

Deep Two-Stream Video Inference for Human Body Pose and Shape Estimation

Ziwen Li¹, Bo Xu¹, Han Huang¹, Cheng Lu² and Yandong Guo^{1,*}
¹OPPO Research Institute, ²Xmotors

yandong.guo@live.com

Abstract

Several video-based 3D pose and shape estimation algorithms have been proposed to resolve the temporal inconsistency of single-image-based methods. However it still remains challenging to have stable and accurate reconstruction. In this paper, we propose a new framework *Deep Two-Stream Video Inference for Human Body Pose and Shape Estimation (DTS-VIBE)*, to generate 3D human pose and mesh from RGB videos. We reformulate the task as a multi-modality problem that fuses RGB and optical flow for more reliable estimation. In order to fully utilize both sensory modalities (RGB or optical flow), we train a two-stream temporal network based on transformer to predict SMPL parameters. The supplementary modality, optical flow, helps to maintain temporal consistency by leveraging motion knowledge between two consecutive frames. The proposed algorithm is extensively evaluated on the Human3.6 and 3DPW datasets. The experimental results show that it outperforms other state-of-the-art methods by a significant margin.

1. Introduction

Considerable amount of research has been done on the 3D human pose and shape estimation from a single RGB image [17, 34, 18, 15]. More recently, some methods try to improve 3D human reconstruction by exploiting temporal information from monocular video [17, 15]. However, those methods still struggle to reconstruct accurate 3D human body when there is complex human joint movement or severe occlusion, largely because of limited sensory modality and training data.

To address this, we revisit the 3D human reconstruction from video with the following two beliefs. First, we argue that RGB feature alone is insufficient to interpret the high degree of freedom of human behavior, and additional sensory information is needed. Second, we argue that a much stronger temporal network should be introduced into 3D human reconstruction, considering the diverse and complex human motion.

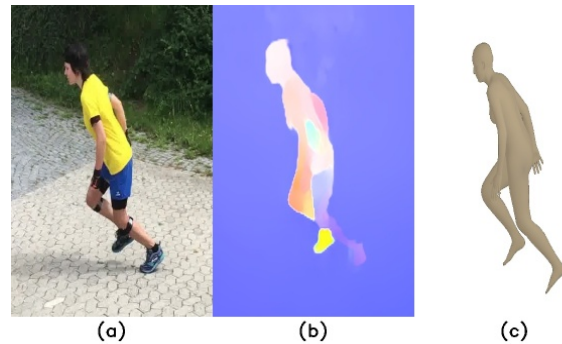


Figure 1: DTS-VIBE learns the complementation of multi-sensory modalities for 3D human mesh reconstruction. Given a RGB video (a), we first estimate its optical flow (b) by auto encoder and predict the human mesh (c) by our two-stream temporal encoder-to-decoder network based on transformer.

Consequently, we design our reconstruction pipeline with careful reconsideration in the following three ways.

First, we propose a new two-stream architecture named **Deep Two-Stream Video Inference for Human Body Pose and Shape Estimation (DTS-VIBE)**, which allows multi-modality fusion for 3D human reconstruction. With this new architecture, we can supplement other sensory modalities, including but not limited to depth, optical flow or other motion cues, to extend RGB feature space for better 3D human pose and shape estimation. To the best of our knowledge, this is the first time such multi-modal architecture is introduced into this area to simultaneously compensate motion instability and improve temporal consistency. This architecture is highly modularized and each of its component is interchangeable, so it can easily accommodate other modalities and their own optimal seq2seq encoder/decoder.

Second, we select optical flow among those candidate modalities to fit the above architecture. The optical flow has proven to be highly effective to understand human behaviour [33] because it can, intuitively, help the two-stream network to bridge each two adjacent frames by understanding the motion. It is worth noting that our optical flow input

does not require actual physical sensor such as event camera, which causes extra cost on data collection, labelling and synchronization with RGB image. Instead, we estimate the optical flow from RGB image sequence as a virtual sensory modality to provide an applicable and inexpensive solution. In other words, optical flow is directly estimated from RGB images without any additional labels.

Third, to correlate video frames in a way that best serves the pose and shape estimation, we build our temporal encoder based on transformer [35] instead of Gated Recurrent Units (GRU) [5] which is used as the temporal network in the state-of-the-art method *VIBE* [17]. Despite being widely applied on various temporal tasks, GRU has its shortcomings that the temporal information is inevitably lost during the recursion. For example, when human body is partially occluded in the leading frames of a sequence, GRU cannot effectively provide reliable temporal information to those frames because referring to prior frames are impossible. However, transformer can alleviate this situation by applying global multi-head attention, particularly to the latter frames that are highly correlated to those leading frames in this case, to better estimate shapes and poses. Similarly, when human body is occluded in the middle of a sequence, transformer can combine short-term and long-term attentions simultaneously to better infer and regularize the motion in the middle. To fully exploit the proposed transformer network, we introduce a new loss named L_{flow} which is inspired by [11]. The loss is designed so that optical flow information is used to regularize the estimated pose by enforcing the trajectory of certain motion. We experimentally show that the combination of transformer network and flow loss significantly improves both accuracy and stability.

Overall, we employ dual convolutional neural networks (CNN) to extract two feature streams from RGB image/video sequence and its optical flows. The extracted image feature streams are fed into a transformer-based temporal encoder, and then combined with the corresponding optical flow streams by late fusion. The fused feature will be finally fed into a transformer-based regressor [35]. Then, we follow [17] to add a discriminator that tries to distinguish between the regressed body poses and samples from the AMASS [15] dataset, which can provide a real/fake label for each sequence.

To justify our solutions, we conduct extensive experiments on multiple datasets. The experimental results show that our proposed method surpasses all state-of-the-art video-based and single-image-based approaches. Overall, the contributions of this paper are:

- To our best knowledge, this is the first end-to-end two-stream architecture which allows multi-modality fusion for video 3D human reconstruction.
- We introduce virtual optical flow which can supple-

ment the corresponding motion information for the RGB domain to predict more accurate pose and shape.

- We propose a transformer-based temporal network to establish more robust temporal correlations and a transformer-based SMPL regressor for better body shape parameterization.
- We introduce a flow loss to regularize the predicted keypoints and reduce the acceleration error.

2. Related works

2.1. Single image-based human reconstruction

Parametric human models [2, 22, 29] are widely used in many prior studies. Pre-trained parametric models such as SMPL [22], STAR [28], SMPL-X [29] are able to recover a human mesh with a few coefficients. Bogo *et al.* [4] propose SMPLify, which first uses CNN to detect the 2D keypoints from a given image, then fits SMPL model to the predicted keypoints. Because of the high cost of 3D human data collection and annotation, people tend to use projected 2D keypoints loss as the less ideal supervision. HMR [14] directly regresses SMPL parameters given an input image and projects 3D joints to 2D using the estimated camera parameters to match the ground truth 2D keypoints. It also proposes a discriminator network to distinguish the predicted and ground truth parameters, generating a more plausible mesh. Kolotouros *et al.* [18] proposed SPIN that iteratively fits the SMPL model to 2D joint and use the current estimation to supervise the network in the next iteration.

Meanwhile, some methods [29, 37] also use intermediate output to help reconstruct human mesh. Pavlako *et al.* [29] first generates 2D joint heatmap and mask, then combines them with input image to predict SMPL parameters. Pengfei *et al.* [37] aligns 3D mesh to 2D image and use UV map to represent the 3D human mesh and uses encoder-decoder model to predict the UV map. Besides, there are plenty of other approaches which directly regress 3D human mesh from a given still image [27, 34, 19, 6, 32]. Gyeongsik *et al.* [27] proposed an Image-to-Lixel network that directly predict three 1D-heatmaps of xyz coordinates of the human mesh vertices. Varolet *et al.* [34] proposes BodyNet, which estimates voxels of human shape in the 3D volumetric space.

2.2. Video-based 3D human reconstruction

Despite the progress of 3D human pose and shape estimation from single image, there are some video based methods [15, 17, 23, 8, 12, 26, 30] which take advantage of temporal information and subsequently achieve impressive outcomes. Hossain *et al.* [12] propose a LSTM model to predict a sequence of 3D joints from given 2D joints.

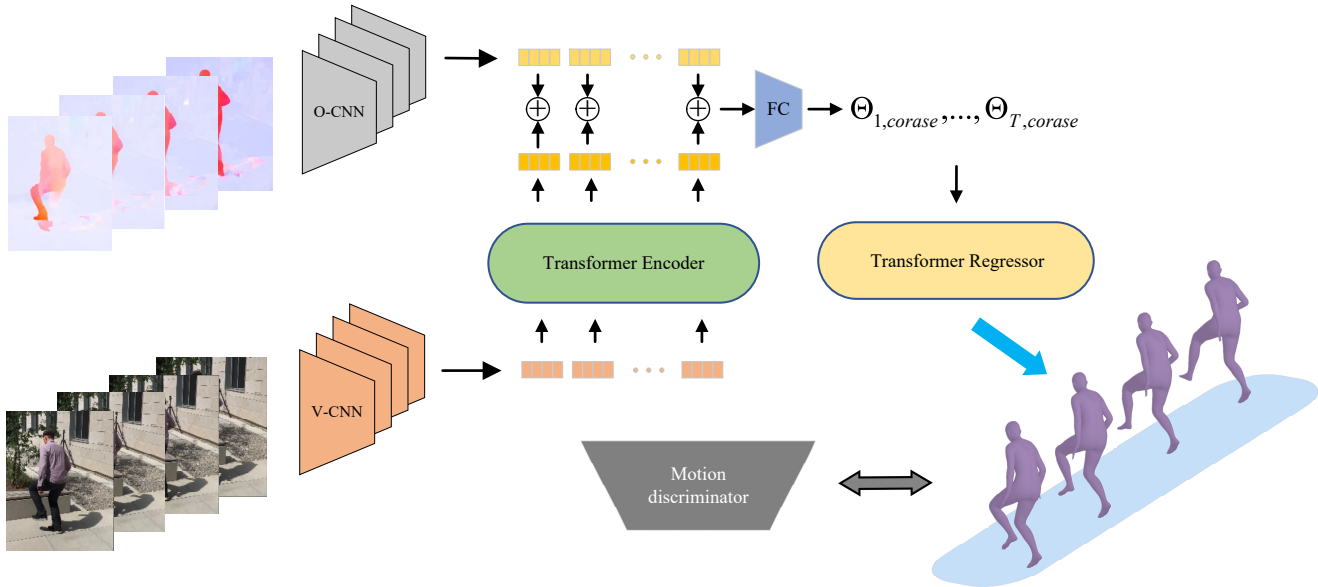


Figure 2: **Architecture of the DTS-VIBE.** The sequence of images and its corresponding flow are sent to the CNN to extract features. The transformer encoder extract temporal features and adds it with the flow feature. A fully connected layer is used to predict coarse SMPL parameters, which is processed by transformer regressor, providing final SMPL parameters. The whole network is trained with a discriminator similar to VIBE.

Pavloet *al.* [30] use a fully-convolutional network to process a sequence of 2D joints to estimate its 3D location, and these predicted 3D joints are used to self-supervise the network on the original 2D joints. [3] first predict 2D joints and SMPL parameters for each frame, then jointly optimize them to reduce the error. Kanazawa *et al.* [15] propose HMMR, which uses 1D CNN as temporal encoder to find features from sequence of images. It predicts 3D poses not only for the target frame but also past and future frames. Such strategy guarantees the smooth estimation. VIBE [17] utilize GRU to encode features from single image into temporal feature and regress SMPL parameters. It also introduces a motion discriminator to guide the generator and encourage it to predict more reasonable poses when compared to the labels in auxiliary dataset. All of them prove to be great success, yet they still lack of temporal consistency in certain degree. In this regard, MEVA [23] feed a sequence of human motion into an auto-encoder framework to first get coarse motion with VME, then refine it with MRR. Its network with encoder-decoder architecture learns the motion of human pose from AMASS[24], which helps to smooth pose estimation at inference.

2.3. Two Stream

Two stream architecture is widely adopted in video action recognition [33, 7, 20, 9]. The feature obtained from optical flow is intuitively helpful when trying to learn the motion of people. Simonyan[33] propose to extract spa-

tial feature from images and extract temporal feature from dense optimal flow. The two types of features are combined by late fusion and sent to the classification module. Meanwhile, optical flow is also used to help 3D reconstruction. [10] combine flow and 2D keypoints to predict poses over the input video and achieve decent performance under occluded scenarios. Similarly, our approach leverages the optical flow to enhance the temporal feature.

3. Architecture

In this section, we describe the deep two-stream video inference network for human body pose and shape estimation (DTS-VIBE). As summarized in Figure 2, the DTS-VIBE model consists of two-stream encoder, transformer-based regressor and motion discriminator.

3.1. Two-stream encoder

The intuition behind using a two-stream encoder to process RGB and optical flow separately is that optical flow can provide additional sensory information of human motion. This is particularly useful for difficult occasions when human is severely occluded or the pose of human is irregular. Optical flow information can also help to maintain the smoothness of human prediction.

Given a sequence of continuous RGB frames v_1, \dots, v_T , we first group every two consecutive video frames to obtain estimated optical flow o_1, \dots, o_T using an auto encoder [21]. The RGB stream v_1, \dots, v_T and optical-flow

stream o_1, \dots, o_T are fed into a dual CNNs (V-CNN and O-CNN), which outputs two-stream feature vectors $f_i^v \in \mathbb{R}^{2048}$ and $f_i^o \in \mathbb{R}^{2048}$, $i \in (1, T)$. We adopt ResNet-50 [7] as the backbone network of V-CNN and O-CNN, followed by the fully connected layers to reduce the dimensions from 2048 to 512. Instead of using Gated Recurrent Units (GRU), which demonstrates decent performance on video-based 3D human reconstruction [17], our temporal network is instead built on transformer [35]. Although GRU is widely applied in various temporal tasks, it shows disadvantages that the temporal information is inevitably lost during the recursion. Transformer can alleviate this situation, because it directly works on every feature simultaneously with global multi-head attention.

The sequence of RGB feature vectors $f_i^v \in \mathbb{R}^{512}$ are fed into our transformer [35]-based temporal encoder that yields latent RGB feature vectors $g_i^v \in \mathbb{R}^{512}$. Then the optical-flow feature vectors $f_i^o \in \mathbb{R}^{512}$ are added on g_i^v and sent to a transformer-based regressor [18] to predict the per-frame SMPL parameters and corresponding camera parameters. Our transformer-based temporal encoder consists of N layers, each of which contains two sub-layers: a multi-head attention layer and a feed forward layer. As shown in Figure 3, given a sequence of input features x_1, \dots, x_T , the multi-head attention layer first gets the *query* (Q), *key*(K), and *value*(V) by using fully connected layer:

$$Q = W^Q X, K = W^K X, V = W^V X \quad (1)$$

Then the output is computed as

$$a = \text{softmax}\left(\frac{QK^T}{\sqrt{d}}\right)V, \quad (2)$$

where d is the dimension of key. The feed forward layer consists of 2 fully connected layer and dropout. The output of transformer is then added with the output of O-CNN, which contains the temporal information between each two frames, followed by a fully connected layer to reduce feature dimension from 512 to 157. The transformer regressor takes the low-dimension feature as input and outputs accurate SMPL parameters $\hat{\theta}_1, \dots, \hat{\theta}_T$. The overall loss of our module is:

$$L_{total} = \lambda_1 L_{3D} + \lambda_2 L_{2D} + \lambda_3 L_{smpl} + \lambda_4 L_{adv} \quad (3)$$

where $\lambda_1, \lambda_2, \lambda_3, \lambda_4$ are weighted coefficients for each loss. Specifically, L_{3D} is the L_2 distance between ground truth 3D joints location $X_{i,3D}$ and 3D joints location $\hat{X}_{i,3D}$ from predicted SMPL model:

$$L_{3D} = \sum_{i=1}^T \|X_{i,3D} - \hat{X}_{i,3D}\|_2^2 \quad (4)$$

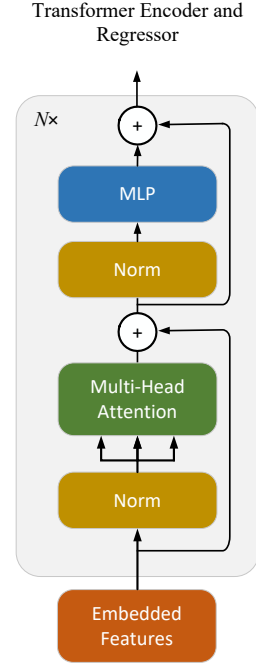


Figure 3: **Architecture of Transformer.** Our temporal encoder use the same transformer structure but with different number of layers or heads.

where L_{2D} is the L_2 distance between ground truth 2D joints location and the weak-perspective projection of predicted 3D joints using corresponding camera parameters:

$$L_{2D} = \sum_{i=1}^T \|X_{i,2D} - \hat{X}_{i,2D}\|_2^2 \quad (5)$$

where L_{smpl} is the L2 distance between ground truth and predicted SMPL parameters:

$$L_{smpl} = \sum_{i=1}^T \|\theta_i - \hat{\theta}_i\|_2^2 + \|\beta_i - \hat{\beta}_i\|_2^2 \quad (6)$$

3.2. Motion discriminator and flow supervision

Motion discriminator. Follow [17], we use a motion discriminator D_M to distinguish the predicted SMPL parameters θ (fake) and the real SMPL parameters $\hat{\theta}$ (real) from AMASS [24]. Motion discriminator helps to produce more feasible real world poses that are aligned with 2D joint locations:

$$L_{adv} = E_{\Theta \sim p_G} [(D_M(\hat{\Theta}) - 1)^2] \quad (7)$$

The objective for the discriminator is:

$$L_D = E_{\Theta \sim p_R} [(D_M(\Theta) - 1)^2] + E_{\Theta \sim p_G} [D_M(\hat{\Theta})^2] \quad (8)$$

Method		3DPW				MPI-INF-3DHP			Human3.6		
		PA-MPJPE↓	MPJPE↓	PVE↓	Accel↓	PA-MPJPE↓	MPJPE↓	Accel↓	PA-MPJPE↓	MPJPE↓	Accel↓
Single image	HMR	76.7	130.0	-	37.4	89.8	124.2	-	56.8	88.0	-
	GraphCMR	70.2	-	-	-	-	-	-	50.1	-	-
	SPIN	59.2	96.9	116.4	29.8	67.5	105.2	-	41.1	-	18.3
	I2L-MeshNet	57.7	93.2	110.1	30.9	-	-	-	41.1	55.7	13.4
	Pose2Mesh	58.3	88.9	106.3	22.6	-	-	-	46.3	64.9	23.9
Video	HMMR	72.6	116.5	139.3	15.2	-	-	-	56.9	-	-
	Sun et al.	69.5	-	-	-	-	-	-	42.4	59.1	-
	SPIN	59.2	96.9	116.4	29.8	67.5	105.2	-	41.1	-	18.3
	MEVA	54.7	86.9	-	11.6	65.4	96.4	11.1	53.2	76.0	15.3
	VIBE	51.9	82.9	99.1	23.4	64.6	96.6	27.3	41.4	65.6	-
	DST-VIBE(Ours)	50.3	76.7	93.5	11.0	62.2	93.4	11.9	39.3	60.5	5.0

Table 1: **Evaluation of state-of-the-art methods on 3DPW, MPI-INF-3DHP, and Human3.6M datasets.** Models of VIBE, MEVA and DST-VIBE use 3DPW train set. And all methods except MEVA use Human3.6 for training. So the most fair comparison is between VIBE and ours. DST-VIBE achieves state-of-the-art results with almost all the metrics.

where p_G is a generated motion sequence and p_R is a real motion sequence from the AMASS dataset.

Flow supervision. Inspired by [11], we introduce a new flow loss to refine our DTS-VIBE network. Since optical flow represents the motion of two adjacent frames, the 2D joints location should move following its corresponding flow. So we use our flow sequence o_1, \dots, o_T to supervise the movement of 2D joints location. We back trace the previous frame 2d joint position $X_{i-1,2D}^f$ using optical flow o_i and current 2d joint position $X_{i,2D}$, the flow loss is calculated as L_2 distance between backward 2D joints and original 2D joints:

$$L_{flow} = \sum_{i=2}^T \|X_{i-1,2D} - X_{i-1,2D}^f\|_2^2 \quad (9)$$

Note that we do not add flow loss at the beginning of the training, we only use it as supervision during refinement.

Among the dual convolutional network we use, the V-CNN is pre-trained on frame-based pose and shape estimation task [18], while the O-CNN is pre-trained on the ImageNet [31] dataset. Similar to [17], we choose sequence length $T = 16$, which also leads to the best results compared with other alternative ones *i.e.* 8, 32, 64 and 128. For the transformer-based temporal encoder, we choose the number of layers to be 6 and number of head to be 8. The transformer-based regressor has 3 layers with 4 heads. We set the four loss term $\lambda_1, \lambda_2, \lambda_3, \lambda_4$ in Eq. 3 to 300, 200, 120 and 60 respectively. We use Adam optimizer [16] with learning rate 5×10^{-5} for the generator and 5×10^{-4} for the motion discriminator. During the refinement with the flow loss L_{flow} as supervision, we remove the discriminator and use Adam optimizer [16] with learning rate 1×10^{-5} .

4. Experiments

In this section, we first list the training/testing datasets and evaluation metrics we use for comparison. Then we compare our method with other state of the art frame-based and video-based methods. We also conduct multiple experiments to illustrate the influence of each component in our model. We provide both qualitative and quantitative results.

4.1. Dataset

As we all know, 2D datasets (annotated by 2D joint labels) contain a large amount of in-the-wild videos and their labels are easy to obtain. 3D datasets (annotated by 3D joint labels) contain less in-the-wild videos and their labels are more difficult to acquire, but can better supervise the 3D construction model than 2D ones. Therefore we conduct hybrid-training to leverage both 2D and 3D datasets. We use InstaVaryety[15], Penn Action [38] and PoseTrack [1] as 2D datasets and 3DPW [36], MPI-INF-3DHP [25] and Human3.6M [13] as 3D datasets for training and evaluation. Human3.6M is a large scale indoor dataset with 2D and 3D annotations. We use subjects S1, S5, S6, S7 and S8 as training data, and test our models on subjects S9 and S11. 3DPW is an outdoor dataset with 2D and 3D annotations. Following previous method, we also train our model on the 3DPW training set before evaluating on 3DPW test set. We also use AMASS [24] dataset for adversarial training to discriminate a real/fake label for each sequence.

4.2. Evaluation Metrics

We use three image-based metrics: Mean per joint position error (MPJPE), Procrustes-aligned mean per joint position error (PA-MPJPE) and per vertex error (PVE). Additionally, we report acceleration error [15], which computes the average difference between the predicted and ground




Figure 4: Attention maps. Visualization of attention values among different attention heads. The x-axis and y-axis correspond to input and target frames respectively. We visualize the attention matrix value of the eight attention head of the first transformer encoder layer.(a) to (h) represent the attention value of eight attention heads separately. Lighter color indicates stronger attention value.

truth acceleration of each joint in mm/s^2 . We use acceleration error as smooth indicator of temporal methods. A lower image-based metric error means a better performance in terms of accuracy of a model. And lower acceleration error means smoother results.

4.3. Comparison results

As shown in Table 1, we compare our result with previous state-of-the-art methods. Our method outperforms VIBE in all image-based metrics on 3DPW and Human3.6M by a large margin. Besides, we achieve comparable result on MPI-INF-3DHP. Meanwhile, our result significantly outperforms all the image-based and video-based 3D reconstruction methods. The improvement of our model indicates that our two stream network helps to fuse spatial and temporal information, leading to more accurate and smoother 3D reconstruction. It also demonstrates that our transformer encoder and regressor are better at capturing temporal human motion information and predicting plausible SMPL parameters.

Table 1 shows the comparison on temporal smoothness between DTS-VIBE and state-of-the-arts video-based methods. We present some visualization results in Figure 6 which demonstrates that our model is better than VIBE when part of the body is self-occluded. The acceleration error of DTS-VIBE is lower than VIBE. As shown in Table 1, we produce smoother result than MEVA, which is state-of-the-art model on temporal smoothness, in terms of the ac-

model	PA-MPJPE	Accel
GRU + TR	51.1	31.6
TE + HR	50.6	17.6
TE + TR(ours)	50.3	11.0

Table 2: **Ablation study on transformer.** We replace our transformer encoder (TE) with GRU and replace Transformer Regressor with HMR regressor

model	PA-MPJPE	Accel
w/o flow	54.7	12.9
w/ flow	52.7	13.2
w/ flow + flow loss	52.9	12.3

Table 3: **Ablation study on flow.** We conduct experiments that removes flow feature and flow loss from our model.

celeration error. We reduce the acceleration error by 67.3% on Human3.6M, while maintaining the comparable result on 3DPW and MPI-INF-3DHP. Figure 5 shows the acceleration error comparison between other methods and ours measured on a sample of 3DPW test set. Compared with VIBE, our model keeps acceleration error at lower level. And our result is slightly better than MEVA except some rare spikes.

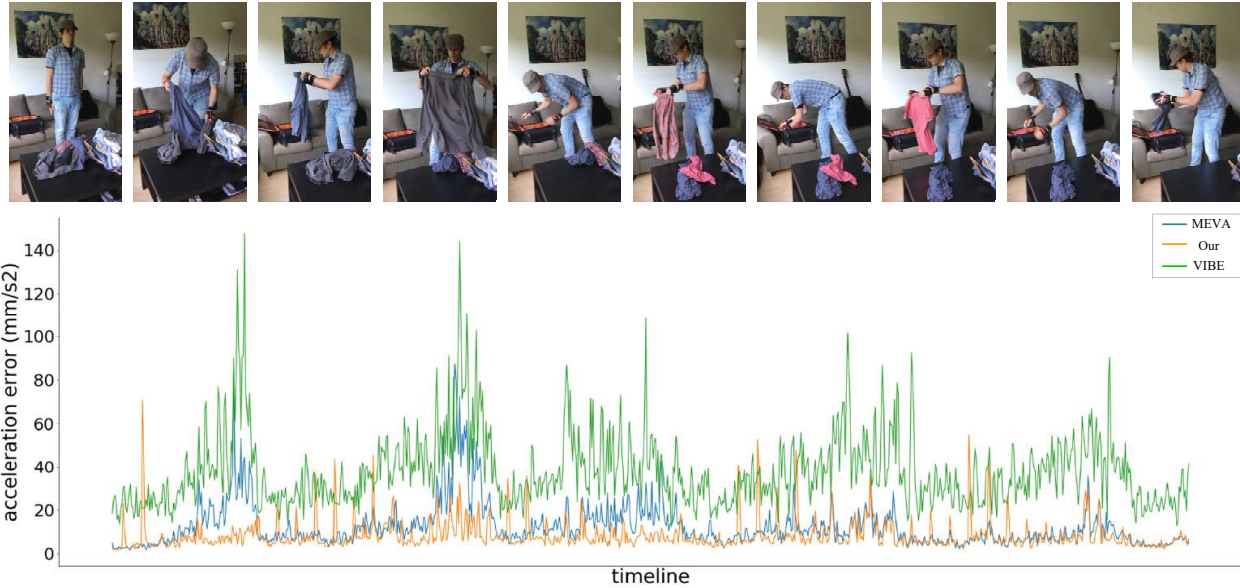


Figure 5: **Comparison of acceleration errors between VIBE, MEVA and our method.** The top of the figure shows some samples from test images and the bottom graph is the acceleration error along the timeline for three models. Overall, our result is slightly smoother than MEVA and significantly better than VIBE.

4.4. Ablation study

Table 3 and Table 2 show the ablation study results with and without each component. For each experiment, we use 3DPW, MPI-INF-3DHP and Human3.6M for training, and 3DPW for evaluation.

Effectiveness of Transformer. To single out the improvement introduced by our transformer based temporal encoder, we revert the temporal encoder into GRU and compare it with our model. Similarly, we also use regressor from [14] and our transformer-based regressor to prove the effectiveness of our transformer regressor. Table 2 shows that the use of transformer as temporal encoder not only improves the accuracy of reconstruction but also provides smoother results. It indicates that transformer can better understand the complexity and variability of real human motion. In addition, our transformer-based regressor significantly outperforms HMR regressor by 0.6% on PA-MPJPE and 37.5% on acceleration error. This demonstrates that our transformer regressor is able to predict more reasonable SMPL parameters.

Effectiveness of optical flow. We perform additional experiments to demonstrate the benefit of dense motion cues - optical flow. Here we remove human3.6M from training set. We also remove the flow-CNN feature and only use image feature to predict SMPL parameters. Table 3 shows the comparison between our model with and without flow feature. The accuracy of our model increases after adding optical flow features while maintaining high degree of smoothness.

Effectiveness of flow loss. We report the results with and without flow loss in Table 3. It demonstrates that our acceleration error is reduced after using flow loss. Overall, the optical flow and flow loss together enhance the performance of model not only on the accuracy but also on the consistency.

Non-local Temporal Interaction. To further understand the effectiveness of our DTS-VIBE model on temporal interaction, we visualize the attention map of transformer encoder. Non-local temporal information from input video are required to stabilize pose estimation while achieving high accuracy. Self attention mechanism of transformer can catch both long-term and short-term temporal information while GRU model from VIBE[17] focuses more on catching short term temporal information.

Figure 4 shows the average attention value of eight attention head of the first transformer encoder layer. We sample 1000 video sequences from 3DPW test set and visualize its self-attention map. Each pixel indicates the self attention value of the target frame with respect to the input frame. Note that different attention head focuses on different input frames. For instance, attention head (d) and (h) tend to catch information from frames far apart while (c) and (g) tend to catch information within nearby frames.

5. Conclusion

In this paper, we introduce a two-stream 3D human reconstruction model named Deep Two-Stream Video Inference for Human Body Pose and Shape Estimation (DTS-



Figure 6: **Comparison visualization.** The first row corresponds to the origin video. The second row is the ground truth body mesh. The third and fourth rows are our predicted results from camera view and alternate view. The fifth and sixth rows are VIBE results from camera view and alternate view.

VIBE), which consists of two-stream encoder, transformer-based regressor and motion discriminator. DTS-VIBE avoids extra inputs other than RGB. Additionally, we introduce a new flow loss to refine our model. Extensive experi-

ments demonstrate the necessity and effectiveness of virtual multi-modality fusion in 3D human reconstruction. Furthermore, our model outperforms current state-of-the-art algorithms on 3D human reconstruction by a significant margin.

References

- [1] M. Andriluka, U. Iqbal, E. Insafutdinov, L. Pishchulin, A. Milan, J. Gall, and B. Schiele. Posetrack: A benchmark for human pose estimation and tracking. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5167–5176, 2018.
- [2] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, and J. Davis. Scape: shape completion and animation of people. In *ACM SIGGRAPH 2005 Papers*, pages 408–416. 2005.
- [3] A. Arnab, C. Doersch, and A. Zisserman. Exploiting temporal context for 3d human pose estimation in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3395–3404, 2019.
- [4] F. Bogo, A. Kanazawa, C. Lassner, P. Gehler, J. Romero, and M. J. Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *European conference on computer vision*, pages 561–578. Springer, 2016.
- [5] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- [6] H. Choi, G. Moon, and K. M. Lee. Pose2mesh: Graph convolutional network for 3d human pose and mesh recovery from a 2d human pose. In *European Conference on Computer Vision*, pages 769–787. Springer, 2020.
- [7] R. Christoph and F. A. Pinz. Spatiotemporal residual networks for video action recognition. *Advances in Neural Information Processing Systems*, pages 3468–3476, 2016.
- [8] R. Dabral, A. Mundhada, U. Kusupati, S. Afaque, A. Sharma, and A. Jain. Learning 3d human pose from structure and motion. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 668–683, 2018.
- [9] A. Diba, V. Sharma, and L. Van Gool. Deep temporal linear encoding networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2329–2338, 2017.
- [10] C. Doersch and A. Zisserman. Sim2real transfer learning for 3d human pose estimation: motion to the rescue. *arXiv preprint arXiv:1907.02499*, 2019.
- [11] X. Dong, S.-I. Yu, X. Weng, S.-E. Wei, Y. Yang, and Y. Sheikh. Supervision-by-registration: An unsupervised approach to improve the precision of facial landmark detectors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 360–368, 2018.
- [12] M. R. I. Hossain and J. J. Little. Exploiting temporal information for 3d human pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 68–84, 2018.
- [13] C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence*, 36(7):1325–1339, 2013.
- [14] A. Kanazawa, M. J. Black, D. W. Jacobs, and J. Malik. End-to-end recovery of human shape and pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7122–7131, 2018.
- [15] A. Kanazawa, J. Y. Zhang, P. Felsen, and J. Malik. Learning 3d human dynamics from video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5614–5623, 2019.
- [16] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [17] M. Kocabas, N. Athanasiou, and M. J. Black. Vibe: Video inference for human body pose and shape estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5253–5263, 2020.
- [18] N. Kolotouros, G. Pavlakos, M. J. Black, and K. Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2252–2261, 2019.
- [19] N. Kolotouros, G. Pavlakos, and K. Daniilidis. Convolutional mesh regression for single-image human shape reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4501–4510, 2019.
- [20] Z. Lan, Y. Zhu, A. G. Hauptmann, and S. Newsam. Deep local video feature for action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 1–7, 2017.
- [21] P. Liu, M. Lyu, I. King, and J. Xu. Selfflow: Self-supervised learning of optical flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [22] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black. Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6):1–16, 2015.
- [23] Z. Luo, S. A. Golestaneh, and K. M. Kitani. 3d human motion estimation via motion compression and refinement. In *Proceedings of the Asian Conference on Computer Vision*, 2020.
- [24] N. Mahmood, N. Ghorbani, N. F. Troje, G. Pons-Moll, and M. J. Black. Amass: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5442–5451, 2019.
- [25] D. Mehta, H. Rhodin, D. Casas, P. Fua, O. Sotnychenko, W. Xu, and C. Theobalt. Monocular 3d human pose estimation in the wild using improved cnn supervision. In *2017 international conference on 3D vision (3DV)*, pages 506–516. IEEE, 2017.
- [26] D. Mehta, O. Sotnychenko, F. Mueller, W. Xu, S. Sridhar, G. Pons-Moll, and C. Theobalt. Single-shot multi-person 3d pose estimation from monocular rgb. In *2018 International Conference on 3D Vision (3DV)*, pages 120–130. IEEE, 2018.
- [27] G. Moon and K. M. Lee. I2l-meshnet: Image-to-lixel prediction network for accurate 3d human pose and mesh estimation from a single rgb image. *arXiv preprint arXiv:2008.03713*, 2020.
- [28] A. A. Osman, T. Bolkart, and M. J. Black. Star: Sparse trained articulated human body regressor. *arXiv preprint arXiv:2008.08535*, 2020.

- [29] G. Pavlakos, V. Choutas, N. Ghorbani, T. Bolkart, A. A. Osman, D. Tzionas, and M. J. Black. Expressive body capture: 3d hands, face, and body from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10975–10985, 2019.
- [30] D. Pavllo, C. Feichtenhofer, D. Grangier, and M. Auli. 3d human pose estimation in video with temporal convolutions and semi-supervised training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7753–7762, 2019.
- [31] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [32] S. Saito, Z. Huang, R. Natsume, S. Morishima, A. Kanazawa, and H. Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2304–2314, 2019.
- [33] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. *arXiv preprint arXiv:1406.2199*, 2014.
- [34] G. Varol, D. Ceylan, B. Russell, J. Yang, E. Yumer, I. Laptev, and C. Schmid. Bodynet: Volumetric inference of 3d human body shapes. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 20–36, 2018.
- [35] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017.
- [36] T. von Marcard, R. Henschel, M. J. Black, B. Rosenhahn, and G. Pons-Moll. Recovering accurate 3d human pose in the wild using imus and a moving camera. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 601–617, 2018.
- [37] P. Yao, Z. Fang, F. Wu, Y. Feng, and J. Li. Densebody: Directly regressing dense 3d human pose and shape from a single color image. *arXiv preprint arXiv:1903.10153*, 2019.
- [38] W. Zhang, M. Zhu, and K. G. Derpanis. From actemes to action: A strongly-supervised representation for detailed action understanding. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2248–2255, 2013.