

MHFormer: Multi-Hypothesis Transformer for 3D Human Pose Estimation

Wenhao Li¹ Hong Liu^{1,*} Hao Tang² Pichao Wang³ Luc Van Gool²

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

²Computer Vision Lab, ETH Zurich ³Alibaba Group

{wenhaoli, hongliu}@pku.edu.cn

{hao.tang, vangool}@vision.ee.ethz.ch pichao.wang@alibaba-inc.com

Abstract

Estimating 3D human poses from monocular videos is a challenging task due to depth ambiguity and self-occlusion. Most existing works attempt to solve both issues by exploiting spatial and temporal relationships. However, those works ignore the fact that it is an inverse problem where multiple feasible solutions (*i.e.*, hypotheses) exist. To relieve this limitation, we propose a Multi-Hypothesis Transformer (MHFormer) that learns spatio-temporal representations of multiple plausible pose hypotheses. In order to effectively model multi-hypothesis dependencies and build strong relationships across hypothesis features, the task is decomposed into three stages: (i) Generate multiple initial hypothesis representations; (ii) Model self-hypothesis communication, merge multiple hypotheses into a single converged representation and then partition it into several diverged hypotheses; (iii) Learn cross-hypothesis communication and aggregate the multi-hypothesis features to synthesize the final 3D pose. Through the above processes, the final representation is enhanced and the synthesized pose is much more accurate. Extensive experiments show that MHFormer achieves state-of-the-art results on two challenging datasets: Human3.6M and MPI-INF-3DHP. Without bells and whistles, its performance surpasses the previous best result by a large margin of 3% on Human3.6M. Code and models are available at <https://github.com/Vegetebird/MHFormer>.

1. Introduction

3D human pose estimation (HPE) from monocular videos is a fundamental vision task with a wide range of applications, such as action recognition [23, 24, 39], human-computer interaction [7], and augmented/virtual reality [31]. This task is typically solved by dividing it into two decoupled subtasks, *i.e.*, 2D pose detection to localize the keypoints

*Corresponding author: hongliu@pku.edu.cn. This work is supported by National Key R&D Program of China (No. 2020AAA0108904), Science and Technology Plan of Shenzhen (No. JCYJ20200109140410340).




Figure 1. Given a frame with occluded body parts (right arm and elbow), a recent state-of-the-art 3D HPE method, PoseFormer [46], outputs a single solution that is inconsistent with the 2D input. In contrast, our MHFormer generates multiple plausible hypotheses (different colors) consistent with the 2D evidence and finally synthesizes a more accurate 3D pose (green). For easy comparison, the input frame is shown at a novel viewpoint.

on the image plane, followed by 2D-to-3D lifting to infer joint locations in the 3D space from 2D keypoints. Despite their impressive performance [4, 9, 29, 34], it remains an inherently ill-posed problem because of self-occlusion and depth ambiguity in 2D representations.

To alleviate such issues, most methods [2, 11, 38, 46] focus on exploring spatial and temporal relationships. They either employ graph convolutional networks to estimate 3D poses with a spatio-temporal graph representation of human skeletons [2, 11, 38] or apply a pure Transformer-based model to capture spatial and temporal information from 2D pose sequences [46]. Yet, the 2D-to-3D lifting from monocular videos is an inverse problem [1] where multiple feasible solutions (*i.e.*, hypotheses) exist due to its ill-posed nature given the missing depth [17]. Those approaches ignore this problem and only estimate a single solution, which often leads to unsatisfactory results, especially when the person is severely occluded (see Figure 1).

Recently, a couple of methods [14, 16, 35, 40] that generate multiple hypotheses have been proposed for the inverse problem. They often rely on the one-to-many mapping by adding multiple output heads to an existing architecture with a shared feature extractor, while failing to build the relationships among the features of different hypotheses. That is an important shortcoming, as such ability is vital to improve the

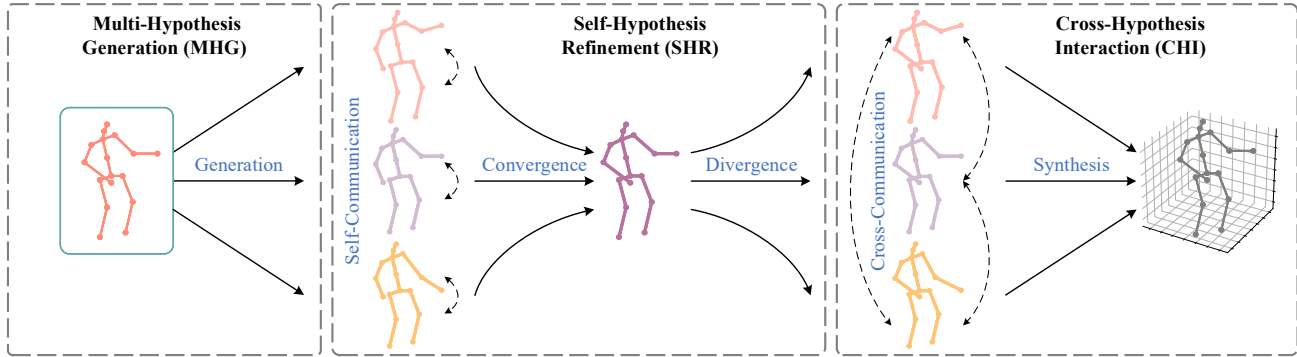


Figure 2. The proposed MHFormer constructs a three-stage framework by starting from generating multiple initial representations and then communicating among them in both independent and mutual ways to synthesize a more precise estimation. For easy illustration, we only show the process of a single-frame 2D pose as input.

expressiveness and the performance of the model. In view of the ambiguous inverse problem of 3D HPE, we argue that it is more reasonable to conduct a one-to-many mapping first and then a many-to-one mapping with various intermediate hypotheses, as this way can enrich the diversity of features and produce a better synthesis for the final 3D pose.

To this end, we present Multi-Hypothesis Transformer (MHFormer), a novel Transformer-based method for 3D HPE from monocular videos. The key insight is to allow the model to learn spatio-temporal representations of diverse pose hypotheses. To accomplish this, we introduce a three-stage framework that starts from generating multiple initial representations and gradually communicates across them to synthesize a more accurate prediction, as shown in Figure 2. This framework more effectively models multi-hypothesis dependencies while also building stronger relationships among hypothesis features. Specifically, in the first stage, a Multi-Hypothesis Generation (MHG) module is built to model the intrinsic structure information of human joints and generate several multi-level features in the spatial domain. Those features contain diverse semantic information in different depths from shallow to deep and hence can be regarded as initial representations of multiple hypotheses.

Next, we propose two novel modules to model temporal consistencies and enhance those coarse representations in the temporal domain, which have not been explored by the existing works that generate multiple hypotheses. In the second stage, a Self-Hypothesis Refinement (SHR) module is proposed to refine every single-hypothesis feature. The SHR consists of two new blocks. The first block is a multi-hypothesis self-attention (MH-SA) which models single-hypothesis dependencies independently to construct self-hypothesis communication, enabling message passing within each hypothesis for feature enhancement. The second block is a hypothesis-mixing multi-layer perceptron (MLP) that exchanges information across hypotheses. The multiple hypotheses are merged into a single converged representation, and then this representation is partitioned into several

diverged hypotheses.

Although those hypotheses are refined by SHR, the connections across different hypotheses are not strong enough since the MH-SA in the SHR only passes intra-hypothesis information. To address this issue, in the last stage, a Cross-Hypothesis Interaction (CHI) module models interactions among multi-hypothesis features. Its key component is the multi-hypothesis cross-attention (MH-CA), which captures mutual multi-hypothesis correlations to build cross-hypothesis communication, enabling message passing among hypotheses for better interaction modeling. Subsequently, a hypothesis-mixing MLP is used to aggregate the multiple hypotheses to synthesize the final prediction.

With the proposed MHFormer, multi-hypothesis spatio-temporal feature hierarchies are explicitly incorporated into Transformer models, where the multiple hypothesis information of body joints can be independently and mutually processed in an end-to-end manner. As a result, the representation ability is potentially enhanced and the synthesized pose is much more accurate. Our contributions are summarized as follows:

- We present a new Transformer-based method, called Multi-Hypothesis Transformer (MHFormer), for 3D HPE from monocular videos. MHFormer can effectively learn spatio-temporal representations of multiple pose hypotheses in an end-to-end manner.
- We propose to communicate among multi-hypothesis features both independently and mutually, providing powerful self-hypothesis and cross-hypothesis message passing, and strong relationships among hypotheses.
- Our MHFormer achieves state-of-the-art performance on two challenging datasets for 3D HPE, significantly outperforming PoseFormer [46] by 3% with 1.3 *mm* error reduction on Human3.6M [12].

2. Related Work

3D Human Pose Estimation. Existing single-view 3D pose estimation methods can be divided into two mainstream

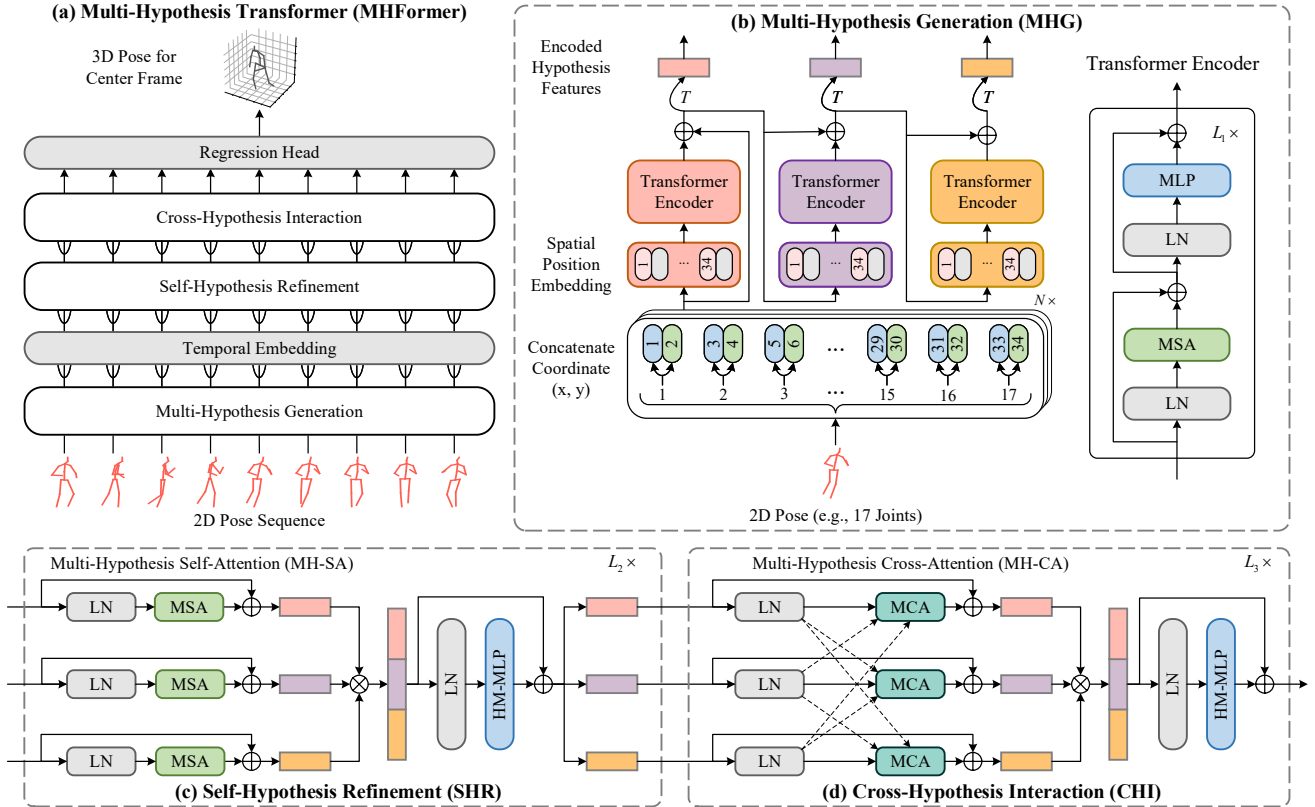


Figure 3. (a) Overview of the proposed Multi-Hypothesis Transformer (MHFormer). (b) Multi-Hypothesis Generation (MHG) module extracts the intrinsic structure information of human joints within each frame and generates multiple hypothesis representations. N is the number of input frames and T is the matrix transposition. (c) Self-Hypothesis Refinement (SHR) module is used to refine single-hypothesis features. (d) Cross-Hypothesis Interaction (CHI) module following SHR enables interactions among multi-hypothesis features.

types: one-stage approaches and two-stage ones. One-stage approaches directly infer 3D poses from input images without intermediate 2D pose representations [18, 28, 33, 36], while two-stage ones first obtain 2D keypoints from pre-trained 2D pose detections and then feed them into a 2D-to-3D lifting network to estimate 3D poses. Benefiting from the excellent performance of 2D human pose estimation, this 2D-to-3D pose lifting method can efficiently and accurately regress 3D poses using detected 2D keypoints. For instance, SimpleBaseline [29] proposes a fully-connected residual network to lift 2D keypoints to 3D joint locations from a single frame. Anatomy3D [4] decomposes the task into bone direction and bone length predictions to ensure temporal consistency over a sequence. Despite the promising results achieved by using temporal correlations from fully convolutional [4, 25, 34] or graph-based [2, 11, 38] architectures, these methods are less efficient in capturing global-context information across frames.

Vision Transformers. Recently, Transformer [37] equipped with the powerfully global self-attention mechanism has received increasingly research interest in the computer vision community [10, 22, 27, 43]. For the basic image classification task, ViT [6] is proposed to apply a standard Transformer ar-

chitecture directly to sequential image patches. For the pose estimation task, PoseFormer [46] applies a pure Transformer to capture human joint correlations and temporal dependencies. Strided Transformer [20] introduces a Transformer-based architecture with strided convolutions to lift a long 2D pose sequence to a single 3D pose. Our work is inspired by them and similarly uses the Transformer as the basic architecture. But we do not just utilize a simple architecture with a single representation; instead, the seminal ideas of multi-hypothesis and multi-level feature hierarchies are connected within Transformers, which makes the model not only expressive but also strong. Besides, a cross-attention mechanism is introduced for effective multi-hypothesis learning.

Multi-Hypothesis Methods. Single-view 3D HPE is ill-posed and therefore assuming only a single solution might be sub-optimal. Several works generate diverse hypotheses for the inverse problem and achieve substantial performance gains [13, 17, 32, 40]. For example, Jahangiri *et al.* [13] generated multiple 3D pose candidates consistent with 2D keypoints via a compositional model and anatomical constraints. Wehrbein *et al.* [40] modeled the posterior distribution of 3D pose hypotheses with normalized flows. Unlike these works that focus on a one-to-many mapping, we learn

a one-to-many mapping first and then a many-to-one mapping, which allows for the effective modeling of different features corresponding to the various hypotheses to improve the representation ability.

3. Multi-Hypothesis Transformer

The overview of the proposed MHFormer is depicted in Figure 3 (a). Given a consecutive 2D pose sequence estimated by an off-the-shelf 2D pose detector from a video, our method aims to reconstruct the 3D pose of the center frame by making full use of spatial and temporal information in the multi-hypothesis feature hierarchies. To achieve our proposed three-stage framework, MHFormer is built upon (i) three major modules: Multi-Hypothesis Generation (MHG), Self-Hypothesis Refinement (SHR), and Cross-Hypothesis Interaction (CHI), and (ii) two auxiliary modules: temporal embedding and regression head.

3.1. Preliminary

In this work, we adopt a Transformer-based architecture since it performs well in long-range dependency modeling. We first give a brief description of the basic components in the Transformer [37], including a multi-head self-attention (MSA) and a multi-layer perceptron (MLP).

MSA. In the MSA, the inputs $x \in \mathbb{R}^{n \times d}$ are linearly mapped to queries $Q \in \mathbb{R}^{n \times d}$, keys $K \in \mathbb{R}^{n \times d}$, and values $V \in \mathbb{R}^{n \times d}$, where n is the sequence length, and d is the dimension. The scaled dot-product attention can be computed by:

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

MSA splits the queries, keys, and values for h times as well as performs the attention in parallel. Then, the outputs of h attention heads are concatenated.

MLP. The MLP consists of two linear layers, which are used for non-linearity and feature transformation:

$$\text{MLP}(x) = \sigma(xW_1 + b_1)W_2 + b_2, \quad (2)$$

where σ denotes the GELU activation function, $W_1 \in \mathbb{R}^{d \times d_m}$ and $W_2 \in \mathbb{R}^{d_m \times d}$ are the weights of the two linear layers respectively, and $b_1 \in \mathbb{R}^{d_m}$ and $b_2 \in \mathbb{R}^d$ are the bias terms.

3.2. Multi-Hypothesis Generation

In the spatial domain, we address the inverse problem by explicitly designing a cascaded Transformer-based architecture to generate multiple features in different depths of the latent space. To this end, MHG is introduced to model the human joint relations and initialize the multi-hypothesis representations (see Figure 3 (b)). Suppose there are M different hypotheses and L_1 layers in the MHG, it takes a sequence of 2D poses $X \in \mathbb{R}^{N \times J \times 2}$ with N video frames

and J body joints as input and outputs multiple hypotheses $[X_{L_1}^1, X_{L_1}^2, \dots, X_{L_1}^M]$, where $X_{L_1}^m \in \mathbb{R}^{(J \cdot 2) \times N}$ is the m -th hypothesis.

More specifically, we concatenate the (x, y) coordinates of joints for each frame to $\bar{X} \in \mathbb{R}^{(J \cdot 2) \times N}$, retain their spatial information of joints via a learnable spatial position embedding $E_{SPos}^m \in \mathbb{R}^{(J \cdot 2) \times N}$, and feed the embedded features into the encoders of the MHG. To encourage gradient propagation, a skip residual connection is applied between the original input and output features from the encoder. These procedures can be formulated as:

$$\begin{aligned} X_0^m &= \text{LN}(X^m) + E_{SPos}^m, \\ X_l^m &= X_{l-1}^m + \text{MSA}^m(\text{LN}(X_{l-1}^m)), \\ X_l^m &= X_l^m + \text{MLP}^m(\text{LN}(X_l^m)), \\ X_{L_1}^m &= X^m + \text{LN}(X_{L_1}^m), \end{aligned} \quad (3)$$

where $\text{LN}(\cdot)$ is the LayerNorm layer, $l \in [1, \dots, L_1]$ is the index of MHG layers, $X^1 = \bar{X}$, and $X^m = X_{L_1}^{m-1}$ ($m > 1$). The outputs of the MHG (*i.e.*, $X_{L_1}^m$) are multi-level features containing diverse semantic information. Therefore, those features can be regarded as initial representations of different pose hypotheses and need to be further enhanced.

3.3. Temporal Embedding

The MHG helps to generate initial multi-hypothesis features in the spatial domain, whereas the capabilities of such features are not strong enough. Considering this limitation, we propose to build relationships across hypothesis features and capture temporal dependencies in the temporal domain with two carefully designed modules: an SHR module followed by a CHI module (see Figure 3 (c) and (d)).

In order to exploit temporal information, we should first convert the spatial domain into the temporal domain. For this purpose, the encoded hypothesis features $X_{L_1}^m$ of each frame are embedded to the high-dimensional features $\tilde{Z}^m \in \mathbb{R}^{N \times C}$ using a transposition operation and a linear embedding, where C is the embedding dimension. Then, a learnable temporal position embedding $E_{TPos}^m \in \mathbb{R}^{N \times C}$ is utilized to retain positional information of frames, which can be formulated as: $\tilde{Z}_0^m = \tilde{Z}^m + E_{TPos}^m$.

3.4. Self-Hypothesis Refinement

In the temporal domain, we first construct the SHR to refine single-hypothesis features. Each SHR layer consists of a multi-hypothesis self-attention (MH-SA) block and a hypothesis-mixing MLP block.

MH-SA. The core of the Transformer model is MSA, through which any two elements can interact with each other, thus modeling long-range dependencies. Instead, our MH-SA aims to capture single-hypothesis dependencies within each hypothesis independently for self-hypothesis communication. Specifically, the embedded features $\tilde{Z}_0^m \in \mathbb{R}^{N \times C}$




Figure 4. **Left:** Multi-head self-attention (MSA). **Right:** Multi-head cross-attention (MCA).

of different hypotheses are fed into several parallel MSA blocks, which can be expressed as:

$$\tilde{Z}_l^m = \tilde{Z}_{l-1}^m + \text{MSA}^m(\text{LN}(\tilde{Z}_{l-1}^m)), \quad (4)$$

where $l \in [1, \dots, L_2]$ is the index of SHR layers. Therefore, the message of different hypothesis features can be passed in a self-hypothesis way for feature enhancement.

Hypothesis-Mixing MLP. The multiple hypotheses are processed independently in the MH-SA, but there is no information exchange across hypotheses. To handle this issue, we add a hypothesis-mixing MLP after the MH-SA. The features of multiple hypotheses are concatenated and fed into the hypothesis-mixing MLP to merge (*i.e.*, converge) themselves. Then, the converged features are evenly partitioned (*i.e.*, diverged) into non-overlapping chunks along the channel dimension, forming refined hypothesis representations. The procedure can be formulated as:

$$\begin{aligned} \tilde{Z}_l' &= \text{Concat}(\tilde{Z}_l^1, \dots, \tilde{Z}_l^M) \in \mathbb{R}^{N \times (C \cdot M)}, \\ \text{Concat}(\tilde{Z}_l^1, \dots, \tilde{Z}_l^M) &= \tilde{Z}_l' + \text{HM-MLP}(\text{LN}(\tilde{Z}_l')), \end{aligned} \quad (5)$$

where $\text{Concat}(\cdot)$ is the concatenation operation and $\text{HM-MLP}(\cdot)$ is the function of hypothesis-mixing MLP which shares the same format as Eq. (2). This process explores the relations among channels of different hypotheses.

3.5. Cross-Hypothesis Interaction

We then model interactions among multi-hypothesis features via the CHI, which contains two blocks: multi-hypothesis cross-attention (MH-CA) and hypothesis-mixing MLP.

MH-CA. The MH-SA lacks connections across hypotheses, which limits its interaction modeling. To capture multi-hypothesis correlations mutually for cross-hypothesis com-

munication, the MH-CA composed of multiple multi-head cross-attention (MCA) elements in parallel is proposed.

The MCA measures the correlation among cross-hypothesis features and has a similar structure to MSA. The common configuration of MCA uses the same input between keys and values [3, 26, 41]. However, an issue with this configuration is that it will result in more blocks (*e.g.*, 6 MCA blocks for 3 hypotheses). Here, we adopt a more efficient strategy, which reduces the number of parameters by using different inputs (only require 3 MCA blocks), as shown in Figure 4 (Right). The multiple hypotheses Z^m are alternately regarded as queries, keys, and values and fed into the MH-CA:

$$Z_l'^m = Z_{l-1}^m + \text{MCA}^m(\text{LN}(Z_{l-1}^{m_1}), \text{LN}(Z_{l-1}^{m_2}), \text{LN}(Z_{l-1}^m)), \quad (6)$$

where $l \in [1, \dots, L_3]$ is the index of CHI layers, $Z_0^m = \tilde{Z}_{L_2}^m$, m_1 and m_2 are the other two corresponding hypotheses, and $\text{MCA}(Q, K, V)$ denotes the function of the MCA. Thanks to the MH-CA, the message passing can be performed in a crossing way to significantly improve modeling power.

Hypothesis-Mixing MLP. The hypothesis-mixing MLP in the CHI serves as the same function as the process in Eq. (5). The outputs of the MH-CA are fed into it:

$$\begin{aligned} Z_l' &= \text{Concat}(Z_l^1, \dots, Z_l^M) \in \mathbb{R}^{N \times (C \cdot M)}, \\ \text{Concat}(Z_l^1, \dots, Z_l^M) &= Z_l' + \text{HM-MLP}(\text{LN}(Z_l')). \end{aligned} \quad (7)$$

In the hypothesis-mixing MLP of the last CHI layer, the partition operation is not used so that the features of all hypotheses are finally aggregated to synthesize a single hypothesis representation $Z_{L_3} \in \mathbb{R}^{N \times (C \cdot M)}$.

3.6. Regression Head

In the regression head, a linear transformation layer is applied on the output Z_{L_3} to perform regression to produce the 3D pose sequence $\tilde{X} \in \mathbb{R}^{N \times J \times 3}$. Finally, the 3D pose of the center frame $\hat{X} \in \mathbb{R}^{J \times 3}$ is selected from \tilde{X} as our final prediction.

3.7. Loss Function

The entire model is trained in an end-to-end manner with a Mean Squared Error (MSE) loss, which is applied to minimize the error between the estimated and ground truth poses:

$$\mathcal{L} = \sum_{n=1}^N \sum_{i=1}^J \left\| Y_i^n - \tilde{X}_i^n \right\|_2, \quad (8)$$

where \tilde{X}_i^n and Y_i^n represent the predicted and ground truth 3D poses of joint i at frame n , respectively.

Table 1. Quantitative comparison with the state-of-the-art methods on Human3.6M under Protocol 1, using detected 2D poses (top) and ground truth 2D poses (bottom) as inputs. (†) - uses temporal information. **Bold**: best; Underlined: second best.

Method	Dir.	Disc	Eat	Greet	Phone	Photo	Pose	Purch.	Sit	SitD.	Smoke	Wait	WalkD.	Walk	WalkT.	Avg.
Fang <i>et al.</i> (AAAI'18) [8]	50.1	54.3	57.0	57.1	66.6	73.3	53.4	55.7	72.8	88.6	60.3	57.7	62.7	47.5	50.6	60.4
GraphSH (CVPR'21) [42]	45.2	49.9	47.5	50.9	54.9	66.1	48.5	46.3	59.7	71.5	51.4	48.6	53.9	39.9	44.1	51.9
MGCN (ICCV'21) [47]	45.4	49.2	45.7	49.4	50.4	58.2	47.9	46.0	57.5	63.0	49.7	46.6	52.2	38.9	40.8	49.4
ST-GCN (ICCV'19) [2] (†)	44.6	47.4	45.6	48.8	50.8	59.0	47.2	43.9	57.9	61.9	49.7	46.6	51.3	37.1	39.4	48.8
VPose (CVPR'19) [34] (†)	45.2	46.7	43.3	45.6	48.1	55.1	44.6	44.3	57.3	65.8	47.1	44.0	49.0	32.8	33.9	46.8
SGNN (ICCV'21) [45] (†)	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	45.7
UGCN (ECCV'20) [38] (†)	<u>41.3</u>	43.9	44.0	<u>42.2</u>	48.0	57.1	42.2	43.2	57.3	61.3	47.0	43.5	47.0	32.6	<u>31.8</u>	45.6
Liu <i>et al.</i> (CVPR'20) [25] (†)	41.8	44.8	41.1	44.9	47.4	54.1	43.4	42.2	56.2	63.6	<u>45.3</u>	43.5	<u>45.3</u>	<u>31.3</u>	32.2	45.1
PoseFormer (ICCV'21) [46] (†)	41.5	44.8	39.8	42.5	<u>46.5</u>	<u>51.6</u>	42.1	<u>42.0</u>	53.3	60.7	45.5	43.3	46.1	31.8	32.2	44.3
Anatomy3D (TCSVT'21) [4] (†)	41.4	<u>43.2</u>	<u>40.1</u>	42.9	46.6	51.9	<u>41.7</u>	42.3	53.9	60.2	45.4	<u>41.7</u>	46.0	31.5	32.7	<u>44.1</u>
MHFormer (Ours) (†)	39.2	43.1	<u>40.1</u>	40.9	44.9	51.2	40.6	41.3	<u>53.5</u>	<u>60.3</u>	43.7	41.1	43.8	29.8	30.6	43.0
Method	Dir.	Disc	Eat	Greet	Phone	Photo	Pose	Purch.	Sit	SitD.	Smoke	Wait	WalkD.	Walk	WalkT.	Avg.
P-LSTM (ECCV'18) [15] (†)	32.1	36.6	34.3	37.8	44.5	49.9	40.9	36.2	44.1	45.6	35.3	35.9	30.3	37.6	35.5	38.4
PoseAug (CVPR'21) [9]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	38.2
VPose (CVPR'19) [34] (†)	35.2	40.2	32.7	35.7	38.2	45.5	40.6	36.1	48.8	47.3	37.8	39.7	38.7	27.8	29.5	37.8
Liu <i>et al.</i> (CVPR'20) [25] (†)	34.5	37.1	33.6	34.2	32.9	37.1	39.6	35.8	40.7	41.4	33.0	33.8	33.0	26.6	26.9	34.7
Anatomy3D (TCSVT'21) [4] (†)	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	32.3
SRNet (ECCV'20) [44] (†)	34.8	32.1	28.5	<u>30.7</u>	31.4	36.9	35.6	30.5	38.9	40.5	32.5	31.0	29.9	<u>22.5</u>	24.5	32.0
PoseFormer (ICCV'21) [46] (†)	<u>30.0</u>	<u>33.6</u>	29.9	31.0	<u>30.2</u>	33.3	<u>34.8</u>	31.4	<u>37.8</u>	38.6	<u>31.7</u>	<u>31.5</u>	29.0	23.3	<u>23.1</u>	<u>31.3</u>
MHFormer (Ours) (†)	27.7	32.1	<u>29.1</u>	28.9	30.0	<u>33.9</u>	33.0	<u>31.2</u>	37.0	<u>39.3</u>	30.0	31.0	<u>29.4</u>	22.2	23.0	30.5

Table 2. Comparison with the methods of generating multiple 3D pose hypotheses on Human3.6M. The number of hypotheses is denoted as M . **Bold**: best; Underlined: second best.

Method	M	MPJPE (mm)
Li <i>et al.</i> (CVPR'19) [16]	5	52.7
Sharma <i>et al.</i> (ICCV'19) [35]	200	46.8
Oikarinen (IJCNN'21) [32]	200	46.2
Wehrbein <i>et al.</i> (ICCV'21) [40]	200	<u>44.3</u>
MHFormer (Ours)	3	43.0

4. Experiments

4.1. Datasets and Evaluation Metrics

We evaluate our method on two widely-used datasets for 3D HPE: Human3.6M [12] and MPI-INF-3DHP [30].

Human3.6M. The Human3.6M dataset [12] is the largest and most representative benchmark for 3D HPE. This dataset consists of 3.6 million images captured from four synchronized cameras at 50 Hz. There are 15 daily activities performed by 11 human subjects in an indoor environment. Following previous works [4, 25, 34, 38], we train a single model on five subjects (S1, S5, S6, S7, S8) and test it on two subjects (S9 and S11). We adopt the most commonly used evaluation protocols: Protocol 1 is the MPJPE which measures the mean Euclidean distance between the ground truth and estimated joints in millimeters; Protocol 2 is the MPJPE after aligning the predicted 3D pose with the ground truth using translation, rotation, and scale (P-MPJPE).

MPI-INF-3DHP. The MPI-INF-3DHP [30] is a large 3D pose dataset in both indoor and outdoor environments. This dataset provides 1.3 million frames, containing more diverse motions than Human3.6M. Following the setting in [4, 21, 30, 46], we report metrics of MPJPE, Percentage of Correct

Keypoint (PCK) with the threshold of 150 mm , and Area Under Curve (AUC) for a range of PCK thresholds.

4.2. Implementation Details

In our implementation, the proposed MHFormer contains $L_1=4$ MHG, $L_2=2$ SHR, and $L_3=1$ CHI layers. The MHFormer model is implemented in PyTorch framework on one GeForce RTX 3090 GPU. We train our model in an end-to-end manner from scratch using Amsgrad optimizer. The initial learning rate is set to 0.001 with a shrink factor of 0.95 applied after each epoch and 0.5 after every 5 epochs. For a fair comparison, the same horizontal flip augmentation is adopted following [2, 4, 34, 46]. We perform the 2D pose detection using cascaded pyramid network (CPN) [5] for Human3.6M following [2, 25, 34] and ground truth 2D pose for MPI-INF-3DHP following [4, 21, 46].

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

Results on Human3.6M. The proposed MHFormer is compared with the state-of-the-art methods on Human3.6M. The results of our model with a receptive field of 351 frames using 2D detected inputs [5] are reported in Table 1 (top). Without bells and whistles, our MHFormer outperforms all previous state-of-the-art methods by a large margin under both Protocol 1 (43.0 mm) and Protocol 2 (34.4 mm , see supplemental material). Compared to the very recent Transformer-based method, *i.e.*, PoseFormer [46], MHFormer noticeably surpasses it by 1.3 mm in MPJPE (relative 3% improvement). Figure 5 shows the qualitative comparison with the PoseFormer and the baseline model (same architecture as ViT [6]) on some challenging poses. To further explore the lower bound of our method, we compared our MHFormer with the state-of-the-art methods with ground truth 2D poses

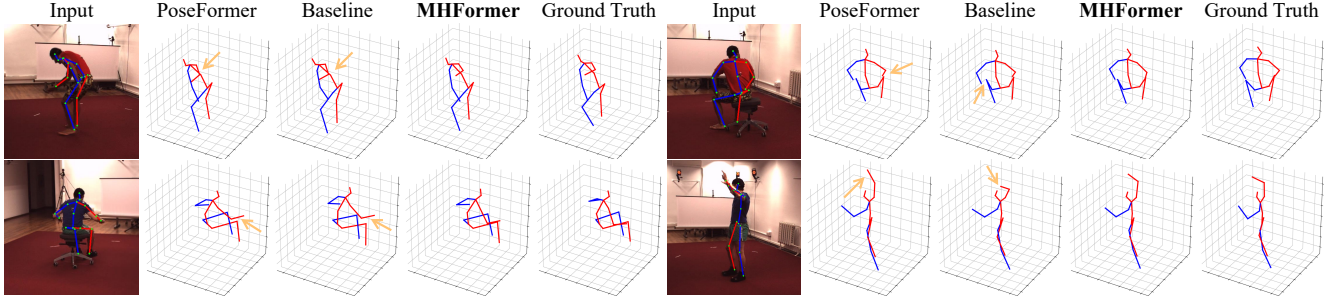


Figure 5. Qualitative comparison among the proposed method (MHFormer), the baseline method, and the previous state-of-the-art method (PoseFormer) [46] on Human3.6M dataset. Wrong estimations are highlighted by yellow arrows.

Table 3. Quantitative comparison with the state-of-the-art methods on MPI-INF-3DHP. Best in bold, second best underlined.

Method	PCK \uparrow	AUC \uparrow	MPJPE \downarrow
Mehta <i>et al.</i> (3DV'17) [30]	75.7	39.3	117.6
Lin <i>et al.</i> (BMVC'19) [21]	83.6	51.4	79.8
VPose (CVPR'19) [34]	86.0	51.9	84.0
Li <i>et al.</i> (CVPR'20) [19]	81.2	46.1	99.7
Anatomy3D (TCSVT'21) [4]	87.9	54.0	78.8
PoseFormer (ICCV'21) [46]	<u>88.6</u>	<u>56.4</u>	<u>77.1</u>
MHFormer (Ours)	93.8	63.3	58.0

Table 4. Ablation study on different receptive fields with MPJPE (mm). CPN - cascaded pyramid network; GT - 2D ground truth.

	9	27	81	243	351
CPN	47.8	45.9	44.5	43.2	43.0
GT	36.6	34.3	32.7	30.9	30.5

as inputs. The results are shown in Table 1 (bottom). It can be seen that our method achieves the best performance (30.5 mm in MPJPE), outperforming all other methods.

Additionally, our method is compared with previous methods of generating multiple 3D pose hypotheses. The results are shown in Table 2. It is noteworthy that these methods report metrics for the best hypothesis due to the adopted one-to-many mapping, while our method reports metrics with a specific solution by learning a deterministic mapping, which is much more practical in reality. Even though we use much fewer hypothesis numbers (3 vs. 200), our proposed method consistently outperforms previous works.

Results on MPI-INF-3DHP. To assess the generalization ability, we evaluate our method on MPI-INF-3DHP dataset. Following [46], we use 2D pose sequences of 9 frames as our model input due to the fewer samples and shorter sequence lengths of this dataset compared to Human3.6M. The results in Table 3 demonstrate that our method achieves the best performance on all metrics (PCK, AUC, and MPJPE). It emphasizes the effectiveness of our MHFormer in improving performance in outdoor scenes.

4.4. Ablation Study

To verify the impact of each component and design in the proposed model, we conduct extensive ablation experiments

Table 5. Ablation study on different parameters of MHG. Here, L_1 is the number of MHG layers and M is the hypothesis number.

M	L_1	Params (M)	FLOPs (G)	MPJPE (mm)
3	2	18.91	1.03	46.4
3	3	18.92	1.03	46.3
3	4	18.92	1.03	45.9
3	5	18.93	1.04	46.1
1	4	6.32	0.34	47.6
2	4	12.61	0.69	46.7
3	4	18.92	1.03	45.9
4	4	25.22	1.38	46.9

on Human3.6M dataset under Protocol 1 with MPJPE.

Impact of Receptive Fields. For the video-based 3D HPE task, a large receptive field is essential for estimation accuracy. Table 4 shows the results of our method with different input frames. It can be seen that our method obtains larger gains with more frames fed into the model. The error has a great decrease of 16.7% from 9-frames to 351-frames with GT input, which indicates the effectiveness of our method in capturing long-range dependencies across frames with a large receptive field. Next, ablations in the following parts are carried out using a receptive field of 27 frames to balance the computation efficiency and performance.

Impact of Parameters in MHG. In the top part of Table 5, we report the results with different numbers of MHG layers. Experiments show that stacking more layers in MHG can slightly improve the performance with few parameter increases, but the gain disappears when the layer number is larger than 4. Moreover, we investigate the influence of using different numbers of hypotheses in MHG. The results are shown in the bottom part of Table 5. Increasing the number of hypotheses can improve the result, but the performance saturates when using 3 hypothesis representations. Notably, our model equipped with 3 hypotheses shows significant gains over the single-hypothesis model with 1.7 mm error reduction. This demonstrates that exploiting different representations of multiple pose hypotheses is helpful to improve the performance of the model, validating our motivation.

Impact of Parameters in SHR and CHI. Table 6 reports how the different parameters of SHR and CHI impact the performance and computation complexity of our model. The




Figure 6. Diverse 3D pose hypotheses generated by MHFormer. For easy illustration, we color-code the hypotheses to show the difference among them. Green colored 3D pose corresponds to the final synthesized estimation of our method.

Table 6. Ablation study on different parameters of SHR and CHI. Here, L_2 and L_3 indicate the number of SHR and CHI layers, respectively. C is the embedding dimension.

L_2	L_3	C	Params (M)	FLOPs (G)	MPJPE (mm)
2	1	256	4.72	0.26	47.2
2	1	384	10.65	0.58	46.4
2	1	512	18.92	1.03	45.9
2	1	768	42.50	2.31	47.4
2	1	512	18.92	1.03	45.9
1	3	512	25.20	1.38	46.7
2	2	512	25.20	1.38	46.8
3	1	512	25.20	1.38	46.3

results show that enlarging the embedding dimension from 256 to 512 can boost the performance, but using dimensions larger than 512 cannot bring further improvements. In addition, we observe no more gains by stacking either more SHR or CHI layers. Therefore, the optimal parameters for our model are $L_2=2$, $L_3=1$, and $C=512$.

Effect of Model Components. In Table 7, we carry out experiments to quantify the influence of our proposed components. Firstly, we compare our method with the baseline model. For a fair comparison, the results of the baseline are reported at the same number of layers as MHFormer with different embedding dimensions, since our hypothesis-mixing MLP in MHFormer takes concatenated hypothesis features as inputs (the dimension is $512 \times 3 = 1536$). The results show that the baseline model is prone to overfitting due to the excessive number of parameters, whereas our method performs well. Additionally, it can be seen that our MHFormer built upon MHG, SHR, and CHI outperforms varying variants of baseline models (1.9 mm improvement). Then, when we incorporate multi-hypothesis representations and SHR or CHI within the baseline, the performance has significant gains (-1.3 mm for MHG-SHR and -1.0 mm for MHG-CHI). Besides, we remove the MHG in MHFormer (SHR-CHI). At this point, the model only captures temporal information and its error heavily increases by 1.3 mm. These ablations indicate that learning multi-hypothesis spatio-temporal representations is significant for 3D HPE, and the different hypothesis representations should be modeled in both independent and mutual ways.

We also explore the use of multi-level features in MHG by simply building the MHG upon several parallel Transformer encoders (MHFormer*). As shown in the table, our

Table 7. Ablation study on different components of our MHFormer. Here, * means no multi-level features in MHG.

Method	MHG	SHR	CHI	MPJPE (mm)
Baseline ($C=256$)	✗	✗	✗	49.9
Baseline ($C=512$)	✗	✗	✗	47.8
Baseline ($C=1536$)	✗	✗	✗	50.4
SHR-CHI	✗	✓	✓	47.2
MHG-SHR	✓	✓	✗	46.5
MHG-CHI	✓	✗	✓	46.8
MHFormer*	✓	✓	✓	46.5
MHFormer (Ours)	✓	✓	✓	45.9

MHFormer equipped with the multi-level features increases the performance, which indicates that the multi-level features can bring valuable information to the final estimation.

5. Qualitative Results

Although our method does not aim to produce multiple 3D pose predictions, for better observation, we add additional regression layers and finetune our model to visualize the intermediate hypotheses. The several qualitative results are shown in Figure 6. It can be seen that our method is able to generate different plausible 3D pose solutions, especially for ambiguous body parts with depth ambiguity, self-occlusion, and 2D detector uncertainty. Moreover, the final 3D pose synthesized by aggregating multi-hypothesis information is more reasonable and accurate.

6. Conclusion

This paper presents Multi-Hypothesis Transformer (MHFormer), a new Transformer-based three-stage framework for the ambiguous inverse problem of 3D HPE from monocular videos. MHFormer first generates initial representations of multiple pose hypotheses in the spatial domain and then communicates across them in both independent and mutual ways in the temporal domain. Extensive experiments show that the proposed MHFormer has a fundamental advantage over single-hypothesis Transformers and achieves state-of-the-art performance on two benchmark datasets. We hope that our approach will foster further research in 2D-to-3D pose lifting considering various ambiguities.

Limitation. One limitation of our method is the relatively larger computational complexity. The excellent performance of Transformers comes at a price of high computational cost.

References

- [1] Christopher M. Bishop. Mixture density networks. Technical report, Aston University, 1994. [1](#)
- [2] Yujun Cai, Liuhao Ge, Jun Liu, Jianfei Cai, Tat-Jen Cham, Junsong Yuan, and Nadia Magnenat Thalmann. Exploiting spatial-temporal relationships for 3D pose estimation via graph convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 2272–2281, 2019. [1](#), [3](#), [6](#)
- [3] Chun-Fu Richard Chen, Quanfu Fan, and Rameswar Panda. CrossViT: Cross-attention multi-scale vision transformer for image classification. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 357–366, 2021. [5](#)
- [4] Tianlang Chen, Chen Fang, Xiaohui Shen, Yiheng Zhu, Zhili Chen, and Jiebo Luo. Anatomy-aware 3D human pose estimation with bone-based pose decomposition. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(1):198–209, 2021. [1](#), [3](#), [6](#), [7](#)
- [5] Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, and Jian Sun. Cascaded pyramid network for multi-person pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7103–7112, 2018. [6](#)
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021. [3](#), [6](#)
- [7] Andrew Errity. Human–computer interaction. In *An Introduction to Cyberpsychology*, pages 263–278. 2016. [1](#)
- [8] Hao-Shu Fang, Yuanlu Xu, Wenguan Wang, Xiaobai Liu, and Song-Chun Zhu. Learning pose grammar to encode human body configuration for 3D pose estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. [6](#)
- [9] Kehong Gong, Jianfeng Zhang, and Jiashi Feng. PoseAug: A differentiable pose augmentation framework for 3D human pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8575–8584, 2021. [1](#), [6](#)
- [10] Shuting He, Hao Luo, Pichao Wang, Fan Wang, Hao Li, and Wei Jiang. TransReID: Transformer-based object Re-Identification. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 15013–15022, 2021. [3](#)
- [11] Wenbo Hu, Changgong Zhang, Fangneng Zhan, Lei Zhang, and Tien-Tsin Wong. Conditional directed graph convolution for 3D human pose estimation. In *Proceedings of the ACM International Conference on Multimedia (ACMMM)*, pages 602–611, 2021. [1](#), [3](#)
- [12] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6M: Large scale datasets and predictive methods for 3D human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, 2013. [2](#), [6](#)
- [13] Ehsan Jahangiri and Alan L Yuille. Generating multiple diverse hypotheses for human 3D pose consistent with 2d joint detections. In *Proceedings of the IEEE International Conference on Computer Vision Workshops (ICCVW)*, pages 805–814, 2017. [3](#)
- [14] Rawal Khirodkar, Visesh Chari, Amit Agrawal, and Amrith Tyagi. Multi-instance pose networks: Rethinking top-down pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 3122–3131, 2021. [1](#)
- [15] Kyoungoh Lee, Inwoong Lee, and Sanghoon Lee. Propagating LSTM: 3D pose estimation based on joint interdependency. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 119–135, 2018. [6](#)
- [16] Chen Li and Gim Hee Lee. Generating multiple hypotheses for 3D human pose estimation with mixture density network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9887–9895, 2019. [1](#), [6](#)
- [17] Chen Li and Gim Hee Lee. Weakly supervised generative network for multiple 3D human pose hypotheses. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2020. [1](#), [3](#)
- [18] Sijin Li and Antoni B Chan. 3D human pose estimation from monocular images with deep convolutional neural network. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, pages 332–347, 2014. [3](#)
- [19] Shichao Li, Lei Ke, Kevin Pratama, Yu-Wing Tai, Chi-Keung Tang, and Kwang-Ting Cheng. Cascaded deep monocular 3D human pose estimation with evolutionary training data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6173–6183, 2020. [7](#)
- [20] Wenhao Li, Hong Liu, Runwei Ding, Mengyuan Liu, Pichao Wang, and Wenming Yang. Exploiting temporal contexts with strided transformer for 3D human pose estimation. *IEEE Transactions on Multimedia*, 2022. [3](#)
- [21] Jiahao Lin and Gim Hee Lee. Trajectory space factorization for deep video-based 3D human pose estimation. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2019. [6](#), [7](#)
- [22] Kevin Lin, Lijuan Wang, and Zicheng Liu. End-to-end human pose and mesh reconstruction with transformers. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1954–1963, 2021. [3](#)
- [23] Mengyuan Liu, Hong Liu, and Chen Chen. Enhanced skeleton visualization for view invariant human action recognition. *Pattern Recognition*, 68:346–362, 2017. [1](#)
- [24] Mengyuan Liu and Junsong Yuan. Recognizing human actions as the evolution of pose estimation maps. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1159–1168, 2018. [1](#)
- [25] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5064–5073, 2020. [3](#), [6](#)

- [26] Xuehu Liu, Pingping Zhang, Chenyang Yu, Huchuan Lu, Xuesheng Qian, and Xiaoyun Yang. A video is worth three views: Trigeminal transformers for video-based person Re-Identification. *arXiv preprint arXiv:2104.01745*, 2021. 5
- [27] 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 *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 10012–10022, 2021. 3
- [28] Xiaoxuan Ma, Jiajun Su, Chunyu Wang, Hai Ci, and Yizhou Wang. Context modeling in 3D human pose estimation: A unified perspective. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6238–6247, 2021. 3
- [29] Julieta Martinez, Rayat Hossain, Javier Romero, and James J Little. A simple yet effective baseline for 3D human pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 2640–2649, 2017. 1, 3
- [30] Dushyant Mehta, Helge Rhodin, Dan Casas, Pascal Fua, Oleksandr Sotnychenko, Weipeng Xu, and Christian Theobalt. Monocular 3D human pose estimation in the wild using improved cnn supervision. In *Proceedings of the International Conference on 3D Vision (3DV)*, pages 506–516, 2017. 6, 7
- [31] Dushyant Mehta, Srinath Sridhar, Oleksandr Sotnychenko, Helge Rhodin, Mohammad Shafiei, Hans-Peter Seidel, Weipeng Xu, Dan Casas, and Christian Theobalt. VNet: Real-time 3D human pose estimation with a single rgb camera. *ACM Transactions on Graphics*, 36(4):1–14, 2017. 1
- [32] Tuomas Oikarinen, Daniel Hannah, and Sohrob Kazerounian. GraphMDN: Leveraging graph structure and deep learning to solve inverse problems. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, pages 1–9, 2021. 3, 6
- [33] Georgios Pavlakos, Xiaowei Zhou, Konstantinos G Derpanis, and Kostas Daniilidis. Coarse-to-fine volumetric prediction for single-image 3D human pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7025–7034, 2017. 3
- [34] Dario Pavlo, Christoph Feichtenhofer, David Grangier, and Michael Auli. 3D human pose estimation in video with temporal convolutions and semi-supervised training. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7753–7762, 2019. 1, 3, 6, 7
- [35] Saurabh Sharma, Pavan Teja Varigonda, Prashast Bindal, Abhishek Sharma, and Arjun Jain. Monocular 3D human pose estimation by generation and ordinal ranking. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 2325–2334, 2019. 1, 6
- [36] Xiao Sun, Bin Xiao, Fangyin Wei, Shuang Liang, and Yichen Wei. Integral human pose regression. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 529–545, 2018. 3
- [37] 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)*, pages 5998–6008, 2017. 3, 4
- [38] Jingbo Wang, Sijie Yan, Yuanjun Xiong, and Dahua Lin. Motion guided 3D pose estimation from videos. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 764–780, 2020. 1, 3, 6
- [39] Pichao Wang, Wanqing Li, Zhimin Gao, Chang Tang, and Philip O Ogunbona. Depth pooling based large-scale 3-d action recognition with convolutional neural networks. *IEEE Transactions on Multimedia*, 20(5):1051–1061, 2018. 1
- [40] Tom Wehrbein, Marco Rudolph, Bodo Rosenhahn, and Bastian Wandt. Probabilistic monocular 3D human pose estimation with normalizing flows. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 11199–11208, 2021. 1, 3, 6
- [41] Tongkun Xu, Weihua Chen, Pichao Wang, Fan Wang, Hao Li, and Rong Jin. CDTrans: Cross-domain transformer for unsupervised domain adaptation. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021. 5
- [42] Tianhan Xu and Wataru Takano. Graph stacked hourglass networks for 3D human pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16105–16114, 2021. 6
- [43] Guanglei Yang, Hao Tang, Mingli Ding, Nicu Sebe, and Elisa Ricci. Transformer-based attention networks for continuous pixel-wise prediction. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 16269–16279, 2021. 3
- [44] Ailing Zeng, Xiao Sun, Fuyang Huang, Minhao Liu, Qiang Xu, and Stephen Lin. SRNet: Improving generalization in 3D human pose estimation with a split-and-recombine approach. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 507–523, 2020. 6
- [45] Ailing Zeng, Xiao Sun, Lei Yang, Nanxuan Zhao, Minhao Liu, and Qiang Xu. Learning skeletal graph neural networks for hard 3D pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 11436–11445, 2021. 6
- [46] Ce Zheng, Sijie Zhu, Matias Mendieta, Taojiannan Yang, Chen Chen, and Zhengming Ding. 3D human pose estimation with spatial and temporal transformers. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 1, 2, 3, 6, 7
- [47] Zhiming Zou and Wei Tang. Modulated graph convolutional network for 3D human pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 11477–11487, 2021. 6