

Exploring Severe Occlusion: Multi-Person 3D Pose Estimation with Gated Convolution

Renshu Gu^{1,3}, Gaoang Wang^{2,*}, Jenq-Neng Hwang³

¹Hangzhou Dianzi University, Hangzhou, Zhejiang 310018, China

²Zhejiang University-University of Illinois at Urbana-Champaign Institute, Haining, Zhejiang 314400, China

³University of Washington, Seattle, 98195, USA

*Email: gaoangwang@intl.zju.edu.cn

Abstract—3D human pose estimation (HPE) is crucial in many fields, such as human behavior analysis, augmented reality/virtual reality (AR/VR) applications, and self-driving industry. Videos that contain multiple potentially occluded people captured from freely moving monocular cameras are very common in real-world scenarios, while 3D HPE for such scenarios is quite challenging, partially because there is a lack of such data with accurate 3D ground truth labels in existing datasets. In this paper, we propose a temporal regression network with a gated convolution module to transform 2D joints to 3D and recover the missing occluded joints in the meantime. A simple yet effective localization approach is further conducted to transform the normalized pose to the global trajectory. To verify the effectiveness of our approach, we also collect a new moving camera multi-human (MMHuman) dataset that includes multiple people with heavy occlusion captured by moving cameras. The 3D ground truth joints are provided by accurate motion capture (MoCap) system. From the experiments on static-camera based Human3.6M data and our own collected moving-camera based data, we show that our proposed method outperforms most state-of-the-art 2D-to-3D pose estimation methods, especially for the scenarios with heavy occlusions.

I. INTRODUCTION

Human pose estimation (HPE) is gaining high attention in recent years, as it is crucial in many areas such as human behavior analysis, augmented reality/virtual reality (AR/VR) applications, and self-driving industry. There has been a great amount of reports on 3D human pose estimation for videos recorded by static cameras in the experimental setting. However, there exist limited reports of quantitative results for natural videos that are captured by moving cameras and contain multi-person interactions.

Occlusion handling, including self-occlusion or inter-person occlusion, is an important issue in human pose estimation. Unfortunately, there is a lack of data with accurate ground truth for multi-person moving camera videos with inter-person occlusions. Multi-person 3D pose estimation with heavy occlusion is still not well-evaluated nor well-addressed in the literature.

In this paper, we aim to address and evaluate the occlusion problem in 3D human pose estimation. We first introduce a new dataset that has commonly encountered multi-person interactions, with various occlusion scenarios captured by a moving camera with ground truth being collected by a motion capture (MoCap) system simultaneously. The dataset includes

several commonly seen actions. To estimate 3D pose from 2D pose, we propose a temporal regression network using gated convolutions, which is inspired by the image inpainting application [1]. The estimated 3D position of each joint is with respect to the root position, and a global trajectory recovery approach is adopted to localize the joints based on the 3D position in the camera coordinate. This simple yet effective global trajectory recovery approach is proposed based on 3D-to-2D back-projection and temporal continuity.

In summary, we claim the following contributions:

- 1) To our best knowledge, we are the first to use temporal gated convolutions to recover missing poses and address the occlusion issues in the pose estimation. The idea of missing pose recovery is inspired by the image inpainting tasks.
- 2) We present a new human pose dataset, which is aimed at multi-person 3D pose estimation with occlusion handling in moving camera environments.
- 3) A simple yet effective approach is proposed to transform normalized poses to the global trajectory into the camera coordinate.
- 4) We provide extensive experiments on videos containing multiple people from moving cameras, which are lacking in the current literature yet critical in today's applications. Our solution provides great new opportunities to understand and predict human behaviors using freely moving cameras.

The organization of the paper is as follows: In Section II, we review some related works of 3D human pose estimation based on monocular cameras. The data collection and the new dataset description are then introduced in Section III. In Section IV, the proposed approach of 3D human pose estimation with temporal gated convolution and global trajectory estimation is discussed in detail. Experimental setup and results are presented in Section V and Section VI, respectively, followed by the conclusion in Section VII.

II. RELATED WORK

Recent efforts of 3D human pose estimation, mainly CNN based, show promising results. However, severe or full occlusion remains a challenging issue, especially in multi-person pose estimation. Since the multi-person 3D pose dataset is rare in the literature, the inter-person occlusion problem is still not well-studied. For the remaining section, we will briefly review

some related works of 3D human pose estimation and survey the related 3D human pose estimation datasets.

A. 3D Human Pose Estimation

One-stage 3D pose estimation. This class of methods, mainly based on convolutional neural networks (CNNs), focuses on end-to-end reconstruction [2]–[10] by directly estimating 3D poses from RGB images without intermediate supervision. Since image cues are directly used, localizing and cropping the subject greatly affects the results. A fine discretization of the 3D space around the subject is proposed in [5], which employs a coarse-to-fine prediction, given that input is a single RGB image. However, despite their attempt to evaluate on real-world scenarios, results are only reported on one sequence, and as they clarify, the available ground truth is limited and not exactly accurate. A weakly-supervised transfer learning method that uses mixed 2D and 3D labels in a unified deep neural network is presented in [6]. Rogez et al. [7], propose an end-to-end architecture LCR-net that contains a pose proposal generator, a classifier that predicts the different anchor-pose labels, and an anchor-pose-specific regressor. They further extend LCR-net to LCR-net++ and adds additional synthetic training data and show better performance on partially occluded or truncated persons in [8]. However, LCR-Net and LCR-net++ do not use any temporal information, thus they cannot handle full-body or severe occlusions. Mehta et al.’s work [9] uses occlusion-robust pose-maps (ORPM) to enable pose inference under partial occlusions within a single frame.

Two-stage 3D pose estimation. The second class of methods first predicts 2D joint positions in image space (keypoints) which are subsequently lifted to 3D [11]–[16]. Recent work [11] shows that given accurate 2D keypoints, predicting 3D poses is in fact relatively straightforward. [2], [12] simply search for a predicted set of 2D keypoints over a large set of 2D keypoints to find the corresponding 3D pose. Some approaches [13] fuse 2D joints, image cues, and 3D joint coordinates along the way. [6] and [17] predicts the depth of 2D joints to obtain the 3D pose. The advantages of two-stage methods include that (1) 2D pose is easy to predict with a large amount of 2D pose estimation data, since 2D ground truth data is much cheaper to obtain compared to 3D ground truth, (2) 2d-to-3d methods are not sensitive to diverse scenarios and environments. However, the two-stage solutions are more sensitive to the 2D pose estimation performance. For challenging scenarios with severe occlusions in the videos, the unreliable 2D keypoints may have a large influence on the 3D estimation performance. How to obtain robust 3D results even with noisy and occluded 2D keypoints as the input remains a challenging task.

3D pose estimation with sequential input. In recent years, researchers try to utilize temporal information to estimate the 3D pose with sequential input. Meanwhile, the occlusion problem is gaining more attention [14], [18]–[20]. A sequence-to-sequence network composed of layer-normalized LSTM units is proposed in [14], which imposes temporal smoothness

constraints during training. Dilated temporal convolutions are employed to capture long-term temporal information in [18]. An occlusion-aware network for video using incomplete 2D keypoints, instead of complete but incorrect 2D keypoints, is proposed in [19]. However, most of the existing works focus on self-occlusion with single person pose estimation. Moreover, the recovered poses are compared with respect to the root position in most experiments. However, poses in the camera coordinate are far more needed in practical applications. Different from existing works, our proposed method handles inter-person occlusions and long-time occlusions, which is quite common in real-world scenarios. Furthermore, our framework can efficiently estimate global trajectories and each subject’s relative positions independent of any dataset.

3D human pose estimation datasets. The most commonly used datasets for 3D human pose estimation include Human3.6M [21], HumanEva [22], TotalCapture [23], MPI-INF-3DHP [24], MuPoTS-3D [9] and 3DPW [25]. Human3.6M is a large-scale dataset featuring a single person at the center of videos recorded by a commercialized MoCap system. MoCap system is considered the most accurate system for obtaining 3D ground truth in the literature. However, the cameras in the MoCap system are static. HumanEva is similarly captured by a MoCap system much earlier and contains simpler poses. TotalCapture captures diverse body motion using a massively multi-view sensor system, where the sensors are also statically installed. MPI-INF-3DHP tries to include more challenging poses in the wild, but still focuses on single persons and static cameras. MuPoTS-3D is a multi-person dataset recorded by a static camera. 3DPW is the first dataset that includes video footage taken from a moving phone camera. However, it uses IMUs to record ground truth, which is reported to possess a 26mm average 3D pose error when verified with the Mocap system.

III. THE MOVING CAMERA MULTI-HUMAN DATASET

Unlike existing 3D human pose estimation datasets, we propose a new dataset, multi-person moving camera (MMHuman) dataset, which is aimed at providing accurate MoCap ground truth for multi-person interaction videos captured by moving cameras, and intentionally includes occlusions and human interactions that are common in real-world scenarios.

A. Data Collection and Design

MMHuman dataset was captured via a commercialized MoCap system and additional mobile cameras. The MoCap system consists of 8 static IR cameras that record synchronized high-resolution 120Hz videos. The additional moving cameras are hand-held commercial smartphones, manually synchronized with the MoCap system. The MoCap system relies on small reflective markers attached to the subjects’ bodies and tracks them over time. Tracking maintains the labeled identity and propagates it through time from an initial pose which is labeled either manually or automatically. A fitting process uses the position and identity of each of the body labels to infer accurate pose parameters. The final 3D ground truth pose is




Fig. 1: The motion capture environment.



Fig. 2: Examples of actions in MMHuman.

provided in the global coordinate, i.e. the calibrated MoCap coordinate. Our MoCap environment is shown in Fig. 1.

We do not see our proposed dataset as an alternative to existing datasets; rather, MMHuman complements existing ones with accurate ground truth of new multi-person interaction sequences shot by freely moving cameras. Our new dataset allows a quantitative evaluation of state-of-the-art approaches for multi-human interaction cases that have heavy and inter-person occlusions.

B. Dataset Description

The dataset consists of 6 actions that involve common human-interactions and other variations, as shown in Fig. 2. Table I shows a summary of our dataset. Summing up all videos recorded by different cameras, the total number of frames available is also provided.

TABLE I: MMHuman Dataset Description

	Shake Hand	Walk Cross	High Five	Pull Up	Hand Over	Kungfu
# Videos	6	6	5	3	5	2
# Frames	11,573	11,379	9,537	5,809	9,570	3,883
Total Frames	51,751					

IV. APPROACH

In this section, we explain our proposed 3D pose estimation approach in detail, including the temporal regression model with modified gated convolution, model architecture, and global pose trajectory estimation.

A. Temporal Gated Convolution

Given a temporal sequence of detected 2D joint locations $\mathbf{x} \in \mathbb{R}^{2N_{jt} \times T}$, our goal is to reconstruct the 3D pose $\mathbf{X} \in \mathbb{R}^{3N_{jt} \times 1}$ of the center frame, where T is the temporal window size, and N_{jt} is the number of joints in the pose model.

When occlusions occur, the 2D joints estimated by human pose detectors are either missing or with low confidence. Such noisy 2D joints adversely affect the accuracy of the 3D pose estimation. Intuitively, we can use attention or soft gated convolution to focus on non-occluded joints in the regression model, which is a well-known technique in image inpainting networks, such as [1], [26], [27]. Similarly, we can also treat 3D pose estimation as a special case of inpainting task when occlusions occur. For the initial input, additional to feed in the sequential 2D poses, we also feed the binary occlusion mask, $\mathbf{M} \in \mathbb{R}^{2N_{jt} \times T}$. Two separate convolution kernels are adopted to obtain the feature maps and soft masks individually. The soft masks can be regarded as soft gating operation or attention on the output features. Specifically, the soft gated convolution is defined as follows,

$$\begin{aligned}
 \mathbf{Y}^l &= \phi \left(\mathbf{X}^{l-1} * \mathbf{W}_f^l \right), \\
 \mathbf{M}^l &= \sigma \left(\mathbf{X}^{l-1} * \mathbf{W}_g^l \right), \\
 \mathbf{X}^l &= \mathbf{Y}^l \otimes \mathbf{M}^l,
 \end{aligned} \tag{1}$$

where \mathbf{X}^{l-1} is the output of the previous layer, \mathbf{W}_f^l are the convolution kernels of the feature map, \mathbf{W}_g^l are the convolution kernels of the soft mask, \mathbf{Y}^l is the feature map of the current layer, \mathbf{M}^l is the mask of the soft gate, \mathbf{X}^l is the output of the current layer, \otimes represents element-wise multiplication, and ϕ and σ are the activation functions related to feature map and soft mask, respectively. The gated convolution is illustrated in Fig. 3(b).

However, in our experiments, the soft gated convolution defined in Eq. (1) does not perform very well. This is reasonable since the soft mask and feature maps are calculated from the same input. It is difficult for the networks to learn different roles of soft gating and feature map extraction at the same time without further supervision. Since the soft gating is




Fig. 3: Temporal convolution layer in the pose estimation model. (a): The plain convolution layer used in [18]. (b): The soft gated convolution proposed in [1] for inpainting task. (c): Modified two-stream soft gated convolution proposed in this paper.

more correlated to the input masks rather than the input feature maps, we separate the soft gating and feature map extraction into two streams as shown in Fig. 3(c). Specifically, we modify Eq. (1) as follows,

$$\begin{aligned} \mathbf{Y}^l &= \phi(\mathbf{X}^{l-1} * \mathbf{W}_f^l), \\ \mathbf{M}^l &= \sigma(\mathbf{M}^{l-1} * \mathbf{W}_g^l), \\ \mathbf{X}^l &= \mathbf{Y}^l \otimes \mathbf{M}^l, \end{aligned} \quad (2)$$

where the soft mask is calculated only from the soft mask of the previous layer. This modification can make the network training benefit from easy convergence.

B. Network Architecture

We follow the same regression architecture used in [18]. Apart from a sequence of 2D joints, we also input the occlusion mask. After the first temporal gated convolution layer, 4 skip blocks are used for regression to estimate the 3D joints. For each layer, we replace the plain convolution as shown in Fig. 3(a) with the modified temporal gated convolution as shown in Fig. 3(c). For each layer, we input both the feature maps and the soft mask from the previous layer. The temporal dilated convolution is conducted separately for input feature maps and soft mask with different kernels, followed by the batch normalization and activation. We use ReLU for feature map extraction and Sigmoid function for the soft gating. For the supervision, we use the mean square error between predicted 3D joints and ground truth as the loss, which is defined as,

$$\mathcal{L} = \frac{1}{N_{jt}} \sum_{i=1}^{N_{jt}} \|\hat{\mathbf{X}}_i - \mathbf{X}_i^*\|^2, \quad (3)$$

where $\hat{\mathbf{X}}$ is the predicted 3D joints and \mathbf{X}^* is the 3D ground truth.

C. Global Pose Trajectory Estimation via Back-Projection

We adopt the same output format as [18], i.e. the output of the temporal regression network is the position of all joints with respect to the root location. In [18], to obtain the global position of each joint, a second network called trajectory model, is trained separately to estimate the trajectory of the human in the camera coordinate. However, this extra effort can be efficiently replaced by a simple optimization approach that utilizes the temporal consistency and 2D-3D back-projection relationship.

For simplicity, we illustrate how to reconstruct the trajectory of one person in the following subsection. For multi-person scenarios, we can estimate all the individual trajectories separately. Denote the root position in the camera coordinate at frame t as \mathbf{C}_t . Then we can define the projection error as follows,

$$\epsilon_t = \sum_{i=1}^{N_{jt}} \|\rho(\mathbf{X}_{i,t} + \mathbf{C}_t) - \mathbf{x}_{i,t}\|^2, \quad (4)$$

where ρ is the projection of the pinhole camera, which maps the 3D coordinate into image coordinate, i.e.,

$$\begin{aligned} \rho_x(\mathbf{X}) &= \frac{fX}{Z} + c_x, \\ \rho_y(\mathbf{X}) &= \frac{fY}{Z} + c_y, \end{aligned} \quad (5)$$

where $\mathbf{X} = [X, Y, Z]^T$ represents the 3D point in the camera coordinate, f is the focal length, $\{c_x, c_y\}$ is the center of the image coordinate.

Notice that all the joints share the same global trajectory \mathbf{C}_t at each frame. Then the root trajectory \mathbf{C}_t can be estimated by minimizing the projection error, i.e.,

$$\hat{\mathbf{C}}_t = \arg \min_{\mathbf{C}_t} \epsilon_t, \quad (6)$$

Since the back-projection is only available for non-occluded joints, the occlusion mask is introduced to constrain the error only on non-occluded joints, i.e.,

$$\epsilon_t = \sum_{i=1}^{N_{jt}} (1 - M_{i,t}) \|\rho(\mathbf{X}_{i,t} + \mathbf{C}_t) - \mathbf{x}_{i,t}\|^2, \quad (7)$$

where $M_{i,t}$ is defined as

$$M_{i,t} = \begin{cases} 1, & \text{if } i \text{ is occluded,} \\ 0, & \text{if } i \text{ is visible.} \end{cases} \quad (8)$$

However, if the human is fully occluded in certain frames t , then the projection error ϵ_t becomes zero, and solving \mathbf{C}_t is impossible. As a result, we adopt the temporal continuity of the person's movement and estimate multi-frame trajectory simultaneously as follows,

$$\begin{aligned} \hat{\mathbf{C}} = \arg \min_{\mathbf{C}} & \sum_{t=1}^F \epsilon_t + \lambda_1 \sum_{t=2}^F \|\mathbf{C}_t - \mathbf{C}_{t-1}\|^2 \\ & + \lambda_2 \sum_{t=2}^{F-1} \|\mathbf{C}_{t-1} + \mathbf{C}_{t+1} - 2\mathbf{C}_t\|^2, \end{aligned} \quad (9)$$

where λ_1 and λ_2 control the first order and second order temporal continuity, respectively, and F is the number of frames used for estimation. From Eq. (9), the global pose trajectory can be estimated even if full-body occlusions occur.

V. EXPERIMENTAL SETUP AND RESULTS

A. Datasets

We evaluate the performance of our proposed method on two motion capture datasets, Human3.6M, and our recorded dataset MMHuman. Following previous work [5], [11], [13], [18], [28]–[32] on Human3.6M, we adopt a 17-joint skeleton, training on five subjects (S1, S5, S6, S7, S8), and testing on two subjects (S9 and S11).

In our experiments, we consider two evaluation protocols: Protocol 1 is the mean per-joint position error after alignment with the ground truth in translation, rotation, and scale [11], [14], [18], [28]–[31]. Protocol 2 aligns predicted poses with the ground-truth only in scale in the camera coordinate with the global trajectory. We do not use the mean per-joint position error without alignment as the evaluation metric since it is sensitive to the camera intrinsic parameters.

B. Implementation Details for 2D Pose Estimation

To fairly evaluate our performance, all the comparison methods use the same 2D inputs. For experiments on Human3.6M, following VideoPose3D [18] we use fine-tuned cascaded pyramid network (CPN) keypoints [33]. For experiments on MMHuman, we use the pre-trained OpenPose [34] to detect 2D keypoints of every frame.



Fig. 4: An example of the generated mask in 2D pose estimation. Horizontal and vertical axes represent frame index and joint index, respectively. The white color represents the occluded joints while the black color represents non-occluded joints.

C. Implementation Details for 3D Pose Estimation with Sequential 2D Joints

For Human3.6M dataset, we randomly generate occlusion masks in both training and testing on the sequential 2D joints to test the effectiveness of temporal gated convolution on handling long-time occlusion. At first, we generate a matrix P with the same size of the 2D input \mathbf{x} with elements sampled from the uniform distribution $\mathcal{U}(0,1)$. Then the occlusion mask M can be set as $M = \mathbb{I}_{\{P > \theta\}}$, where \mathbb{I} is the indicator function that binarizes P according to the threshold θ . However, it is found that such settings cannot well represent long-time occlusion. The occlusion time span is usually short, and the 2D missing joints can be recovered well enough just by linear interpolation. Instead of using the uniformly sampled matrix for binarization directly, we convolve the matrix P with the normalized all-one kernel $\mathbf{1}$ for each joint dimension separately to mimic long-time occlusion, i.e.,

$$M = \mathbb{I}_{\{\hat{P} > \theta\}}, \quad (10)$$

where

$$\hat{P} = \frac{1}{K} P * \mathbf{1}, \quad (11)$$

where $*$ represents convolution operation and K is the temporal kernel size. The larger the kernel size K is, the stronger relationship across the temporal domain there is. As a result, we can change K to represent the intensity of long-time occlusion. An example of the generated mask is shown in Fig. 4. We can see that the generated mask can have both partial occlusion and full occlusion.

For MMHuman, since occlusions naturally occur in the video recording, we use the originally detected 2D keypoints without generating occlusion masks using the strategy mentioned above. If the confidence of the detected 2D keypoint is below a threshold (we set 0.3 in our experiments), then the joint is treated as under occlusion. Different from Human3.6M dataset, our MMHuman contains multi-person in the video, therefore we apply multi-object tracking as a pre-processing step for all the comparison methods to solve the association problem. We adopt the TrackletNet tracking (TNT) [35] for human tracking since it is proven highly effective when handling complicated multi-object tracking cases with various occlusions for both static and moving cameras' recording. The details of TNT can be found in [35].

For the comparison approaches listed in the following section, missing joints from the inputs could cause trouble

TABLE II: 3D Pose Error (Protocol 1, Errors in mm)

	Dir.	Disc.	Eat	Greet	Phone	Photo	Pose	Purch.
Hossain <i>et al.</i> [14]	50.8	60.3	59.6	60.1	75.0	67.9	53.7	59.8
Pavlo <i>et al.</i> [18]	46.3	53.9	49.3	56.8	48.2	57.7	51.2	51.5
Lin <i>et al.</i> [20]	51.3	52.5	49.7	55.0	47.1	64.9	47.8	50.2
Proposed	42.9	43.5	41.5	44.7	41.9	49.4	42.0	41.0
	Sit	SitD.	Smoke	Wait	WalkD.	Walk	WalkT.	Avg.
Hossain <i>et al.</i> [14]	71.1	84.7	62.5	59.2	68.1	69.9	72.0	65.0
Pavlo <i>et al.</i> [18]	53.0	63.6	50.8	45.5	56.3	57.0	56.2	53.2
Lin <i>et al.</i> [20]	53.1	61.7	51.0	47.6	60.1	65.7	62.1	53.9
Proposed	50.4	59.8	43.9	41.3	46.5	35.1	32.6	43.8

in the inference. Therefore, we use linear interpolation in the temporal domain to fill out the missing inputs instead of directly feeding incomplete data. This gives us the best performances we could possibly get using the comparison approaches.

D. Hyper Parameter Setting

For the temporal gated convolution model, we use Amsgrad [36] as the optimizer and train for 80 epochs. We set the batch size as 1024 in the training. For Human3.6M, we adopt an exponentially decaying learning rate schedule, starting from $\eta = 0.001$ with a shrink factor $\alpha = 0.95$ applied on each epoch.

For the global pose trajectory estimation, we set $\lambda_1 = 1$, $\lambda_2 = 1$ and use $F = 100$ frames in the optimization.

VI. EXPERIMENTAL RESULTS

A. Temporal Gated Convolution Model

We compare our method with recent state-of-the-art temporal based models [14], [18], [20], where both open source codes with pre-trained models provided. To explore severe occlusion cases, we randomly generate binary masks with occlusion ratio equaling to 50% in our experiments, i.e., 50% input joints are treated as occluded joints. Table II shows the 3D pose error following protocol 1 on Human3.6M dataset. From the table, we can see our proposed method achieves promising results. It shows the effectiveness of using gated convolution to handle heavy occlusion in the 3D pose estimation.

To further analyze the impact of occlusions in the input 2D pose, We also report the pose reconstruction errors on Human3.6M under protocol 1 with different occlusion ratios in the generated mask in Table III. The results meet our expectations. As the occlusion ratio increases, all methods have higher errors in the 3D pose estimation. When there is 50% percent of missing joints, our method achieves significantly better performance, reporting an average error of 43.8 mm, which is 9.4 mm better than the second-best. Moreover, our proposed method does not have a large fluctuation in the reconstruction errors given different settings of occlusion ratios. It shows the robustness of the proposed method facing different occlusion situations. Note that for the first column

TABLE III: 3D Pose Error on Human3.6M with Different Occlusion Ratios (Errors in mm)

Occ. Ratio	0%	25%	50%
Hossain <i>et al.</i> [14]	44.1	57.2	65.0
Pavlo <i>et al.</i> [18]	36.5	40.3	53.2
Lin <i>et al.</i> [20]	36.8	41.9	53.9
Proposed	37.2	39.2	43.8

in the table when no occlusion occurs, the reconstruction error of the proposed method is slightly higher (about 0.7 mm) than VideoPose3D [18]. This is because since we add a large number of challenging examples with occlusions during training, the data distribution is slightly different between the training set and the testing set if no occlusion occurs in the testing time. This will lead to a slightly higher error than [18]. Besides that, when no occlusion occurs, the input mask becomes a zero-matrix, and the advantage of gated convolution is not well utilized.

Table IV shows the results with the occlusion ratio equaling to 50% following protocol 2 on Human3.6M dataset. The errors are compared in the camera coordinate. Since no global pose trajectory is provided for methods [14], [20], we only compare with [18]. From the table, we can see that the proposed method also outperforms VideoPose3D [18] which uses an additional network to obtain the global trajectory. It shows the effectiveness of the proposed global pose trajectory estimation method.

Table V shows the results on our recorded dataset MMHuman. This dataset, as explained in Section III, includes more inter-person occlusions during human interactions. For this dataset, each detected 2D keypoint by OpenPose is associated with a confidence score. We treat it as occlusion if the confidence below 0.3. To test the generalization of the methods, we use the pre-trained model on Human3.6M and not fine-tuned on MMHuman dataset. From Table V, we can see that our proposed method achieves the best performance. It further proves that our pose estimation method is more robust facing diverse real-world scenarios.

TABLE IV: 3D Pose Error (Protocol 2, Errors in mm)

	Dir.	Disc.	Eat	Greet	Phone	Photo	Pose	Purch.
Pavlo <i>et al.</i> [18]	70.9	92.3	69.0	107.9	66.7	98.1	69.8	73.8
Proposed	53.5	60.3	59.9	87.2	66.7	79.1	59.1	64.8
	Sit	SitD.	Smoke	Wait	WalkD.	Walk	WalkT.	Avg.
Pavlo <i>et al.</i> [18]	68.1	115.4	69.7	93.8	89.8	87.5	77.9	83.4
Proposed	63.9	99.7	57.2	73.1	85.5	107.5	80.0	73.2

TABLE V: 3D Pose Error on MMHuman (Errors in mm)

	ShakeH.	WalkC.	HighF.	PullUp	HandO.	Kungfu	Avg.
Hossain <i>et al.</i> [14]	119.3	119.1	119.1	120.2	115.2	119.3	118.5
Pavlo <i>et al.</i> [18]	87.7	86.1	82.8	88.8	79.1	78.7	84.2
Lin <i>et al.</i> [20]	110.5	116.3	123.2	116.3	107.7	118.4	114.9
Proposed	83.7	74.8	78.6	83.9	75.1	76.7	78.3

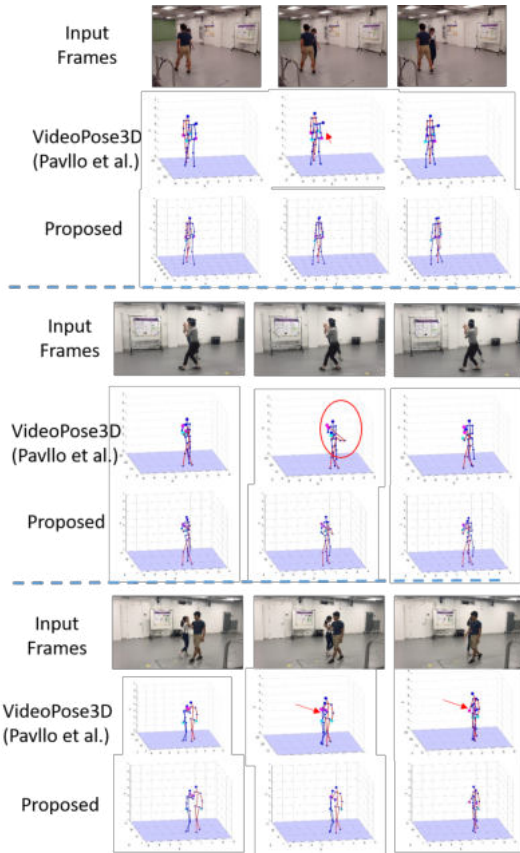


Fig. 5: Qualitative results on MMHuman.

B. Qualitative Results for Occlusion Handling

Fig. 5 shows some qualitative results of our proposed method. The recovered 3D poses using our proposed method as well as state-of-the-art method VideoPose3D [18] are displayed in the camera coordinate. Each example has either severe partial occlusion or full-body occlusion, which is a common problem in human pose estimation the pro-

posed method is targeted at. The first sequence (belonging to "PullUp") shows that being at the end of a inter-person occlusion period, our proposed method demonstrates reliable performance while [18] gets adversely affected. The second sequence ("Kungfu") shows under severe inter-person occlusion, the proposed method still outputs reasonable result, whereas [18] messes up. For the third ("HighFive") sequence, [18] gives unnatural limb estimations, whereas the proposed method predicts natural poses that are consistent with temporal context. It demonstrates that the proposed method is clearly more reliable than other competing methods in occlusion cases.

VII. CONCLUSION

In this paper, we propose a temporal gated convolution model for 3D human pose estimation to address the occlusion issues in the real-world scenario. We also introduce a new 3D human pose dataset MMHuman that features multi-person heavy occlusions and moving cameras, which facilitates evaluation and future research in estimating 3D human poses in real-world scenarios. Meanwhile, the global pose trajectory is efficiently and effectively estimated via temporal back-projection between 2D and 3D joint sequences. We outperform several state-of-the-art 3D pose estimation methods with temporal models on Human3.6M and our self-recorded dataset MMHuman. For future work, we will combine the recent adversarial learning models to enhance the robustness of the 3D pose reconstruction with occlusion.

REFERENCES

- [1] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. S. Huang, "Free-form image inpainting with gated convolution," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 4471–4480.
- [2] H. Jiang, "3d human pose reconstruction using millions of exemplars," in *2010 20th International Conference on Pattern Recognition*. IEEE, 2010, pp. 1674–1677.
- [3] S. Li and A. B. Chan, "3d human pose estimation from monocular images with deep convolutional neural network," in *Asian Conference on Computer Vision*. Springer, 2014, pp. 332–347.

- [4] B. Tekin, A. Rozantsev, V. Lepetit, and P. Fua, "Direct prediction of 3d body poses from motion compensated sequences," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 991–1000.
- [5] G. Pavlakos, X. Zhou, K. G. Derpanis, and K. Daniilidis, "Coarse-to-fine volumetric prediction for single-image 3d human pose," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 7025–7034.
- [6] X. Zhou, Q. Huang, X. Sun, X. Xue, and Y. Wei, "Towards 3d human pose estimation in the wild: a weakly-supervised approach," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 398–407.
- [7] G. Rogez, P. Weinzaepfel, and C. Schmid, "Lcr-net: Localization-classification-regression for human pose," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 3433–3441.
- [8] —, "Lcr-net++: Multi-person 2d and 3d pose detection in natural images," *IEEE transactions on pattern analysis and machine intelligence*, 2019.
- [9] D. Mehta, O. Sotnychenko, F. Mueller, W. Xu, S. Sridhar, G. Pons-Moll, and C. Theobalt, "Single-shot multi-person 3d pose estimation from monocular rgb," in *2018 International Conference on 3D Vision (3DV)*. IEEE, 2018, pp. 120–130.
- [10] A. Zanfir, E. Marinou, M. Zanfir, A.-I. Popa, and C. Sminchisescu, "Deep network for the integrated 3d sensing of multiple people in natural images," in *Advances in Neural Information Processing Systems*, 2018, pp. 8410–8419.
- [11] J. Martinez, R. Hossain, J. Romero, and J. J. Little, "A simple yet effective baseline for 3d human pose estimation," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 2640–2649.
- [12] C.-H. Chen and D. Ramanan, "3d human pose estimation= 2d pose estimation+ matching," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 7035–7043.
- [13] B. Tekin, P. Márquez-Neila, M. Salzmann, and P. Fua, "Learning to fuse 2d and 3d image cues for monocular body pose estimation," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 3941–3950.
- [14] M. Rayat Imtiaz Hossain and J. J. Little, "Exploiting temporal information for 3d human pose estimation," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 68–84.
- [15] R. Gu, G. Wang, and J.-N. Hwang, "Efficient multi-person hierarchical 3d pose estimation for autonomous driving," in *2019 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*. IEEE, 2019, pp. 163–168.
- [16] R. Gu, G. Wang, Z. Jiang, and J.-N. Hwang, "Multi-person hierarchical 3d pose estimation in natural videos," *IEEE Transactions on Circuits and Systems for Video Technology*, 2019.
- [17] G. Moon, J. Y. Chang, and K. M. Lee, "Camera distance-aware top-down approach for 3d multi-person pose estimation from a single rgb image," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 10133–10142.
- [18] D. Pavlo, C. Feichtenhofer, D. Grangier, and M. Auli, "3d human pose estimation in video with temporal convolutions and semi-supervised training," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 7753–7762.
- [19] Y. Cheng, B. Yang, B. Wang, W. Yan, and R. T. Tan, "Occlusion-aware networks for 3d human pose estimation in video," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 723–732.
- [20] J. Lin and G. H. Lee, "Trajectory space factorization for deep video-based 3d human pose estimation," in *BMVC*, 2019.
- [21] C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu, "Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments," *IEEE transactions on pattern analysis and machine intelligence*, vol. 36, no. 7, pp. 1325–1339, 2013.
- [22] L. Sigal and M. J. Black, "Human3.6m: Synchronized video and motion capture dataset for evaluation of articulated human motion," *Brown University TR*, vol. 120, 2006.
- [23] D. Xiang, H. Joo, and Y. Sheikh, "Monocular total capture: Posing face, body, and hands in the wild," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10965–10974.
- [24] D. Mehta, H. Rhodin, D. Casas, P. Fua, O. Sotnychenko, W. Xu, and C. Theobalt, "Monocular 3d human pose estimation in the wild using improved cnn supervision," in *2017 International Conference on 3D Vision (3DV)*. IEEE, 2017, pp. 506–516.
- [25] T. von Marcard, R. Henschel, M. J. Black, B. Rosenhahn, and G. Pons-Moll, "Recovering accurate 3d human pose in the wild using imus and a moving camera," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 601–617.
- [26] G. Liu, F. A. Reda, K. J. Shih, T.-C. Wang, A. Tao, and B. Catanzaro, "Image inpainting for irregular holes using partial convolutions," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 85–100.
- [27] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. S. Huang, "Generative image inpainting with contextual attention," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 5505–5514.
- [28] X. Sun, J. Shang, S. Liang, and Y. Wei, "Compositional human pose regression," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 2602–2611.
- [29] H.-S. Fang, Y. Xu, W. Wang, X. Liu, and S.-C. Zhu, "Learning pose grammar to encode human body configuration for 3d pose estimation," in *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [30] G. Pavlakos, X. Zhou, and K. Daniilidis, "Ordinal depth supervision for 3d human pose estimation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7307–7316.
- [31] W. Yang, W. Ouyang, X. Wang, J. Ren, H. Li, and X. Wang, "3d human pose estimation in the wild by adversarial learning," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 5255–5264.
- [32] D. C. Luvizon, D. Picard, and H. Tabia, "2d/3d pose estimation and action recognition using multitask deep learning," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 5137–5146.
- [33] 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 conference on computer vision and pattern recognition*, 2018, pp. 7103–7112.
- [34] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh, "Realtime multi-person 2d pose estimation using part affinity fields," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 7291–7299.
- [35] G. Wang, Y. Wang, H. Zhang, R. Gu, and J.-N. Hwang, "Exploit the connectivity: Multi-object tracking with trackletnet," in *Proceedings of the 27th ACM International Conference on Multimedia*, 2019, pp. 482–490.
- [36] S. J. Reddi, S. Kale, and S. Kumar, "On the convergence of adam and beyond," *arXiv preprint arXiv:1904.09237*, 2019.