

# Human Mesh Recovery from Monocular Images via a Skeleton-disentangled Representation

Yu Sun<sup>1\*</sup>, Yun Ye<sup>2\*</sup>, Wu Liu<sup>2†\*</sup>, Wenpeng Gao<sup>1</sup>, YiLi Fu<sup>1</sup>, and Tao Mei<sup>2</sup>

<sup>1</sup>Harbin Institute of Technology <sup>2</sup>JD AI Research  
yusun@stu.hit.edu.cn yun.ye@intel.com liuwu@live.cn  
{wpgao, meylfu}@hit.edu.cn tmei@live.com

## Abstract

We describe an end-to-end method for recovering 3D human body mesh from single images and monocular videos. Different from the existing methods try to obtain all the complex 3D pose, shape, and camera parameters from one coupling feature, we propose a skeleton-disentangling based framework, which divides this task into multi-level spatial and temporal granularity in a decoupling manner. In spatial, we propose an effective and pluggable “disentangling the skeleton from the details” (DSD) module. It reduces the complexity and decouples the skeleton, which lays a good foundation for temporal modeling. In temporal, the self-attention based temporal convolution network is proposed to efficiently exploit the short and long-term temporal cues. Furthermore, an unsupervised adversarial training strategy, temporal shuffles and order recovery, is designed to promote the learning of motion dynamics. The proposed method outperforms the state-of-the-art 3D human mesh recovery methods by 15.4% MPJPE and 23.8% PA-MPJPE on Human3.6M. State-of-the-art results are also achieved on the 3D pose in the wild (3DPW) dataset without any fine-tuning. Especially, ablation studies demonstrate that skeleton-disentangled representation is crucial for better temporal modeling and generalization.

## 1. Introduction

Different from traditional 3D pose estimation that usually predicts the location of 14/17 skeleton joints [5, 24], 3D human body mesh recovery from the monocular images is a more complex task, which tries to estimate the more detailed 3D shape and joint angles. In detail, it needs to estimate more than 85 parameters, which controls 6890 ver-

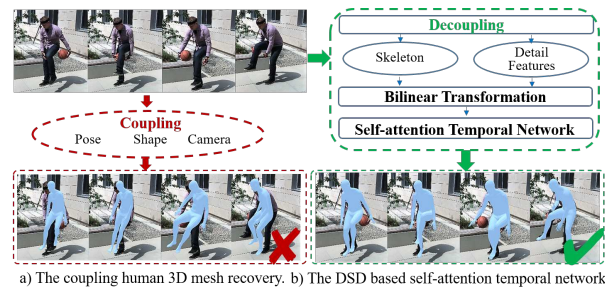


Figure 1. Human 3D mesh recovery from monocular video based on skeleton disentangling and temporal coherence.

tices [22] that form the surface of 3D body mesh. Moreover, the information loss from the 3D scene to a 2D image, inherent ambiguity, and complex changes in human body shape and pose further increase the complexity of this task. Therefore, although the 3D body mesh recovery is important in computer vision, motion/event analysis [10, 11, 20, 23, 19], and virtual try-on [7, 6], it is still a frontier challenge. In this paper, we try to solve this problem via a skeleton-disentangled representation in multi-level spatial and temporal granularity.

Most of the existing methods try to recover the human 3D mesh from single images. Previous multi-stage approaches [17, 26, 27] first extract human body information (e.g. 2D pose and segmentation) and then estimate the 3D model parameters from them, which is typically not optimal. Recently, HMR [15] provides an end-to-end solution to learn a mapping from image pixels directly to model parameters, which shows a significant performance advantage over the two-stage methods. However, as shown in Figure 1, severe coupling problem makes the prediction unstable. In HMR, 3D pose, body shape, and camera parameters are derived directly from the same feature vector using the fully connected layers. Without any decoupling measurements, high complexity makes features of different targets

\*This work is done when Yu Sun is an intern at JD AI Research.

†Corresponding author: Wu Liu. \* Equal contribution.

tightly coupled. Moreover, existing datasets with 3D annotations are collected in a constrained environment with limited motion and shape patterns. Therefore, many previous methods have employed 2D in-the-wild pose datasets, like MPII [2], to learn richer poses. It may be sufficient for 3D pose estimation. While, for 3D mesh recovery, 2D pose is far from enough for recovering the complex human 3D shape and pose details. Lacking supervision makes the predictions of the details vulnerable, which further exacerbates the coupling problem. The coupling strategy makes the model trained on these datasets cannot well generalize to the complex environments and various human states.

To solve this problem, we propose a lightweight and pluggable DSD module to decouple different factors in an end-to-end manner. The main idea of DSD module is to disentangle the skeleton from the 3D mesh details with bilinear transformation. Firstly, information of the pose skeleton and the rest details (e.g., body shape, detailed pose information) are extracted independently. Furthermore, the bilinear transformation is employed to aggregate two pieces of information while keeping their decoupling in the new feature space. Finally, the network is trained end-to-end to keep the global optimal. In the evaluations, we demonstrate that the DSD module outperforms the state-of-the-art methods [16] by 13% PA-MPJPE on Human3.6M dataset. Moreover, the evaluations also demonstrate that the proposed portable DSD module can be easily plugged into other 2D/3D pose estimation network for recovering 3D human mesh.

Recovering human 3D mesh from single images may suffer from the inherent ambiguity as multiple 3D poses and shapes can be mapped to the same situation in a 2D image. To tackle this problem, we propose a self-attention temporal network (SATN) for efficiently optimizing the coherence between predictions of adjacent frames. SATN is the combination of self-attention module and TCN. TCN is employed for its efficient parallel computation and excellent short-term modeling ability. However, the inherently hierarchical structure of convolution layers limits the representation learning of the long-term sequence. More specifically, if the distance of two frames is larger than the kernel size of a single convolution layer, the model requires a stack of convolution layers to construct a long-term connection to relate them. As a result, the associations of the entire sequence cannot be established until the upper layer, which is inefficient. Therefore, we need a more efficient network to establish associations as early as possible. Based on this thought, self-attention is employed to associate the temporal features before TCN. Valuable connections among all frames can be established within just one self-attention layer. In this manner, associations can be efficiently established at the bottom of the temporal network, which greatly helps TCN efficiently learning the short and long-term temporal coherence.

Given a video clip centered at frame  $t$ , SATN is developed to estimate human 3D mesh of frame  $t$ . Correspondingly, the supervision is only for the frame  $t$ . It is hard to determine whether SATN has learned the long-term temporal correlation. In other words, we lack supervision for guiding the temporal representation learning. Therefore, we propose an unsupervised adversarial training strategy for learning the temporal correlation from the order of video frames. In detail, frames of a video clip are first shuffled and then re-sorted using the self-attention and sequence sorting module. By recovering the correct temporal order of the motion sequence in the video, a strong supervision signal is generated for learning the motion dynamics. Besides, considering that the motion orders are reversible and the adjacent poses are similar, the target of sequence sorting is well designed to meet these properties.

Compared with the DSD network, we report an additional 4.2%/7.3% improvements brought by our temporal model, in terms of PA-MPJPE on Human3.6M and 3DPW respectively. The proposed approach outperform the state-of-the-art methods [16, 37] that predict 3D mesh by 15.4% MPJPE and 23.8% PA-MPJPE on Human3.6M. Besides, state-of-the-art results are also achieved on 3DPW without any fine-tuning. Especially, using features reorganized by DSD module for temporal modeling relatively improves 11.1% and 27.5% PA-MPJPE on 3DPW and Human3.6M respectively. It demonstrates that skeleton-disentangled representation is critical for better temporal motion modeling and generalization.

In summary, the following contributions are made:

1. An effective and portable DSD module is proposed to disentangle the skeleton from the rest part of the human 3D mesh. It reduces the complexity and lays a good foundation for temporal modeling.
2. Self-attention temporal network is proposed to learn the short and long-term temporal coherence of 3D human body efficiently.
3. An unsupervised adversarial training strategy is proposed to guide the representation learning of motion dynamics in the video.
4. Our entire framework is trained in an end-to-end manner. We outperform previous approaches that output 3D meshes [15, 16, 27, 37] in terms of 3D joint error.

## 2. Related Work

**Recovering 3D human mesh from single images.** Most of the existing approaches formulated this problem as recovering the parameters of a statistical human body model, SMPL [22]. Recently proposed ConvNet-based methods [15, 26, 32, 27, 37, 17, 39, 36] have shown impressive

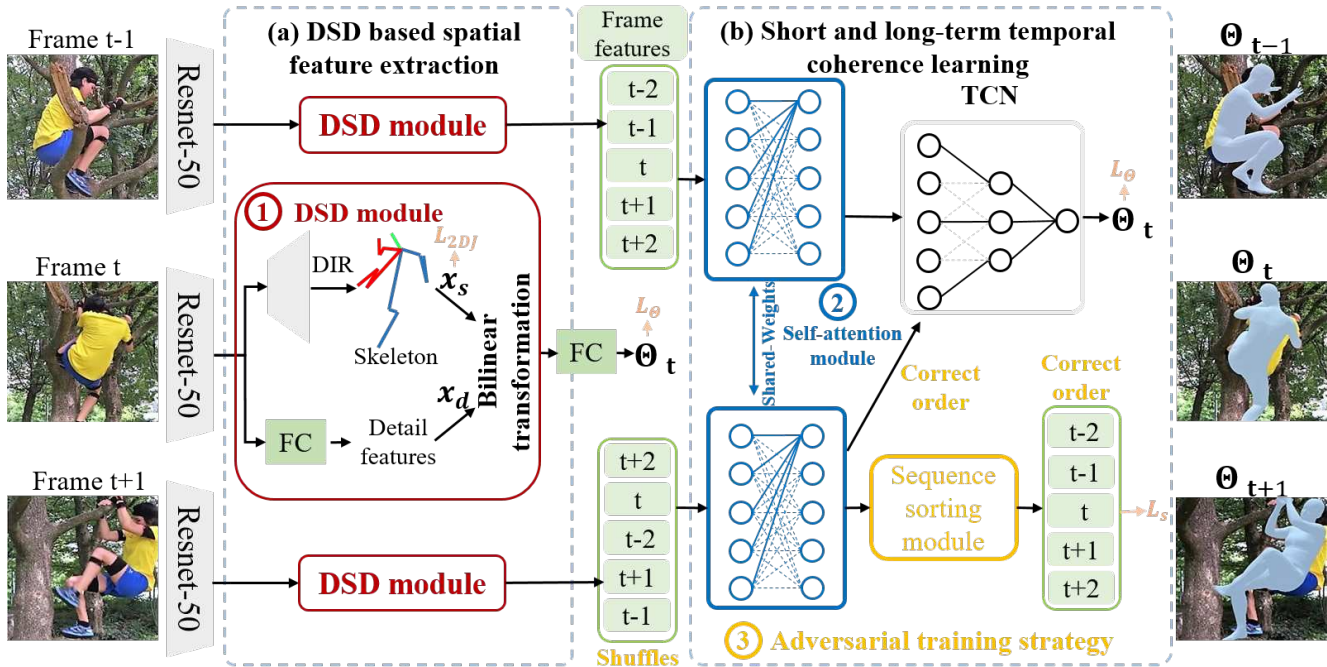


Figure 2. Overview of the proposed skeleton-disentangling based self-attention temporal network.

performance. They can be split into two categories: two stages, direct estimation.

The two-stage methods [26, 27, 17] first predict intermediate results, like human parsing [26], and then predict the SMPL parameters from them. For instance, Pavlakos *et al.* [27] developed two individual networks to infer pose and shape from silhouettes and keypoint locations separately. The two-stage methods are robust to the domain shift, but throw away the important details of human body.

Some other methods [15, 32] directly estimate SMPL parameters from images in an end-to-end manner. In particular, HMR [15] designed a discriminator to distinguish the authenticity of predicted SMPL parameters. In this generative adversarial manner, HMR could be trained with only 2D pose annotations. HMR greatly outperformed all two-stage methods. It may suggest that the end-to-end and non-disruptive feature extraction process is important in this task. Besides, Tung *et al.* [32] developed various supervision approaches, like 2D re-projection of joints and segmentation, to better utilize all available annotations. However, as we introduced before, these methods suffer from severe coupling problem.

To tackle the coupling problem, we propose a DSD module to disentangle the skeleton from the rest details. It reduces network complexity and improves the accuracy of 3D pose recovery. In this process, the detailed information of 2D images is well preserved. The proposed DSD module based single-frame network outperforms all these methods.

**Recovering 3D pose from monocular video.** We focus on 3D pose estimation from monocular video using deep networks, which is the most similar to ours. Various networks are developed to exploit temporal information. For example, TP-Net [5] and Martinez *et al.* [24] trained a fully connected (FC) network to exploit the temporal coherence from the 3D/2D pose sequence. Similarly, Hossain *et al.* [13] employed LSTM to predict 3D poses in a sequence-to-sequence manner. Considering the computation efficiency of RNN-based methods is limited, temporal convolution network (TCN) [3] is proposed and widely used for temporal modeling. For example, Pavllo *et al.* [28] employed the TCN to predict 3D pose from a consecutive 2D pose sequence. They also conduct back-projection of estimated poses from 3D to 2D for supervision using the labels generated by the state-of-the-art 2D pose detector. Currently, as we state before, TCN-based methods still lack temporal supervision. We propose an adversarial training strategy to use the temporal order as supervision.

**Recovering 3D mesh from monocular video.** Few methods are proposed to recover human 3D mesh from monocular video. Generally, we can split the existing methods into two categories: optimization-based and CNN-based. For example, SFV [29] first predicts human 3D meshes of each frame by CNN and then smooths the single-frame results via traditional optimization-based post-processing. However, compared with CNN, the efficiency of optimization-based methods is limited. Most recently,

HMR-video [16] employed the TCN to exploit the temporal information. They proposed to learn the motion dynamic by predicting the actions before and after the current frame. The proposed method outperforms their methods on both 3DPW and Human3.6M.

### 3. Method

The proposed method is to estimate the 3D human mesh parameters  $\Theta$  from single images and monocular videos. As shown in *Figure 2*, the proposed framework has two parts: a) spatially, extracting features of each frame using the proposed DSD module (red) and b) temporally, learning the short and long-term temporal coherence of adjacent frames.

For single images, the most straightforward method is to estimate  $\Theta$  directly from features extracted using Resnet-50 [12]. We follow this simple pipeline and insert a DSD module before the final fully connected layer. The proposed DSD module can re-organize the coupled features and disentangle skeleton from the rest details in feature space.

For monocular videos, we propose a temporal network (SATN) to estimate  $\Theta$  from a tuple of features (light green) re-organized by DSD module. More specifically, from  $n$  ( $n = 5$  in *Figure 2*) frame features centered at frame  $t$ , SATN is built to predict the  $\Theta_t$ . SATN consists of self-attention module (blue) and TCN (gray).

Besides, an adversarial training strategy (gold) is designed to help the self-attention module learning motion dynamics in the video. In the middle of *Figure 2*, features of adjacent frames are shuffled. Self-attention module along with the sequence sorting module is trained to recover the correct temporal order. This strategy is performed in a multi-task learning manner. The correct order of frame features is recovered after the self-attention module. Features in correct temporal order are sent to TCN for predicting the human 3D mesh  $\Theta_t$  as usual.

#### 3.1. 3D Human Body Representation

A parametric statistical 3D human body model, SMPL, is employed to encode the 3D mesh into low-dimensional parameters. SMPL disentangles the shape and pose of a human body. It establishes an efficient mapping  $M(\beta, \theta; \Phi) : \mathbb{R}^{|\theta| \times |\beta|} \mapsto \mathbb{R}^{3 \times 6890}$  from shape  $\beta$  and pose  $\theta$  to a triangulated mesh with 6890 vertices, where  $\Phi$  represents the statistical prior of human body. The shape parameter  $\beta \in \mathbb{R}^{10}$  is the linear combination weights of 10 basic shape. The pose parameter  $\theta \in \mathbb{R}^{3 \times 23}$  represents relative 3D rotation of 23 joints in axis-angle representation. To unify the pose formats of different datasets, 14 common joints of LSP [14] are selected. A linear regressor  $P_{3d}$  is developed to derive these 14 joints from 6890 vertices of human body mesh. The linear combination operation of this regressor guarantees that joints location is differentiable with respect to

shape  $\beta$  and pose  $\theta$  parameters.

A weak-perspective camera model is employed in this task to establish the mapping from 3D space to 2D image plane, in convenience of supervising 3D mesh with 2D pose labels. Finally, a 85 dimensional vector  $\Theta = \{\theta, \beta, \mathbf{R}, \mathbf{t}, \mathbf{s}\}$  is adopted to represent a 3D human body in camera coordinate, where  $\mathbf{R} \in \mathbb{R}^3$  is the global rotation in axis-angle representation,  $\mathbf{t} \in \mathbb{R}^2$  and  $\mathbf{s} \in \mathbb{R}$  represents translation and scale in image plane, respectively. The projection of  $M(\beta, \theta; \Phi)$  is

$$\mathbf{x} = \mathbf{s}\Pi(\mathbf{R}M(\beta, \theta; \Phi)) + \mathbf{t}, \quad (1)$$

where  $\Pi$  is an orthographic projection.

#### 3.2. DSD Module

First of all, we need to extract the 3D mesh information from each frame. As we introduced before, the existing single-frame methods suffer from severe coupling problem. To overcome this problem, we propose the DSD module to disentangle the skeleton from the rest part of the human 3D mesh, including body 3D shape and the detailed poses like the orientation of head, hands, and feet. The DSD module brings two main advantages: a) decoupling reduces the complexity by meticulously task dividing; and b) skeleton-disentangled representation is a better foundation for exploring the temporal motion dynamics.

As shown in the red part of *Figure 2*, a two-branch structure is designed to extract the skeleton and the rest detailed features separately. We follow [31] to estimate joint coordinates of the 2D/3D skeleton with minor modifications. Three de-convolution layers are stacked on top of the backbone and followed by DIR (differential integral regression) to convert the normalized heatmaps  $H$  into 2D/3D joint coordinates. The  $DIR(H)$  is the integration of all location indices  $p$  in the heatmaps, weighted by their probabilities. For instance, the 3D coordinate values of the  $k$ -th joint  $J_k$  could be derived from the  $k$ -th 3D heatmap  $H_k$  by

$$J_k = \sum_{p_z=1}^D \sum_{p_y=1}^H \sum_{p_x=1}^W p H_k(p), \quad (2)$$

where  $D$ ,  $H$ , and  $W$  are depth, height and width of the  $k$ -th heatmap  $H_k$  respectively. For 2D heatmaps,  $D = 1$ .

After independent feature extraction, we need to find a proper way to aggregate them. For the two-factor problems, bilinear transformation is well-known for its strong ability of decoupling the style-content factors [9], such as the identity and head orientation in face recognition; or accent and word class in speech recognition. Therefore, we employ it to disentangle the skeleton from the rest details. Given the 2D skeleton coordinates  $\mathbf{x}_s \in \mathbb{R}^{N \times 28}$  and corresponding detailed features  $\mathbf{x}_d \in \mathbb{R}^{N \times 512}$ , their bilinear transforma-

tion  $\mathbf{y} \in \mathbb{R}^{N \times 512}$  is

$$\mathbf{y} = \mathbf{x}_s \mathbf{A} \mathbf{x}_d^T, \quad (3)$$

where  $\mathbf{A} \in \mathbb{R}^{512 \times 28 \times 512}$  is the learnable weights.

Moreover, with high modularity and the simple design, it is convenient to transplant DSD module to other networks. For instance, DSD module could be easily plugged into the existing 2D/3D pose estimation network for predicting 3D human body meshes. As shown in *Figure 2*, two fully connected layers, and a bilinear transformation layer are all we need for this conversion.

### 3.3. Self-attention Temporal Network

Secondly, we propose a self-attention temporal network (SATN) for learning long and short-term temporal coherence in the video. Given skeleton-disentangled features from DSD module, SATN predicts smoother  $\Theta$  in temporal. TCN performs well in modeling the short-term patterns and efficient parallel computation. But due to its inherent hierarchical structure, it is inefficient to model the long-term correlation in the video. To solve this problem, we propose to build the TCN on top of the self-attention module. As illustrated in *Figure 2*, given a tuple of frame features  $\mathbf{x}_f \in \mathbb{R}^{N \times 512}$  centered at  $t$ , self-attention module is employed to relate different frames and establish the long-term sequence representations. Then, the outputs of the self-attention module are fed into the TCN for predicting the human 3D mesh  $\Theta_t$ .

In particular, the positional encoding is added to each input feature vector for injecting information about their absolute position in sequence. Multi-head attention (MHA) [33] is adopted to accomplish the self-attention mechanism, which is a variance of typical scaled dot-product attention (SDPA) [33]. In this work, we employ 8 heads for MHA. In detail,  $\mathbf{x}_f$  is first linearly projected to  $\{\mathbf{x}_f^i \in \mathbb{R}^{N \times 64}, i = 1, \dots, 8\}$  via 8 fully connected (FC) layers. Then SDPA is performed on each  $\{\mathbf{x}_f^i\}$  in parallel. The outputs are concatenated and linearly projected to  $\mathbf{y}_f \in \mathbb{R}^{N \times 512}$  using a FC layer. The SDPA is derived as

$$SDPA(\mathbf{x}_f) = \text{softmax}\left(\frac{\mathbf{x}_f \mathbf{x}_f^T}{\sqrt{d}}\right) \mathbf{x}_f, \quad (4)$$

where  $d=512$  is the dimension of inputs and serves as the scaling factor. In *Equation 4*, the output is the weight sum of  $\mathbf{x}_f$ . The weight matrix represents the relation between each two frames. The paths length between different frames is reduced to constants. With the assistance of the self-attention module, the bottom layers of TCN could reach to the information of entire sequence. The receptive field is relatively expanded. Short and long-term temporal coherence can be learned more efficiently.

Although we predict from multi-frame features in temporal, the supervision is still in single-frame level. We lack

the temporal supervision for the long-term representation learning. To tackle this problem, we propose an unsupervised adversarial training strategy. Similar to recent methods [25, 30, 8, 18], we also use the temporal order of frames as the supervision for the representation learning. However, there are some differences in the formulation of the problem. Most of the previous methods [25, 30, 8] formulated the problem as the binary classification that verifies the correct/incorrect temporal order. Lee *et al.* [18] developed it to predict  $n!/2$  combinations for each  $n$ -tuple of frames. In this work, we further develop it to directly estimate the original position indices of the shuffled frames for richer supervision. Besides, two adaptive changes of loss function are made to meet the special properties of the motion sequence.

As illustrated in the *Figure 2*, frame features are first shuffled before input. As before, the input is added up with the positional coding, which stays unchanged to avoid disclosure of the order information. After going through the self-attention module, the outputs are sent to the TCN and the sequence sorting module separately. Before fed into TCN, the shuffled order is retrieved to avoid the interference with the normal representation learning of the TCN. The sequence sorting module predicts the position indices of the correct order. Specifically, for each tuple of 9 shuffled frames, we need to predict 9 indices, indicating their positions in the original sequence. Besides, two adaptive modifications of sorting target are made. Firstly, considering the reasonable orders of some actions (e.g., standing up/sitting down) are reversible, the minor loss between the predicted and the forward/backward ground truth orders will be chosen. Secondly, sometimes, the differences between adjacent frames are too small to be properly preserved after single-frame convolution encoding. In this situation, the context order of adjacent frames becomes ambiguous. But the typical one-hot label will make equal punishments on this kind of failures without considering the adjacent ambiguity. To tackle this problem, we replace the hard one-hot label with the soft Gaussian-like label for proper supervision.

### 3.4. Loss Functions

As we mentioned in *Section 1*, full 3D annotations of 2D image available are limited to the experimental environment. Models trained on these data cannot generalize well to the in-the-wild images. To make full use of existing 2D/3D data, the estimated 3D human body parameters are optimized with

$$L_{\Theta} = w_{pm} L_{pm} + w_{3d} L_{3d} + w_{2d} L_{2d} + w_r L_r \quad (5)$$

where  $w_{pm}, w_{3d}, w_{2d}, w_r$  are weights of these loss items. For images with motion capture,  $L_{pm}$  is employed to supervise the the body parameter  $\theta$ , acquired by moshing [21], with L2 loss directly. In particular, inspired by [15, 27], angular pose  $\theta \in \Theta$  is converted from rotation vectors to






Figure 3. Some qualitative results of DSD network.

$3 \times 3$  rotation matrices via the Rodrigues formula for stable training. For images with 2D/3D pose annotations,  $L_{2d}$  and  $L_{3d}$  are employed to supervise the skeleton pose  $J^{2d}/J^{3d}$  of estimated body mesh using L1/L2 loss respectively.  $J^{3d}$  and  $J^{2d}$  are derived from the predicted  $\Theta = \{\theta, \beta, \mathbf{R}, \mathbf{t}, \mathbf{s}\}$  by

$$J^{3d} = P_{3d}(M(\beta, \theta; \Phi)) \quad (6)$$

and

$$J^{2d} = \mathbf{s}\Pi(\mathbf{R}J^{3d}) + \mathbf{t}, \quad (7)$$

where  $M$  represents a mapping from shape  $\beta$  and pose  $\theta$  to a 3D human body mesh,  $P_{3d}$  represents a linear mapping from a 3D mesh to 3D joints coordinates and  $\Pi$  is an orthographic projection, as we defined in *Section 3.1*.

Besides, it is important to establish rational constraints of joint angles and body shape, especially for learning from images with only 2D pose annotations. Therefore, we follow [15] and employ a discriminator to provide rationality loss  $L_r$  using Mocap data [1]. Besides,  $L_{2DJ}$  is employed to supervise the 2D skeleton coordinates  $x_s$  in the DSD module using L1 loss. The sorting results of sequence sorting module are supervised with L2 loss of  $L_s$ .

## 4. Experiments

### 4.1. Datasets

**Human3.6M** is the only dataset with 3D annotations we used to train. It contains videos of multiple actors performing 17 activities, which are captured in a controlled environment. We downsample all videos from 50fps to 10fps for removing the redundancy. The proposed method is evaluated on common protocol following [15] for a more comprehensive comparison. In details, we train on 5 subjects (S1, S5, S6, S7, S8) and test on subject S9 and S11. We report both the mean per joint position error (MPJPE) and Procrustes

Aligned MPJPE (PA-MPJPE), which is MPJPE after rigid alignment of predicted pose with ground truth, in millimeters. Except for the common MPJPE and PA-MPJPE, mean per joint velocity and acceleration error (MPJVE/MPJAE), in  $mm/s$  and  $mm/s^2$  respectively, are adopted to evaluate the smoothness and stability of predictions over time.

**3DPW** [34] is a recent challenge dataset that contains 60 video sequences (24 train, 24 test, 12 validation) of richer activities, such as climbing, golfing, relaxing on the beach, etc. They leverage video and IMU to obtain accurate 3D pose despite the complexity of scenes. For a fair comparison, none of the approaches get trained on 3DPW. We report PA-MPJPE on all splits (train/test/val) of this dataset.

**2D in-the-wild datasets.** We use MPII [2], LSP [14], AICH [35], and Penn Action [38], which are 2D pose datasets without 3D annotations, to train our single-frame DSD network for better generalization. We follow [4, 17, 15] to use 14 LSP joints as the skeleton.

### 4.2. Implement Details

**Architecture:** The framework has spatial and temporal parts. In spatial, features are extracted using ResNet-50, pre-trained on MPII, and further decoupled using DSD module. In temporal, features of video clips are gathered and put into SATN, which compose of self-attention module and TCN. The self-attention module contains 2 Transformer [33] blocks. TCN follows the design of [28, 3] but only contains 2 convolution blocks. Since the Human3.6M is down-sampled, we take a receptive field of 9 frames for the TCN, equivalent to 45 frames in the original video. The sequence sorting module consists of 3 convolution blocks followed by 3 FC layers.

**Training details:** We train both single-frame DSD network and SATN for 40 epochs. Considering the efficiency of training, we pre-compute the single-frame features and train SATN with them directly. The SGD is adopted as the optimizer with momentum=0.9. The learning rate and batch size are set to  $1e-4$  and 16 respectively. The weights of loss items are set as  $w_{pm} = 20$ ,  $w_r = 0.6$ ,  $w_{2d} = 10$ ,  $w_{3d} = 60$ . Besides, considering the domain gap between the train and the test set, we discard the hyper-parameters of the batch normalization layers during the evaluation. Besides, to enhance the stability of occlusion, we enlarge the scale augmentation to generate some samples with half body.

### 4.3. Comparisons to State-of-the-art Approaches

The results in *Table 1* shows that the proposed methods achieve the state-of-the-art results on 3DPW. Again, all approaches are directly tested without any fine-tuning and Human3.6M is the only 3D train set. It demonstrates that the proposed methods perform well in generalization. The most similar method to ours is HMR-video [16], which also learns motion dynamic with TCN for temporal op-

| Method               | Train       | Test        | val         |
|----------------------|-------------|-------------|-------------|
| Simple-baseline [24] | -           | 157.0       | -           |
| SMPLify [4]          | -           | 106.8       | -           |
| TP-Net [5]           | -           | 92.2        | -           |
| HMR-video [16]       | -           | 80.1        | -           |
| HMR-video-L [16]     | 75.9        | 72.6        | 76.7        |
| <b>DSD</b>           | 75.0        | 75.0        | 78.0        |
| <b>DSD+SATN</b>      | <b>68.3</b> | <b>69.5</b> | <b>71.8</b> |

Table 1. Comparisons to state-of-the-art methods on entire 3DPW in terms of PA-MPJPE without any fine-tuning.

| Method          | MPJPE↓      | PA-MPJPE↓   |
|-----------------|-------------|-------------|
| HMR [15]        | 87.9        | 58.1        |
| STN [37]        | 69.9        | 61.4        |
| HMR-video       | -           | 57.8        |
| HMR-video-L     | -           | 55.7        |
| <i>Direct</i>   | 110.0       | 62.8        |
| <i>Concat</i>   | 91.8        | 63.1        |
| <b>DSD</b>      | 60.1        | 44.3        |
| <b>DSD+SATN</b> | <b>59.1</b> | <b>42.4</b> |

Table 2. Comparisons to state-of-the-art methods that predict 3D meshes on Human3.6M.

| Method                 | MPJPE↓      | PA-MPJPE↓   |
|------------------------|-------------|-------------|
| Backbone+TCN           | 87.8        | 58.5        |
| Backbone+SATN          | 80.6        | 58.6        |
| DSD+TCN                | 82.3        | 52.0        |
| DSD+TCN+Self-attention | 59.6        | 43.4        |
| <b>DSD+SATN</b>        | <b>59.1</b> | <b>42.4</b> |

Table 3. Ablation study of the components in SATN.

timization. We compared with their two evaluation settings, HMR-video and HMR-video-L. The main difference between two settings is that HMR-video is trained with datasets of similar scale with ours, while HMR-video-L is trained with their internet dataset which is nearly 20x larger than ours. Our approach outperforms HMR-video in terms of PA-MPJPE by 13.2% on 3DPW test set. On all splits of 3DPW, we evaluate HMR-video-L model they released. As shown in *Table 1*, the proposed method superior to HMR-video-L on all splits.

State-of-the-art results are also achieved on Human3.6M. In *Table 2*, the proposed method outperforms the HMR-video and HMR-video-L by 26.6% and 23.8% in terms of PA-MPJPE respectively. Note that, in *Table 1*, TP-Net and Simple-baseline, also trained on Human3.6M, perform extremely well on Human3.6M, while showing poor generalization ability on 3DPW. Besides, the proposed method greatly outperforms SMPLify on 3DPW, which indicates that the improvements are not brought by the different 3D pose representation.

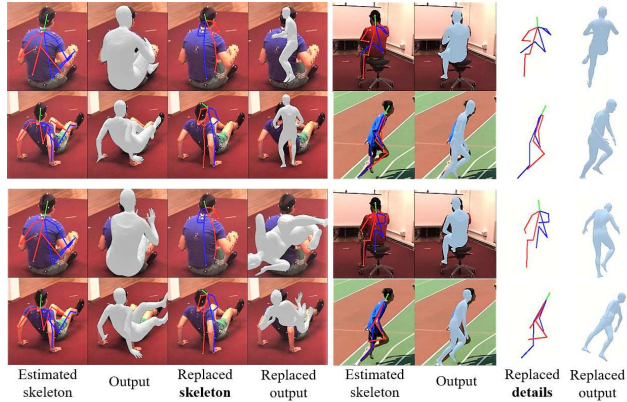


Figure 4. Decoupling effectiveness comparison between DSD module (the two rows upper) and *Concat*(the two rows lower) by replacing the skeleton (shown in the left part) or the detailed features (shown in the right part).

#### 4.4. Evaluation of the DSD module

**1) The ablation study of the DSD module.** Two baseline methods are compared in *Table 2* for evaluating the DSD module. *Direct* is to directly estimate the  $\Theta$  without DSD module. *Concat* is to replace the bilinear transformation layer of DSD module with the concatenation operation. Compared with *Direct*, *Table 2* shows that *Concat* performs even worse in PA-MPJPE, but adding DSD module brings significant improvement (by 45.3% MPJPE and 29.4% PA-MPJPE). The results demonstrate that decoupling skeleton improves the accuracy of 3D pose recovery. Next, we will validate that the skeleton is sufficiently disentangled from the rest details by the DSD module.

**2) Decoupling effectiveness.** We set up a visual experiment to compare the decoupling effectiveness of DSD module and *Concat*. For verifying whether the skeleton is sufficiently decoupled, we deliberately replace the estimated skeleton with the randomly selected one. The left part of *Figure 4* with gray human mesh shows the different responses from DSD module (the upper two rows) and *Concat* (the lower two rows) to this replacement. As we can see, the body postures of 3D mesh predicted by the DSD module are correctly changed together with the replaced skeleton, while the outputs from *Concat* are completely messed up. The experimental results show that the skeleton is successfully disentangled from the rest details. Likewise, a similar conclusion can also be drawn from the right part of *Figure 4* with blue human mesh, which is the results when we replace the detailed features and keep skeleton unchanged.

**3) Stability.** In addition, *Table 4* and *5* show their MPJVE and MPJAE. Note that *Concat* has much larger MPJVE and MPJAE than the rest, even including *Direct*. It indicates that the predictions derived from the concatenated features are more unstable. By contrast, the DSD module

| Method        | Dir. | Disc. | Eat  | Greet | Phone | Pose | Purch. | Sit  | SitD. | Smoke | Photo | Wait | Walk | WalkD. | WalkT. | Avg.↓ |
|---------------|------|-------|------|-------|-------|------|--------|------|-------|-------|-------|------|------|--------|--------|-------|
| <i>Direct</i> | 14.5 | 15.6  | 14.9 | 15.6  | 15.4  | 15.0 | 15.9   | 15.4 | 16.1  | 15.6  | 15.7  | 15.7 | 15.8 | 16.0   | 16.1   | 16.1  |
| <i>Concat</i> | 20.0 | 20.9  | 19.4 | 19.7  | 19.2  | 18.8 | 19.5   | 19.1 | 20.1  | 19.4  | 19.7  | 19.6 | 19.7 | 20.4   | 20.4   | 20.5  |
| DSD           | 12.3 | 12.8  | 12.2 | 12.6  | 12.4  | 12.1 | 12.1   | 11.8 | 12.3  | 12.0  | 12.0  | 12.0 | 12.2 | 12.4   | 12.5   | 12.5  |
| DSD+SATN      | 8.8  | 9.1   | 8.5  | 8.9   | 8.7   | 8.6  | 8.7    | 8.3  | 8.4   | 8.3   | 8.3   | 8.3  | 8.4  | 8.6    | 8.7    | 8.7   |

Table 4. MPJVE( $mm/s$ ): Velocity error over the poses of the predicted 3D human meshes.

| Method        | Dir. | Disc. | Eat  | Greet | Phone | Pose | Purch. | Sit  | SitD. | Smoke | Photo | Wait | Walk | WalkD. | WalkT. | Avg.↓ |
|---------------|------|-------|------|-------|-------|------|--------|------|-------|-------|-------|------|------|--------|--------|-------|
| <i>Direct</i> | 13.1 | 14.2  | 13.5 | 14.3  | 14.0  | 13.7 | 14.7   | 14.0 | 14.3  | 14.0  | 14.0  | 14.0 | 14.3 | 14.6   | 14.8   | 14.8  |
| <i>Concat</i> | 17.1 | 18.0  | 16.6 | 17.1  | 16.5  | 16.2 | 16.8   | 16.2 | 16.7  | 16.1  | 16.4  | 16.4 | 16.7 | 17.3   | 17.4   | 17.4  |
| DSD           | 11.5 | 12.0  | 11.3 | 11.8  | 11.5  | 11.2 | 11.3   | 10.8 | 11.0  | 10.8  | 10.8  | 10.8 | 11.1 | 11.3   | 11.5   | 11.5  |
| DSD+SATN      | 7.0  | 7.4   | 6.8  | 7.2   | 6.9   | 6.8  | 6.9    | 6.5  | 6.5   | 6.3   | 6.4   | 6.3  | 6.5  | 6.7    | 6.8    | 6.8   |

Table 5. MPJAE( $mm/s^2$ ): Acceleration error over the poses of the predicted 3D human meshes.

improves stability with smoother predictions.

**4) Comparisons to the state-of-the-art single-frame methods.** In *Table 2*, we compare proposed DSD module with HMR and STN. We tightly follow the same evaluation protocol and use the same backbone, ResNet-50, as stated in their articles. Compared with HMR, adding DSD module reduces the error by 31.6% MPJPE and 23.7% PA-MPJPE, which further proves the effectiveness of DSD module. We also compare to STN [37], which is the most recent proposed method that also predicts angular pose and outputs 3D meshes. In *Table 2*, DSD network outperforms the STN by 14% MPJPE and 27.9% PA-MPJPE. Some qualitative results of DSD module are illustrated in Figure 3.

#### 4.5. Evaluation of SATN

**1) The ablation study of SATN.** *Table 2* shows that compared with the single-frame DSD network, adding SATN brings an additional 4.2%/7.3% improvement in terms of MPJPE/PA-MPJPE on Human3.6M. In addition, *Table 1* shows that adding SATN brings a higher improvement (7.3% PA-MPJPE) on 3DPW. Besides, except for reducing the error of 3D pose recovery, as shown in *Table 4* and 5, SATN significantly improves the smoothness of the predictions.

**2) The ablation study of components in SATN.** *Table 3* shows that performance gets steadily improved by adding the proposed components, including the self-attention module and the adversarial training strategy. Note that making temporal modeling simply with the TCN is not enough. Compared with the results of the single-frame DSD network in *Table 2*, *Table 3* shows that DSD+TCN leads to the degradation of performance. Similar conclusions can also be drawn from HMR-video [16]. As shown in *Table 3*, involving self-attention module greatly improves the performance of the temporal network and reverses the degradation. This phenomenon demonstrates that it is vital for TCN to establish associations of the entire input sequence

as early as possible and self-attention module performs well in this task.

#### 4.6. Role of the DSD Module for SATN

The results in *Table 3* shows that skeleton-disentangled representation is of great importance for effective temporal modeling. Backbone+TCN in *Table 3* is to train the TCN with features from the backbone directly without DSD module. As shown in *Table 2*, the results of Backbone+TCN are comparative to the HMR-video [16], which has similar architecture. By contrast, involving DSD module (DSD+TCN in *Table 3*) brings a 11.1% improvement in terms of PA-MPJPE. Besides, when we replace TCN with SATN, the performances get further improvement. Especially, in *Table 3*, the combination of DSD and SATN (DSD+TCN+Self-attention+adversarial training) causes a magic reaction and outperforms all backbone-based methods by 26.6% MPJPE and 27.5% PA-MPJPE. However, PA-MPJPE of Backbone+TCN and Backbone+SATN in *Table 3* are nearly equal, which indicates that skeleton-disentangled representation is essential for motion modeling.

### 5. Conclusion

We propose an end-to-end framework for recovering 3D human mesh from single images and monocular videos via a skeleton-disentangled representation. Decoupling skeleton reduces the complexity and the 3D pose error. The proposed DSD module could be regarded as an efficient bridge between 2D/3D pose estimation and 3D mesh recovery. Moreover, SATN is well designed to explore long and short-term temporal coherence. Massive evaluations demonstrate that we provide an attractive and efficient baseline method for related problems, including human motion analysis, 3D virtual try-on and multi-person 3D meshes recovery.



## References

- [1] CMU graphics lab motion capture database.
- [2] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2014.
- [3] Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv:1803.01271*, 2018.
- [4] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it smpl: Automatic estimation of 3D human pose and shape from a single image. In *European Conference on Computer Vision, ECCV*, 2016.
- [5] Rishabh Dabral, Anurag Mundhada, Uday Kusupati, Safeer Afaque, Abhishek Sharma, and Arjun Jain. Learning 3d human pose from structure and motion. In *European Conference on Computer Vision, ECCV*, 2018.
- [6] Haoye Dong, Xiaodan Liang, Xiaohui Shen, Bochao Wang, Hanjiang Lai, Jia Zhu, Zhiting Hu, and Jian Yin. Towards multi-pose guided virtual try-on network. In *IEEE International Conference on Computer Vision, ICCV*, 2019.
- [7] Haoye Dong, Xiaodan Liang, Xiaohui Shen, Bowen Wu, Bing-Cheng Chen, and Jian Yin. Flow-navigated warping gan for video virtual try-on. In *IEEE International Conference on Computer Vision, ICCV*, 2019.
- [8] Basura Fernando, Hakan Bilen, Efstratios Gavves, and Stephen Gould. Self-supervised video representation learning with odd-one-out networks. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2017.
- [9] William T Freeman and Joshua B Tenenbaum. Learning bilinear models for two-factor problems in vision. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 1997.
- [10] Chuang Gan, Naiyan Wang, Yi Yang, Dit-Yan Yeung, and Alex G Hauptmann. Devnet: A deep event network for multimedia event detection and evidence recounting. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2015.
- [11] Chuang Gan, Yi Yang, Linchao Zhu, Deli Zhao, and Yueting Zhuang. Recognizing an action using its name: A knowledge-based approach. *International Journal of Computer Vision*, 2016.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2016.
- [13] Mir Rayat Imtiaz Hossain and James J Little. Exploiting temporal information for 3d human pose estimation. In *European Conference on Computer Vision, ECCV*, 2018.
- [14] Sam Johnson and Mark Everingham. Clustered pose and nonlinear appearance models for human pose estimation. In *Proceedings of the British Machine Vision Conference*, 2010.
- [15] Angjoo Kanazawa, Michael J. Black, David W. Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2018.
- [16] Angjoo Kanazawa, Jason Zhang, Panna Felsen, and Jitendra Malik. Learning 3d human dynamics from video. *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2019.
- [17] Christoph Lassner, Javier Romero, Martin Kiefel, Federica Bogo, Michael J Black, and Peter V Gehler. Unite the people: Closing the loop between 3d and 2d human representations. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2017.
- [18] Hsin-Ying Lee, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Unsupervised representation learning by sorting sequences. In *IEEE International Conference on Computer Vision, ICCV*, 2017.
- [19] Kun Liu, Wu Liu, Chuang Gan, Mingkui Tan, and Huadong Ma. T-c3d: Temporal convolutional 3d network for real-time action recognition. In *Thirty-second AAAI conference on artificial intelligence*, 2018.
- [20] Wu Liu, Yongdong Zhang, Sheng Tang, Jinhui Tang, Richang Hong, and Jintao Li. Accurate estimation of human body orientation from rgb-d sensors. *IEEE Transactions on cybernetics*, 2013.
- [21] Matthew Loper, Naureen Mahmood, and Michael J Black. Mosh: Motion and shape capture from sparse markers. *ACM Transactions on Graphics, TOG*, 2014.
- [22] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Transactions on Graphics, TOG*, 2015.
- [23] Huadong Ma and Wu Liu. A progressive search paradigm for the internet of things. *IEEE MultiMedia*, 2017.
- [24] Julieta Martinez, Rayat Hossain, Javier Romero, and James J. Little. A simple yet effective baseline for 3d human pose estimation. In *IEEE International Conference on Computer Vision, ICCV*, 2017.
- [25] Ishan Misra, C Lawrence Zitnick, and Martial Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *European Conference on Computer Vision, ECCV*, 2016.
- [26] 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 *IEEE International Conference on 3D Vision, 3DV*, 2018.
- [27] Georgios Pavlakos, Luyang Zhu, XiaoWei Zhou, and Kostas Daniilidis. Learning to estimate 3D human pose and shape from a single color image. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2018.
- [28] Dario Pavllo, Christoph Feichtenhofer, David Grangier, and Michael Auli. 3d human pose estimation in video with temporal convolutions and semi-supervised training. *arXiv preprint arXiv:1811.11742*, 2018.
- [29] Xue Bin Peng, Angjoo Kanazawa, Jitendra Malik, Pieter Abbeel, and Sergey Levine. SfV: Reinforcement learning of physical skills from videos. In *ACM SIGGRAPH Asia 2018 Technical Papers*, 2018.
- [30] Omer Sumer, Tobias Dencker, and Bjorn Ommer. Self-supervised learning of pose embeddings from spatiotempo-

- ral relations in videos. In *IEEE International Conference on Computer Vision, ICCV*, 2017.
- [31] Xiao Sun, Bin Xiao, Fangyin Wei, Shuang Liang, and Yichen Wei. Integral human pose regression. In *European Conference on Computer Vision, ECCV*, 2018.
- [32] Hsiao-Yu Tung, Hsiao-Wei Tung, Ersin Yumer, and Katerina Fragkiadaki. Self-supervised learning of motion capture. In *Advances in Neural Information Processing Systems, NIPS*, 2017.
- [33] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems, NIPS*, 2017.
- [34] Timo von Marcard, Roberto Henschel, Michael Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3d human pose in the wild using imus and a moving camera. In *European Conference on Computer Vision, ECCV*, 2018.
- [35] Jiahong Wu, He Zheng, Bo Zhao, Yixin Li, Baoming Yan, Rui Liang, Wenjia Wang, Shipai Zhou, Guosen Lin, Yanwei Fu, et al. Ai challenger: A large-scale dataset for going deeper in image understanding. *arXiv preprint arXiv:1711.06475*, 2017.
- [36] Weipeng Xu, Avishek Chatterjee, Michael Zollhöfer, Helge Rhodin, Dushyant Mehta, Hans-Peter Seidel, and Christian Theobalt. Monoperfcap: Human performance capture from monocular video. *ACM Transactions on Graphics, TOG*, 2018.
- [37] Yusuke Yoshiyasu, Ryusuke Sagawa, Ko Ayusawa, and Akihiko Murai. Skeleton transformer networks: 3d human pose and skinned mesh from single rgb image. *arXiv preprint arXiv:1812.11328*, 2018.
- [38] Weiyu Zhang, Menglong Zhu, and Konstantinos G Derpanis. From actemes to action: A strongly-supervised representation for detailed action understanding. In *IEEE International Conference on Computer Vision, ICCV*, 2013.
- [39] Xiaowei Zhou, Menglong Zhu, Georgios Pavlakos, Spyridon Leonardos, Konstantinos G Derpanis, and Kostas Daniilidis. Monocap: Monocular human motion capture using a cnn coupled with a geometric prior. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.