 3d pose estimation. In *ACM MM*, pages 3521–3529, 2022. 2
- [41] Zhongwei Qiu, Qiansheng Yang, Jian Wang, and Dongmei Fu. Ivt: An end-to-end instance-guided video transformer for 3d pose estimation. In *ACM MM*, pages 6174–6182, 2022. 1
- [42] Dahu Shi, Xing Wei, Liangqi Li, Ye Ren, and Wenming Tan. End-to-end multi-person pose estimation with transformers. In *CVPR*, pages 11069–11078, 2022. 2, 3
- [43] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *CVPR*, pages 5693–5703, 2019. 3
- [44] Yu Sun, Qian Bao, Wu Liu, Yili Fu, Michael J Black, and Tao Mei. Monocular, one-stage, regression of multiple 3d people. In *ICCV*, pages 11179–11188, 2021. 1, 2, 6, 7
- [45] Yu Sun, Wu Liu, Qian Bao, Yili Fu, Tao Mei, and Michael J Black. Putting people in their place: Monocular regression of 3d people in depth. In *CVPR*, pages 13243–13252, 2022. 1, 2, 6, 7, 8
- [46] 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. 2, 4
- [47] Márton Véges and A Lőrincz. Temporal smoothing for 3d human pose estimation and localization for occluded people. In *ICONIP*, pages 557–568. Springer, 2020. 1
- [48] 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. 6, 7, 8
- [49] 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. 1, 2, 3, 7
- [50] 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. 1, 2, 3, 6, 7, 8
- [51] Ye Yuan, Umar Iqbal, Pavlo Molchanov, Kris Kitani, and Jan Kautz. Glamr: Global occlusion-aware human mesh recovery with dynamic cameras. In *CVPR*, pages 11038–11049, 2022. 1, 2, 3, 7
- [52] Andrei Zanfir, Elisabeta Marinoiu, and Cristian Sminchisescu. Monocular 3d pose and shape estimation of multiple people in natural scenes—the importance of multiple scene constraints. In *CVPR*, pages 2148–2157, 2018. 7
- [53] Andrei Zanfir, Elisabeta Marinoiu, Mihai Zanfir, Alin-Ionut Popa, and Cristian Sminchisescu. Deep network for the integrated 3d sensing of multiple people in natural images. *NeurIPS*, 31, 2018. 7
- [54] Mihai Zanfir, Andrei Zanfir, Eduard Gabriel Bazavan, William T Freeman, Rahul Sukthankar, and Cristian Sminchisescu. Thundr: Transformer-based 3d human reconstruction with markers. In *ICCV*, pages 12971–12980, 2021. 2
- [55] Hongwen Zhang, Yating Tian, Xinchu Zhou, Wanli Ouyang, Yebin Liu, Limin Wang, and Zhenan Sun. Pymaf: 3d human pose and shape regression with pyramidal mesh alignment feedback loop. In *ICCV*, pages 11446–11456, 2021. 2, 7
- [56] Jianan Zhen, Qi Fang, Jiaming Sun, Wentao Liu, Wei Jiang, Hujun Bao, and Xiaowei Zhou. Smap: Single-shot multi-person absolute 3d pose estimation. In *ECCV*, pages 550–566. Springer, 2020. 6, 7
- [57] 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. 2