

Multi-view Pose Fusion for Occlusion-Aware 3D Human Pose Estimation

Laura Bragagnolo[✉], Matteo Terreran[✉], Davide Allegro[✉], and Stefano Ghidoni[✉]

Intelligent Autonomous Systems Lab, Department of Information Engineering,
University of Padova, 35131 Padua, Italy
{bragagnolo, terreran, allegrodav, ghidoni}@dei.unipd.it

Abstract. Robust 3D human pose estimation is crucial to ensure safe and effective human-robot collaboration. Accurate human perception, however, is particularly challenging in these scenarios due to strong occlusions and limited camera viewpoints. Current 3D human pose estimation approaches are rather vulnerable in such conditions. In this work we present a novel approach for robust 3D human pose estimation in the context of human-robot collaboration. Instead of relying on noisy 2D features triangulation, we perform multi-view fusion on 3D skeletons provided by absolute monocular methods. Accurate 3D pose estimation is then obtained via reprojection error optimization, introducing limbs length symmetry constraints. We evaluate our approach on the public dataset Human3.6M and on a novel version Human3.6M-Occluded, derived adding synthetic occlusions on the camera views with the purpose of testing pose estimation algorithms under severe occlusions. We further validate our method on real human-robot collaboration workcells, in which we strongly surpass current 3D human pose estimation methods. Our approach outperforms state-of-the-art multi-view human pose estimation techniques and demonstrates superior capabilities in handling challenging scenarios with strong occlusions, representing a reliable and effective solution for real human-robot collaboration setups.

Keywords: 3D Human Pose Estimation · Multi-view Fusion · Human-Robot Collaboration

1 Introduction

Three-dimensional human pose estimation is a fundamental task in computer vision, relevant for many applications, such as virtual reality [16], sport analysis [2] and human action recognition [23]. Accurate 3D human pose estimation is crucial for human-robot collaboration. Reliable 3D human perception is essential for robots to understand the exact position and motion of humans, effectively preventing collisions and ensuring safety. Human pose monitoring also helps in assessing and improving worker ergonomics, enhancing overall workplace health and productivity. Moreover, robust 3D human pose estimation is critical for precise and prompt recognition of human actions and intentions, that can enable smooth and efficient human-robot collaboration [28, 30].

In industrial scenarios, accurate human pose estimation is a real challenge. The typical setting includes a multi-camera setup that frames the human-robot collaboration area—cameras are usually very tilted and often not synchronized. Moreover, the human body is often only partially visible, as surrounding obstacles and equipment often occlude the subject. Such conditions are very under-represented in 3D human pose estimation datasets, which are usually collected in very controlled environments.

As multiple viewpoints are often available in human-robot collaboration workcells, most commonly adopted solutions are multi-view human pose estimation methods, that can accurately estimate 3D human joint position when camera calibration parameters are known. Most state-of-the-art multi-view approaches are, however, based on triangulation of 2D keypoints [3, 11, 37], so they cannot handle significant amounts of occlusions and clutter, as 2D pose estimation quality quickly degrades under these conditions and good predictions cannot be recovered. Recent monocular 3D human pose estimation methods have achieved remarkable results on public benchmarks [7, 34] and exhibit strong generalization to new and diverse scenarios [13, 35]. On the other hand, when huge parts of the body are not visible, they tend to struggle with precise absolute scale prediction and reliable 3D joint configuration. For these reasons, current human pose estimation systems are unsuitable for real industrial scenarios [1].

In this work, we propose a novel multi-view 3D human pose estimation method based on the fusion of 3D keypoints from multiple viewpoints to achieve robust and reliable human pose prediction in human-robot collaboration settings, characterized by the presence of heavy occlusions and few camera viewpoints. While many current approaches are based on multi-view fusion of 2D features and significantly suffer when human joints are obscured, we take as input 3D human pose predictions given by absolute monocular methods. To mitigate joints position artifacts produced in presence of large occlusions, we leverage multiple viewpoints and introduce a novel multi-view 3D keypoint fusion module to merge single-view 3D pose predictions. Considering per-joint reprojection errors on all camera views, we can effectively correct body pose hallucinations and scale recovery errors. Accurate 3D pose localization is then achieved thanks to an optimization procedure on the fused pose, considering reprojection errors on 2D views. The final prediction is further refined considering body symmetry constraints during optimization. With respect to previous work [3, 8] we do not require any prior information or statistics about human limbs length.

We validate our method on data collected from real human-robot collaboration workcells and on Human3.6M-Occluded, a novel benchmark derived from the public dataset Human3.6M [10] adding synthetic occlusions on camera views.

To summarise, the main contributions of this paper are:

1. a multi-view 3D human pose estimation method based on per-joint multi-view fusion and optimization, designed to tackle common problems in complex real scenarios, such as occlusions;
2. limbs length symmetry constraints to refine 3D human pose, that do not depend on any prior information or statistics;

3. an analysis on robustness and generalization capabilities of multi-view 3D human pose estimation methods in real human-robot collaboration workcells;
4. the novel benchmark Human3.6M-Occluded, to test pose estimation algorithms under severe occlusions;
5. extensive evaluation and ablation studies considering multi-view 3D human pose estimation methods on the public dataset Human3.6M and on Human3.6M-Occluded.

2 Related Work

Multi-view 3D Human Pose Estimation Multi-view human pose estimation methods leverage multiple camera viewpoints to provide accurate human pose predictions in 3D space. They typically include two steps: first they extract 2D features from each view independently, and then reconstruct 3D information via triangulation [9, 11, 37]. If triangulation-based approaches are very precise when provided with good quality 2D predictions, occluded environments are traditionally challenging for these methods, as 2D estimated poses tend to be very noisy. Some approaches try to refine 2D predictions and triangulation incorporating epipolar geometry constraints: the epipolar transformer [9] learns attention weights to fuse pixels along the epipolar line in neighboring views; Ma *et al.* [17] introduce Transfusion, a transformer framework for multi-view 3D pose estimation, aiming at directly improving 2D predictions integrating information from different views via geometry position encoding. Similarly, [19, 32] propose to leverage geometric correlations between views to refine 2D pose estimations. Zhang *et al.* [37] present AdaFuse, which tackles human pose estimation in presence of occlusions by learning an adaptive fusion weight to reduce the impact of poor quality views. In [8] and [3] constraints on human limbs length are exploited in the human pose estimation process in order to boost accuracy and robustness. They, however, require knowledge about subject’s bones length or bone length statistics computed on the dataset of interest, making it difficult to apply such methods in unconstrained environments.

These methods typically do not generalize well to new scenarios and camera arrangements without extensive re-training and still struggle when dealing with strong occlusion. Therefore, they are not suitable to industrial multi-camera environments, which are prone to significant occlusions and in which collecting huge amounts of data for training may be impractical or even impossible.

Absolute 3D Human Pose Estimation Recently, many 3D human pose estimation approaches have focused on single-view images or videos [7, 15, 34]. This is an ill-posed problem, affected by depth ambiguities and occlusions. Most works [14, 36] mainly focus on root-relative human pose estimation, which provide human pose predictions with respect to a reference frame centered in the pelvis joint (root). Nevertheless, for real-world applications such as human-robot collaboration, it is crucial to accurately estimate the human body pose in the

global 3D space [23, 30]. Pavlakos et al. [22] used a simple anthropometric approach, optimizing the absolute person distance based on 2D pixel positions and root-relative depth estimates from volumetric heatmaps, ensuring that the back-projected skeleton’s bone lengths matched the average bone lengths from the training set. This method, however, has limited generalization power, as it requires pre-calculated statistics about human bone lengths. Zhan *et al.* [33] use an intrinsic-parameter-invariant representation to convert the 2D keypoints from pixel space to 3D rays in a normalized 3D space; such rays are then fused from consecutive frames using temporal convolution. Sáráandi *et al.* introduce MeTRAbs [24, 26], tackling scale and distance estimation of 3D poses while also addressing body occlusions. They train a fully-convolutional network to output location and depth heatmaps for 2D joints, providing accurate metric-scale depth information. Still, when huge parts of the body are not visible, predicted scale and pose may be unreliable.

Occlusion-Aware Human Pose Estimation Occlusion-aware human pose estimation has been primarily addressed by root-relative single-view methods. To boost resilience to missing joints, synthetically occluded data can be exploited during training, as suggested in [25]. Cheng *et al.* [4] make use of 2D confidence heatmaps of keypoints and introduce an optical-flow consistency constraint, to filter out the unreliable estimations. More recently, generative models have been successfully employed to retrieve occlusion-robust 3D pose, as in [36]. Several single-view datasets representing scenarios with occlusions, such as 3DPW-Occ [35] and 3DOH50K [35] have also been proposed. Such works tackle root-relative 3D human pose estimation from monocular images and videos and mainly focus on everyday life scenarios.

Our work addresses the gaps in the current literature by proposing an occlusion-aware 3D human pose estimation method from multi-view images, targeting human-robot collaboration environments. Differently from current multi-view approaches, mostly relying on 2D feature triangulation, we aggregate single-view 3D predictions from absolute human pose estimation algorithms, to achieve robustness against strong occlusions.

3 Method

The proposed method consists of three main parts: (i) single-view 3D pose prediction; (ii) multi-view 3D pose aggregation; (iii) 3D pose refinement based on reprojection error minimization. In the first step, for each camera in the network, a candidate 3D pose is estimated by means of a monocular 3D human pose estimation algorithm, such as [24, 33]. The 3D poses obtained from each view are then aggregated together by means of a dedicated multi-view fusion module, which aims to combine all the independent single-view predictions in a same 3D pose. Finally, the aggregated 3D pose is further refined in an optimization module based on constrained reprojection error minimization.

3.1 Single-view 3D Pose Estimation

As a first step, our approach takes as input multiple images captured simultaneously by C cameras with known calibration parameters. Three-dimensional human poses are extracted from 2D images following a top-down approach: (i) a people detector finds all the possible patches containing a person, (ii) for each patch, a CNN predicts the 3D location of the human body joints with respect to the camera reference frame. Our method assumes that all the cameras in the network are calibrated with respect to the same reference frame (e.g., robot base frame), so that the rotation R_k and translation t_k of each camera k with respect to the reference frame are known. For each view, the final output is a tuple of 3D poses, 2D bounding boxes and confidence scores.

3.2 Multi-view Fusion Module

Monocular 3D human pose estimation methods typically provide accurate and plausible body joint predictions when considering relative joint positions. However, they typically struggle in estimating the absolute distance of each body joints from the camera, especially in presence of large body occlusions. Moreover, when portions of the human body are not visible, there exists an ambiguity between distance and size: small objects close to the camera may look similar to large objects farther away from the camera [26], leading to incorrect 3D pose prediction in metric space. In order to alleviate this problem, we propose to leverage information from multiple views, from the intuition that averaging the absolute positions of the single-view estimations should give a good approximation of the true location of the human body.

Consider a network of C cameras and denote as $\{\hat{P}_{j,i}\}$ the set of 3D joints predicted from camera i . The j^{th} joint of the fused skeleton can be expressed as the weighted average of the j^{th} joint for all single-view detections:

$$\bar{P}_j = \frac{\sum_{i=1}^C w_i^j \hat{P}_{j,i}}{\sum_{i=1}^C w_i^j} . \quad (1)$$

The use of weights w_i allows to penalize the views framing the subject from a higher distance or the detections with lower confidence. When the human body is partially or heavily occluded, in fact, the 3D pose obtained from monocular 3D pose estimation methods is not fully reliable: if some parts of the body are not visible, the pose estimation algorithm uses clues from the visible parts to estimate the positions of the missing joints. Significant occlusions reduce evidence for the obscured joints' positions, leading to incorrect final predictions compared to the true pose. For instance, if a person's legs are occluded by a table, the model lacks evidence to determine whether the person is standing with crossed, bent or straight legs.

To better incorporate all the available information, we propose the use of per-joint weights: if a 3D joint (e.g. the right ankle) is incorrectly predicted in one view, the error given by projecting that joint over all the other views will

be higher with respect to the reprojection error computed for that same joint correctly predicted from another view.

In particular, given the set $\{\hat{P}_{j,i}\}$ of 3D joints predicted from camera i , a single joint $\hat{P}_{j,i}$ can be easily projected onto the image plane of a camera k by means of the camera intrinsic parameters K_k and extrinsic parameters R_k and t_k . Averaging the reprojection error between the 2D joints and the predictions on each frame $\{p_{j,i}\}$, allows to define the per-joint weight $w_{j,i}$ as follows:

$$w_{j,i} = \frac{1}{e_{j,i}}, \quad e_{j,i} = \frac{1}{C} \sum_{k=1}^C \left\| [K_k | \mathbf{0}] [R_k | t_k] \hat{P}_{j,i} - p_{j,k} \right\|^2. \quad (2)$$

3.3 Optimization Module

The fusion module merges multi-view 3D predictions, leveraging multiple view-points to mitigate artifacts in joint positions caused by occluded body parts. However, when significant occlusions occur in multiple camera views, averaging predictions may not suffice, as estimated pose and camera distance can be slightly incorrect in all views. To address this issue, we formulate an optimization problem to refine the fused 3D pose \bar{P}_j , minimizing the per-joint reprojection error across all C cameras available:

$$\operatorname{argmin}_{\bar{P}_j} \sum_{k=1}^C \sum_{j=1}^J \left\| [K_k | \mathbf{0}] [R_k | t_k] \bar{P}_{j,i} - p_{j,k} \right\|^2. \quad (3)$$

Without additional constraints, the optimization process can change the position of all body joints; however, there are symmetries in the human body that can be exploited to guide optimization. In particular, we design a limb length symmetry constraint considering four pairs of bones: left and right upper arms, lower arms, upper legs and lower legs. For each pair we constrain the limb length to be the same, that is, we enforce left and right limbs to have consistent measures. In particular, we compute the left (L) and right (R) 3D limb lengths as the euclidean distance $d(\bar{P}_i, \bar{P}_j)$ between the corresponding left and right joints. If u and v are a pair of corresponding limb joints, the final cost for the symmetry constraint is given by:

$$C_{sym} = \sum_{u=1}^J \sum_{v=1}^J \mathbf{1}(u, v) \cdot \left\| d(\bar{P}_{u_L}, \bar{P}_{v_L}) - d(\bar{P}_{u_R}, \bar{P}_{v_R}) \right\|^2, \quad (4)$$

$$\text{where } \mathbf{1}(u, v) = \begin{cases} 1 & \text{if } (u, v) \text{ is upper/lower arm, upper/lower leg} \\ 0 & \text{otherwise.} \end{cases}$$

Note that, with respect to previous work [3, 8, 22], we do not require any prior information on human limbs length or pre-computed statistics on datasets' 3D annotations. Finally, the optimization problem is defined as follows:

$$\operatorname{argmin}_{\bar{P}_j} \sum_{k=1}^C \sum_{j=1}^J \left\| [K_k | \mathbf{0}] [R_k | t_k] \bar{P}_{j,i} - p_{j,k} \right\|^2 + C_{sym}. \quad (5)$$

4 Experiments

In order to evaluate the capabilities of our method to handle significant occlusions and clutter, following the work of [25] and [35], we propose Human3.6M-Occluded, a novel benchmark derived from Human3.6M [10], adding synthetic occlusions on camera views. The robustness and reliability of our method in challenging human-robot environments is then inspected considering data sequences of human-robot collaboration tasks acquired in a real industrial workcell and in a laboratory environment. We finally conduct ablation studies to analyse the behaviour of our framework in presence of few occluded viewpoints and camera synchronization errors, typical in industrial multi-camera setups.

4.1 Experimental Settings and Metrics

Human3.6M and Human3.6M-Occluded The Human3.6M dataset is currently one of the largest multi-view datasets providing 3D human pose annotations. It consists of 3.6 million frames captured by 4 synchronized 50 Hz cameras. 3D and 2D keypoint ground truth is obtained by a marker-based motion capture system composed of 10 IR-cameras. The dataset includes 11 different subjects performing 15 daily actions, such as eating, walking, and giving directions. Each scene features only one subject. Following the traditional protocol, subjects S9 and S11 are considered for evaluation. To provide fair comparisons with previous work [11], [17], sequences of subject S9 with incorrect 3D annotations (parts of ‘Greeting’, ‘SittingDown’ and ‘Waiting’) were ignored for final results calculations.

Motivated by the lack of human body occlusions in Human3.6M and of public multi-view human pose estimation datasets including real occlusions in cluttered scenarios, we introduce the novel benchmark Human3.6M-Occluded. This is derived from the Human3.6M dataset by synthetically adding significant body occlusions, by taking objects and animals images from the Pascal VOC 2012 dataset [5] as occluders, as done in [25]. Considering the four available camera views for each scene in the dataset, we partially cover the subject’s body on three out of four views. This kind of scenario is very difficult to deal with, as each body joint is not guaranteed to be visible by at least two views. An occluded view includes two random objects overlapping the human bounding box. Object size and location inside the box are chosen at random. Sample images from the Human3.6M-Occluded benchmark are shown in Figure 1.

Differently from previous works presenting similar approaches, we make reproducible code available to generate the dataset, to encourage further research on occlusion-aware human pose estimation and to facilitate comparison with other methods. Code is available here ¹.

Human-Robot Workcell To test the robustness of our human pose estimation method in real industrial environments, we conducted experimental validation

¹ <https://github.com/laurabragagnolo/human3.6m-occluded>




Fig. 1: Sample sequence from the Human3.6M-Occluded dataset, with occlusions on three of four views.

on data acquired in two different human-robot collaboration scenarios: an industrial workcell developed as use case for the DrapeBot project [6, 29], and a laboratory collaborative workcell. The former has been designed for collaborative transportation of flexible materials and covers an area of about 20 m². It was equipped with four Intel RealSense Depth D455 sensors surrounding an ABB industrial manipulator at an average distance of 4 m from the base of the robot. The latter, equipped with four Intel RealSense Depth D455 cameras and a Franka Emika Panda collaborative manipulator, has been designed for the collaborative assembly of furniture, (*e.g.*, wooden chair). The industrial workcell provides a very challenging testing scenario, with a lot of clutter in the scene and very tilted cameras. The laboratory workcell represents a simpler scenario, with less occluders. In both settings cameras are not perfectly synchronized. Sample images extracted from acquired sequences are shown in Figure 2.

In all sequences collected, participants were equipped with the Xsens MVN Awinda motion capture system [27] to record whole-body kinematics. The system consists of 17 wearable inertial measurement units (IMUs) placed all over the body, providing an accurate full-body 3D pose not affected by occlusions that has been used as a ground truth information. Since the 3D pose obtained from inertial sensors suffers from drifting errors in the absolute position over long period of time, we correct this effect using the information provided by the multi-camera pose tracking solution OpenPTrack [20], by aligning the inertial-based 3D pose with the 3D position of the the neck joint detected by OpenPTrack, which has been empirically proven to be the most robust to occlusions among all body joints.

Evaluation Metrics The 3D pose estimation accuracy is measured considering the Mean Per Joint Position Error (MPJPE), that is, the Euclidean distance between predicted 3D position and ground truth, for each joint. We consider both the absolute and the root-relative MPJPE, that requires aligning the pelvis joints coordinates of the estimated and ground truth pose. While the absolute error takes into account the accuracy of the global localization of the prediction, the relative error focuses on the accuracy of the human keypoints configuration




Fig. 2: Sample images from the sequences acquired in real human-robot collaboration scenarios. a) Industrial workcell; b) Laboratory workcell.

with respect to the pelvis joint, also known as the root joint. The MPJPE is expressed in millimeters, where lower values represent higher accuracy.

Implementation Details The single-view 3D predictions for multi-camera fusion are obtained using MeTRAbs [24]. The considered model has been trained considering multiple skeleton conventions, *e.g.* H36M [10], 3DHP [18], SMPL-X [21], and several 3D human pose estimation datasets, including Human3.6M [10], CMU Panoptic [12] and TotalCapture [31]. For each of the experiments presented below, we do not perform any fine-tuning.

4.2 Results on Human3.6M and Human3.6M-Occluded

The results of the evaluation of our method and state-of-the-art multi-view human pose estimation algorithms on the Human3.6M dataset is reported in Table 1 and Table 2. Considering the root-relative MPJPE (Table 1), in scenarios with very few occlusions and in which the human body is well visible from all the four viewpoints, our method performs slightly worse with respect to current approaches. For multi-view methods, however, the absolute position error is much more interesting to look at, since it represents the quality of the 3D global localization of the human body prediction. Results for the absolute MPJPE are reported in Table 2. Our framework surpasses most recent approaches and gives comparable results to strong baselines such as algebraic triangulation [11]. Looking the state of the art, it is worth noticing that our method and TransFusion [17] yield very similar results, while AdaFuse [37] produces the strongest performance on this benchmark, considering both relative and absolute errors.

The loss of accuracy with respect to the state of the art is mainly due to the fact that competing approaches rely on heavy geometrical constraints, such as triangulation and epipolar lines, to perform 2D feature fusion. When the human body is fully visible and 2D features are very reliable, exploiting geometry gives extremely precise results that, however, poorly adapt to scenarios with occlusions, in which 2D joint predictions are not as reliable. Differently, our

Table 1: The results of evaluation on the Human3.6M dataset. The table presents the MPJPE relative to pelvis for state-of-the-art multi-view human pose estimation methods.

MPJPE relative to pelvis, mm																
Methods	Dir.	Disc.	Eat	Greet	Phone	Photo	Pose	Purch.	Sit	SitD.	Smoke	Wait	WalkD.	Walk	WalkT.	Avg
RANSAC (as in [37])	18.7	19.5	19.5	20.9	22.1	17.2	19.6	24.8	31.6	22.5	21.7	22.0	17.4	22.3	17.3	21.7
T-Iskakov <i>et al.</i> [11]	20.4	22.6	20.5	19.7	22.1	20.6	19.5	23.0	25.8	33.0	23.0	21.6	20.7	23.7	21.3	22.6
V-Iskakov <i>et al.</i> [11]	19.9	20.0	18.9	18.5	20.5	19.4	18.4	22.1	22.5	28.7	21.2	20.8	19.7	22.1	20.2	20.8
Zhang <i>et al.</i> [37]	17.0	18.1	17.3	21.1	19.1	16.4	18.1	19.2	26.3	19.4	19.1	21.0	16.8	19.6	16.9	19.4
Ma <i>et al.</i> [17]	25.2	27.8	24.8	24.3	27.2	28.8	24.3	25.9	63.4	32.6	27.4	25.1	27.0	25.2	24.7	29.4
Wan <i>et al.</i> [32]	19.5	20.9	19.5	18.3	21.1	20.0	17.9	21.3	23.9	30.1	21.6	19.9	18.9	22.8	19.5	21.1
Ours	25.7	27.2	27.2	27.2	31.4	26.1	29.3	33.8	36.2	29.3	29.7	28.5	25.7	30.6	26.0	29.0

Table 2: The results of evaluation on the Human3.6M dataset. The table presents the absolute MPJPE for state-of-the-art multi-view human pose estimation methods.

MPJPE absolute, mm																
Methods	Dir.	Disc.	Eat	Greet	Phone	Photo	Pose	Purch.	Sit	SitD.	Smoke	Wait	WalkD.	Walk	WalkT.	Avg
RANSAC (as in [37])	18.3	19.8	20.0	18.4	22.1	17.3	19.8	23.6	28.1	22.8	23.9	17.9	18.0	21.2	17.6	21.5
T-Iskakov <i>et al.</i> [11]	21.7	23.7	22.2	20.4	26.7	24.2	19.9	22.6	31.2	35.6	26.8	21.2	20.9	24.6	21.1	24.5
V-Iskakov <i>et al.</i> [11]	18.0	18.3	16.5	16.1	17.4	18.2	16.5	18.5	19.4	20.1	18.2	17.4	17.2	19.2	16.6	17.9
Zhang <i>et al.</i> [37]	17.4	18.3	17.8	17.8	19.4	16.8	18.5	19.2	22.1	19.4	21.4	16.7	17.4	19.6	17.4	19.4
Ma <i>et al.</i> [17]	24.4	26.4	23.4	21.1	25.2	23.2	24.7	33.8	29.8	26.4	26.8	24.2	23.2	26.1	23.3	25.8
He <i>et al.</i> [9]	25.7	27.7	23.7	24.8	26.9	31.4	24.9	26.5	28.8	31.7	28.2	26.4	23.6	28.3	23.5	26.9
Moliner <i>et al.</i> [19]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	26.0
Ours	23.8	24.6	23.4	24.1	26.9	24.0	25.3	28.3	31.7	25.9	26.9	24.2	23.0	27.3	24.2	25.6

method uses information about the geometry of the scene, *i.e.*, rotation and translation between cameras, for the sole 3D joint reprojection and then finds the optimal pose that minimizes reprojection errors on multiple views. This enforces looser geometrical constraints on the final pose, which lead to slightly worse results on the standard Human3.6M with respect to the other methods, but more robust performance in presence of occlusions.

To validate this, we now test our approach in presence of significant occlusions, on the proposed Human3.6M-Occluded benchmark. For comparisons we select AdaFuse [37] and TransFusion [17], which have a specific focus on human pose estimation in unconstrained environments with occlusions and provide official codes and training weights on Human3.6M. We also compare with RANSAC (implementation from [37]) and the algebraic triangulation [11], taken as baselines. For this experiment all the models are trained on the original Human3.6M dataset, without any fine-tuning on the Human3.6M-Occluded data.

Results for the relative and absolute MPJPE on this benchmark are shown in Table 3 and Table 4. Our system provides the best performances with respect to both evaluation metrics, proving its superior capabilities in handling challenging scenes with occlusions with respect to competing methods relying on 2D features and strong geometrical constraints.

Table 3: The results of evaluation on the proposed Human3.6M-Occluded benchmark. The table presents the MPJPE relative to pelvis for our method and state-of-the-art multi-view human pose estimation algorithms.

MPJPE relative to pelvis, mm																
Methods	Dir.	Disc.	Eat	Greet	Phone	Photo	Pose	Purch.	Sit	SitD.	Smoke	Wait	WalkD.	Walk	WalkT.	Avg
RANSAC (as in [37])	302.8	69.7	57.8	48.1	49.2	94.1	67.5	132.4	137.7	58.0	66.3	54.8	38.1	267.4	56.9	97.0
T-Iskakov <i>et al.</i> [11]	69.4	75.1	83.3	65.4	76.2	91.0	106.3	192.9	632.4	117.7	91.9	105.0	49.5	101.9	70.7	134.9
Zhang <i>et al.</i> [37]	32.5	41.3	33.9	35.0	37.7	38.0	36.4	50.5	83.0	47.4	47.9	38.7	28.9	40.3	37.6	43.5
Ma <i>et al.</i> [17]	48.8	65.8	51.9	48.7	61.7	57.1	62.3	75.1	384.8	224.8	74.1	50.6	122.9	97.6	35.9	97.3
Ours	34.1	40.6	36.2	35.7	42.9	35.9	40.0	60.8	68.1	42.7	41.3	35.5	30.8	43.0	34.3	41.9

Table 4: The results of evaluation on the proposed Human3.6M-Occluded benchmark. The table presents the absolute MPJPE for our method and state-of-the-art multi-view human pose estimation algorithms.

MPJPE absolute, mm																
Methods	Dir.	Disc.	Eat	Greet	Phone	Photo	Pose	Purch.	Sit	SitD.	Smoke	Wait	WalkD.	Walk	WalkT.	Avg
RANSAC (as in [37])	305.5	61.3	51.9	47.2	46.6	130.0	62.2	117.5	117.5	53.7	65.5	55.0	35.2	60.2	51.6	80.7
T-Iskakov <i>et al.</i> [11]	67.1	71.3	80.4	63.4	74.8	88.5	100.9	188.7	611.2	111.7	92.2	110.6	47.1	99.7	69.8	127.4
Zhang <i>et al.</i> [37]	31.5	39.5	32.4	32.3	36.3	36.6	34.7	46.4	71.7	45.4	47.8	37.0	28.0	37.8	35.3	41.1
Ma <i>et al.</i> [17]	48.1	65.2	50.9	47.9	60.8	56.1	61.6	74.1	384.2	224.2	73.4	49.9	122.3	96.5	35.8	96.5
Ours	32.0	36.9	32.9	31.5	37.9	32.9	36.2	53.7	64.9	37.5	38.9	32.8	27.8	38.5	31.4	37.8

4.3 Results on the Human-Robot Workcell Scenario

We now inspect the robustness of our human pose estimation approach on real industrial environments. For this, we consider the data sequences collected in the industrial workcell and in our laboratory, presented earlier in this section. Results are shown in Table 5 and Table 6. Sequences acquired in the industrial workcell include the participation of two different subjects (performing two sequences each) and are numbered from 1 to 4; sequences collected in our laboratory include only one subject and are numbered from 5 to 7.

Our framework demonstrates a great advantage with respect to the other approaches, producing reliable predictions even on such challenging scenes, with an average absolute joint position error of about 14 cm. On the other hand, the performance of the other methods dramatically decreases in the same context, leading to errors of several meters with respect to the ground truth reference, as shown in Figure 3 and Figure 4. Differently, our framework produces consistently better results on both scenarios, displaying excellent generalization power and adaptability to real multi-camera setups.

While results on Human3.6M-Occluded already confirmed the superior performance of our method in presence of significant occlusions, results on the real human-robot collaboration scenario really highlight the shortcomings of current methods and the true advantage of our framework. Nevertheless, the discrepancy that we observe between our results and those provided by the literature is due to conditions that characterize industrial multi-camera setups, such as camera synchronization errors and large occlusions on all viewpoints. These factors

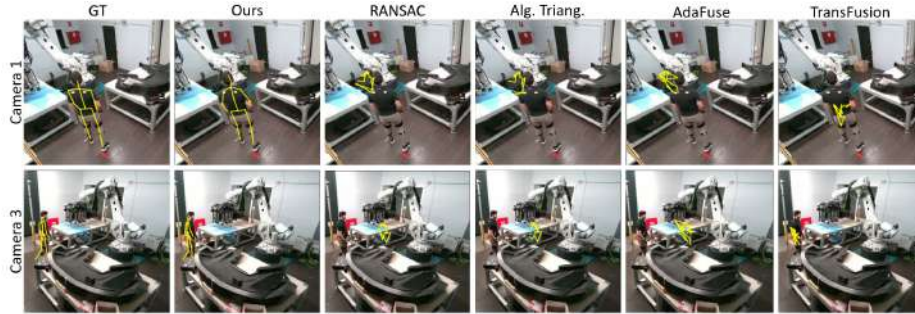


Fig. 3: Qualitative comparison between our approach and the state of the art on the industrial workcell scenario.



Fig. 4: Qualitative results of our approach and state-of-the-art multi-view methods on the laboratory collaborative workcell.

are not considered in public datasets [10, 12], and thus, not handled by most approaches in the literature.

4.4 Ablation Studies

In the remainder of this section we consider the Human3.6M and the Human3.6M-Occluded benchmarks to carry out ablation studies on different aspects typical of real industrial scenarios: (i) imperfect synchronization of cameras and (ii) heavy occlusions with few viewpoints available.

Camera Synchronization While public datasets are acquired in very constrained and controlled environments employing reliable camera synchronization mechanisms, such conditions are hardly met in industrial environments. To assess the advantages of our method, we now isolate the effect of slight errors in camera synchronization on multi-view human pose estimation methods. Considering the Human3.6M dataset, we study the variation of the mean joint position error as we simulate synchronization errors on an increasing number of cameras.

Table 5: The results of evaluation on data sequences collected in real human-robot workcells. The table presents the MPJPE relative to pelvis for our method and state-of-the-art multi-view human pose estimation algorithms.

MPJPE relative to pelvis, mm									
Methods	Seq 1	Seq 2	Seq 3	Seq 4	Seq 5	Seq 6	Seq 7	Avg	
RANSAC (as in [37])	541.0	265.2	328.7	312.9	261.1	287.7	260.4	322.4	
T-Iskakov <i>et al.</i> [11]	439.5	257.5	315.8	307.0	270.7	288.1	257.8	305.2	
Zhang <i>et al.</i> [37]	580.7	370.0	402.1	447.0	347.2	344.9	321.7	401.9	
Ma <i>et al.</i> [17]	505.7	453.7	446.9	481.4	412.5	430.5	410.9	448.8	
Ours	90.4	98.7	87.9	88.8	100.9	114.0	109.8	98.6	

Table 6: The table presents the absolute MPJPE on data collected in real human-robot workcells for our method and state-of-the-art multi-view human pose estimation algorithms.

MPJPE absolute, mm									
Methods	Seq 1	Seq 2	Seq 3	Seq 4	Seq 5	Seq 6	Seq 7	Avg	
RANSAC (as in [37])	8981.5	7295.3	8109.3	7336.4	2800.5	2817.2	2813.5	5736.2	
T-Iskakov <i>et al.</i> [11]	8784.7	7291.9	8112.4	7327.2	2796.3	2813.9	2808.4	5705.0	
Zhang <i>et al.</i> [37]	9000.7	7276.1	8116.3	7308.6	2754.0	2765.7	2773.4	5713.5	
Ma <i>et al.</i> [17]	8907.5	7370.7	7296.0	8211.2	2788.6	2785.8	2786.8	5735.2	
Ours	110.7	106.1	171.7	154.8	116.7	198.1	140.9	142.7	

At first, three of the four cameras will give the correct frame corresponding to time t while we simulate a synchronization error on the fourth one. The de-synchronized camera will give a frame corresponding to a wrong timestamp t_e , defined according to Equation (6).

$$t_e = t + e \quad \text{with} \quad e \sim \text{Uniform}(\{-2, -1, +1, +2\}) \quad (6)$$

Then, we simulate a synchronization error on two of the four cameras and, finally, on three out of four cameras. Figure 5 shows how the absolute and relative MPJPE vary over the three experiments. Our method proves to be the most robust to camera synchronization errors in terms of both relative and absolute MPJPE, which increase of only 1 and 2 mm when three cameras out of four are unsynchronized.

Number of Occluded Viewpoints To study the robustness of multi-view pose estimation approaches in presence of few occluded viewpoints, we consider Human3.6-Occluded and progressively reduce the number of available cameras to three and to two. It is worth noticing that in Human3.6M-Occluded, for each scene, three out of four views contain occlusions and the occluded cameras are selected at random. This means that when removing a camera, it can happen that for some scenes all views are occluded.

Figure 6 shows the relative and absolute MPJPE with reduced number of viewpoints for the compared methods. For a fair comparison, we do not consider

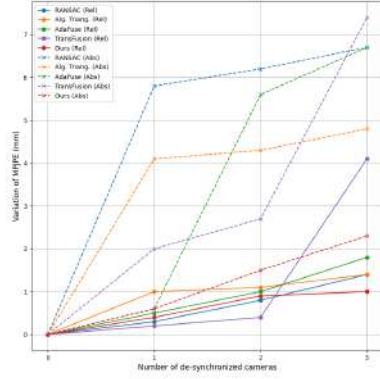


Fig. 5: Variation of relative and absolute MPJPE with increasing number of de-synchronized cameras. Solid lines represent variation of the relative error, dashed lines represent variation of the absolute errors.




Fig. 6: Relative and absolute MPJPE gradually removing the number of view-points, with occlusions. Solid lines represent variation of the relative error, dashed lines represent variation of the absolute errors.

TransFusion [17] for this ablation study, as this approach requires to be re-trained when the number of available cameras change. One can easily notice that our framework is the most robust to significant occlusions with few viewpoints, yielding a relative and absolute joint position error of about 6 cm when tested on scenes with only two cameras and occlusions on at least one view.

5 Conclusion

In this work, we proposed a novel multi-view 3D pose fusion approach for robust and reliable 3D human pose estimation in human-robot collaboration scenarios, generally characterized by significant occlusions and limited camera viewpoints. While most methods rely on 2D keypoints triangulation, we aggregate single-view 3D joint predictions, proposing a per-joint weighting depending on the mean reprojection error on all views, to mitigate the effect of incorrect predictions. The fused 3D pose is further refined via reprojection error optimization, introducing limb length symmetry constraints to enhance estimation accuracy.

Our framework was evaluated on challenging scenarios with sever occlusions, considering data sequences from real human-robot collaboration workcells and the proposed benchmark Human3.6M-Occluded. Our method outperforms state-of-the-art techniques on heavy occlusions and exhibits strong generalization capabilities to unseen environments, making it a suitable solution for 3D human pose estimation in challenging human-robot collaboration settings.

Acknowledgements

The research leading to these results has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 101006732 (DrapeBot).

References

1. Boldo, M., De Marchi, M., Martini, E., Aldegheri, S., Quaglia, D., Fummi, F., Bombieri, N.: Real-time multi-camera 3d human pose estimation at the edge for industrial applications. *Expert Syst. Appl.* **252**, 124089 (2024)
2. Bridgeman, L., Volino, M., Guillemaut, J.Y., Hilton, A.: Multi-person 3d pose estimation and tracking in sports. In: *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.* pp. 0–0 (2019)
3. Chen, Z., Zhao, X., Wan, X.: Structural triangulation: A closed-form solution to constrained 3d human pose estimation. In: *Eur. Conf. Comput. Vis.* pp. 695–711. Springer (2022)
4. Cheng, Y., Yang, B., Wang, B., Yan, W., Tan, R.T.: Occlusion-aware networks for 3d human pose estimation in video. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 723–732 (2019)
5. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>
6. Ghidoni, S., Terreran, M., Evangelista, D., Eitzinger, C., Zambal, S., Villagrossi, E., Pedrocchi, N., Castaman, N., Malecha, M.: A smart workcell for human-robot cooperative assembly of carbon fiber parts. In: *Atti della 4^a Conferenza Italiana di Robotica e Macchine Intelligenti* (2021). <https://doi.org/10.5281/zenodo.6367920>
7. Gong, J., Foo, L.G., Fan, Z., Ke, Q., Rahmani, H., Liu, J.: Diffpose: Toward more reliable 3d pose estimation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 13041–13051 (2023)
8. Guidolin, M., Tagliapietra, L., Menegatti, E., Reggiani, M.: Hi-ros: Open-source multi-camera sensor fusion for real-time people tracking. *Comput. Vis. Image Und.* **232**, 103694 (2023)
9. He, Y., Yan, R., Fragkiadaki, K., Yu, S.I.: Epipolar transformer for multi-view human pose estimation. In: *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.* pp. 1036–1037 (2020)
10. Ionescu, C., Papava, D., Olaru, V., Sminchisescu, C.: Human3. 6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Trans. Pattern Anal. Mach. Intell.* **36**(7), 1325–1339 (2013)
11. Isakov, K., Burkov, E., Lempitsky, V., Malkov, Y.: Learnable triangulation of human pose. In: *Int. Conf. Comput. Vis.* pp. 7718–7727 (2019)
12. Joo, H., Liu, H., Tan, L., Gui, L., Nabbe, B., Matthews, I., Kanade, T., Nobuhara, S., Sheikh, Y.: Panoptic studio: A massively multiview system for social motion capture. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 3334–3342 (2015)
13. Joo, H., Neverova, N., Vedaldi, A.: Exemplar fine-tuning for 3d human model fitting towards in-the-wild 3d human pose estimation. In: *3DV.* pp. 42–52. IEEE (2021)

14. Kocabas, M., Huang, C.H.P., Hilliges, O., Black, M.J.: Pare: Part attention regressor for 3d human body estimation. In: *Int. Conf. Comput. Vis.* pp. 11127–11137 (2021)
15. Li, W., Liu, H., Tang, H., Wang, P., Van Gool, L.: Mhformer: Multi-hypothesis transformer for 3d human pose estimation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 13147–13156 (2022)
16. Lin, H.Y., Chen, T.W.: Augmented reality with human body interaction based on monocular 3d pose estimation. In: *International Conference on Advanced Concepts for Intelligent Vision Systems.* pp. 321–331. Springer (2010)
17. Ma, H., Chen, L., Kong, D., Wang, Z., Liu, X., Tang, H., Yan, X., Xie, Y., Lin, S.Y., Xie, X.: Transfusion: Cross-view fusion with transformer for 3d human pose estimation. *arXiv preprint arXiv:2110.09554* (2021)
18. Mehta, D., Sotnychenko, O., Mueller, F., Xu, W., Sridhar, S., Pons-Moll, G., Theobalt, C.: Single-shot multi-person 3d pose estimation from monocular rgb. In: *3DV.* pp. 120–130. IEEE (2018)
19. Moliner, O., Huang, S., Åström, K.: Geometry-biased transformer for robust multi-view 3d human pose reconstruction. In: *FG.* pp. 1–8. IEEE (2024)
20. Munaro, M., Basso, F., Menegatti, E.: Openprtrack: Open source multi-camera calibration and people tracking for rgb-d camera networks. *RAS* **75**, 525–538 (2016)
21. Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3d hands, face, and body from a single image. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 10975–10985 (2019)
22. Pavlakos, G., Zhu, L., Zhou, X., Daniilidis, K.: Learning to estimate 3d human pose and shape from a single color image. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 459–468 (2018)
23. Peng, K., Roitberg, A., Yang, K., Zhang, J., Stiefelhagen, R.: Delving deep into one-shot skeleton-based action recognition with diverse occlusions. *IEEE Trans. Multimedia* **25**, 1489–1504 (2023)
24. Sáráandi, I., Hermans, A., Leibe, B.: Learning 3d human pose estimation from dozens of datasets using a geometry-aware autoencoder to bridge between skeleton formats. In: *WACV.* pp. 2956–2966 (2023)
25. Sáráandi, I., Linder, T., Arras, K.O., Leibe, B.: How robust is 3d human pose estimation to occlusion? *arXiv preprint arXiv:1808.09316* (2018)
26. Sáráandi, I., Linder, T., Arras, K.O., Leibe, B.: Metrabs: metric-scale truncation-robust heatmaps for absolute 3d human pose estimation. *IEEE Trans. Biom. Behav. Identity Sci.* **3**(1), 16–30 (2020)
27. Schepers, M., Giuberti, M., Bellusci, G., et al.: Xsens mvn: Consistent tracking of human motion using inertial sensing. *Xsens Technol* **1**(8), 1–8 (2018)
28. Terreran, M., Barcellona, L., Ghidoni, S.: A general skeleton-based action and gesture recognition framework for human–robot collaboration. *RAS* **170**, 104523 (2023)
29. Terreran, M., Ghidoni, S., Menegatti, E., Villagrossi, E., Pedrocchi, N., Castaman, N., Gottardi, A., Eitzinger, C., Vescovi, L., Salemi, G., et al.: A smart workcell for cooperative assembly of carbon fiber parts guided by human actions. In: *Atti della 4^a Conferenza Italiana di Robotica e Macchine Intelligenti* (2022)
30. Terreran, M., Lazzaretto, M., Ghidoni, S.: Skeleton-based action and gesture recognition for human-robot collaboration. In: *International Conference on Intelligent Autonomous Systems.* pp. 29–45. Springer (2022)
31. Trumble, M., Gilbert, A., Malleson, C., Hilton, A., Collomosse, J.P.: Total capture: 3d human pose estimation fusing video and inertial sensors. In: *Brit. Mach. Vis. Conf.* vol. 2, pp. 1–13. London, UK (2017)

32. Wan, X., Chen, Z., Zhao, X.: View consistency aware holistic triangulation for 3d human pose estimation. *Comput. Vis. Image Und.* **236**, 103830 (2023)
33. Zhan, Y., Li, F., Weng, R., Choi, W.: Ray3d: ray-based 3d human pose estimation for monocular absolute 3d localization. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 13116–13125 (2022)
34. Zhang, J., Tu, Z., Yang, J., Chen, Y., Yuan, J.: Mixste: Seq2seq mixed spatio-temporal encoder for 3d human pose estimation in video. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 13232–13242 (2022)
35. Zhang, T., Huang, B., Wang, Y.: Object-occluded human shape and pose estimation from a single color image. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 7376–7385 (2020)
36. Zhang, Y., Ji, P., Wang, A., Mei, J., Kortylewski, A., Yuille, A.: 3d-aware neural body fitting for occlusion robust 3d human pose estimation. In: *Int. Conf. Comput. Vis.* pp. 9399–9410 (2023)
37. Zhang, Z., Wang, C., Qiu, W., Qin, W., Zeng, W.: Adafuse: Adaptive multiview fusion for accurate human pose estimation in the wild. *Int. J. Comput. Vis.* **129**, 703–718 (2021)