

GLA-GCN: Global-local Adaptive Graph Convolutional Network for 3D Human Pose Estimation from Monocular Video

Bruce X.B. Yu¹ Zhi Zhang¹ Yongxu Liu¹
Sheng-hua Zhong² Yan Liu¹ Chang Wen Chen^{1‡}

¹The Hong Kong Polytechnic University ²Shenzhen University

{bruce.xb.yu, zhi271.zhang, yongxu.liu}@connect.polyu.hk csszhong@szu.edu.cn
{yan.liu, changwen.chen}@polyu.edu.hk

Abstract

3D human pose estimation has been researched for decades with promising fruits. 3D human pose lifting is one of the promising research directions toward the task where both estimated pose and ground truth pose data are used for training. Existing pose lifting works mainly focus on improving the performance of estimated pose, but they usually underperform when testing on the ground truth pose data. We observe that the performance of the estimated pose can be easily improved by preparing good quality 2D pose, such as fine-tuning the 2D pose or using advanced 2D pose detectors. As such, we concentrate on improving the 3D human pose lifting via ground truth data for the future improvement of more quality estimated pose data. Towards this goal, a simple yet effective model called Global-local Adaptive Graph Convolutional Network (GLA-GCN) is proposed in this work. Our GLA-GCN globally models the spatiotemporal structure via a graph representation and backtraces local joint features for 3D human pose estimation via individually connected layers. To validate our model design, we conduct extensive experiments on three benchmark datasets: Human3.6M, HumanEva-I, and MPI-INF-3DHP. Experimental results show that our GLA-GCN¹ implemented with ground truth 2D poses significantly outperforms state-of-the-art methods (e.g., up to 3%, 17%, and 14% error reductions on Human3.6M, HumanEva-I, and MPI-INF-3DHP, respectively).

1. Introduction

3D Human Pose Estimation (HPE) in videos aims to predict the pose joint locations of the human body in 3D space, which can facilitate plenty of applications such as video surveillance, human-robot interaction, and physio-

therapy [53]. 3D human poses can be directly retrieved from advanced motion sensors such as motion capture systems, depth sensors, or stereotyped cameras [51, 73]. The 3D HPE task can be performed under either multi-view or monocular view settings. Although state-of-the-art multi-view methods [30, 52, 78, 26] generally show superior performance than monocular ones [28, 77], ordinary RGB monocular cameras are much cheaper than these off-the-shelf motion sensors and more widely applied in real-world surveillance scenarios. Hence, 3D HPE from a monocular video is an important and challenging task, which has been attracting increasing research interest. Recent monocular view works can be grouped into model-based and model-free methods [13]. Model-based methods [15, 23] incorporate parametric body models such as kinematic [67], planar [60], and volumetric models [2] for 3D HPE. Model-free methods can be further grouped into single-stage and 2D to 3D lifting methods. Single-stage methods estimate the 3D pose directly from images in an end-to-end manner [36, 46, 12, 65, 43, 81]. 2D to 3D lifting methods have an intermediate 2D pose estimation layer [44, 50, 41, 61]. Among these methods, 2D to 3D lifting methods implemented with ground truth 2D poses achieved better performance.

The advantages of 2D to 3D lifting methods can be summarized as two main points: allowing make use of advances in 2D human pose detection and exploiting temporal information along multiple 2D pose frames [50, 31]. For the 2D human pose detection, it has achieved remarkable progress via detectors such as Mask R-CNN (MRCNN) [24], Cascaded Pyramid Network (CPN) [14], Stacked Hourglass (SH) detector [47], and HR-Net [58]. The intermediate 2D pose estimation stage via these 2D pose detectors significantly reduces the data volume and complexity of the 3D HPE task. For the temporal information, existing mainstream methods [50, 41, 61, 28, 37, 77] gained noticeable improvements by feeding a long sequence of 2D pose

¹Code is available: <https://github.com/bruceyo/GLA-GCN>

[‡]Corresponding author.

frames to their models, among which [77] achieved the state-of-the-art performance via ground truth 2D poses. Recent methods [77, 84] simply fine-tuned these 2D pose detectors on the target datasets and achieved great improvements in the performance of estimated 2D pose data but remain far behind the results of using ground truth 2D pose, which motivates us to concentrate on improving the 3D HPE via ground truth 2D pose data for potential improvements via future more quality estimated 2D pose data.

Given the promising performance and advantages of 2D to 3D lifting methods, our work contributes to the literature in this direction. For 2D to 3D lifting approaches, since [44] proposed Fully Connected Network (FCN), recent advanced models have three main groups: Temporal Convolutional Network (TCN)-based [50, 41], Graph Convolutional Network (GCN)-based [79, 61, 28], and Transformer-based ones [38, 37, 77]. On the one hand, we observe that existing TCN- and Transformer-based methods can receive large receptive fields (i.e., a long 2D pose sequence) with strided convolutions. However, it can be difficult to make further intuitive designs to backtrace local joint features based on the pose structure, since the 2D pose sequence is flattened and fed to the model. Meanwhile, the estimation of different pose joints relies on the same fully connected layer, which lacks considering the independent characteristic of different pose joints. On the other hand, GCN-based models can explicitly reserve the structure of 2D and 3D human pose during convolutional propagation. However, this advantage of GCN remains under-explored. Existing GCN-based methods [79, 61] also utilized a fully connected layer for the estimation of different 3D pose joints, which does not consider the structural features of GCN representations.

To this end, we propose Global-local Adaptive GCN (GLA-GCN) for 2D to 3D human pose lifting. Our GLA-GCN contains two modules: global representation and local 3D pose estimation. In the global representation, we use an adaptive Graph Convolutional Network (GCN) to reconstruct the global representation of an intermediate 3D human pose sequence from its corresponding 2D sequence. For the local 3D pose joint estimation, we temporally shrink the global representation optimized by the reconstructed 3D pose sequence with a strided design. Then, an individual connected layer is proposed to locally estimate the 3D human pose joints from the shrunken global representation. Our contributions can be threefold as follows:

- We propose a global-local learning architecture that leverages the global spatiotemporal representation and local joint representation in the GCN-based model for 3D human pose estimation.
- We are the first to introduce an individually connected layer that has two components to divide joint nodes and input the joint node representation for 3D pose joint estimation instead of based on pooled features.

- Our GLA-GCN model performs better than corresponding state-of-the-art methods [37, 77] with considerable margins e.g., up to 3% and 17% error reductions on Human3.6M [29] and HumanEva-I [57], respectively.

2. Related Work

2D to 3D Lifting. 3D HPE is a traditional vision problem that has been studied for decades [18, 75, 8, 63, 22, 32, 64, 66, 62, 42]. Existing works of 3D HPE from a monocular view usually target two main scenarios: single person and multi-person [13]. This work aims to improve the performance of single-person 3D HPE. [33, 1, 70] represent early efforts that attempt to infer 3D position from 2D projections. They usually rely on manually chosen parameters based on assumptions about pose joint mobility. Methods [23, 76] estimating 3D pose from fewer frames or even a single frame has shown great progress but can be a lack of considering temporal information. Recent advances in 2D human pose estimation [47, 24, 14] enable 2D to 3D lifting approaches to achieve remarkable performance over other counterparts. Inspired by [44], there have been more well-designed learning architectures being proposed to improve the performance, in particular, by utilizing temporal information. These methods are also known as 2D to 3D lifting, which can be grouped into three directions: TCN-, GCN-, and Transformer-based architectures [50, 41, 79, 61, 28, 37, 77].

TCN-based methods [50, 41] successfully push the performance of 2D to 3D lifting forward with a strided design for their learning architectures built upon 1D CNN layers. The strided design is on the temporal dimension of the input, which allows the features to shrink from a 2D pose sequence to a feature embedding for the 3D pose estimation via a final fully connected layer. The number of channels for the fully connected layer is conventionally set to 1024, which is shared to predict the 3D positions of all pose joints. While varied numbers of input 2D pose frames have been extensively investigated, which shows input 2D pose frames with reasonable length can benefit the 3D pose reconstruction. The strided design can effectively reduce the feature size by shrinking the number of temporal frames along the propagation of several TCN blocks. Using this strided structure, Transformer-based methods [37, 77] show promising performance, especially [77] that takes advantage of weighted and temporal loss functions and helps it outperform the GCN-based methods optimized with an additional motion loss [61, 28]. The motion loss was shown not very effective in [77]. These observations compel us to explore effective models in the direction of GCN-based models with inspiring designs in mind but without relying on various novel loss functions.

Graph Convolutional Network. A popular method representing the pose data with GCN is Spatial Temporal GCN

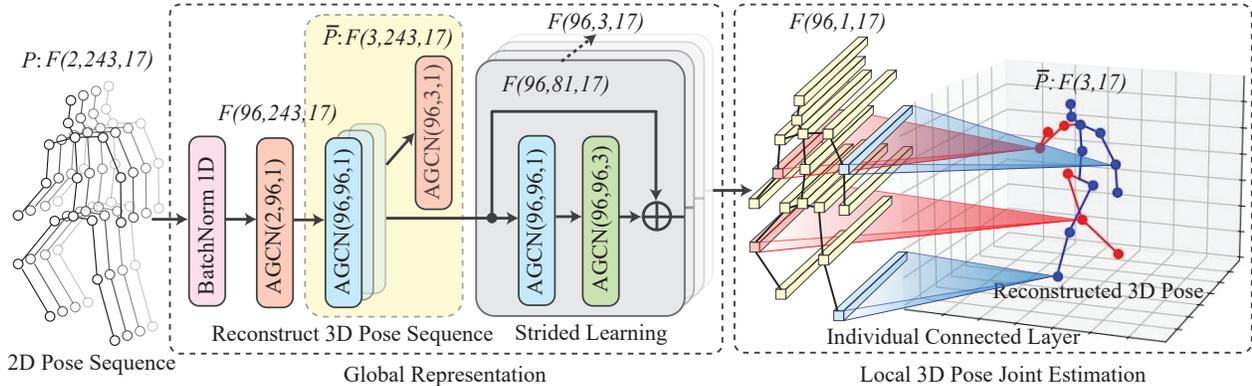


Figure 1. Learning architecture of our GLA-GCN. $AGCN(C_{in}, C_{out}, S)$ represents AGCN blocks with the specific values of the input channel, output channel, and stride length. $F(C', T', N')$ represents the size of a feature map. The individual connected layer shows the prediction process of four pose joint examples that use separate 1D CNN layers.

(ST-GCN) [69], which is originally proposed to model large receptive fields for skeleton-based action recognition. Following ST-GCN, advanced GCN models have been proposed to advance 3D HPE [17, 79, 16, 61].

Regarding GCN-based models for 3D HPE, Ci et al. [17] proposed Locally Connected Network (LCN) that takes advantage of FCN [44] and GCN [19]. LCN has a similar design for the convolutional filters to ST-GCN [69], which defines a neighbor set for a node based on the distance to perform the convolutional operation. Zhao et al. [79] proposed an architecture called SemGCN that stacks GCN layers by flattening output to a fully connected layer. The optimization of SemGCN is based on both joint positions and bone vectors. Choi et al. [16] also proposed to use GCN to recover 3D human pose and mesh from a 2D human pose. Liu et al. [40] investigated how weight-sharing schemes in GCNs affect the pose lifting task, which shows the pre-aggregation method leads to relatively better performance. The architecture in [40] is similar with that of SemGCN. The above-mentioned GCN-based methods achieved good performance via a single pose frame input but they did not take advantage of temporal information in a 2D pose sequence.

Taking multiple 2D pose frames as input, U-shaped Graph Convolution Networks (UGCN) [61, 28] further improves the performance of GCN-based methods by paying attention to the temporal characteristics of a pose motion. Specifically, UGCN utilizes spatial-temporal GCN [69] to predict a 3D pose sequence from a 2D pose sequence for the reconstruction of a single 3D pose frame. A motion loss term that regulates the temporal trajectory of pose joints based on the prediction of a 3D pose sequence and its corresponding ground truth 3D pose sequence. Despite the improvements gained with novel loss terms in works such as SemGCN and UGCN, we aim to contribute to the literature

of 2D-3D lifting by using the consistent loss term used in [50, 41]. In our model design, we propose to incorporate the strided convolutions to a GCN-based model that represents global information of a 2D pose sequence. Based on the structure of GCN representation, we explicitly utilize the structured features of different pose joints to locally predict their corresponding 3D pose locations.

3. Method

Given the temporal information of a 2D human pose sequence estimated from a video $P = \{p_{t,i} \in \mathbb{R}^2 | t = 1, \dots, T; i = 1, \dots, N\}$, where T is the number of pose frames and N is the number of pose joints, we aim to utilize this 2D pose sequence P to reconstruct 3D coordinates of pose joints $\bar{P} = \{\bar{p}_i \in \mathbb{R}^3 | i = 1, \dots, N\}$. Figure 1 shows the learning architecture of our GLA-GCN, which uses AGCN layers to globally represent the 2D pose sequence and locally estimate the 3D pose via an individual connected layer. In the following of this section, we introduce the detailed design of our GLA-GCN.

3.1. Global Representation

Adaptive Graph Convolutional Network. An AGCN block [35, 56] is based on the GCN with an adaptive design that improves the flexibility of a typical ST-GCN block [69]. Let us represent the 2D pose sequence P as a spatial-temporal graph $\mathcal{G} = \{v_t, \varepsilon_t | t = 1, \dots, T\}$, where $v_t = \{v_{t,i} | i = 1, \dots, N\}$ represents the pose joints and ε_t represents the corresponding pose bones. To implement a basic ST-GCN block, a neighbor set \mathcal{B}_i is first defined to indicate the spatial graph convolutional filter for a specific pose joint $v_{t,i}$. Specifically, for the graph convolutional filter of a vertex node, we apply three distance neighbor subsets: the vertex itself, the centripetal subset, and the centrifugal subset. The definitions of centripetal and centrifugal subsets

are based on the pose frame’s gravity center (i.e., the average coordinate of all pose joints). Centripetal and centrifugal subsets represent nodes that are closer and farther to the average distance from the gravity center, respectively. Empirically, similar with 2D convolution, we set the kernel size K to 3, which will lead to 3 subsets in \mathcal{B}_i . To implement the subsets, a mapping $h_{t,i} \rightarrow \{0, \dots, K - 1\}$ is used to index each subset with a numeric label, where centripetal and centrifugal subsets are respectively labeled as 1 and 2. Subsets that have the average distance to the gravity center are indexed to 0. This graph convolutional operation can be written as

$$f_{out}(v_{t,i}) = \sum_{v_{t,j} \in \mathcal{B}_i} \frac{1}{Z_{t,j}} f_{in}(v_{t,j}) W(h_{t,i}(v_{t,j})) \quad (1)$$

where $f_{in} : v_{t,j} \rightarrow \mathbb{R}^2$ is a mapping that gets the attribute features of joint node $v_{t,j}$ and $Z_{t,j}$ is a normalization term that equals to the subset’s cardinality. $W(h_{t,i}(v_{t,j}))$ is a weight function $W(v_{t,i}, v_{t,j}) : \mathcal{B}_i \rightarrow \mathbb{R}^2$ implemented by indexing a $(2, K)$ tensor. For a pose frame, the determined graph convolution of a sampling strategy (e.g., centripetal and centrifugal subsets) can be implemented by an $N \times N$ adjacency matrix. Specifically, with K spatial sampling strategies $\sum_{k=1}^K \mathbf{A}_k$ and the adaptive design, Equation 1 can be transformed into

$$\mathbf{f}_{out}(v_t) = \sum_{k=1}^K (\mathbf{A}_k + \mathbf{B}_k + \mathbf{C}_k) \mathbf{f}_{in} \mathbf{W}_k \quad (2)$$

where $\Lambda_k^{-\frac{1}{2}} \bar{\mathbf{A}}_k \Lambda_k^{-\frac{1}{2}}$ is a normalized adjacency matrix of $\bar{\mathbf{A}}_k$ with its elements indicating whether a vertex $v_{t,j}$ is included in the neighbor set. $\Lambda_k^{ii} = \sum_j (\bar{\mathbf{A}}_k^{ij}) + \alpha$ is a diagonal matrix with α set to 0.001 to prevent empty rows. \mathbf{W}_k denotes the weighting function of Equation 1, which is a weight tensor of the 1×1 convolutional operation. Unlike \mathbf{A}_k that represents the physical structure of a human pose, \mathbf{B}_k represents learnable parameters that indicate the connection strength between pose joints, which is implemented with an $N \times N$ adjacency matrix initialized to 0. \mathbf{C}_k performs the similar function of \mathbf{B}_k , which is implemented by the dot product of two feature maps calculated by embedding functions (i.e., θ and ϕ) to calculate the similarity between pose joints. Calculation of \mathbf{C}_k can be represented as

$$\mathbf{C}_k = \text{SoftMax}(\mathbf{f}_{in}^T \mathbf{W}_{\theta k}^T \mathbf{W}_{\phi k} \mathbf{f}_{in}) \quad (3)$$

where \mathbf{W}_{θ} and \mathbf{W}_{ϕ} are learnable parameters of the two embedding functions, which are initialized as 0.0. Then an AGCN block is realized with a $1 \times \Gamma$ classical 2D convolutional layer (Γ is the temporal kernel size that we set to 9) and the defined adaptive graph convolution $\mathbf{f}_{out}(v_t)$, which are both followed by a batch normalization layer and a ReLU layer and a dropout layer in between them. Meanwhile, a residual connection [25] is added to the AGCN block.

Reconstruct 3D Pose Sequence. Taking the inspiration of recent works [61, 28, 38, 37], the introduced AGCN block is then used to extract the spatiotemporal structural information in the global graph representation, which is supervised by estimating the 3D pose sequence from the corresponding 2D sequence (see Figure 1 [Reconstruct 3D Pose Sequence]). Here, each AGCN block has three key parameters: the number of input channels C_{in} , the number of output channels C_{out} , and the stride S of the temporal convolution, while the other parameters are kept consistent (e.g., the temporal convolution kernel size is three). Given an input C_{in} -dim pose representation $F(C_{in}, T_{in}, N)$, the AGCN block derives the output C_{out} -dim pose $F(C_{out}, T_{out}, N)$ via convolution on the pose structure sequence, where T_{out} depends on N_{in} and S . To reconstruct the 3D pose sequence, we first use $AGCN(2, 96, 1)$ to convert the 2D pose sequence $F(2, T, N)$ into a 96D pose representation $F(96, T, N)$. Following the settings of related work, we set T to 243 and N to 17 for the Human3.6M dataset. That is, the input 2D pose sequence of $F(2, 243, 17)$ is converted into a 96D pose sequence of $F(96, 243, 17)$. Then, we stack iterative layers of $AGCN(96, 96, 1)$ to construct the deep spatiotemporal structural representation of the 96D pose sequence. The output of the last AGCN block is fed into an $AGCN(96, 3, 1)$ to estimate the 3D pose sequence based on the 96D joint representation and derive $F(3, 243, 17)$. Then, we let $\ddot{p}_{t,i} \in \mathbb{R}^3$ be the 3D position of the i -th joint at time t , and minimize the difference between the estimated 3D pose sequence and the ground truth 3D pose sequence:

$$\mathcal{L}_{global} = \frac{1}{T} \frac{1}{N} \sum_{t=1}^T \sum_{i=1}^N \|\ddot{p}_{t,i} - p_{t,i}\|_2 \quad (4)$$

Strided Learning Architecture. Inspired by the TCN-based approaches [50, 41], we further adapt the strided learning architecture to the AGCN model, using strided convolution to reduce long time sequences and aggregate temporal information near time t for pose estimation. The gray module in Figure 1 (Strided Learning) illustrates the design of the strided AGCN modules. Each strided AGCN module has two consecutive AGCN blocks, which are surrounded by residual connections [25]. We perform strided convolutions at the second AGCN block of each strided AGCN module to gradually shrink the feature size at the temporal dimension. The input of the first strided AGCN module is the intermediate output in 3D pose sequence reconstruction, i.e., the extracted $F(96, 243, 17)$. After the propagation through the first strided AGCN module, the 96D pose sequence will be shrunken to $F(96, 81, 17)$. Then, we repetitively perform subsequence AGCN layers until the feature size is shrunken to the size of $96 \times 1 \times 17$. In this way, the pattern of the temporal neighbor in the pose sequence will be aggregated for subsequent local 3D

pose joint estimation to estimate the 3D pose of the centric timestep.

3.2. Local 3D Pose Joint Estimation

Based on the above-mentioned strided AGCN modules, the input 2D pose sequence represented as $F(96, 243, 17)$ can be transformed into a feature map $F(96, 1, 17)$. The next step is to estimate the 3D position of joint nodes based on the feature map.

Individually Connected Layers. Existing TCN- and GCN-based methods [50, 41, 79, 61] usually flatten the derived feature maps and use a global skeleton representation consisting of all joint nodes to estimate every single joint, neglecting the matching information between joints and corresponding vectors in feature maps. Unlike existing works, we believe the global knowledge of the temporal and spatial neighborhoods has been aggregated via the proposed global representation. Thus, it is crucial to scope at the spatial information of the corresponding joint node to infer its 3D position. Based on this idea, this paper first proposes an individual connected layer to estimate the 3D position of every single joint based on the corresponding joint node feature $F(96, 1, 1)$, instead of the pooled representation of all joint nodes $F(96, 1, 17)$. Mathematically, the individual connected layer can be denoted as:

$$\dot{p}_i^{(unshared)} = v_i \mathbf{W}_i + \mathbf{b}_i \quad (5)$$

where the estimated 3D position of joint i is denoted by \dot{p}_i and v_i represents the flattened features of $F(96, 1, i)$ joint node i . The weight parameters of the individual connected layer is represented by \mathbf{W}_i and $\mathbf{W}_i \in \mathbb{R}^{96 \times 3}$, whose bias parameter is \mathbf{b}_i and $\mathbf{b}_i \in \mathbb{R}^{1 \times 3}$.

Due to the weight \mathbf{W}_i and bias \mathbf{b}_i are not shared between joints, we name the above individually connected layers as unshared individually connected layers. On top of that, we find that individually connected layers in the unshared fashion may ignore the shared rules between joints in 2D to 3D pose lifting, resulting in overfitting joint-specific distribution. Therefore, we further designed shared individually connected layers:

$$\dot{p}_i^{(shared)} = v_i \mathbf{W}_s + \mathbf{b}_s \quad (6)$$

The weight parameters of the shared individual connected layer are represented by \mathbf{W}_s and $\mathbf{W}_s \in \mathbb{R}^{96 \times 3}$, whose bias parameter is \mathbf{b}_s and $\mathbf{b}_s \in \mathbb{R}^{1 \times 3}$. Then, the 3D pose estimation of each joint can be formulated as the weighted average of the estimated results from the shared and unshared individually connected layers:

$$\bar{p}_i = \lambda \dot{p}_i^{(unshared)} + (1 - \lambda) \dot{p}_i^{(shared)} \quad (7)$$

Here, λ is the parameter that weighs the shared individual connected layer and the unshared individual connected

layer. When λ is 0.0, the model uses only the shared individual connected layer for estimation, and when λ is 1.0, the model uses only the unshared individual connected layer for prediction. Especially, for convenience, the connected layers are implemented via a 1D CNN layer in this paper. Finally, we wish to minimize the difference between the estimated joint pose \bar{p}_i and the ground truth joint pose p_i via L_{local} :

$$\mathcal{L}_{local} = \frac{1}{N} \sum_{i=1}^N \|\bar{p}_i - p_i\|_2 \quad (8)$$

During the training process, we optimize \mathcal{L}_{global} and \mathcal{L}_{local} in two stages. In the first stage, we minimize $\mathcal{L}_{global} + \mathcal{L}_{local}$ to optimize the model using globally supervised signal guidance. In the second stage, we minimize \mathcal{L}_{local} to improve the 3D pose estimation performance.

4. Experiments

4.1. Datasets and Evaluation

Our experiments are based on three public datasets: Human3.6M [29], HumanEva-I [57], and MPI-INF-3DHP [45]. With respect to Human3.6M, the data of subjects S1, S5, S6, S7, and S8 are applied for training, while that of S9 and S11 are used for testing, which is consistent with the training and validation settings of existing works [50, 41, 79, 61]. In terms of HumanEva-I, following [44] and [41], data for actions “walk” and “jog” from subjects S1, S2, and S3 are used for training and testing. For MPI-INF-3DHP, we follow the experimental setting of the recent state-of-the-art [54] for a fair comparison.

Standard evaluation protocols: Mean Per-Joint Position Error (MPJPE) and Pose-aligned MPJPE (P-MPJPE), respectively known as *Protocol#1* and *Protocol#2*, are used for both datasets. The calculation of MPJPE is based on the mean Euclidean distance between the predicted 3D pose joints aligned to root joints (i.e., pelvis) and the ground truth 3D pose joints collected via motion capture, which follows [83, 59, 49]. Compared with MPJPE, P-MPJPE is also based on the mean Euclidean distance but has an extra post-processing step with rigid alignments (e.g., scale, rotation, and translation) to the predicted 3D pose. P-MPJPE leads to smaller differences with the ground truth and it follows [44, 27, 21].

4.2. Implementation Details

We introduce the implementation detail of our GLA-GCN from three main perspectives: 2D pose detections, model setting, and hyperparameters for the training process. For a fair comparison, we follow the 2D pose detections of Human3.6M [29] and HumanEva-I [57] used in [50, 41], which are detected by CPN [14] and MRCNN [24], respectively. The CPN’s 2D pose detection has 17 joints while

Method	Dir.	Disc.	Eat.	Greet	Phone	Photo	Pose	Purch.	Sit	SitD.	Smoke	Wait	WalkD.	Walk	WalkT.	Avg.
Martinez et al. [44] (ICCV'17)	51.8	56.2	58.1	59.0	69.5	78.4	55.2	58.1	74.0	94.6	62.3	59.1	65.1	49.5	52.4	62.9
Fang et al. [20] (AAAI'18)	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
Pavlakos et al. [48] (CVPR'18)	48.5	54.4	54.4	52.0	59.4	65.3	49.9	52.9	65.8	71.1	56.6	52.9	60.9	44.7	47.8	56.2
Lee et al. [34] (ECCV'18) †	40.2	49.2	47.8	52.6	50.1	75.0	50.2	43.0	55.8	73.9	54.1	55.6	58.2	43.3	43.3	52.8
Zhao et al. [79] (CVPR'19)	47.3	60.7	51.4	60.5	61.1	49.9	47.3	68.1	86.2	55.0	67.8	61.0	<u>42.1</u>	60.6	45.3	57.6
Ci et al. [17] (ICCV'19)	46.8	52.3	44.7	50.4	52.9	68.9	49.6	46.4	60.2	78.9	51.2	50.0	54.8	40.4	43.3	52.7
Pavlo et al. [50] (CVPR'19) †	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
Cai et al. [9] (ICCV'19) †	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
Pavlo et al. [50] (CVPR'19) †	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
Xu et al. [67] (CVPR'20) †	37.4	43.5	42.7	42.7	46.6	59.7	41.3	45.1	<u>52.7</u>	60.2	45.8	43.1	47.7	33.7	37.1	45.6
Liu et al. [41] (CVPR'20) †	41.8	44.8	41.1	44.9	47.4	54.1	43.4	42.2	56.2	63.6	45.3	43.5	45.3	<u>31.3</u>	32.2	45.1
Zeng et al. [74] (ECCV'20) †	46.6	47.1	43.9	41.6	45.8	49.6	46.5	40.0	53.4	61.1	46.1	42.6	43.1	31.5	32.6	44.8
Xu and Takano [68] (CVPR'21)	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
Zhou et al. [82] (PAMI'21) †	<u>38.5</u>	45.8	<u>40.3</u>	54.9	39.5	45.9	39.2	43.1	49.2	71.1	41.0	53.6	44.5	33.2	34.1	45.1
Li et al. [38] (CVPR'22) †	39.2	<u>43.1</u>	40.1	40.9	<u>44.9</u>	51.2	<u>40.6</u>	41.3	53.5	60.3	<u>43.7</u>	41.1	43.8	29.8	30.6	<u>43.0</u>
Shan et al. [54] (ECCV'22) †	38.9	42.7	40.4	<u>41.1</u>	45.6	49.7	40.9	39.9	55.5	<u>59.4</u>	44.9	42.2	42.7	29.4	29.4	42.8
Our GLA-GCN (T=243, CPN) †	41.3	44.3	40.8	41.8	45.9	54.1	42.1	<u>41.5</u>	57.8	62.9	45.0	<u>42.8</u>	45.9	29.4	29.9	44.4
Martinez et al. [44] (ICCV'17)	37.7	44.4	40.3	42.1	48.2	54.9	44.4	42.1	54.6	58.0	45.1	46.4	47.6	36.4	40.4	45.5
Lee et al. [34] (ECCV'18) †	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
Zhao et al. [79] (CVPR'19)	37.8	49.4	37.6	40.9	45.1	41.4	40.1	48.3	50.1	42.2	53.5	44.3	40.5	47.3	39.0	43.8
Ci et al. [17] (ICCV'19)	36.3	38.8	29.7	37.8	34.6	42.5	39.8	32.5	<u>36.2</u>	<u>39.5</u>	34.4	38.4	38.2	31.3	34.2	36.3
Liu et al. [41] (CVPR'20) †	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
Xu and Takano [68] (CVPR'21)	35.8	38.1	31.0	35.3	35.8	43.2	37.3	31.7	38.4	45.5	35.4	36.7	36.8	27.9	30.7	35.8
Zheng et al. [80] (ICCV'21) †	30.0	33.6	29.9	31.0	30.2	33.3	34.8	31.4	37.8	38.6	31.7	31.5	29.0	23.3	23.1	31.3
Li et al. [38] (CVPR'22) †	<u>27.7</u>	32.1	<u>29.1</u>	28.9	30.0	33.9	33.0	31.2	37.0	39.3	30.0	31.0	29.4	22.2	23.0	30.5
Shan et al. [54] (ECCV'22) †	28.5	<u>30.1</u>	28.6	<u>27.9</u>	<u>29.8</u>	<u>33.2</u>	<u>31.3</u>	<u>27.8</u>	36.0	37.4	29.7	<u>29.5</u>	<u>28.1</u>	<u>21.0</u>	<u>21.0</u>	<u>29.3</u>
Our GLA-GCN (T=243, GT) †	26.5	27.2	29.2	25.4	28.2	31.7	29.5	26.9	37.8	39.9	<u>29.9</u>	27.0	27.3	20.5	20.8	28.5
Wang et al. [61] (ECCV'20) †*	23.0	25.7	22.8	22.6	24.1	30.6	24.9	24.5	31.1	35.0	25.6	24.3	25.1	19.8	18.4	25.6
Li et al. [37] (TMM'22) †*	27.1	29.4	26.5	27.1	28.6	33.0	30.7	26.8	38.2	34.7	29.1	29.8	26.8	19.1	19.8	28.5
Hu et al. [28] (MM'22) †*	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	22.7
Zhang et al. [77] (CVPR'22) †*	<u>21.6</u>	<u>22.0</u>	<u>20.4</u>	<u>21.0</u>	20.8	24.3	<u>24.7</u>	<u>21.9</u>	26.9	24.9	<u>21.2</u>	<u>21.5</u>	<u>20.8</u>	<u>14.7</u>	<u>15.7</u>	<u>21.6</u>
Our method (T=243, GT) †*	20.1	21.2	20.0	19.6	<u>21.5</u>	<u>26.7</u>	23.3	19.8	<u>27.0</u>	<u>29.4</u>	20.8	20.1	19.2	12.8	13.8	21.0

Table 1. *Protocol#1*: Reconstruction error with MPJPE (mm) on Human3.6M. Top-table: input 2D pose sequences are detected by (CPN) - cascaded pyramid network. Bottom-table: input 2D pose sequences with ground truth (GT). Best in bold, second best underlined, the lower the better. † indicates using temporal information. * indicates reconstructing an intermediate 3D pose sequence.

Method	Walk			Jog			Avg
	S1	S2	S3	S1	S2	S3	
Martinez et al. [44] (ICCV'17)	19.7	17.4	46.8	26.9	18.2	18.6	24.6
Fang et al. [20] (AAAI'18)	19.4	16.8	37.4	30.4	17.6	16.3	23.0
Pavlakos et al. [48] (CVPR'18)	18.8	12.7	29.2	23.5	15.4	14.5	19.0
Lee et al. [34] (ECCV'18) †	18.6	19.9	30.5	25.7	16.8	17.7	21.5
Pavlo et al. [50] (CVPR'19) †	13.9	10.2	46.6	20.9	13.1	13.8	19.8
Liu et al. [41] (CVPR'20) †	13.1	<u>9.8</u>	<u>26.8</u>	16.9	<u>12.8</u>	<u>13.3</u>	<u>15.5</u>
Zheng et al. [80] (ICCV'21) †	14.4	10.2	46.6	22.7	13.4	13.4	20.1
Li et al. [37] (TMM'22) †*	14.0	10.0	32.8	19.5	13.6	14.2	17.4
Zhang et al. [77] (CVPR'22) †*	<u>12.7</u>	10.9	17.6	22.6	15.8	17.0	16.1
Ours (T=27, MRCNN) †	12.5	9.1	26.9	<u>18.5</u>	12.7	12.8	15.4
Li et al. [37] (TMM'22) †*	<u>9.7</u>	<u>7.6</u>	15.8	<u>12.3</u>	<u>9.4</u>	<u>11.2</u>	<u>11.1</u>
Ours (T=27, GT) †	8.7	6.8	11.5	10.1	8.2	9.9	9.2

Table 2. Results of *Protocol#2* for HumanEva-I. † uses temporal information. Best in bold, second best underlined. * indicates reconstructing an intermediate 3D pose sequence.

the MRCNN's 2D pose detection has 15 joints. Besides, we also conduct experiments for the ground truth (GT) 2D pose detections of the two datasets.

Based on the specific structure of 2D pose, we implement the graph convolutional operation filters of AGCN blocks, e.g., the sizes of \mathbf{A}_k , \mathbf{B}_k , \mathbf{C}_k are set to 17×17 , 15×15 , and 17×17 for Human3.6M, HumanEva-I, and MPI-INF-3DHP, respectively. The designed model has some key parameters that can be adjusted to get better performance. For this part, we conduct ablation experiments

Method	PCK↑	AUC↑	MPJPE↓
Mehta et al. [45] (3DV'17, T=1)	75.7	39.3	117.6
Pavlo et al. [50] (CVPR'19, T=81) †	86.0	51.9	84.0
Lin et al. [39] (BMVC'19, T=25) †	83.6	51.4	79.8
Wang et al. [61] (ECCV'20, T=96) †*	86.9	62.1	68.1
Zheng et al. [80] (ICCV'21, T=9) †	88.6	56.4	77.1
Chen et al. [10] (TCSVT'21, T=81) †	87.9	54.0	78.8
Hu et al. [28] (MM'22, T=96) †*	97.9	69.5	42.5
Shan et al. [54] (ECCV'22, T=81) †	97.9	75.8	32.2
Our GLA-GCN (T=27) †	<u>98.19</u>	<u>76.53</u>	<u>31.36</u>
Our GLA-GCN (T=81) †	98.53	79.12	27.76

Table 3. Results of *Protocol#1* for MPI-INF-3DHP. † uses temporal information. Best in bold, second best underlined. * indicates reconstructing an intermediate 3D pose sequence.

with different numbers of channels and 2D pose frames (i.e., C_{out} and T , respectively) on Human3.6M. To verify the proper design of the proposed model regarding the strided design and individual connected layer, we perform further ablation experiments on both datasets.

In terms of the hyperparameters, we respectively set the batch size to 512, 256, and 256 for Human3.6M, HumanEva-I, and MPI-INF-3DHP. Being consistent with [41], we adopt the ranger optimizer and train the model with the MPJPE loss for 80 and 1000 epochs for Human3.6M and HumanEva-I, respectively, using an initial learning rate of 0.01. Meanwhile, we set the dropout rate to 0.1. For both

Method	Dir.	Disc.	Eat.	Greet	Phone	Photo	Pose	Purch.	Sit	SitD.	Smoke	Wait	WalkD.	Walk	WalkT.	Avg.
Martinez et al. [44] (ICCV'17)	39.5	43.2	46.4	47.0	51.0	56.0	41.4	40.6	56.5	69.4	49.2	45.0	49.5	38.0	43.1	47.7
Fang et al. [20] (AAAI'18)	38.2	41.7	43.7	44.9	48.5	55.3	40.2	38.2	54.5	64.4	47.2	44.3	47.3	36.7	41.7	45.7
Pavlakos et al. [48] (CVPR'18)	34.7	39.8	41.8	38.6	42.5	47.5	38.0	36.6	50.7	56.8	42.6	39.6	43.9	32.1	36.5	41.8
Lee et al. [34] (ECCV'18) †	34.9	35.2	43.2	42.6	46.2	55.0	37.6	38.8	50.9	67.3	48.9	35.2	31.0	50.7	34.6	43.4
Cai et al. [9] (ICCV'19) †	35.7	37.8	36.9	40.7	39.6	45.2	37.4	34.5	46.9	50.1	40.5	36.1	41.0	29.6	33.2	39.0
Pavlo et al. [50] (CVPR'19) †	34.1	36.1	34.4	37.2	36.4	42.2	34.4	33.6	45.0	52.5	37.4	33.8	37.8	25.6	27.3	36.5
Xu et al. [67] (CVPR'20) †	<u>31.0</u>	<u>34.8</u>	34.7	34.4	36.2	43.9	31.6	33.5	<u>42.3</u>	49.0	37.1	33.0	39.1	26.9	31.9	36.2
Chen et al. [14] (ICCV'20) †	32.9	35.2	35.6	34.4	36.4	42.7	<u>31.2</u>	32.5	45.6	50.2	37.3	32.8	36.3	26.0	23.9	35.5
Liu et al. [41] (CVPR'20) †	32.3	35.2	33.3	35.8	35.9	41.5	33.2	32.7	44.6	50.9	37.0	32.4	37.0	25.2	27.2	35.6
Shan et al. [55] (MM'21) †	32.5	36.2	33.2	35.3	35.6	42.1	32.6	31.9	42.6	<u>47.9</u>	36.6	<u>32.1</u>	34.8	24.2	25.8	35.0
Shan et al. [54] (ECCV'22) †	31.3	35.2	32.9	<u>33.9</u>	35.4	<u>39.3</u>	32.5	<u>31.5</u>	44.6	48.2	<u>36.3</u>	<u>32.9</u>	<u>34.4</u>	<u>23.8</u>	<u>23.9</u>	<u>34.4</u>
Zhang et al. [77] (CVPR'22) †*	30.8	33.1	30.3	31.8	33.1	39.1	31.1	30.5	42.5	44.5	34.0	30.8	32.7	22.1	22.9	32.6
Our GLA-GCN (T=243, CPN) †	32.4	35.3	<u>32.6</u>	34.2	<u>35.0</u>	42.1	32.1	31.9	45.5	49.5	36.1	32.4	35.6	23.5	24.7	34.8
Martinez et al. [44] (ICCV'17)	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	37.1
Ci et al. [17] (ICCV'19)	<u>24.6</u>	<u>28.6</u>	<u>24.0</u>	<u>27.9</u>	<u>27.1</u>	<u>31.0</u>	<u>28.0</u>	<u>25.0</u>	<u>31.2</u>	<u>35.1</u>	<u>27.6</u>	<u>28.0</u>	<u>29.1</u>	<u>24.3</u>	<u>26.9</u>	<u>27.9</u>
Our GLA-GCN (T=243, GT) †	20.2	21.9	21.7	19.9	21.6	24.7	22.5	20.8	28.6	33.1	22.7	20.6	20.3	15.9	16.2	22.0
Our GLA-GCN (T=243, GT) †*	16.6	18.1	16.2	17.0	17.0	22.2	19.0	17.1	22.4	25.9	17.5	16.4	16.3	10.8	11.6	17.6

Table 4. *Protocol#2*: Reconstruction error after rigid alignment with P-MPJPE (mm) on Human3.6M. Top-table: input 2D joints are acquired by detection (CPN) - cascaded pyramid network. Bottom-table: input 2D joints with (GT) - ground truth. † indicates using temporal information. * indicates reconstructing an intermediate 3D pose sequence. Best in bold, second best underlined.

the training and testing phases, data augmentation is applied by horizontally flipping the pose data. All experiments are conducted with two GeForce GTX 3090 GPUs.

4.3. Comparison with State-of-the-Art

Tables 1 and 4 show the comparison on Human3.6M and HumanEva-I with state-of-the-art methods under *Protocol#1* and *Protocol#2*, respectively. Based on the implementation via GT 2D pose respectively optimized with or without loss of reconstructing the intermediate 3D pose sequence (defined in Equation 4), our GLA-GCN outperforms the state-of-the-art method [77] in terms of averaged results of two evaluation protocols. Figure 3 shows the training process of our GLA-GCN on Human3.6M, which indicates our model converges quickly without observable overfitting. For the HumanEva-I dataset, the results of *Protocol#2* (see Table 2) also show that our method is superior to state-of-the-art methods by just using the MPJPE loss. We also conduct a qualitative comparison with the state-of-the-art method that does not have a 3D pose sequence reconstruction module [41]. Figure 2 shows the visualized improvements over [41]. For example, in the ‘‘S11 WalkT.’’ action, the visualizations of right-hip and right-hand joints estimated with our method and the ground truth 3D pose are clearly separate but those of [41] are connected to each other. Moving on to the MPI-INF-3DHP dataset in Table 3, we can see a significant decline in MPJPE with our model. Compared with the state-of-the-art method, P-STMO[54], the MPJPE of our model decreases from 32.2mm to 27.76mm, representing an error reduction of approximately 14%.

Comparing the performance by using estimated 2D pose (i.e., CPN or HR-Net pose data) is also regarded as important by most existing works. However, models such as [38, 54] can perform well on relatively low-quality esti-

mated 2D pose data but fail to well generalize the good performance to high-quality 2D data (i.e., ground truth poses). We note that our method underperforms on relatively low-quality estimated 2D pose when compared with some recent methods [38, 54, 77]. In the following, we conduct an in-depth discussion on this issue.

Discussion: the Effect of 2D Pose Quality. Tracing back to the first work on 3D Pose lifting, Martinez et al. [44] used the SH 2D pose detector fine-tuned on the Human3.6M dataset to improve the 3D HPE (significantly from 67.5mm average MPJPE to 62.9mm), indicating that the quality of 2D pose can be essential for the 3D HPE. Recent works [61, 38, 77] took advantage of advanced 2D pose detector HR-Net and achieved better performance (e.g., 39.8mm average MPJPE). Zhu et al. [84] also successfully improved the result to 37.5mm average MPJPE by fine-tuning the SH network [47] on Human3.6M, which remains far behind the results implemented with GT 2D pose.

A similar observation is also applicable to the HumanEva-I and MPI-INF-3DHP datasets. As shown in Table 2, our method yields a remarkable 40% drop in P-MPJPE on the HumanEva-I dataset. Given the GT 2D pose, the P-MPJPE goes from 15.4mm to 9.2mm compared with the best state-of-the-art algorithm. While on MPI-INF-3DHP, the MPJPE goes from 32.2mm to 27.76mm.

Hence, improving the performance of the estimated pose purely relies on preparing quality 2D pose data, which can be easily achieved by either using an advanced 2D pose detector that can generate pose data similar to the GT 2D pose or just arbitrarily fine-tuning the existing pose detectors. On the other hand, it remains unclear for what scenario the reconstructed 3D pose with advanced pose detectors can be beneficial. One scenario is 3D human pose estimation in the wild, which is usually evaluated with qualitative visualization [37]. However, whether the 3D pose reconstructed

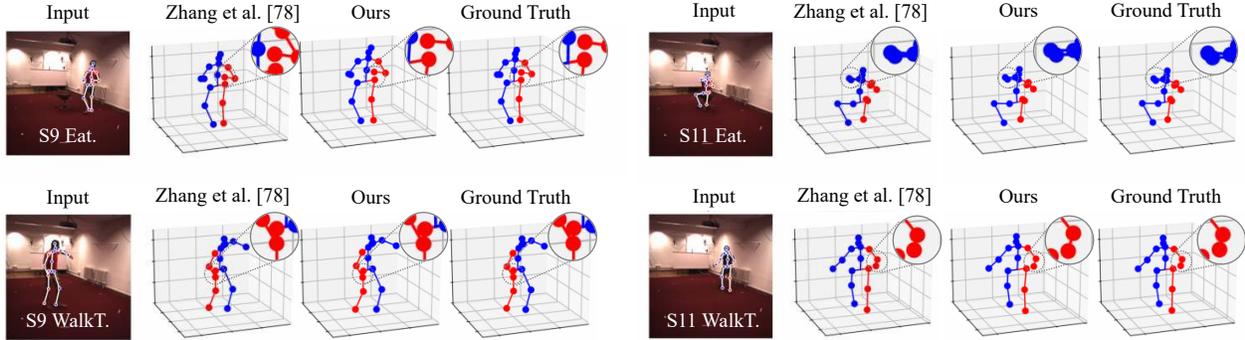


Figure 2. Qualitative comparison with Zhang et al. [77] for S9 and S11 on two actions of Human3.6M. Noticeable improvements are enlarged. For most cases, our GLA-GCN gets better results, while it occasionally generates worse results (e.g., the bottom right case).

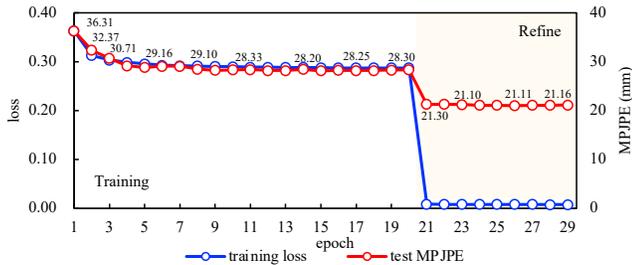


Figure 3. Loss on the training set and MPJPE on the test set.

from the estimated 2D pose can contribute to pose-based tasks remains under-explored. Given that how to improve the performance of the estimated 2D pose is straightforward and its usage remains lack of good applicable scenario, we argue that the comparison based on the GT 2D pose can more properly reflect a model’s 3D HPE ability than comparisons based on the estimated 2D pose.

4.4. Ablation Studies

In the following, we ablate our model design gradients (i.e., AGCN layers, strided design, and individual connected layer). To validate the properness of using AGCN layers, we compare our model with the version implemented with ST-GCN [69] blocks, which leads to the ablation of AGCN. As shown in row #1 of Table 6, the results of *Protocol#2* on two datasets consistently indicate that using AGCN blocks can achieve better performances.

For the ablation of strided design, we perform average pooling on the second (i.e., temporal) dimension of the feature map. Results in row #2 of Table 6 indicate that it is not as effective as the strided design. Without the strided design, it will not only lead to a larger feature map representation, i.e., increased from $F(C_{out}, 1, N)$ to $F(C_{out}, T, N)$ but also affects the 3D HPE.

To verify the design of our individual connected layer, we compare it with the implementation of a fully connected layer that takes the expanded feature map as its input. The results in row #3 of Table 6 indicate that our individual connected layer can make better use of the structured repre-

Method	Frames	No. of Parameters	MPJPE (mm)
Pavlo et al. [50] (CVPR’19) †	$T = 27$	8.56M	40.6
Liu et al. [41] (CVPR’20) †	$T = 27$	5.69M	38.9
Li et al. [38] (CVPR’22) †*	$T = 27$	18.92M	34.3
Our GLA-GCN †	$T = 27$	0.84M	34.4
Pavlo et al. [50] (CVPR’19) †	$T = 81$	12.75M	38.7
Liu et al. [41] (CVPR’20) †	$T = 81$	8.46M	36.2
Li et al. [38] (CVPR’22) †*	$T = 81$	$\geq 18.92M$	32.7
Our GLA-GCN †	$T = 81$	1.10M	31.5
Pavlo et al. [50] (CVPR’19) †	$T = 243$	16.95M	37.8
Liu et al. [41] (CVPR’20) †	$T = 243$	11.25M	34.7
Our GLA-GCN †	$T = 243$	1.35M	28.5
Wang et al. [61] (ECCV’20) †*	$T = 96$	1.69M	25.6
Our GLA-GCN ($C_{out=96}$) †*	$T = 243$	1.88M	24.5
Hu et al. [28] (MM’22) †*	$T = 96$	3.42M	22.7
Li et al. [38] (CVPR’22) †*	$T = 351$	$\geq 18.92M$	30.5
Zhang et al. [77] (CVPR’22) †*	$T = 243$	33.70M	21.6
Our GLA-GCN ($C_{out=512}$) †*	$T = 243$	48.64M	21.0

Table 5. Comparison with state-of-the-art methods on Human3.6M implemented with different receptive fields of ground truth 2D pose. Results of *Protocol#1* are reported. * indicates reconstructing an intermediate 3D pose sequence.

sentation of GCN and thus significantly improves the performance. The differences of features before the prediction layers (i.e., individually and fully connected layers) are respectively visualized in the upper and lower rows of Figure 4. The visualization indicates that our individual connected layer can make predictions based on more interpretable features that cannot be traced by using a fully connected layer. For example, the arm and leg joints show relatively higher independence than other joints for actions “eating” and “walking”, respectively. Feeding these independent features of all joints to a fully connected layer will interfere with the prediction of a specific joint.

We further verify the advantage of this structured representation of GCN by swapping the left and right limbs of 2D pose input data, leading to the break of pose structure. Results in row #4 of Table 6 show that breaking the pose structure will affect the 3D pose estimation. This observation, in turn, further indicates the proper design of our individual connected layer.

Discussion: Limitation on Model Size. Similar to state-of-the-art methods [28, 38, 77], we note that our method is

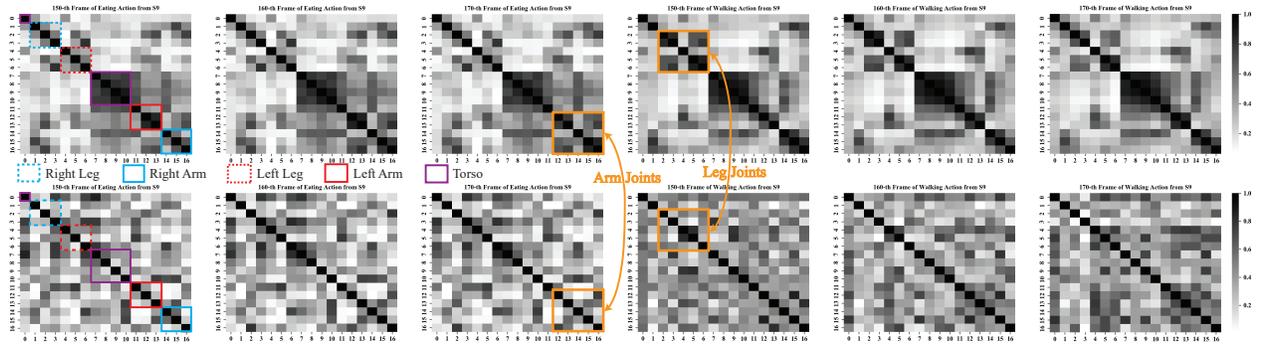


Figure 4. Visualizations of inter-joint feature cosine similarity from actions: “eating” (first 3 columns) and “walking” (last three columns) of Human3.6M. **Upper row** uses an individual connected layer; **lower row** uses a fully connected layer (please zoom in for a better view).

#	Method	Human3.6M		HumanEva-I	
		CPN	GT	MRCNN	GT
1	w/o AGCN	39.1	28.0	18.2	11.7
2	w/o strided design	41.2	30.6	22.6	12.9
3	w/o individual connected layer	39.0	27.7	17.6	12.4
4	Swap the structure of input 2D pose	38.3	28.1	16.4	12.4
5	GLA-GCN (T=27) †	37.8	25.8	15.4	9.2

Table 6. Ablation study on key designs of our GLA-GCN. The results are based on the average value of *Protocol#2* implemented with 27 receptive fields for various 2D pose detections of the Human3.6M and HumanEva-I datasets.

faced with the issue of model size. Specifically, the lower table of Table 5 shows that our model can achieve better performance than state-of-the-art methods [61, 77] but uses slightly more model parameters. We aim to tackle this issue in the future by using techniques such as pruning.

5. Conclusion

This paper proposes a GCN-based method utilizing the structured representation for 3D HPE in the 2D to 3D lifting paradigm. The proposed GLA-GCN globally represents the 2D pose sequence and locally estimates the 3D pose joints via an individual connected layer. Results show that our GLA-GCN outperforms corresponding state-of-the-art methods implemented with GT 2D poses on datasets Human3.6M, HumanEva-I, and MPI-INF-3DHP. We verify the properness of model design with extensive ablation studies and visualizations. In the future, we will tackle the issue of parameter efficiency of our model via tuning techniques [71, 72]. Meanwhile, we will consider its effect on application scenarios such as human behavior understanding [6, 5, 4, 7, 3] and aim to improve the results of the estimated 2D pose by preparing high-quality 2D pose data via fine-tuned 2D pose detectors (e.g., SH detector [47] and HR-Net [58]), and investigate the effects of other loss terms (e.g., based on bone features [11] and motion trajectory [61]).

References

- [1] Mykhaylo Andriluka, Stefan Roth, and Bernt Schiele. Pictorial structures revisited: People detection and articulated pose estimation. In *2009 IEEE conference on computer vision and pattern recognition*, pages 1014–1021. IEEE, 2009.
- [2] Anurag Arnab, Carl Doersch, and Andrew Zisserman. Exploiting temporal context for 3d human pose estimation in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3395–3404, 2019.
- [3] XB Bruce, Yan Liu, and Keith CC Chan. Skeleton-based detection of abnormalities in human actions using graph convolutional networks. In *2020 Second International Conference on Transdisciplinary AI (TransAI)*, pages 131–137. IEEE, 2020.
- [4] XB Bruce, Yan Liu, and Keith CC Chan. Multimodal fusion via teacher-student network for indoor action recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 3199–3207, 2021.
- [5] XB Bruce, Yan Liu, Keith CC Chan, Quintai Yang, and Xiaoying Wang. Skeleton-based human action evaluation using graph convolutional network for monitoring alzheimer’s progression. *Pattern Recognition*, 119:108095, 2021.
- [6] Yu Bruce X.B., Yan Liu, Xiang Zhang, Sheng-hua Zhong, and Keith CC Chan. Mmnet: A model-based multimodal network for human action recognition in rgb-d videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3):3522–3538, 2022.
- [7] Yu Bruce X.B., Liu Yan, Zhang Xiang, Chen Gong, and Chan Keith C.C. EgcN: An ensemble-based learning framework for exploring effective skeleton-based rehabilitation exercise assessment. *EGCN: An Ensemble-based Learning Framework for Exploring Effective Skeleton-based Rehabilitation Exercise Assessment*, pages 3681–3687, 2022.
- [8] Yujun Cai, Lihao 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/CVF International Conference on Computer Vision (ICCV)*, pages 2272–2281, 2019.
- [9] Yujun Cai, Lihao Ge, Jun Liu, Jianfei Cai, Tat-Jen Cham, Junsong Yuan, and Nadia Magnenat Thalmann. Exploit-

- ing spatial-temporal relationships for 3d pose estimation via graph convolutional networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2272–2281, 2019.
- [10] 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.
- [11] 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*, 2021.
- [12] Xipeng Chen, Pengxu Wei, and Liang Lin. Deductive learning for weakly-supervised 3d human pose estimation via uncalibrated cameras. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 1089–1096, 2021.
- [13] Yucheng Chen, Yingli Tian, and Mingyi He. Monocular human pose estimation: A survey of deep learning-based methods. *Computer Vision and Image Understanding*, 192:102897, 2020.
- [14] 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*, pages 7103–7112, 2018.
- [15] Yu Cheng, Bo Yang, Bo Wang, and Robby T Tan. 3d human pose estimation using spatio-temporal networks with explicit occlusion training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 10631–10638, 2020.
- [16] Hongsuk Choi, Gyeongsik Moon, and Kyoung Mu Lee. Pose2mesh: Graph convolutional network for 3d human pose and mesh recovery from a 2d human pose. In *European Conference on Computer Vision*, pages 769–787. Springer, 2020.
- [17] Hai Ci, Chunyu Wang, Xiaoxuan Ma, and Yizhou Wang. Optimizing network structure for 3d human pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2262–2271, 2019.
- [18] Qi Dang, Jianqin Yin, Bin Wang, and Wenqing Zheng. Deep learning based 2d human pose estimation: A survey. *Tsinghua Science and Technology*, 24(6):663–676, 2019.
- [19] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. *Advances in neural information processing systems*, 29:3844–3852, 2016.
- [20] 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.
- [21] 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.
- [22] Erik Gärtner, Mykhaylo Andriluka, Hongyi Xu, and Cristian Sminchisescu. Trajectory optimization for physics-based reconstruction of 3d human pose from monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13106–13115, 2022.
- [23] Kehong Gong, Jianfeng Zhang, and Jiashi Feng. Poseaug: A differentiable pose augmentation framework for 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8575–8584, 2021.
- [24] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
- [25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [26] Yihui He, Rui Yan, Katerina Fragkiadaki, and Shoou-I Yu. Epipolar transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7779–7788, 2020.
- [27] Mir Rayat Imtiaz Hossain and James J Little. Exploiting temporal information for 3d human pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 68–84, 2018.
- [28] 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 29th ACM International Conference on Multimedia*, pages 602–611, 2021.
- [29] 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.
- [30] Karim Iskakov, Egor Burkov, Victor Lempitsky, and Yury Malkov. Learnable triangulation of human pose. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7718–7727, 2019.
- [31] Xiaopeng Ji, Qi Fang, Juntao Dong, Qing Shuai, Wen Jiang, and Xiaowei Zhou. A survey on monocular 3d human pose estimation. *Virtual Reality & Intelligent Hardware*, 2(6):471–500, 2020.
- [32] Jogendra Nath Kundu, Siddharth Seth, Pradyumna YM, Varun Jampani, Anirban Chakraborty, and R Venkatesh Babu. Uncertainty-aware adaptation for self-supervised 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20448–20459, 2022.
- [33] Hsi-Jian Lee and Zen Chen. Determination of 3d human body postures from a single view. *Computer Vision, Graphics, and Image Processing*, 30(2):148–168, 1985.
- [34] 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.

- [35] Ruoyu Li, Sheng Wang, Feiyun Zhu, and Junzhou Huang. Adaptive graph convolutional neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [36] Sijin Li, Weichen Zhang, and Antoni B Chan. Maximum-margin structured learning with deep networks for 3d human pose estimation. In *Proceedings of the IEEE international conference on computer vision*, pages 2848–2856, 2015.
- [37] 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.
- [38] Wenhao Li, Hong Liu, Hao Tang, Pichao Wang, and Luc Van Gool. Mhformer: Multi-hypothesis transformer for 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13147–13156, 2022.
- [39] 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*, 2019.
- [40] Kenkun Liu, Rongqi Ding, Zhiming Zou, Le Wang, and Wei Tang. A comprehensive study of weight sharing in graph networks for 3d human pose estimation. In *European Conference on Computer Vision*, pages 318–334. Springer, 2020.
- [41] 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 5064–5073, 2020.
- [42] Zhiwei Liu, Xiangyu Zhu, Lu Yang, Xiang Yan, Ming Tang, Zhen Lei, Guibo Zhu, Xuetao Feng, Yan Wang, and Jinqiao Wang. Multi-initialization optimization network for accurate 3d human pose and shape estimation. In *Proceedings of the ACM International Conference on Multimedia*, pages 1976–1984, 2021.
- [43] 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 6238–6247, 2021.
- [44] 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*, pages 2640–2649, 2017.
- [45] 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 *International Conference on 3D Vision*, pages 506–516, 2017.
- [46] Gyeongsik Moon and Kyoung Mu Lee. I2l-meshnet: Image-to-lixel prediction network for accurate 3d human pose and mesh estimation from a single rgb image. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VII 16*, pages 752–768. Springer, 2020.
- [47] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European conference on computer vision*, pages 483–499. Springer, 2016.
- [48] Georgios Pavlakos, Xiaowei Zhou, and Kostas Daniilidis. Ordinal depth supervision for 3d human pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7307–7316, 2018.
- [49] 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*, pages 7025–7034, 2017.
- [50] Dario Pavllo, 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*, pages 7753–7762, 2019.
- [51] Liliana Lo Presti and Marco La Cascia. 3d skeleton-based human action classification: A survey. *Pattern Recognition*, 53:130–147, 2016.
- [52] N Dinesh Reddy, Laurent Guigues, Leonid Pishchulin, Jayan Eledath, and Srinivasa G Narasimhan. Tesseract: End-to-end learnable multi-person articulated 3d pose tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15190–15200, 2021.
- [53] Nikolaos Sarafianos, Bogdan Boteanu, Bogdan Ionescu, and Ioannis A Kakadiaris. 3d human pose estimation: A review of the literature and analysis of covariates. *Computer Vision and Image Understanding*, 152:1–20, 2016.
- [54] Wenkang Shan, Zhenhua Liu, Xinfeng Zhang, Shanshe Wang, Siwei Ma, and Wen Gao. P-stmo: Pre-trained spatial temporal many-to-one model for 3d human pose estimation. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part V*, pages 461–478. Springer, 2022.
- [55] Wenkang Shan, Haopeng Lu, Shanshe Wang, Xinfeng Zhang, and Wen Gao. Improving robustness and accuracy via relative information encoding in 3d human pose estimation. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 3446–3454, 2021.
- [56] Lei Shi, Yifan Zhang, Jian Cheng, and Hanqing Lu. Two-stream adaptive graph convolutional networks for skeleton-based action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12026–12035, 2019.
- [57] Leonid Sigal, Alexandru O Balan, and Michael J Black. Humaneva: Synchronized video and motion capture dataset and baseline algorithm for evaluation of articulated human motion. *International journal of computer vision*, 87(1-2):4, 2010.
- [58] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5693–5703, 2019.
- [59] Bugra Tekin, Artem Rozantsev, Vincent Lepetit, and Pascal Fua. Direct prediction of 3d body poses from motion compensated sequences. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 991–1000, 2016.

- [60] Raquel Urtasun and Pascal Fua. 3d human body tracking using deterministic temporal motion models. In *European conference on computer vision*, pages 92–106. Springer, 2004.
- [61] Jingbo Wang, Sijie Yan, Yuanjun Xiong, and Dahua Lin. Motion guided 3d pose estimation from videos. In *European Conference on Computer Vision*, pages 764–780. Springer, 2020.
- [62] Junjie Wang, Zhenbo Yu, Zhengyan Tong, Hang Wang, Jinxian Liu, Wenjun Zhang, and Xiaoyan Wu. Ocr-pose: Occlusion-aware contrastive representation for unsupervised 3d human pose estimation. In *Proceedings of the ACM International Conference on Multimedia*, pages 5477–5485, 2022.
- [63] Luyang Wang, Yan Chen, Zhenhua Guo, Keyuan Qian, Mude Lin, Hongsheng Li, and Jimmy S Ren. Generalizing monocular 3d human pose estimation in the wild. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 0–0, 2019.
- [64] Zitian Wang, Xuecheng Nie, Xiaochao Qu, Yunpeng Chen, and Si Liu. Distribution-aware single-stage models for multi-person 3d pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13096–13105, 2022.
- [65] Tom Wehrbein, Marco Rudolph, Bodo Rosenhahn, and Bastian Wandt. Probabilistic monocular 3d human pose estimation with normalizing flows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11199–11208, 2021.
- [66] Wen-Li Wei, Jen-Chun Lin, Tyng-Luh Liu, and Hong-Yuan Mark Liao. Capturing humans in motion: temporal-attentive 3d human pose and shape estimation from monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13211–13220, 2022.
- [67] Jingwei Xu, Zhenbo Yu, Bingbing Ni, Jiancheng Yang, Xiaokang Yang, and Wenjun Zhang. Deep kinematics analysis for monocular 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 899–908, 2020.
- [68] Tianhan Xu and Wataru Takano. Graph stacked hourglass networks for 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16105–16114, 2021.
- [69] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *32nd AAAI conference on artificial intelligence*, 2018.
- [70] Yi Yang and Deva Ramanan. Articulated pose estimation with flexible mixtures-of-parts. In *CVPR 2011*, pages 1385–1392. IEEE, 2011.
- [71] Bruce XB Yu, Jianlong Chang, Lingbo Liu, Qi Tian, and Chang Wen Chen. Towards a unified view on visual parameter-efficient transfer learning. *arXiv preprint arXiv:2210.00788*, 2022.
- [72] Bruce XB Yu, Jianlong Chang, Haixin Wang, Lingbo Liu, Shijie Wang, Zhiyu Wang, Junfan Lin, Lingxi Xie, Haojie Li, Zhouchen Lin, et al. Visual tuning. *arXiv preprint arXiv:2305.06061*, 2023.
- [73] Bruce XB Yu, Yan Liu, and Keith CC Chan. A survey of sensor modalities for human activity recognition. 1:282–294, 2020.
- [74] 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 *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*, pages 507–523. Springer, 2020.
- [75] 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/CVF International Conference on Computer Vision (ICCV)*, pages 11436–11445, 2021.
- [76] Yu Zhan, Fenghai Li, Renliang Weng, and Wongun Choi. Ray3d: ray-based 3d human pose estimation for monocular absolute 3d localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13116–13125, 2022.
- [77] Jinlu Zhang, Zhigang Tu, Jianyu Yang, Yujin Chen, and Junsong Yuan. Mixste: Seq2seq mixed spatio-temporal encoder for 3d human pose estimation in video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13232–13242, 2022.
- [78] Zhe Zhang, Chunyu Wang, Weichao Qiu, Wenhui Qin, and Wenjun Zeng. Adafuse: Adaptive multiview fusion for accurate human pose estimation in the wild. *International Journal of Computer Vision*, 129(3):703–718, 2021.
- [79] Long Zhao, Xi Peng, Yu Tian, Mubbasir Kapadia, and Dimitris N Metaxas. Semantic graph convolutional networks for 3d human pose regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3425–3435, 2019.
- [80] 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/CVF International Conference on Computer Vision*, pages 11656–11665, 2021.
- [81] Kun Zhou, Xiaoguang Han, Nianjuan Jiang, Kui Jia, and Jiangbo Lu. Hemlets pose: Learning part-centric heatmap triplets for accurate 3d human pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2344–2353, 2019.
- [82] Kun Zhou, Xiaoguang Han, Nianjuan Jiang, Kui Jia, and Jiangbo Lu. Hemlets posh: Learning part-centric heatmap triplets for 3d human pose and shape estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [83] Xiaowei Zhou, Menglong Zhu, Spyridon Leonardos, Konstantinos G Derpanis, and Kostas Daniilidis. Sparseness meets deepness: 3d human pose estimation from monocular video. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4966–4975, 2016.
- [84] Wentao Zhu, Xiaoxuan Ma, Zhaoyang Liu, Libin Liu, Wayne Wu, and Yizhou Wang. Motionbert: Unified pretraining for human motion analysis. *arXiv preprint arXiv:2210.06551*, 2022.