

Graph and Temporal Convolutional Networks for 3D Multi-person Pose Estimation in Monocular Videos

Yu Cheng,¹ Bo Wang,² Bo Yang,² Robby T. Tan^{1,3}

¹National University of Singapore,

²Tencent Game AI Research Center,

³Yale-NUS College

e0321276@u.nus.edu, {bohawkwang, brandonyang}@tencent.com, robby.tan@nus.edu.sg

Abstract

Despite the recent progress, 3D multi-person pose estimation from monocular videos is still challenging due to the commonly encountered problem of missing information caused by occlusion, partially out-of-frame target persons, and inaccurate person detection. To tackle this problem, we propose a novel framework integrating graph convolutional networks (GCNs) and temporal convolutional networks (TCNs) to robustly estimate camera-centric multi-person 3D poses that does not require camera parameters. In particular, we introduce a human-joint GCN, which unlike the existing GCN, is based on a directed graph that employs the 2D pose estimator's confidence scores to improve the pose estimation results. We also introduce a human-bone GCN, which models the bone connections and provides more information beyond human joints. The two GCNs work together to estimate the spatial frame-wise 3D poses, and can make use of both visible joint and bone information in the target frame to estimate the occluded or missing human-part information. To further refine the 3D pose estimation, we use our temporal convolutional networks (TCNs) to enforce the temporal and human-dynamics constraints. We use a joint-TCN to estimate person-centric 3D poses across frames, and propose a velocity-TCN to estimate the speed of 3D joints to ensure the consistency of the 3D pose estimation in consecutive frames. Finally, to estimate the 3D human poses for multiple persons, we propose a root-TCN that estimates camera-centric 3D poses without requiring camera parameters. Quantitative and qualitative evaluations demonstrate the effectiveness of the proposed method. Our code and models are available at <https://github.com/3dpose/GnTCN>.

Introduction

Significant progress has been made in 3D human pose estimation in recent years, e.g. (Sun et al. 2019a; Pavlo et al. 2019; Cheng et al. 2019, 2020). In general, existing methods can be classified as either top-down or bottom-up. Top-down approaches use human detection to obtain the bounding box of each person, and then perform pose estimation for every person. Bottom-up approaches are human-detection free and can estimate the poses of all persons simultaneously. Top-down approaches generally demonstrate more superior performance in pose estimation accuracy, and are suitable for

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

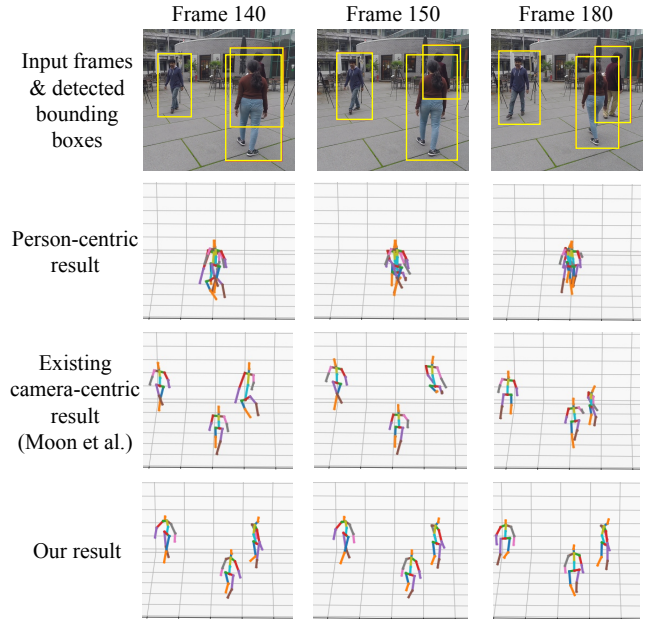


Figure 1: Incorrect 3D multi-person pose estimation caused by person-centric pose estimation or occlusions. The person-centric estimation loses the location of each person in the scene (2nd row) and existing method suffers from missing information due to occlusion (3rd row).

many applications that require high pose estimation precision (Pavlo et al. 2019; Cheng et al. 2020); while bottom-up approaches are better in efficiency (Cao et al. 2017, 2019). In this paper, we aim to further improve 3D pose estimation accuracy, and thus push forward the frontier of the top-down approaches.

Most top-down methods focus on single person and define a 3D pose in a person-centric coordinate system (e.g., pelvis-based origin), which cannot be extended to multiple persons. Since for multiple persons, all the estimated skeletons need to reside in a single common 3D space in correct locations. The major problem here is that by applying the person-centric coordinate system, we lose the location of each person in the scene, and thus we do not know where to put them, as shown in Fig. 1, second row. Another major

problem of multiple persons is the missing information of the target persons, due to occlusion, partially out-of-frame, inaccurate person detection, etc. For instance, inter-person occlusion may confuse human detection (Lin and Lee 2020; Sárárdi et al. 2020), causing erroneous pose estimation (Li et al. 2019; Umer et al. 2020), and thus affect the 3D pose estimation accuracy (as shown in Fig. 1, third row). Addressing these problems is critical for multi-person 3D pose estimation from monocular videos.

In this paper, we exploit the use of the visible human joints and bone information spatially and temporally utilizing GCNs (Graph Convolutional Networks) and TCNs (Temporal Convolutional Networks). Unlike most existing GCNs, which are based on undirected graphs and only consider the connection of joints, we introduce a directed graph that can capture the information of both joints and bones, so that the more reliably estimated joints/bones can influence the unreliable ones caused by occlusions (instead of treating them equally as in undirected graphs). Our human-joint GCN (in short, joint-GCN) employs the 2D pose estimator’s heatmap confidence scores as the weights to construct the graph’s edges, allowing the high-confidence joints to correct low-confidence joints in our 3D pose estimation. While our human-bone GCN (in short, bone-GCN) makes use of the confidence scores of the part affinity field (Cao et al. 2019) to provide complementary information to the joint GCN. The features produced by the joint- and bone-GCNs are concatenated and fed into our fully connected layers to estimate a person-centric 3D human pose.

Our GCNs focus on recovering the spatial information of target persons in a frame-by-frame basis. To increase the accuracy across the input video, we need to put more constraints temporally, both in terms of the smoothness of the motions and the correctness of the whole body dynamics (i.e., human dynamics). To achieve this, we first employ a joint-TCN that takes a sequence of the 3D poses produced by the GCN module as input, and estimate the person-centric 3D pose of the central frame. The joint-TCN imposes a smoothness constraint in its prediction and also imposes the constraints of human dynamics. However, the joint-TCN can estimate only person-centric 3D poses (not camera-centric). Also, the joint-TCN is not robust to occlusion. To resolve the problems, we introduce two new types of TCNs: root-TCN and velocity-TCN.

Relying on the output of the joint-TCN, our root-TCN produces the camera-centric 3D poses, where the X, Y, Z coordinates of the person center, i.e. the pelvis, are in the camera coordinate system. The root-TCN is based on the weak perspective camera model, and does not need to be trained with large variation of camera parameters, since it estimates only the relative depth, Z/f . Our velocity-TCN takes the person-centric 3D poses and the velocity from previous frames as input, and estimates the velocity at the current frame. Our velocity-TCN estimates the current pose based on the previous frames using motion cues. Hence, it is more robust to missing information, such as in the case of occlusion. The reason is because the joint-TCN focuses on the correlations between past and future frames regardless of the trajectory, while the velocity-TCN focuses on the motion

prediction, and thus makes the estimation more robust.

In summary, our contributions are listed as follows.

- Novel directed graph-based joint- and bone-GCNs to estimate 3D poses that can predict human 3D poses even though the information of the target person is incomplete due to occlusion, partially out-of-frame, inaccurate human detection, etc.
- Root-TCN that can estimate the camera-centric 3D poses using the weak perspective projection without requiring camera parameters.
- Combination of velocity- and joint-TCNs that utilize velocity and human dynamics for robust 3D pose estimation.

Related Works

3D human pose estimation in video Recent 3D human pose estimation methods utilize temporal information via recurrent neural network (RNN) (Hossain and Little 2018; Lee, Lee, and Lee 2018; Chiu et al. 2019) or TCN (Pavlo et al. 2019; Cheng et al. 2019; Sun et al. 2019b; Cheng et al. 2020) improve the temporal consistency and show promising results on single-person video datasets such as HumanEva-I, Human3.6M, and MPI-INF-3DHP (Sigal, Balan, and Black 2010; Ionescu et al. 2014; Mehta et al. 2017a), but they still suffer from the inter-person occlusion issue when applying to multi-person videos. Although a few works take occlusion into account (Ghiasi et al. 2014; Charles et al. 2016; Belagiannis and Zisserman 2017; Cheng et al. 2019, 2020), in a top-down framework, it is difficult to reliably estimation 3D multi-person human poses in videos due to erroneous detection and occlusions. Moreover, none of these method estimate camera-centric 3D human poses.

Monocular 3D human pose estimation Earlier approaches that tackle camera-centric 3D human pose from monocular camera require camera parameters as input or assume fixed camera pose to project the 2D posture into camera-centric coordinate (Mehta et al. 2017b, 2019; Pavlo et al. 2019). As a result, these methods are inapplicable for wild videos where camera parameters are not available. Removing the requirement of camera parameters has drawn researcher’s attention recently. Moon et al. (Moon, Chang, and Lee 2019) first propose to learn a correction factor for a person’s root depth estimation from a single image. Several recent works (Li et al. 2020; Lin and Lee 2020; Zhen et al. 2020) show improved performance compared with (Moon, Chang, and Lee 2019). Li et al. (Li et al. 2020) develop an integrated method for detection, person-centric pose, and depth estimation from a single image. Lin et al. (Lin and Lee 2020) propose to formulate the depth regression as a bin index estimation problem. Zhen et al. (Zhen et al. 2020) propose to estimate 2.5D representation of body parts first and then reconstruct 3D human pose. Unlike their approach, our method is video-based where temporal information is utilized by TCN on top of GCN output, which leads to improved 3D pose estimation.

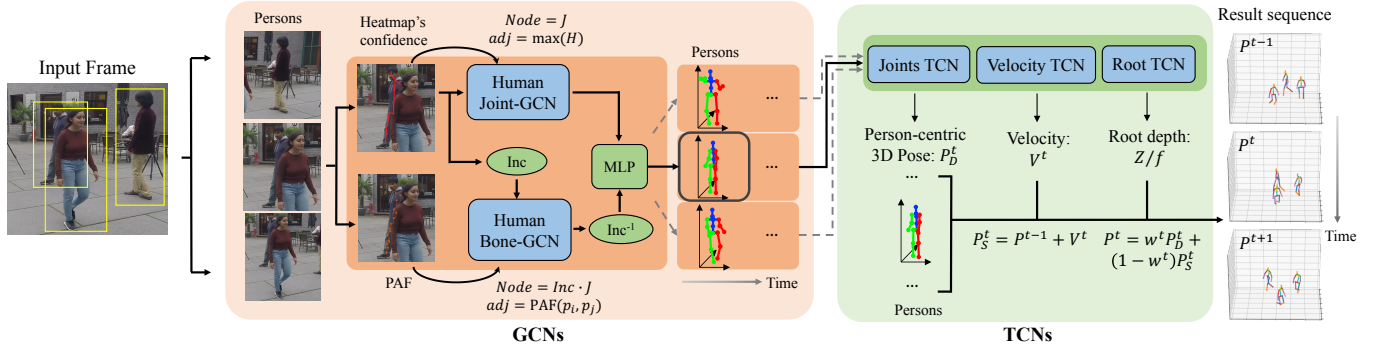


Figure 2: The framework of our approach. The 2D poses and part affinity field for each bounding box are fed into our joint- and bone-GCNs to obtain the full 3D poses (left). After obtaining all poses in the video, they are grouped by IDs which is provided by pose tracker, and fed into the joint-, root- and velocity-TCN to obtain the camera-centric 3D pose estimation (right).

GCN for pose estimation Graph convolutional network (GCN) has been applied to 2D or 3D human pose estimation in recent years (Zhao et al. 2019; Cai et al. 2019; Ci et al. 2019; Qiu et al. 2020). Zhao et al. (Zhao et al. 2019) propose a graph neural network architecture to capture local and global node relationship and apply the proposed GCN for single-person 3D pose estimation from image. Ci et al. (Ci et al. 2019) explore different network structures by comparing fully connected network and GCN and develop a locally connected network to improve the representation capability for single-person 3D human pose estimation from image as well. Cai et al. (Cai et al. 2019) construct an undirected graph to model the spatial-temporal dependencies between different joints for single-person 3D pose estimation from video data. Qiu et al. (Qiu et al. 2020) develop a dynamic GCN framework for multi-person 2D pose estimation from a image. Our method is different from all these methods in terms of we propose to use directed graph to incorporate heatmap and part affinity field confidence in graph construction, which brings the benefit of overcoming the limitation of human detection on top-down pose estimation methods.

Method

The overview of our framework is shown as Fig. 2. Having obtained the 2D poses from the 2D pose estimator, the poses are normalized so that they are centered at the root point, which is at the hip of human body. Each pose is then fed into our joint- and bone-GCNs to obtain its 3D full pose, despite the input 2D pose might be incomplete. Finally, a 3D full pose sequence is fed into the joint-TCN, root-TCN, and velocity-TCN to obtain the camera-centric 3D human poses that have smooth motion and comply with natural human dynamics.

Joint-GCN and Bone GCN

Existing top-down methods are erroneous when the target human bounding box is incorrect, due to missing information (occlusion, partially out-of-frame, blur, etc.). To address this common problem, we introduce joint-GCN and bone-GCN that can correct the 3D poses from inaccurate 2D pose estimator. These GCNs work on a frame-by-frame basis.

Following the structure of the human body, we assign the coordinates (x_i, y_i) of the human joints from the 2D pose estimator to each vertex of our graph, and establish connections between each pair of the joints. Unlike most GCNs, which are based on an undirected graph, we propose a GCN based on a directed graph. The directed graph allows us to propagate information more from high-confident joints to low-confident ones, and thus reduces the risk of propagating erroneous information (e.g., occluded joints or missing joints) in the graph. In other words, the low-confident joints contribute less to the message propagation than the high-confident ones. Details of the directed graph are available in the supplementary material.

The joint-GCN uses the 2D joints as the vertices and the confidence scores of the 2D joints as the edge weights, while the bone-GCN uses the confidence scores of part affinity field (Cao et al. 2017)) as the edge weights. The features produced by the two GCNs are concatenated together and fed to a Multi Layer Perceptron to obtain the person-centric 3D pose estimation.

In GCNs, the message is propagated according to adjacent matrix, which indicates the edge between each pair of vertices. The adjacency matrix is formed by the following rule:

$$A_{i,j} = \{ \max(H_i) e^{-\text{order}(i,j)} (i \neq j) \max(H_i) (i = j) \}, \quad (1)$$

where H is the heatmap from the 2D pose estimator. $\text{order}(i, j)$ stands for the number of the order of neighboring vertices, which means the number of hops required to reach vertex j from vertex i . This formation of adjacency imposes more weight for close vertices and less for distance vertices.

The forward propagation of each GCN layer can be expressed as:

$$h_i = \sigma(F(h_{i-1})W_i), \quad (2)$$

where F is the feature transformation function, and W is the learnable parameter of layer i . To learn a network with strong generalization ability, we follow the idea of Graph

SAGE (Hamilton, Ying, and Leskovec 2017) to learn a generalizable aggregator, which is formulated as:

$$F(h_i) = Ah_i \oplus h_i, \quad (3)$$

where h_i is the output of layer i in the GCN and \oplus stands for the concatenation operation. A is the normalized adjacency matrix. Since our method is based on a directed graph, which uses a non-symmetric adjacency matrix, the normalization is $A_{i,j} = \frac{A_{i,j}}{D_j}$ instead of $A_{i,j} = \frac{A_{i,j}}{\sqrt{D_i D_j}}$ in (Kipf and Welling 2016), D_i and D_j are the indegree of vertices i and j , respectively. This normalization ensures that the indegree of each vertex sums to 1, which prevents numerical instability.

Our joint-GCN considers only human-joints and does not include the information of bones, which can be critical for the cases when the joints are missing due to occlusion or other reasons. To exploit the bone information, we created a bone-GCN. First, we construct the incidence matrix I_n of shape $[\#bones, \#joints]$ to represent the bone connections, where each row represents an edge and the columns represent vertices. For each bone, the parent joint is assigned with -1 and the child joint is assigned with 1 . Second, the incidence matrix I_n is multiplied with the joint matrix J to obtain the bone matrix B , which will be further fed into our bone-GCN.

In joint matrix J , each row stands for the 2D coordinate (x, y) of a joint. Unlike our joint-GCN, where the adjacency matrix is drawn from the joint heatmap produced by 2D pose estimator, our human-bone GCN utilizes the confidence scores from the part affinity field, following the method of (Cao et al. 2017), as the adjacency. Finally, the outputs from our human-joint GCN and human-bone GCN are concatenated together and fed into an MLP (Multi-layer Perceptron). The loss function we use is the $L2$ loss between GCN 3D joints output P_{GCN} and 3D ground-truth skeleton P , which is $L_{GCN} = \|P - P_{GCN}\|_2^2$.

In the training stage, to obtain sufficient variation and to increase the robustness of our GCNs, we use not only the results from our 2D pose estimator, but also augmented data from our ground-truths. Each joint is assigned with a random confidence score and random noise.

Root-TCN

In most of the videos, the projection can be modelled as weak perspective:

$$\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = 1/Z f 0 c_x 0 f c_y 0 0 1 \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}, \quad (4)$$

where x and y are the image coordinates, X, Y and Z are the camera coordinates. f, c_x, c_y stands for the focal length and camera centers, respectively. Thus we have:

$$X = \frac{Z}{f}(x - c_x) \quad Y = \frac{Z}{f}(y - c_y). \quad (5)$$

By assuming (c_x, c_y) as the image center, which is applicable for most cameras, the only parameters we need to estimate is depth Z and focal length f . To be more practical, we

jointly estimate Z/f , instead of estimating them separately. This enables our method to be able to take wild videos that the camera parameters are unknown.

According to the weak perspective assumption, the scale of a person in a frame indicates the depth in the camera coordinates. Hence, we propose a network, root Temporal Convolutional Network (root-TCN), to estimate the Z/f from 2D pose sequences. We first normalize each 2D pose by scaling the average joint-to-pelvis distance to 1, using a scale factor s . Then we concatenate the normalized pose p , scale factor s , as well as the person’s center in the frame as c , and feed a list of such concatenated features in a local temporal window into the TCN for depth estimation in the camera coordinates.

As directly learning Z/f is not easy to converge, we transform this regression problem into a classification problem. For each video, we divide the depth into N discrete ranges, set to 60 in our experiments, and our root-TCN outputs a vector with length N as $\{x_1, \dots, x_N\}$, where x_i indicates the probability that Z/f is within the i th discrete range. Then, we apply *Soft-argmax* to this vector to get the final continuous estimation of the depth as:

$$\left[\frac{Z}{f}\right]^t = \text{Soft-argmax} (f_R(p^{t-n:t+n}, c^{t-n:t+n}, s^{t-n:t+n})), \quad (6)$$

where t is the time stamp, and n is half of the temporal window’s size. This improves the training stability and reduces the risk of large errors.

The loss function for the depth estimation is defined as the mean squared error between the ground truth and predictions, expressed as $L_{Root} = (\frac{Z}{f} - \frac{\hat{Z}}{\hat{f}})^2$, where Z/f is the predicted value, and \hat{Z}/\hat{f} denotes the ground truth. According to Eq.(5), we can calculate the coordinates for the person’s center as P_D^t .

Joint-TCN and Velocity-TCN

To increase the accuracy of the 3D poses across the input video, we impose temporal constraints, by employing a temporal convolutional network (TCN) (Cheng et al. 2020) that takes a sequence of consecutive 3D poses as input. We call this TCN a joint-TCN, which is trained using various 3D poses and their augmentation, and hence capture human dynamics. The joint-TCN outputs the person-centric 3D pose, P_D . The TCN utilizes temporal information to interpolate the poses of occluded frames with temporal information.

However, when persons get close and occlude each other, there may be fewer visible joints belonging to a person and more distracting joints from other persons. To resolve the problem, in addition to the joint-TCN, we propose a velocity-based estimation network, velocity-TCN, which takes the 3D joints and their velocities as input, and predicts the velocity of all joints as:

$$V^t = (v_x^t, v_y^t, v_z^t) = \text{TCN}_v(p^{t-n:t-1}, V^{t-n:t-1}), \quad (7)$$

where p stands for the 2D pose and V^t denotes the velocity at time t . TCN_v is the velocity-TCN. The velocity here

is proportional to $1/f$ according to Eq. (5). We normalize the velocity both in training and testing. With estimated V^t , we can obtain the coordinate $P_S^t = P^{t-1} + V^t$, where P_S^t and P^{t-1} are estimated coordinates at time t and $t - 1$. The calculation of P^{t-1} is discussed later in Eq.(8).

The joint-TCN predicts the joints by interpolating the past and future poses, while our velocity-TCN predicts the future poses using motion cues. Both of them are able to handle the occlusion frames, but the joint-TCN focuses on the connection between past and future frames regardless of the trajectory, while the velocity-TCN focuses on the motion prediction, which can handle a motion drift. To leverage the benefits of both, we introduce an adaptive weighted average of their estimated coordinates P_D^t and P_S^t .

We utilize the 2D pose tracker (Umer et al. 2020) to detect and track human poses in the video. In the tracking procedure, we regard the heatmaps with less than 0.5 confidence value as occluded joints, and the pose with less than 30% non-occluded joints as the occluded pose. By doing this, we obtain the occlusion information for both joints and frames. Note that the values here are obtained empirically through our experiments. Suppose we find an occlusion duration from T_{occ}^{start} to T_{occ}^{end} for a person, then we generate the final coordinates as:

$$P^t = w^t P_D^t + (1 - w^t) P_S^t, \quad (8)$$

where $w_v^t = e^{-\min(t-T_{occ}^{start}, T_{occ}^{end}-t)}$. For frames that are closer to occlusion duration boundaries, we trust P_D^t more; for those are far from occlusion boundaries, we trust P_S^t more. The velocity-TCN loss is the $L2$ loss between the predicted 3D points at t and ground-truth 3D points.

Experiments

MuPoTS-3D is a 3D multi-person testing set with both indoor and outdoor scenes (Mehta et al. 2018). The ground-truth 3D pose of each person in the video is obtained from multi-view markerless capture, which is suitable for evaluating 3D multi-person pose estimation performance in both person-centric and camera-centric coordinates. Unlike previous methods (Moon, Chang, and Lee 2019) using the training set (MuCo-3DHP) to train their models and then do evaluation on MuPoTS-3D, we use MuPoTS-3D for testing only without fine-tuning.

3DPW is an outdoor multi-person dataset for 3D human pose reconstruction (von Marcard et al. 2018). Following previous methods (Kanazawa et al. 2019; Sun et al. 2019b), we use 3DPW for testing only without any fine-tuning. The ground-truth of 3DPW is SMPL 3D mesh model (Loper et al. 2015), where the definition of joints differs from the one commonly used in 3D human pose estimation (skeleton-based) like Human3.6M (Tripathi et al. 2020), so it is unfair to evaluate skeleton-based methods on it even after joint adaption or scaling. To perform a fair comparison, we select an occlusion subset from the 3DPW test set (please refer to the supplementary material for details). And the performance change of a method between the full test set and the subset indicates how well the method can handle the missing information problem caused by occlusions.

Method	AP_{25}^{root}	AUC_{rel}	PCK	PCK_{abs}
Baseline	24.1	32.9	74.4	29.8
Baseline (GT box)	28.5	34.2	78.9	31.2
Baseline + GCNs	35.4	39.7	83.2	35.1
Baseline + TCNs	38.4	43.1	85.3	38.7
Full model	45.2	48.9	87.5	45.7

Table 1: Ablation study on MuPoTS-3D dataset. Best in bold, second best underlined.

Method	AP_{25}^{root}	AUC_{rel}	PCK	PCK_{abs}
Joint* GCN	24.1	27.3	73.1	25.6
Joint GCN	28.5	30.1	76.8	29.0
Joint + Bone* GCN	28.4	31.9	78.1	29.7
Joint + Bone GCN	33.4	37.9	82.6	34.3
Joint + Bone + Aug.	35.4	39.7	83.2	35.1
Joint TCN	43.1	45.8	86.2	42.6
Joint + Velocity	45.2	48.9	87.5	45.7

Table 2: Ablation study on our proposed Joint and Bone GCNs and TCNs. * stands for the GCN structure with undirected graph. We keep the GCN as the best one (joint + bone + aug.) to perform an ablation study on TCN.

Human3.6M is a widely used dataset and benchmark for 3D human pose estimation (Ionescu et al. 2014). It contains 3.6 million single-person indoor images captured by the MoCap system, which is suitable for evaluation of single-person pose estimation and camera-centric coordinates prediction. Following previous works (Hossain and Little 2018; Pavllo et al. 2019; Wandt and Rosenhahn 2019), the subject 1,5,6,7,8 are used for training, and 9 and 11 for testing.

Evaluation and Implementation MPJPE, PA-MPJPE, PCK, and AUC_{rel} are used for person-centric pose estimation evaluation. AP_{25}^{root} and PCK_{abs} are used for camera-centric pose estimation evaluation. Each GCN and TCN is trained for 100 epochs with initial learning rate $1e-3$, more details are available in the supplementary material.

Ablation Studies In Table 1, we provide the results of an ablation study to validate the major components of the proposed framework. MuPoTS-3D is used as it has been used for 3D multi-person pose evaluation in person-centric and camera-centric coordinates (Moon, Chang, and Lee 2019). AUC_{rel} and PCK metrics are used to evaluate person-centric 3D pose estimation performance, AP_{25}^{root} and PCK_{abs} metrics are used to evaluate camera-centric 3D pose (i.e., camera-centric coordinate) estimation following (Moon, Chang, and Lee 2019).

In particular, we use the joint-TCN with time window 1 plus a root-TCN with time window 1 as a baseline for both person-centric and camera-centric coordinate estimation as shown at the 1st row in Table 1. We use the baseline with ground-truth bounding box (i.e., perfect 2D tracking) as a second baseline to illustrate even with perfect detection bounding box the baseline still performs poorly because it cannot deal with occlusion and distracting joints from other persons. On the contrary, we can see significant performance (e.g., 18% \sim 29% improvement against the baseline in

Group	Method	PCK	PCK _{abs}
Person-centric	Mehta et al. (2018)	65.0	n/a
	Rogez et al., (2019)	70.6	n/a
	Cheng et al. (2019)	74.6	n/a
	Cheng et al. (2020)	80.5	n/a
Camera-centric	Moon et al. (2019)	82.5	31.8
	Lin et al. (2020)	<u>83.7</u>	35.2
	Zhen et al. (2020)	80.5	38.7
	Li et al. (2020)	82.0	<u>43.8</u>
	Our method	87.5	45.7

Table 3: Quantitative evaluation on multi-person 3D dataset, MuPoTS-3D. Best in bold, second best underlined.

PCK_{abs}) improvements after adding the proposed GCN and TCN modules as shown in row 3 - 5 in Table 1. The benefits from the TCNs are slightly larger than those from the joint and bone GCNs as temporal information is used by the TCNs while GCNs only use frame-wise information. Lastly, our full model shows the best performance, with 53% improvement against the baseline in terms of PCK_{abs}.

We perform a second ablation study to break down the different pieces in our GCN and TCN modules to show the effectiveness of each individual component in Table 2. We observe the undirected graph-based GCN performs the worst as shown in the 1st row. Our joint-GCN, bone-GCN, and data augmentation (applying random cropping and scaling) show improved performance in row 2 - 4. On top of the full GCN module, we also show the contribution of the joint-TCN and velocity-TCN in row 5 - 6 (row 6 is our full model). Similar to Table 1 we can see the TCN module brings more improvement compared with the GCN module as temporal information is used.

Quantitative Results To compare with the state-of-the-art (SOTA) methods in 3D multi-person human pose estimation, we perform evaluations on MuPoTS-3D as shown in Table 3. Please note our network is trained only on Human3.6M to have a fair comparison with other methods (Cheng et al. 2019, 2020). As the definition of keypoints in MuPoTS-3D is different from the one where our model is trained on, we use joint adaptation (Tripathi et al. 2020) to transfer the definition of keypoints. Among the methods in Table 3, (Moon, Chang, and Lee 2019; Lin and Lee 2020; Li et al. 2020) are fine-tuned on 3D training set MuCo-3DHP.

Regarding to the performance on MuPoTS-3D, our camera-centric pose estimation accuracy beat the SOTA (Li et al. 2020) by 4.3% on PCK_{abs}. A few papers reported their results on AP_{25}^{root} , where (Moon, Chang, and Lee 2019) is 31.0, (Lin and Lee 2020) is 39.4, and our result is 45.2, where we beat the SOTA (Lin and Lee 2020) by 14.7%. We also compare with other methods on person-centric 3D pose estimation, and get improvement of 4.5% on PCK against the SOTA (Lin and Lee 2020). Please note we do not fine-tune on MuCo-3DHP like others (Moon, Chang, and Lee 2019; Lin and Lee 2020; Li et al. 2020), which is the training dataset for evaluation on MuPoTS-3D. Moreover, the SOTA method (Lin and Lee 2020) on person-centric metric PCK shows poor performance on PCK_{abs} (ours vs theirs:

Dataset	Method	PA-MPJPE	δ
Original	Dabral et al. (2018)	92.2	n/a
	Doersch et al. (2019)	74.7	n/a
	Kanazawa et al. (2019)	72.6	n/a
	Cheng et al. (2020)	71.8	n/a
	Sun et al. (2019b)	69.5	n/a
	Kolotouros et al. (2019)*	<u>59.2</u>	n/a
	Kocabas et al., (2020)*	51.9	n/a
	Our method	64.2	n/a
Subset	Cheng et al. (2020)	96.1	+24.1
	Sun et al. (2019b)	94.1	+24.6
	Kolotouros et al. (2019)*	88.9	+29.7
	Kocabas et al., (2020)*	82.5	+30.6
	Our method	<u>85.7</u>	+21.5

Table 4: Quantitative evaluation using PA-MPJPE in millimeter on original 3DPW test set and its occlusion subset. * denotes extra 3D datasets were used in training. Best in bold, second best underlined.

45.7 - 35.2, 29.8% improvement), and the SOTA method (Li et al. 2020) on camera-centric metric PCK_{abs} has mediocre performance on PCK (ours vs theirs: 87.5 vs 82.0, 6.7% improvement). All of these results clearly show that our method not only surpasses all existing methods, but also is the only method that is well-balanced in both person-centric and camera-centric 3D multi-person pose estimation.

3DPW dataset (von Marcard et al. 2018) is a new 3D multi-person human pose dataset that contains multi-person outdoor scenes for person-centric pose estimation evaluation. Following previous works (Kanazawa et al. 2019; Sun et al. 2019b), 3DPW is only used for testing and the PA-MPJPE values on test set are shown in Table 4. As discussed in the Datasets section, the ground-truth definitions are different between 3D pose reconstruction and estimation where the ground-truth of 3DPW is SMPL mesh model, even we follow (Tripathi et al. 2020) to perform joint adaptation to transform the estimated joints but still have a disadvantage, and the PA-MPJPE values cannot objectively reflect the performance of skeleton-based pose estimation methods.

As aforementioned, we select a subset out of the original test set with the largest detection errors, and run the code of the top-performing methods in Table 4 on this subset for comparison. Table 4 shows that even with the disadvantage of different definition of joints, our method is the 3rd best on the original testing test, and becomes the 2nd best on the subset where the difference to the best one (Kocabas, Athanasiou, and Black 2020) is greatly shrunk. More importantly, the δ of PA-MPJPE between the original testing set and the subset in the 4th column in Table 4, our method shows the least error increase compared with all other top-performing methods. In the particular, the two best methods (Kolotouros et al. 2019; Kocabas, Athanasiou, and Black 2020) on the original testing set show 29.7 and 30.6 mm error increase while our method shows only 21.5 mm. The performance change of PA-MPJPE between the original testing set and the subset clearly demonstrates that our method is the best in terms of solving the missing information problem which is critical for 3D multi-person pose estimation.

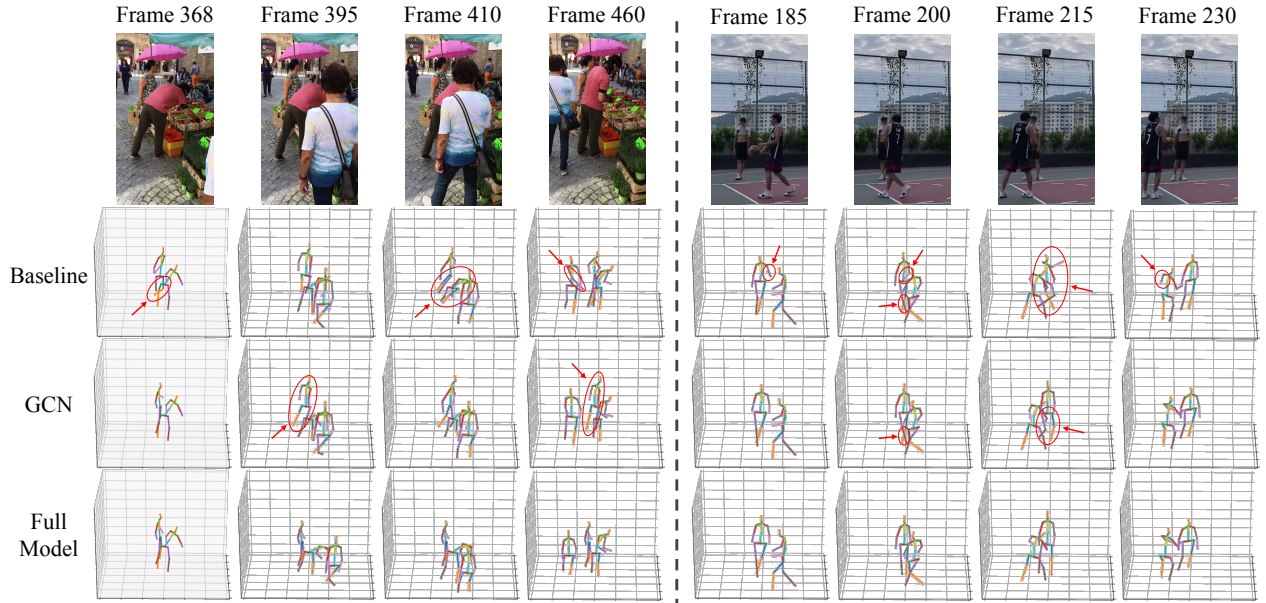


Figure 3: Examples of results from our whole framework compared with different baseline results. First row shows the images from two video clips; second row shows the results from the baseline described in Ablation Studies; third row shows the result of the GCN module; last row shows the results of the whole framework. Wrong estimations are labeled with red circles.

Group	Method	MPJPE	PA-MPJPE
Person-centric	Hossain et al., (2018)	51.9	42.0
	Wandt et al., (2019)*	50.9	38.2
	Pavlo et al., (2019)	46.8	36.5
	Cheng et al., (2019)	42.9	32.8
	Kocabas et al., (2020)	65.6	41.4
	Kolotouros et al. (2019)	<u>41.1</u>	n/a
Camera-centric	Moon et al., (2019)	54.4	35.2
	Zhen et al., (2020)	54.1	n/a
	Li et al., (2020)	48.6	<u>30.5</u>
	Ours	40.9	30.4

Table 5: Quantitative evaluation on Human3.6M for normalized and camera-centric 3D human pose estimation. * denotes ground-truth 2D labels are used. Best in bold, second best underlined.

In order to further illustrate the effectiveness of both person-centric and camera-centric 3D pose estimation of our method, we perform evaluations on the widely used single-person dataset, Human3.6M. To evaluate camera-centric pose estimation, we use mean root position error (MPRE), a new evaluation metric proposed by (Moon, Chang, and Lee 2019). Our result is 88.1 mm, the result of (Moon, Chang, and Lee 2019) is 120 mm, the result of (Lin and Lee 2020) is 77.6 mm. Our method outperforms the result of (Moon, Chang, and Lee 2019) by a large margin: 31.9 mm error reduction and 26% improvement. Although depth estimation focused method (Lin and Lee 2020) shows better camera-centric performance on this single-person dataset Human3.6M, their camera-centric result on multi-person dataset MuPoTS-3D is much worse than ours (ours vs. theirs

in PCK_{abs} : 45.7 - 35.2, 29.8% improvement). Camera-centric 3D human pose estimation is for multi-person pose estimation, good performance only on single-person dataset is not enough to solve the problem.

To compare with most of the existing methods that evaluate person-centric 3D pose estimation on Human3.6M using MPJPE and PA-MPJPE, we report our results using the same metrics in Table 5. As Human3.6M contains only single-person videos, we do not expect to our method bring much improvement. It is observed that our method is comparable with the SOTA methods. In addition, although our method shows improved performance over others that use kinematic constraints (Wandt and Rosenhahn 2019; Cheng et al. 2019) because of our GCNs and TCNs, adding kinematic constraints could potentially improve our performance further (Akhter and Black 2015; Kundu et al. 2020).

Qualitative Results As shown in Figure 3, our full model can better handle occlusions and incorrect detection compared with the baselines and the relative positions among all persons are well captured without camera parameters. More comparisons against SOTA methods and qualitative results on wild videos are available in the supplementary material.

Conclusion

We propose a new framework to unify GCNs and TCNs for camera-centric 3D multi-person pose estimation. The proposed method successfully handles missing information due to occlusion, out-of-frame, inaccurate detections, etc., in videos and produces continuous pose sequences. Experiments on different datasets validate the effectiveness of our framework as well as our individual modules.

Acknowledgements

This research/project is supported by the National Research Foundation, Singapore under its Strategic Capability Research Centres Funding Initiative. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

References

- Akhter, I.; and Black, M. J. 2015. Pose-conditioned joint angle limits for 3D human pose reconstruction. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1446–1455.
- Belagiannis, V.; and Zisserman, A. 2017. Recurrent Human Pose Estimation. In *International Conference on Automatic Face and Gesture Recognition (FG)*. IEEE.
- Cai, Y.; Ge, L.; Liu, J.; Cai, J.; Cham, T.-J.; Yuan, J.; and Thalmann, N. M. 2019. Exploiting spatial-temporal relationships for 3d pose estimation via graph convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, 2272–2281.
- Cao, Z.; Hidalgo, G.; Simon, T.; Wei, S.-E.; and Sheikh, Y. 2019. OpenPose: Realtime Multi-Person 2D Pose Estimation using Part Affinity Fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence* doi:10.1109/TPAMI.2019.2929257.
- Cao, Z.; Simon, T.; Wei, S.-E.; and Sheikh, Y. 2017. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7291–7299.
- Charles, J.; Pfister, T.; Magee, D.; Hogg, D.; and Zisserman, A. 2016. Personalizing Human Video Pose Estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Cheng, Y.; Yang, B.; Wang, B.; and Tan, R. T. 2020. 3D Human Pose Estimation using Spatio-Temporal Networks with Explicit Occlusion Training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 8118–8125.
- Cheng, Y.; Yang, B.; Wang, B.; Yan, W.; and Tan, R. T. 2019. Occlusion-Aware Networks for 3D Human Pose Estimation in Video. In *Proceedings of the IEEE International Conference on Computer Vision*, 723–732.
- Chiu, H.-k.; Adeli, E.; Wang, B.; Huang, D.-A.; and Niebles, J. C. 2019. Action-agnostic human pose forecasting. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 1423–1432.
- Ci, H.; Wang, C.; Ma, X.; and Wang, Y. 2019. Optimizing network structure for 3d human pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, 2262–2271.
- Dabral, R.; Mundhada, A.; Kusupati, U.; Afaq, S.; Sharma, A.; and Jain, A. 2018. Learning 3d human pose from structure and motion. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 668–683.
- Doersch, C.; and Zisserman, A. 2019. Sim2real transfer learning for 3D human pose estimation: motion to the rescue. In *Advances in Neural Information Processing Systems*, 12929–12941.
- Ghiasi, G.; Yang, Y.; Ramanan, D.; and Fowlkes, C. C. 2014. Parsing Occluded People. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Hamilton, W.; Ying, Z.; and Leskovec, J. 2017. Inductive representation learning on large graphs. In *Advances in neural information processing systems*, 1024–1034.
- Hossain, M. R. I.; and Little, J. J. 2018. Exploiting temporal information for 3D human pose estimation. In *ECCV*, 69–86. Springer.
- Ionescu, C.; Papava, D.; Olaru, V.; and Sminchisescu, C. 2014. Human3.6M: Large Scale Datasets and Predictive Methods for 3D Human Sensing in Natural Environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* 36(7): 1325–1339.
- Kanazawa, A.; Zhang, J. Y.; Felsen, P.; and Malik, J. 2019. Learning 3D Human Dynamics From Video. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Kipf, T. N.; and Welling, M. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Kocabas, M.; Athanasiou, N.; and Black, M. J. 2020. VIBE: Video inference for human body pose and shape estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5253–5263.
- Kolotouros, N.; Pavlakos, G.; Black, M. J.; and Daniilidis, K. 2019. Learning to Reconstruct 3D Human Pose and Shape via Model-fitting in the Loop. In *Proceedings of the IEEE International Conference on Computer Vision*, 2252–2261.
- Kundu, J. N.; Seth, S.; Rahul, M.; Rakesh, M.; Radhakrishnan, V. B.; and Chakraborty, A. 2020. Kinematic-Structure-Preserved Representation for Unsupervised 3D Human Pose Estimation. In *AAAI*, 11312–11319.
- Lee, K.; Lee, I.; and Lee, S. 2018. Propagating lstm: 3d pose estimation based on joint interdependency. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 119–135.
- Li, J.; Wang, C.; Liu, W.; Qian, C.; and Lu, C. 2020. HMOR: Hierarchical Multi-Person Ordinal Relations for Monocular Multi-Person 3D Pose Estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Li, J.; Wang, C.; Zhu, H.; Mao, Y.; Fang, H.-S.; and Lu, C. 2019. Crowdpose: Efficient crowded scenes pose estimation and a new benchmark. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 10863–10872.
- Lin, J.; and Lee, G. H. 2020. HDNet: Human Depth Estimation for Multi-Person Camera-Space Localization. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Loper, M.; Mahmood, N.; Romero, J.; Pons-Moll, G.; and Black, M. J. 2015. SMPL: A skinned multi-person linear model. *ACM transactions on graphics (TOG)* 34(6): 248.
- Mehta, D.; Rhodin, H.; Casas, D.; Fua, P.; Sotnychenko, O.; Xu, W.; and Theobalt, C. 2017a. Monocular 3d human pose estimation in the wild using improved cnn supervision. In *2017 International Conference on 3D Vision (3DV)*, 506–516. IEEE.
- Mehta, D.; Sotnychenko, O.; Mueller, F.; Xu, W.; Elgharib, M.; Fua, P.; Seidel, H.-P.; Rhodin, H.; Pons-Moll, G.; and Theobalt, C. 2019. XNect: Real-time Multi-person 3D Human Pose Estimation with a Single RGB Camera. *arXiv preprint arXiv:1907.00837*.
- Mehta, D.; Sotnychenko, O.; Mueller, F.; Xu, W.; Sridhar, S.; Pons-Moll, G.; and Theobalt, C. 2018. Single-shot multi-person 3d pose estimation from monocular rgb. In *2018 International Conference on 3D Vision (3DV)*, 120–130. IEEE.

- Mehta, D.; Sridhar, S.; Sotnychenko, O.; Rhodin, H.; Shafiei, M.; Seidel, H.-P.; Xu, W.; Casas, D.; and Theobalt, C. 2017b. Vnect: Real-time 3d human pose estimation with a single rgb camera. *ACM Transactions on Graphics (TOG)* 36(4): 44.
- Moon, G.; Chang, J.; and Lee, K. M. 2019. Camera Distance-aware Top-down Approach for 3D Multi-person Pose Estimation from a Single RGB Image. In *The IEEE Conference on International Conference on Computer Vision (ICCV)*.
- Pavlo, D.; Feichtenhofer, C.; Grangier, D.; and Auli, M. 2019. 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*, 7753–7762.
- Qiu, Z.; Qiu, K.; Fu, J.; and Fu, D. 2020. DGCN: Dynamic Graph Convolutional Network for Efficient Multi-Person Pose Estimation. In *AAAI*, 11924–11931.
- Rogez, G.; Weinzaepfel, P.; and Schmid, C. 2019. Lcr-net++: Multi-person 2d and 3d pose detection in natural images. *IEEE transactions on pattern analysis and machine intelligence*.
- Sáráandi, I.; Linder, T.; Arras, K. O.; and Leibe, B. 2020. MeTRAbs: Metric-Scale Truncation-Robust Heatmaps for Absolute 3D Human Pose Estimation. *arXiv preprint arXiv:2007.07227*.
- Sigal, L.; Balan, A. O.; and Black, M. J. 2010. 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.
- Sun, K.; Xiao, B.; Liu, D.; and Wang, J. 2019a. Deep High-Resolution Representation Learning for Human Pose Estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Sun, Y.; Ye, Y.; Liu, W.; Gao, W.; Fu, Y.; and Mei, T. 2019b. Human mesh recovery from monocular images via a skeleton-disentangled representation. In *Proceedings of the IEEE International Conference on Computer Vision*, 5349–5358.
- Tripathi, S.; Ranade, S.; Tyagi, A.; and Agrawal, A. 2020. PoseNet3D: Unsupervised 3D Human Shape and Pose Estimation. *arXiv preprint arXiv:2003.03473*.
- Umer, R.; Doering, A.; Leibe, B.; and Gall, J. 2020. Self-supervised Keypoint Correspondences for Multi-Person Pose Estimation and Tracking in Videos. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- von Marcard, T.; Henschel, R.; Black, M. J.; Rosenhahn, B.; and Pons-Moll, G. 2018. Recovering Accurate 3D Human Pose in The Wild Using IMUs and a Moving Camera. In *European Conference on Computer Vision (ECCV)*, 614–631.
- Wandt, B.; and Rosenhahn, B. 2019. RepNet: Weakly Supervised Training of an Adversarial Reprojection Network for 3D Human Pose Estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Zhao, L.; Peng, X.; Tian, Y.; Kapadia, M.; and Metaxas, D. N. 2019. Semantic graph convolutional networks for 3D human pose regression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3425–3435.
- Zhen, J.; Fang, Q.; Sun, J.; Liu, W.; Jiang, W.; Bao, H.; and Zhou, X. 2020. SMAP: Single-Shot Multi-Person Absolute 3D Pose Estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*.