

Part-aware Measurement for Robust Multi-View Multi-Human 3D Pose Estimation and Tracking

Hau Chu^{1,3} Jia-Hong Lee^{1,3} Yao-Chih Lee^{2,5} Ching-Hsien Hsu^{2,6}
Jia-Da Li^{1,3} Chu-Song Chen^{1,4}
¹National Taiwan University ²Academia Sinica, Taiwan
³{hauchu8733, jiahonghenrylee, jiadali}@ntu.edu.tw, ⁴chusong@csie.ntu.edu.tw
⁵ycllee1231@iis.sinica.edu.tw, ⁶applychristw@gmail.com

Abstract

This paper introduces an approach for multi-human 3D pose estimation and tracking based on calibrated multi-view. The main challenge lies in finding the cross-view and temporal correspondences correctly even when several human pose estimations are noisy. Compare to previous solutions that construct 3D poses from multiple views, our approach takes advantage of temporal consistency to match the 2D poses estimated with previously constructed 3D skeletons in every view. Therefore cross-view and temporal associations are accomplished simultaneously. Since the performance suffers from mistaken association and noisy predictions, we design two strategies for aiming better correspondences and 3D reconstruction. Specifically, we propose a part-aware measurement for 2D-3D association and a filter that can cope with 2D outliers during reconstruction. Our approach is efficient and effective comparing to state-of-the-art methods; it achieves competitive results on two benchmarks: 96.8% on Campus and 97.4% on Shelf. Moreover, we extend the length of Campus evaluation frames to be more challenging and our proposal also reach well-performed result. The code will be available at <https://git.io/JO4KE>.

1. Introduction

Multi-human 3D pose estimation and tracking based on multi-view streaming videos have many applications, including marker-less motion capture [20, 12, 35, 25], sports analysis [4], and video surveillance [7]. Recently, plenty of nicely performed 2D pose estimation approaches [8, 37, 6] have been developed; they have then been extended to the estimation of 3D poses from a monocular view [14, 16]. Although great progress has been made on inferring the 3D human poses in a single view, the noisy predictions caused by large pose variations and partial occlusions remain to be

demanding. To problems, address these issues, constructing 3D human poses from multi-camera views becomes a promising way. However, there are still several challenges to be tackled, such as heavy occlusions, high computational complexity, and noisy predictions.

In general, given synced multi-view video input, three main tasks should be dealing with properly: *2D skeleton extraction*, *cross-view association*, and *temporal association*. Human skeletons are estimated for each view initially, which is often achieved via a 2D human pose estimation method with convolution neural networks (CNNs) [8, 24, 37]. The 3D skeletons are then reconstructed according to the associated 2D human skeletons in different views. In the last, associations between the 2D or 3D skeletons are established with those in the next frame in video streams.

Leveraging 2D poses, recent studies [4, 32, 11, 7] follow an initialization-and-tracking framework for 3D pose inference. The framework assumes a streaming mode, i.e., the outputs of multi-view 3D poses are obtained on-line per input frame, and the previously generated outputs cannot be altered. To perform on-line 3D pose inference, initial 3D human skeletons are computed given the first views. This is often achieved via epi-polar-line distance and/or person re-id matching; then, the initial 3D joints are established with multi-view stereo. As for maintaining the flexibility of use, most works [32, 4, 7, 39] do not need fine-tuning on the data collected in the testing environment; pre-trained person re-id models are thus not effective enough due to the testing domain shift in the usual. Hence, epi-polar-line distances are mostly used for initializing the 3D poses. Once initialized, the obtained 3D skeletons of the individuals serve as a guide for human tracking based on the 2D poses of all views detected in future frames. The 3D human skeletons are then updated via multi-view reconstruction according to the new correspondences of the 2D human skeletons obtained when tracking is finished.

In the processing flow of on-line 3D pose inference, tem-

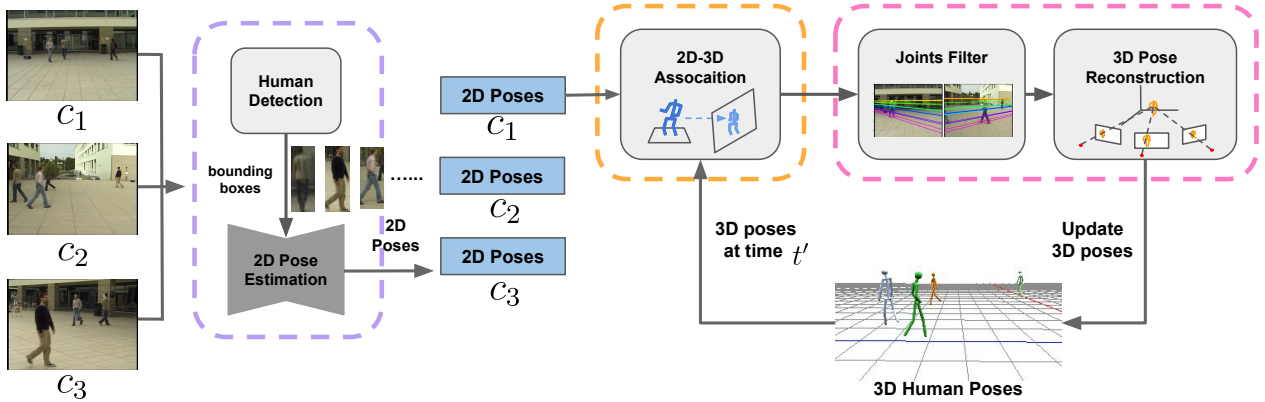


Figure 1: Overview of our 3D pose tracking mechanism. First, initial 3D poses are given via the reconstruction with cross-view 2D correspondences. Then, given new frames as shown in the figure, we apply 2D-3D association to match estimated 2D poses with previously tracked 3D skeletons. While finishing the association of all views, for each matched 2D poses and 3D skeleton, *joints filter* is employed to remove 2D outliers of estimated body joints and update the 3D pose.

poral association plays a main role. Previous approaches such as [11, 32, 36] construct the 3D skeletons from multiple current views at first, and then the 3D skeletons obtained are smoothed temporally for each individual. However, since matching of C views involves $O(C^2)$ pairs of the epi-polar-line distance evaluation, the computation cost could be considerably increased with the number of people and cameras. To alleviate the computation burden, we introduce a 2D-3D matching mechanism to enforce the temporal consistency, which leverages the previously tracked 3D skeletons for finding the correspondences among views. The complexity is $O(C)$ as only the projections of the previous 3D skeletons on the C views is needed. Besides, multi-person interactions in crowded scenes would increase the difficulty of matching and identifying people as individuals across views. To exploit the projected 3D skeletons for finding better correspondences among views, we propose a part-aware measurement to compute the affinity. We also design a greedy algorithm called *joints filter* to remove noisy points and construct a robust 3D pose by taking advantage of epi-polar constraints. Unlike previous works such as [7], the proposed method can handle wrong key-point estimation caused by occlusion or motion blur during association. We compare our method with previous works and show that our solution achieves competitive results on two benchmark datasets: Campus and Shelf.

In the following contents, Section 2 reviews the related approaches of 3D human pose estimation and tracking. Section 3 presents the details of our new approach. Section 4 presents the results comparing the accuracy and efficiency between our approaches and competitive methods on the Campus [3] and Shelf [3] datasets. We also demonstrate the design analysis of our approach in ablation studies. Finally, conclusions are given in Section 5.

2. Related Work

This section reviews the previous approaches to 3D human pose estimation and their tracking techniques. In Section 2.1, we review the methods of 3D human pose estimation, which infer the 3D poses based on isolated images. In Section 2.2, we review the techniques of 3D human pose tracking, where temporal consistency in 3D space is utilized for performance enhancement.

2.1. 3D Human Pose Estimation

Depending on the number of input cameras, 3D human pose estimation methods are divided into a monocular camera for taking single-view video [2, 23, 31, 14, 21, 10, 22, 16, 38] and multiple cameras for taking multi-view videos synchronously [3, 13, 4, 32, 11, 26, 7, 36, 39, 35].

Due to the difficulty of multi-person 3D poses reconstruction in monocular view, most of the single-view approaches are developed to construct a single person's 3D poses [23, 31, 10, 38], where the predicted pose does not include absolute locations in the environment. Therefore, it will limit these approaches applying in different practical surveillance scenarios. Despite a great achievement of multi-human 3D pose estimation in a single view [2, 22, 14], there is still a large deviation when applying these techniques in different practical surveillance scenarios. In particular, the motion blur and occlusions occur in images.

To retrieve absolute location and handle occlusions, the studies of multi-view 3D pose estimation attract more attention recently. It can be applied in various applications, such as sports analysis, video surveillance, animation, and healthcare. Most previous approaches [3, 25, 27] for single-person 3D pose estimation are developed based on the 3D Pictorial Structure model (3DPS) [5], which discretizes the

3D-space as a grid and assigns each joint to one assumed location in the grid. Depending on the cross-view association, 3D pose reconstruction can be solved by minimizing the geometric error [1] between assumed 3D poses and the 3D poses generated from multi-view 2D images. Since CNN has a great performance in the human detector and 2D human pose estimation, most multi-person 3D pose estimation approaches [13, 11, 36] utilize sophisticated CNN-based 2D human pose estimation techniques at the beginning. It detects the human in the image and applies 2D pose estimation in the cropped image to obtain feature maps of present humans in these multi-view images. Feature maps from different views will serve various fusion, matching, or clustering techniques to achieve multi-person 3D poses estimation.

Ershadi-N. *et al.* [13] utilize DeeperCut [15] to detect the human body’s 2D part poses in each image for constructing the initial 3D joints of human body poses in 3D state space. They develop a clustering algorithm to separate all 3D candidate joints into multiple individual 3D human poses before refining these poses using a fully connected pairwise conditional random field (CRF). Dong *et al.* [11] utilize Faster R-CNN [29] with lightweight backbone network and Cascaded Pyramid Network [8] to detect humans’ location and their 2D poses in multi-view images. They develop a multi-way matching method with circle consistency to identify the same person across cameras via human affinity before constructing 3D human poses. The human affinity is calculated from CamStyle [40] generated features and geometric compatibility. Tu *et al.* [36] utilize HRNet [37] to obtain 2D pose heatmaps in each camera view. To fuse projected 2D pose heatmaps into 3D space, they develop a Cuboid Proposal Network and Pose Regression Network to localize people and regress cuboid proposals to the detailed 3D poses. Their novel solution performs well on both benchmarks, yet retraining the model is necessary when it comes to new scenes. Although these approaches can better perform 3D human pose estimation, the computational cost of cross-view matching and 3D human pose estimation is still too high for real-time applications.

2.2. 3D Human Pose Tracking

To process synchronized multiple camera streaming videos, tracking plays an important role in practice by leveraging temporal information to smooth 3D motion capture, handle current pose ambiguity, and accelerate the system’s processing time. Recently, Taylor *et al.* [34] combine Conditional Restricted Boltzmann Machine [33] with Bayesian filtering for tracking single-person 3D poses. With an increasing number of people in the scene, dealing with frequent occlusions and ambiguities becomes crucial. Liu *et al.* [20] combine image segmentation with articulated template models for accurately capturing and tracking 3D articulated skeleton motion of each person. Like the purpose of

Liu *et al.* [20], Elhayek *et al.* [12] and Li *et al.* [19] combine the CNN-based joint detection method with a model-based motion or spatio-temporal tracking algorithm in a unified 3D pose optimization for tracking and capturing 3D articulated skeleton motion of each person.

While in larger indoor/outdoor environments with more people and cameras, most approaches [4, 32, 7, 39] focus on reducing the computation cost while obtaining better performance. Tanke *et al.* [32] utilize a 2D human pose detector to obtain multiple 2D estimated human poses from multiple views and solve the k-partite matching problem using epipolar geometry to build associations among these multiple 2D estimated human poses across multiple views. They thus construct 3D human pose of each person, followed by a greedy algorithm to match and track iteratively across frames. Like Tanke *et al.* [32], Bridgeman *et al.* [4] applies a greedy fashion to seek the best correspondence between views. Estimated 3D skeletons are also exploited as input to improve the tracking quality. Recently, Chen *et al.* [7] develop a novel 3D human pose tracking technique that speeds up tracking tasks in their large-scale camera systems. Compare to Tanke’s work [32], their cross-view tracking with geometric affinity can track and construct 3D human poses of each person across views and frames synchronously. To attain better performance, Zhang *et al.* [39] develop a novel 4D (2D spatial, 1D viewpoint, and 1D temporal) association graph algorithm to enhance the accuracy of finding associations among 2D and 3D human poses from views and frames. Our approach is developed by referring to the inspiration of these approaches [7, 32].

3. Methodology

Given multiple cameras \mathbf{C} from different views in a scene, coupling with their known camera intrinsic \mathbf{K} and extrinsic (rotation \mathbf{R} and translation \mathbf{o}) information, we aim to find out an unknown number of humans’ 3D locations in an area. The proposed 3D reconstruction and tracking framework (as shown in Figure 1) can be divided in to three stages: 2D human pose estimation, 2D-3D association and 3D human pose reconstruction. The overall procedure is shown in Algorithm 1.

3.1. 2D Human Pose Estimation

To reconstruct 3D human pose via multiple views, the potential 2D human poses \mathbf{p} in each view are first extracted via an off-the-shelf 2D human pose estimator. It is worth noting that the error of 2D pose estimation may easily affect the following association and 3D pose reconstruction processes, yielding mis-matching or inaccurate 3D pose. We adopt YOLOv3 [28] coupled with HRNet [37] as our top-down 2D pose estimator. Since HRNet [37] conducted multi-scale fusions by fusing high and low resolution representations, leading to a potentially more accurate and spa-

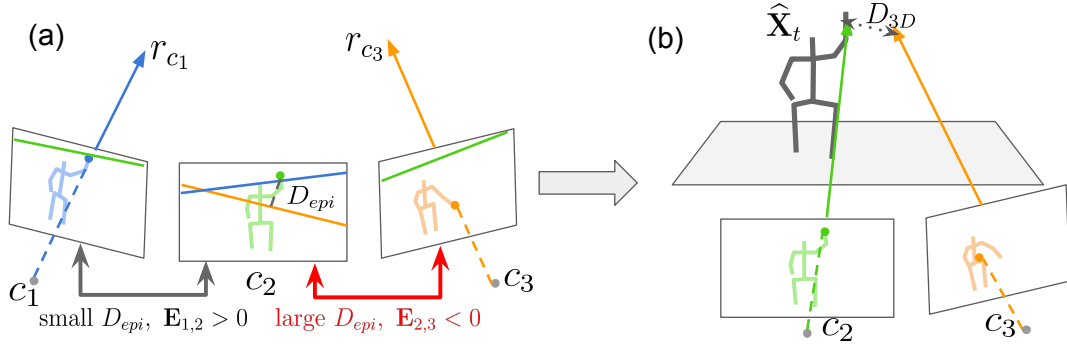


Figure 2: Illustration of *joints filter* algorithm. (a) The blue and orange lines on image plane of c_2 are epipolar lines of back-project-ray r_{c_1} , r_{c_3} , respectively. The joints in views c_2 and c_3 are considered as outlier candidates due to the large epipolar distance thus be examined in the next step. (b) We remove the outlier with the larger distance between back-project-ray of 2D joint and estimated 3D skeleton $\hat{\mathbf{X}}_t$.

tially more precise heatmap predictions, which is different from earlier works that recovered high-resolution representations from lower-resolution representations [9]. This combination performs a more efficient and accurate estimation. Yet the comparisons of these off-the-shelf 2D pose estimators are beyond the scope of this paper.

3.2. Part-aware 2D-3D Association

After extracting 2D human poses \mathbf{p} from each view, each 2D pose $\mathbf{x}_{t,c} \in \mathbf{p}$ is associated with previously reconstructed 3D poses $\mathbf{X}_{t'} \in \mathbf{P}$, which are initialized across views and updated at t' (described in Section 3.3.) The affinity $\mathcal{G}(\mathbf{x}_{t,c}, \mathbf{X}_{t'})$ of a 2D pose $\mathbf{x}_{t,c}$ to a 3D pose $\mathbf{X}_{t'}$ is measured with the geometric constraints of projection difference between $\mathbf{x}_{t,c}$ and re-projecting $\mathbf{X}_{t'}$ onto view c . The geometry affinity of the n^{th} joint \mathcal{G}^n is computed by the equation below:

$$\mathcal{G}^n(\mathbf{x}_{t,c}^n, \mathbf{X}_{t'}^n) = \left(1 - \frac{\|\mathbf{x}_{t,c}^n - \tilde{\mathbf{x}}_{t',c}^n\|}{\alpha_{2D}(t-t')}\right) \cdot e^{-\lambda_\alpha(t-t')}. \quad (1)$$

where $\tilde{\mathbf{x}}_{t',c}^n$ is the 2D projection of $\mathbf{X}_{t'}^n$ on camera c . α_{2D} , λ_α are constants of 2D velocity threshold and penalty rate of time interval respectively.

Common approaches conduct body-aware affinity measurement to acquire a 2D pose affinity to 3D pose by averaging all the joint errors. However, the potential noisy predictions due to occlusion or motion blur may yield mismatching results. To reduce the effect of noisy estimation of joints, we propose a part-aware measurement approach that can handle those noisy joints (shown in Figure 4). In contrast to directly averaging all joint affinities, we only take joints that have positive affinities into consideration, since the outlier joints may yield negative affinities that influence associations.

Therefore, the affinity $\mathcal{G}(\mathbf{x}_{t,c}, \mathbf{X}_{t'})$ is the mean of the joint affinities with positive values, indicating the similar-

ity of $\mathbf{x}_{t,c}^n$ and $\mathbf{X}_{t'}^n$. And for those number of joints with positive affinities smaller than ε , we set $\mathcal{G}(\mathbf{x}_{t,c}, \mathbf{X}_{t'})$ to 0.

Once the affinities of all pairs of 2D poses and 3D poses are computed, an affinity matrix $\mathbf{A} \in \mathbb{R}^{|\mathbf{P}| \times |\mathbf{P}|}$ can be formed to associate 2D poses to 3D skeletons. This becomes a weighted bipartite problem and can easily be solved by Hungarian algorithm [18]. While finishing the associations of all views, cross-view association is accomplished implicitly, where the corresponding 2D poses across views are those assigned to the same 3D skeleton.

3.3. 3D Pose Reconstruction

As long as a tracked 3D skeleton $\mathbf{X}_{t'}$ is matched with 2D poses after association, the new 3D skeleton can be reconstructed with the associated 2D poses from multiple views. However, the association may sometimes give only single view of 2D pose thus cannot perform reconstruction from multiple views. Therefore, we gather a set $\mathbb{P} = \{\mathbf{x}_{t',c}^n | c \in \mathbf{C}; 0 \leq t - t'_c < \tau\}$, where t'_c denotes the time of the last matched 2D pose from camera c and $t'_c \leq t$, to retrieve latest associated 2D pose in each camera within a short time interval τ for triangulation, and we set τ as 75-100ms empirically in practice. Furthermore, we propose *joints filter* as pre-processing of reconstruction to handle noisy joint estimations.

Joints Filter is also a part-aware processing. In some cases, noisy predictions occur because of partial occlusion or motion blur. Epipolar constraint is utilized to remove outliers by computing the distance between epipolar line and the corresponding point. Here we handle each joint independently. Therefore, an epipolar affinity matrix of the n^{th} joint $\mathbf{E} \in \mathbb{R}^{|\mathbf{P}| \times |\mathbf{P}|}$ can be computed.

$$\mathbf{E}_{i,j}(\mathbf{x}_i^n, \mathbf{x}_j^n) = 1 - \frac{d(\mathbf{x}_i^n, L_j) + d(\mathbf{x}_j^n, L_i)}{2\alpha_{epi}}. \quad (2)$$

where L , d denote the epipolar line and the point-to-line distance function, respectively. α_{epi} represents a threshold of acceptable distance error. Whenever there is a negative value at $\mathbf{E}_{i,j}$, at least one of \mathbf{x}_i^n and \mathbf{x}_j^n is regarded as an outlier. Benefit from the known 3D skeleton, we can infer its location in 3D space for correlation measurement. As shown in Figure 2, We design a greedy method to remove it by comparing them to $\hat{\mathbf{X}}_t^n$, which is estimated by \mathbf{X}_t^n and motion model (*e.g.*, linear motion, 3D Kalman Filter [30].) Afterward, the back-project-ray \mathbf{r} of joint \mathbf{x}_i^n can be obtained from the following formula:

$$\mathbf{r}_{i,c}^n = (\mathbf{KR})_c^+ \cdot \mathbf{x}_{i,c}^n + \mathbf{o}_c \quad (3)$$

where $\mathbf{K}, \mathbf{R} \in \mathbb{R}^{3 \times 3}$ are intrinsic and rotation matrix, respectively, and $+$ stands for the inverse matrix. \mathbf{o}_c is the center of camera c with respect to global coordinate system. We then compute the 3D point-to-line distance between $\hat{\mathbf{X}}_t^n$ and \mathbf{r}_i^n . At last, we remove the joint with larger distance as the outlier of 3D skeleton to get a robust reconstruction.

3D Pose Reconstruction. In our method, we use Direct Linear Transformation (DLT) for triangulation. In some cases, only one 2D pose is associated to the 3D skeleton yet triangulation needs at least two points to process. To avoid this problem, a set \mathbb{P} of 2D poses is collected from a small range of time τ . However, traditional triangulation is based on the condition that all the 2D points across views are from the same time, a penalty is set for the time interval, which is referenced from [7]. But sometimes, all the joints are wrong predicted, often happen at lower arms and legs, we regard it as a missing joint $\mathbf{X}_{t,miss}^n$. Here we compensate $\mathbf{X}_{t,miss}^n$ with $\hat{\mathbf{X}}_t^n$. Though *joints filter* is adopted, slight noisy estimation still remains. We conduct temporal information to smooth the 3D pose trajectory with Gaussian filter for further enhancing the performance.

Initialization of 3D pose. After association process in Section 3.2, each extracted 2D pose is labeled as either matched or unmatched. For those unmatched 2D poses from different views possibly forms a new 3D skeleton. Hence, to initialize a new 3D pose, we utilize epipolar constraint to associate unmatched 2D poses across views. Assuming that there is a set \mathbb{U}_c of unmatched 2D poses in camera c , we measure the correlation of unmatched 2D poses from other views and \mathbb{U}_c by Eq. 2. The final affinity matrix $\mathbf{E} = \sum_{n=1}^N \mathbf{E}^n$ can also be solved by Hungarian algorithm [18]. The process is iterative, meaning that we handle cameras one by one. After each association, there might still remain unmatched poses \mathbb{U}'_c that aren't captured in camera c but captured in others. We then add \mathbb{U}'_c into \mathbb{U}_c . While completing association of all views, the *joints filter* is also applied to get a robust reconstruction for initial 3D skeleton with a slight difference. Since there is no former 3D skeleton to reference, we have to solve it another way. The epipolar matrix of the n^{th} joint \mathbf{E}^n is still required, and an

Algorithm 1: Estimation and Tracking procedure

```

Input:
Cameras  $c \in \mathbf{C}$ 
New 2D poses  $\mathbf{x}_t \in \mathbf{p}$ 
Previously tracked 3D skeletons  $\mathbf{X}_{t'} \in \mathbf{P}$  at time  $t'$ 
Output:
Tracked 3D skeletons with new 3D poses  $\mathbf{X}_t \in \mathbf{P}$  at time  $t$ 
1 Initialization:  $\mathbf{P} \leftarrow \emptyset$ ;  $\mathbf{A} \in \mathbb{R}^{|\mathbf{P}| \times |\mathbf{P}|}$ 
  /* cross-view association */
2 foreach  $c \in \mathbf{C}$  do
3   foreach  $i, \mathbf{X}_{t'} \in \mathbf{P}$  do
4     foreach  $j, \mathbf{x}_{t,c} \in \mathbf{p}$  do
5        $\mathbf{A}(i, j) \leftarrow \text{PartAwareAssociation}(\mathbf{x}_{t,c}, \mathbf{X}_{t'})$ 
6     end
7   end
8    $\text{Match}(\mathbf{X}_{t'}, \mathbf{x}_{t,c}) \leftarrow \text{HungarianAlgorithm}(\mathbf{A})$ 
9 end
10  $\tau \leftarrow \text{TimeInterval}$ 
  /* update 3D skeleton */
11 foreach  $\mathbf{X}_{t'} \in \mathbf{P}$  do
12    $\mathbb{P} \leftarrow \{\mathbf{x}_{t',c} | c \in \mathbf{C}; 0 \leq t - t' < \tau\}$ 
13    $\hat{\mathbf{X}}_t \leftarrow \text{CurrentPoseEstimation}(\mathbf{X}_{t'})$ 
14    $\hat{\mathbb{P}} \leftarrow \text{JointsFilter}(\mathbb{P}, \hat{\mathbf{X}}_t)$ 
15    $\mathbf{X}_t \leftarrow \text{3DReconstruction}(\hat{\mathbb{P}}, \hat{\mathbf{X}}_t)$ 
16 end

```

outlier exist while $\mathbf{E}_{i,j}^n$ is a negative value. We then compute the $\sum_{j=0} \mathbf{E}_{i,j}^n$ of \mathbf{x}_i^n and \mathbf{x}_j^n , and remove the joint with smaller sum, which has a weaker affinity with others. Finally, the initial 3D pose can be reconstructed via Sec 3.3.

4. Experimental Results

4.1. Benchmark Datasets

Campus dataset [3]. A small dataset includes three actors and two pedestrians walking in an outdoor environment and captured by three calibrated cameras. There are 2000 frames in the video. The 350 ~ 470th and 650 ~ 750th frames were used for evaluation; others were used for training. It also provides an evaluation metric, percentage of correct parts (PCP), to quantify the accuracy of detected 3D keypoints on six parts of the human body, *i.e.* head, torso, upper arm, lower arm, upper leg, and lower leg. We noted that most of the actors in standard evaluation set videos are well captured and insufficient to test the present approach hardly. Therefore, the additional 471 ~ 649th frames were added into the evaluation set. These frames are more demanding while more people were walking around.

Shelf dataset [3]. In comparison with Campus, it's a more complex dataset captured by five calibrated cameras in an indoor scenario where up to four people were dissembling a shelf at the same time. As a result, people will be shaded by the shelf or each other frequently in the video and increase the difficulty of 3D pose reconstruction. The video contains 3200 frames in total. The 300 ~ 600th frames were used

Campus Dataset, PCP(%)									
Method	Belagiannis <i>et al.</i> [3]	Ershadi-N <i>et al.</i> [13]	Bridgeman <i>et al.</i> [4]	Tanke <i>et al.</i> [32]	Dong <i>et al.</i> [11]	Chen <i>et al.</i> [7]	Tu <i>et al.</i> [36]	Zhang <i>et al.</i> [39]	Ours
Actor1	93.5	94.2	91.8	98	97.6	97.1	97.6	-	98.37
Actor2	75.7	92.9	92.7	91	94.1	94.1	93.8	-	93.76
Actor3	85.4	84.6	93.2	29.2	98.6	98.6	98.8	-	98.26
Average	84.5	90.6	92.6	95.7	96.6	96.6	96.7	-	96.79

Shelf Dataset, PCP(%)									
Method	Belagiannis <i>et al.</i> [3]	Ershadi-N <i>et al.</i> [13]	Bridgeman <i>et al.</i> [4]	Tanke <i>et al.</i> [32]	Dong <i>et al.</i> [11]	Chen <i>et al.</i> [7]	Tu <i>et al.</i> [36]	Zhang <i>et al.</i> [39]	Ours
Actor1	75.3	93.3	99.7	99.8	98.8	99.6	99.3	99.0	99.14
Actor2	69.7	75.9	92.8	90	94.1	93.2	94.1	96.2	95.41
Actor3	87.6	94.8	97.7	98	97.8	97.5	97.6	97.6	97.64
Average	77.5	88.0	96.7	95.9	96.9	96.8	96.0	97.6	97.39

Table 1: Quantitative comparison on the Campus and Shelf dataset. The results of other methods are taken from respective papers. “red” and “blue” represent the first and second highest scores, respectively.

Campus	PCP(%)			
	Ua	La	UI	LI
Tanke <i>et al.</i> [32]	97.7	84.3	99.3	99.0
Chen <i>et al.</i> [7]	98.6	84.6	-	-
[†] Chen <i>et al.</i> [7]	99.8	84.2	100.0	100.0
ours	99.8	84.3	100.0	100.0
Shelf	Ua	La	UI	LI
Tanke <i>et al.</i> [32]	98.0	86.3	100.0	100.0
Chen <i>et al.</i> [7]	98.7	87.7	-	-
[†] Chen <i>et al.</i> [7]	98.8	89.9	100.0	99.9
ours	98.8	90.5	100.0	100.0

Table 2: PCP scores of human parts on different approaches. “U, L, a, l” represent upper, lower, arms and legs. [†] stands for our re-implementation of [7] with same 2D pose estimation method as ours.

for evaluation; others were used for training. We follow the same evaluation protocol as the Campus dataset to quantify the accuracy of detected 3D keypoints.

CMU Panoptic [17]. In contrast to Campus and Shelf, the Panoptic dataset is a large dataset released by CMU Perceptual Computing Lab. They built an enclosed studio with 480 VGA cameras and 31 HD cameras installed on it. The videos contain multi-person engaging in several social activities. For the lack of the groundtruth of 3D poses, only qualitative video and image results are presented on this dataset. The qualitative results are shown in supplementary.

4.2. Comparison with other approaches

We compare our proposed method with the following approaches. Belagiannis *et al.* [3] develop a 3DPS-based

Campus Extended Testing Set, PCP(%)			
Method	Dong <i>et al.</i> [11]	[†] Chen <i>et al.</i> [7]	Ours
Actor1	97.55	98.37	98.37
Actor2	93.33	95.31	95.65
Actor3	98.19	98.13	98.41
Average	96.36	97.27	97.47

Table 3: Quantitative comparison on the Campus extended testing set. The results of other methods are taken from respective papers. “red” indicates the highest scores.

model with temporal consistency thereafter become a baseline on Campus and Shelf datasets. Ershadi-N. *et al.* [13] develop a clustering algorithm to reduce the 3D state space by optimally clustering 3D candidate joints. Bridgeman *et al.* [4] develop a 2D pose association algorithm with error correction to match 2D poses in multiply views more robustly after constructing 3D poses. Tanke *et al.* [32] develop an iterative greedy matching to remove outliers and fill-in missing joints by temporal 3D pose averaging. Dong *et al.* [11] develop a multi-way matching with circle consistency to match 2D poses in multiple views optimally after constructing 3D poses. Chen *et al.* [7] develop an efficient algorithm by leveraging temporal consistency to find out 2D correspondences across views. An incremental reconstruction that can deal with unsynchronized streaming video and fast-moving limbs is also designed. Tu *et al.* [36] develop a cuboid proposal network to localize 3D poses by integrating multi-view projected 2D pose heatmaps. Zhang *et al.* [39] develop a 4D association graph to calculate relevance be-

tween 2D and 3D joints across views and frames to obtain the best matching connection when constructing 3D poses.

Table 1 demonstrates the PCP score of our proposed method and the above approaches evaluated on Campus and Shelf datasets. We reach state-of-the-art on Campus and competitive score on Shelf. In particular, our approach significantly improves Actor1’s results on Campus, and as for Actor2 in Shelf, who suffering from severe occlusion, our method can better tackle the condition than most previous works. A further comparison was demonstrated in Table 2. We compared four human parts to prove that our proposed reconstruction method is valid. Lower arms that have larger motion amplitudes often suffer from motion blur or occlusion. In most works, the PCP score of lower arms performs under 90%; however, our proposal successfully reach the highest accuracy up to 90.5% on Shelf.

Furthermore, we also evaluate some approaches [11, 7] on the extended Campus testing set. With three actors and a pedestrian walking around not only enhance the difficulties of associations but also increase the frequency of partial occlusions, Dong *et al.* [11] method slightly degrade the performance(96.36%). On the other hand, ours and Chen *et al.* [7], which is our re-implementation, maintain considerable results on the extended frames(97.47% and 97.27% respectively). We conjecture that with the help of part-aware measurement and *joints filter*, our approach achieves 0.2% improvement. The result is shown in Table 3. Thus, we can confirm that our approach can effectively and robustly construct 3D pose of people, who suffer from severe occlusion.

4.3. Ablation Studies

In this section, we give some ablation studies to verify that each component is beneficial to our approach. We compare to the baseline with body-aware measurement in association and without *joints filters*.

Measure in body or part-aware? To verify that measuring in part-aware is effective, we only change the method building associations. The other modules remain the same as the baseline. In some cases, even though most joints have good predictions, there were still seldom joints with an enormous error that significantly affect the affinity and causing unmatched consequence(an example is given in Figure 4). As demonstrated in Table 3, singly applying part-aware measurement enhanced approximately 2.1% on Campus, which is a large improvement. We investigate that measuring in part-aware leads to fewer unmatched results, representing that for each 3D skeleton, there are more matched 2D poses across views. Hence, with more cross-view 2D correspondences, a better 3D skeleton can be constructed.

Effectiveness of Joints Filter? As described in Section 3.3, we regard *joints filter* as an important process to handle noisy points that will degrade the reconstruction accuracy. For comparison, we select the torso, lower arms to verify

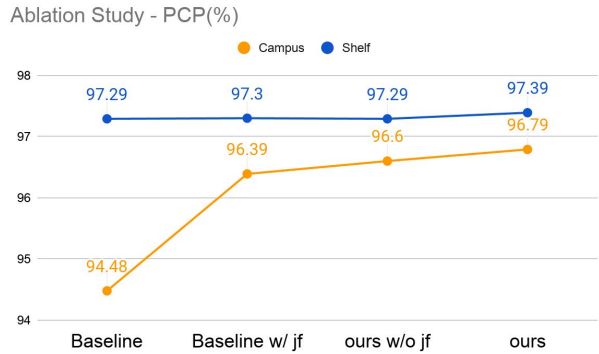


Figure 3: Ablation study on Campus and Shelf datasets. “jf” means *joints filter*, and “ours” indicates that part-aware measurement and *joints filter* are employed. As shown, applying either part-aware measurement or *joints filter* performs better and by using both of them further improved the PCP result.

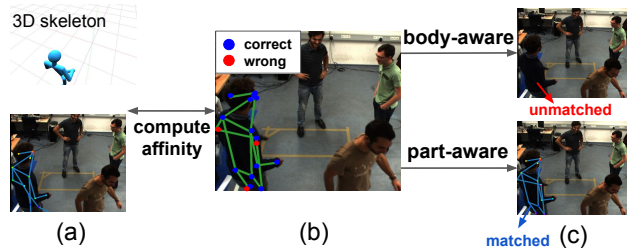


Figure 4: Visual comparison of body-aware and part-aware association on Shelf dataset. (a) a 3D skeleton and its re-projection (b) the corresponding estimated 2D pose. Due to the out-of-view of lower body, the enormous 2D estimation error of ankles leads to weak ankle affinities and affects the affinity of entire body to mis-match in body-aware association. (c) Therefore, we adopt part-aware association by neglecting the few wrong predictions but only comparing affinities of others to reach a robust matching result.

the impact of the *joints filter*. Torso performs well (100%) in every method due to the small range of motion. For the lower arms, which are considered the most demanding limb among whole body parts, the *joints filter* method improves about 0.5~1.0% and PCP score. Some examples on Shelf are shown in Figure 5. Additionally, to justify that *joints filter* not only works in our proposal, we add it to the re-implementation of Chen *et al.* [7]. We can improve the PCP score even from 97.28% to 97.43%, which is a 0.15% enhancement.




Figure 5: A visualization of ablation studies. From left to right indicate the camera’s order, and the last column is the 3D skeleton from c_3 ’s viewpoint. As shown in red bounding boxes, the black-shirt man’s lower body is out of c_3 ’s FOV, leading to wrong prediction. (a) body-aware association failed to associate the 2D pose in c_3 to the corresponding 3D skeleton. (b) altering association with part-aware measurement succeeds. Yet due to the outliers of his lower body, the legs of 3D skeleton perform weird poses (orange box.) (c) further improves the 3D pose reconstruction of (b) with joint filter.

Method	Processing Speed on Shelf (ms)				fps
	Affi.	Asso.	Recon.	Init.	
Bridgeman <i>et al.</i> [4]	-	-	-	-	110
Dong <i>et al.</i> [11]	25	20	60	-	9.5
Chen <i>et al.</i> [7]	-	-	-	-	325
ours w/o jf	1		3.5	1	181
ours	1		8	1	100

Table 4: Runtime comparison on Shelf. Stages ‘Affi.’, ‘Asso.’, ‘Recon.’ and ‘Init.’ stand for affinity computation, association, reconstruction and 3D skeleton initialization, respectively.

4.4. Processing time

In this section, we report the runtime of our algorithm on different datasets. Considering 2D pose estimation is beyond the scope, only tracking and reconstruction time are demonstrated.

Shelf. For the pose estimation, our model takes about 600ms to produce all the 2D poses. Since most of the multi-view 3D human pose estimation methods focus on the reconstruction process, we compare the result with other’s proposals on this part. Dong *et al.* [11] test their experiment on GeForce 1080Ti GPU, which spent an average of 25ms to compute affinities and 20 ms for finding the cross-view association, and 60 ms for reconstruction. In contrast, our implementation runs on Intel CPU 3.70GHz. It takes

about 2 ms for finding the 2D-3D association, 8 ms for reconstruction and 1 ms for the cross-view initialization.

Campus contains fewer people and cameras, the processing speed is faster than the Shelf. It takes about 1 ms for association, 5 ms for reconstruction, and 0.5 ms for initialization.

5. Conclusion

We propose an efficient and robust multi-person 3D pose estimation and tracking from multiple views by directly associating 2D poses estimation to a 3D skeleton in each view and updating 3D poses with a part-aware approach. Our proposed part-aware association is able to achieve better matching in contrast to common body-aware methods. Moreover, joint filter also plays an important role in dealing with noisy joint estimation to get a more accurate reconstruction. In experiments, our approach achieved a competitive result on benchmark datasets without any additional training. And the ablation study verified our part-aware method is effective in partial occlusion.

Acknowledgement. This work was partially supported by Ministry of Science and Technology in Taiwan (MOST 110-2634-F-002-04).

References

- [1] Alex M Andrew. Multiple view geometry in computer vision. *Kybernetes*, 2001.

- [2] Mykhaylo Andriluka, Stefan Roth, and Bernt Schiele. Monocular 3d pose estimation and tracking by detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 623–630. IEEE, 2010.
- [3] Vasileios Belagiannis, Sikandar Amin, Mykhaylo Andriluka, Bernt Schiele, Nassir Navab, and Slobodan Ilic. 3d pictorial structures revisited: Multiple human pose estimation. *IEEE transactions on pattern analysis and machine intelligence*, 38(10):1929–1942, 2015.
- [4] Lewis Bridgeman, Marco Volino, Jean-Yves Guillemaut, and Adrian Hilton. Multi-person 3d pose estimation and tracking in sports. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2019.
- [5] Magnus Burenius, Josephine Sullivan, and Stefan Carlsson. 3d pictorial structures for multiple view articulated pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3618–3625, 2013.
- [6] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Openpose: realtime multi-person 2d pose estimation using part affinity fields. *IEEE transactions on pattern analysis and machine intelligence*, 43(1):172–186, 2019.
- [7] Long Chen, Haizhou Ai, Rui Chen, Zijie Zhuang, and Shuang Liu. Cross-view tracking for multi-human 3d pose estimation at over 100 fps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3279–3288, 2020.
- [8] 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/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7103–7112, 2018.
- [9] Y. Chen, Z. Wang, Y. Peng, Z. Zhang, G. Yu, and J. Sun. Cascaded pyramid network for multi-person pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7103–7112, 2018.
- [10] Yu Cheng, Bo Yang, Bo Wang, Wending Yan, and Robby T Tan. Occlusion-aware networks for 3d human pose estimation in video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 723–732, 2019.
- [11] Junting Dong, Wen Jiang, Qixing Huang, Hujun Bao, and Xiaowei Zhou. Fast and robust multi-person 3d pose estimation from multiple views. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7792–7801, 2019.
- [12] Ahmed Elhayek, Edilson de Aguiar, Arjun Jain, Jonathan Thompson, Leonid Pishchulin, Mykhaylo Andriluka, Christoph Bregler, Bernt Schiele, and Christian Theobalt. Marconi—convnet-based marker-less motion capture in outdoor and indoor scenes. *IEEE transactions on pattern analysis and machine intelligence*, 39(3):501–514, 2016.
- [13] Sara Ershadi-Nasab, Erfan Noury, Shohreh Kasaei, and Esmaeil Sanaei. Multiple human 3d pose estimation from multiview images. *Multimedia Tools and Applications*, 77(12):15573–15601, 2018.
- [14] Riza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7297–7306, 2018.
- [15] Eldar Insafutdinov, Leonid Pishchulin, Bjoern Andres, Mykhaylo Andriluka, and Bernt Schiele. Deeppercut: A deeper, stronger, and faster multi-person pose estimation model. In *European Conference on Computer Vision (ECCV)*, pages 34–50. Springer, 2016.
- [16] Wen Jiang, Nikos Kolotouros, Georgios Pavlakos, Xiaowei Zhou, and Kostas Daniilidis. Coherent reconstruction of multiple humans from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5579–5588, 2020.
- [17] Hanbyul Joo, Tomas Simon, Xulong Li, Hao Liu, Lei Tan, Lin Gui, Sean Banerjee, Timothy Godisart, Bart Nabbe, Iain Matthews, et al. Panoptic studio: A massively multiview system for social interaction capture. *IEEE transactions on pattern analysis and machine intelligence*, 41(1):190–204, 2017.
- [18] H. W. Kuhn and Bryn Yaw. The hungarian method for the assignment problem. *Naval Res. Logist. Quart.*, pages 83–97, 1955.
- [19] Kun Li, Nianhong Jiao, Yebin Liu, Yangang Wang, and Jingyu Yang. Shape and pose estimation for closely interacting persons using multi-view images. In *Computer Graphics Forum*, volume 37, pages 361–371. Wiley Online Library, 2018.
- [20] Yebin Liu, Juergen Gall, Carsten Stoll, Qionghai Dai, Hans-Peter Seidel, and Christian Theobalt. Markerless motion capture of multiple characters using multiview image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 35(11):2720–2735, 2013.
- [21] Dushyant Mehta, Oleksandr Sotnychenko, Franziska Mueller, Weipeng Xu, Srinath Sridhar, Gerard Pons-Moll, and Christian Theobalt. Single-shot multi-person 3d pose estimation from monocular rgb. In *2018 International Conference on 3D Vision (3DV)*, pages 120–130. IEEE, 2018.
- [22] Gyeongsik Moon, Ju Yong Chang, and Kyoung Mu Lee. Camera distance-aware top-down approach for 3d multi-person pose estimation from a single rgb image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10133–10142, 2019.
- [23] Francesc Moreno-Noguer. 3d human pose estimation from a single image via distance matrix regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2823–2832, 2017.
- [24] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. volume 9912, pages 483–499, 10 2016.
- [25] Georgios Pavlakos, Xiaowei Zhou, Konstantinos G Derpanis, and Kostas Daniilidis. Harvesting multiple views for marker-less 3d human pose annotations. In *Proceedings of*

- the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6988–6997, 2017.
- [26] Aleksis Pirinen, Erik Gärtner, and Cristian Sminchisescu. Domes to drones: Self-supervised active triangulation for 3d human pose reconstruction. *Advances in Neural Information Processing Systems*, 32, 2019.
- [27] Haibo Qiu, Chunyu Wang, Jingdong Wang, Naiyan Wang, and Wenjun Zeng. Cross view fusion for 3d human pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4342–4351, 2019.
- [28] Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7263–7271, 2017.
- [29] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence*, 39(6):1137–1149, 2016.
- [30] Y. Salih and A. S. Malik. 3d object tracking using three kalman filters. In *2011 IEEE Symposium on Computers Informatics*, pages 501–505, 2011.
- [31] Xiao Sun, Bin Xiao, Fangyin Wei, Shuang Liang, and Yichen Wei. Integral human pose regression. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 529–545, 2018.
- [32] Julian Tanke and Juergen Gall. Iterative greedy matching for 3d human pose tracking from multiple views. In *German Conference on Pattern Recognition*, 2019.
- [33] Graham W Taylor, Geoffrey E Hinton, and Sam T Roweis. Modeling human motion using binary latent variables. In *Advances in neural information processing systems*, pages 1345–1352. Citeseer, 2007.
- [34] Graham W Taylor, Leonid Sigal, David J Fleet, and Geoffrey E Hinton. Dynamical binary latent variable models for 3d human pose tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 631–638. IEEE, 2010.
- [35] Denis Tome, Matteo Toso, Lourdes Agapito, and Chris Russell. Rethinking pose in 3d: Multi-stage refinement and recovery for markerless motion capture. In *2018 international conference on 3D vision (3DV)*, pages 474–483. IEEE, 2018.
- [36] Hanyue Tu, Chunyu Wang, and Wenjun Zeng. Voxelpose: Towards multi-camera 3d human pose estimation in wild environment. In *European Conference on Computer Vision (ECCV)*, 2020.
- [37] J. Wang, K. Sun, T. Cheng, B. Jiang, C. Deng, Y. Zhao, D. Liu, Y. Mu, M. Tan, X. Wang, W. Liu, and B. Xiao. Deep high-resolution representation learning for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2020.
- [38] Jianfeng Zhang, Xuecheng Nie, and Jiashi Feng. Inference stage optimization for cross-scenario 3d human pose estimation. *Advances in Neural Information Processing Systems*, 33, 2020.
- [39] Yuxiang Zhang, Liang An, Tao Yu, xiu Li, Kun Li, and Yebin Liu. 4d association graph for realtime multi-person motion capture using multiple video cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [40] Zhun Zhong, Liang Zheng, Zhedong Zheng, Shaozi Li, and Yi Yang. Camera style adaptation for person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5157–5166, 2018.