

Body Size and Depth Disambiguation in Multi-Person Reconstruction from Single Images

Nicolas Ugrinovic Adria Ruiz Antonio Agudo Alberto Sanfeliu Francesc Moreno-Noguer
Institut de Robòtica i Informàtica Industrial, CSIC-UPC, Barcelona, Spain
{nugrinovic, aruiz, aagudo, asanfeliu, fmoreno}@iri.upc.edu

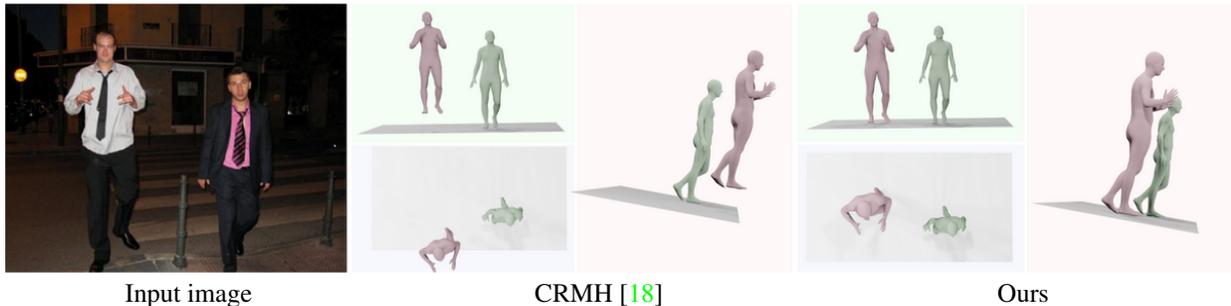


Figure 1: We present a multi-person reconstruction approach that solves the inherent body size and depth ambiguity problem. The figure shows the results of a state-of-the-art approach [18] (center), and the results of our approach (right). Note that while the frontal view of both approaches is very similar, the side view in [18] depicts a large error. Our main contribution consist in proposing an optimization scheme that enforces the people to have their feet on the ground, which allows disambiguating body size from depth and producing accurate spatial arrangements.

Abstract

We address the problem of multi-person 3D body pose and shape estimation from a single image. While this problem can be addressed by applying single-person approaches multiple times for the same scene, recent works have shown the advantages of building upon deep architectures that simultaneously reason about all people in the scene in a holistic manner by enforcing, e.g., depth order constraints or minimizing interpenetration among reconstructed bodies. However, existing approaches are still unable to capture the size variability of people caused by the inherent body scale and depth ambiguity. In this work, we tackle this challenge by devising a novel optimization scheme that learns the appropriate body scale and relative camera pose, by enforcing the feet of all people to remain on the ground floor. A thorough evaluation on MuPoTS-3D and 3DPW datasets demonstrates that our approach is able to robustly estimate the body translation and shape of multiple people while retrieving their spatial arrangement, consistently improving current state-of-the-art, especially in scenes with people of very different heights. Code can be found at: https://github.com/nicolasugrinovic/size_depth_disambiguation

1. Introduction

While recent works on single-person 3D reconstruction have shown impressive results [6, 13, 21, 24, 25, 27, 37, 40, 20, 3, 4, 48, 39, 8, 56], the problem of simultaneously reconstructing multiple humans is still in its infancy. The straightforward solution consists in regarding different people as independent instances and estimating the body shapes and poses one by one using a single-person approach. This strategy, however, may result in inconsistent spatial arrangements and erroneous poses of the reconstructed people. A few works have improved the coherence of the reconstructions by simultaneously reasoning about all people in the image, either for estimating their body poses [14, 34, 53], poses and shape [58, 59, 18, 61, 10], or poses and motion [2, 1]. For instance, Jiang *et al.* [18], have considered global constraints accounting for inter-person occlusions, interpenetration between meshes and depth ordering.

Nevertheless, while showing very promising results, previous approaches are still prone to fail in situations where people in the images have different heights. This is because the constraints considered so far are not able to handle the inherent body size / depth ambiguity, and result in consistent body 2D reprojections, but wrong relative scales and

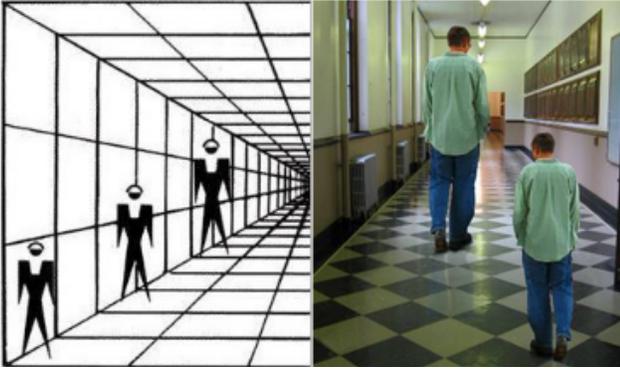


Figure 2: Two examples of the Ponzo illusion [43] on the perception of the human height. In both cases, the contextual information of the ground floor helps to perceive the height and depth of each person.

depths (see Fig. 1-center).

In this paper we present a novel approach for multi-person reconstruction that handles the depth/body size ambiguity. For this purpose, we will consider a novel constraint that estimates body scale and relative camera-body translation while enforcing the feet of all people to remain on the ground. This observation, does indeed have psychological groundings on the geometrical-optical Ponzo illusion [43], that suggests that the human mind judges an object’s size based on its background, in our case the ground floor, and the preconceived prior that humans have their feet on the ground (see Fig. 2).

In order to implement this constraint we present a simple but effective coarse-to-fine strategy where we first use [18, 20] to extract an initial estimation of the shape and pose of all people in the image. This estimate may contain gross errors, like people floating in the air or with large discrepancies in their true size (see again Fig. 1-center). We then extract the ground plane from the image using an off-the-shelf depth [45] and semantic segmentation estimation network [54], and compute the 3D normal direction of that plane. Finally, we devise a learning scheme, in which translation and scale parameters per person are optimized so as to minimize the feet-to-plane 3D distance and the 2D reprojection error. We observed that with only these two terms, more computationally demanding losses preventing interpenetrability of the 3D meshes were not necessary.

We exhaustively evaluate the proposed approach on the MuPoTS-3D [31] and 3DPW [52] datasets, and show a consistent improvement over a depth ordering metric in comparison to recent state-of-the-art. To further stress the significance of our results, we propose two additional metrics that assess the accuracy of the retrieved body sizes and the inter-person distances. In both these metrics we also obtain remarkably better results than previous methods. Fig. 1-right, shows an example output of our approach where we

are able to estimate correct body shape and position of two people with very different sizes.

2. Related work

3D Human Pose Estimation. Estimating the 3D pose of human joints has received growing interest mostly for its importance in real-world applications and by the release of large-scale public MoCap and multi-camera datasets [17, 32, 19, 44]. A lot of important studies have been done on this topic for single-person [28, 51, 7, 33, 65, 36, 41, 49, 57, 64] and multi-person [46, 47, 35, 26, 31]. 3D pose estimation for single-person is done directly from the image or lifting 2D joints to 3D joints and for multi-person either by using top-down or bottom-up approaches. To face the shortcomings of loss of information due to reprojection in 3D pose estimation, parametric models of the human body have been extensively used. Examples of these are SMPL [29], Frank [20] or GHUM [55]. Although it is very difficult to annotate large-scale datasets with these models, there are optimization-based methods such as [7, 40, 9] that allow to fit the parameters very well using only 2D and 3D joint annotations. Due to its simplicity and robustness several works [21, 23, 42, 38, 40, 25, 5, 24, 58, 12, 63, 60] build upon the SMPL model to reconstruct human bodies from single images. The SMPL model can be used to recover the full body shape (as a 3D mesh) and joints.

Multi-person 3D Pose and Shape Estimation. Recently there have been some works that extend 3D pose/shape estimation to a multi-person setting. The most relevant works in this area are [58, 59, 18, 61, 9, 50].

Zanfir *et al.* [59] are the first ones to use a bottom-up model to estimate the pose and shape of multiple-people. In a concurrent but radically different work, Zanfir *et al.* [58] propose a top-down multi-staged approach. They first estimate initial pose for each person and then jointly optimize to absolute poses for multiple people in the image. They also include a ground plane constraint. However, they estimate it using initial 3D joint predictions and require a video sequence to do so. In contrast to them, we cannot assume that our initial 3D joints estimations are correct and we work with a single image. So we take a very different approach for estimating the ground plane normal. While these two works have been one of the firsts to address this multi-person problem, subsequent methods such as [18] show improved results over them.

Very recent approaches [61, 50] propose a novel mesh representation to account for cases of occlusions and also reason about the most probable depth location for each person in the scene while presenting bottom-up architectures. Jiang *et al.* [18] leverage the architecture proposed by HMR [21] to get multi-person shape/pose estimates. They

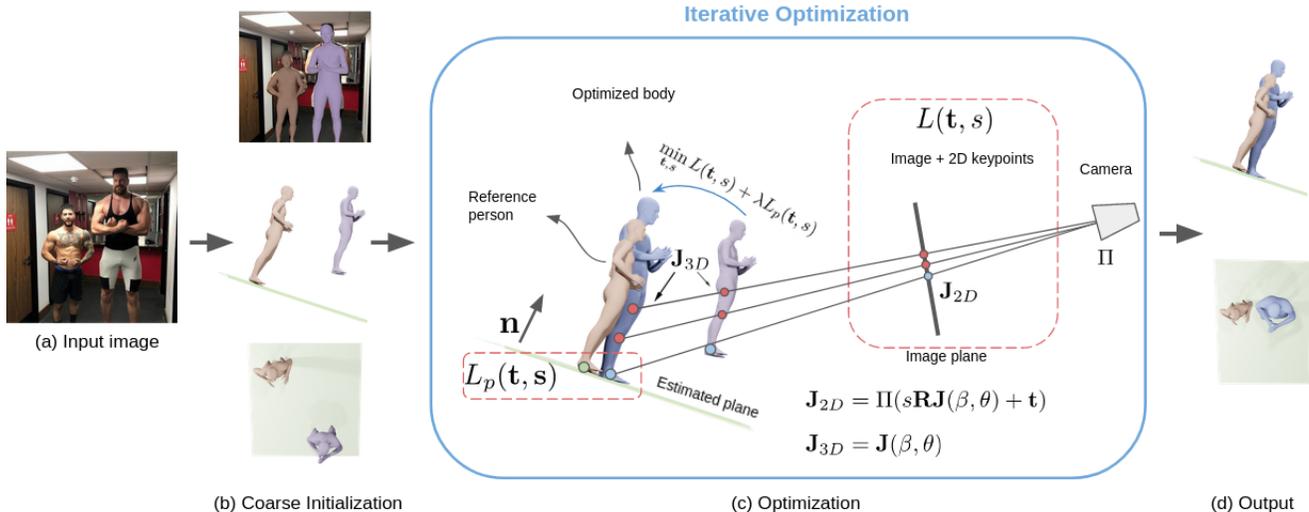


Figure 3: **Overview of our approach.** Given an input image (a), we first estimate the 2D keypoints and SMPL parameters as the initialization (b). Then we estimate the ground plane normal and minimize the reprojection error along with the ground plane constraint for reconstruction (c). We are able to correct the spatial arrangement in the reconstructed 3D scene by forcing all body meshes to touch the ground with their feet (d). Notice the difference with respect to the results from the initial estimation where people are floating in the air (b).

include two novel losses: one that penalizes collisions between people and another depth ordering aware loss that uses segmentation masks to enforce a correct depth order. Due to the nature of this loss, the depth order can be misestimated when no people overlap with each other. As we will see, our method is capable of facing this limitation.

Human-Scene Interaction. Recently, increasing attention has been given to the context of the scene when estimating the body pose and shape along with the spatial arrangement of humans in the scene. So far very few works [62, 10, 15, 16] have been done in this area. For example, PHOSA [62], which is closer to our approach, seeks to model several people in a scene and their interactions with object on images on-the-wild. They propose a two stage approach where they first estimate each person and object in the scene and then globally optimize for spatial arrangement. Instead of objects, we use the ground plane which unifies all people in the scene.

3. Method

We next describe our approach (see Figure 3 for an overview). First, we define the mesh representation adopted (Sec. 3.1). Then our problem formulation and objective function are presented in Sec. 3.2. Finally, we explain our feet-to-ground constraint (Sec. 3.3) and optimization scheme (Sec. 3.4).

3.1. SMPL and Pose/Shape Estimation

We adopt SMPL [30] to represent the 3D body geometry. This model is specially appealing in our context given its ability to represent a large number of shape and pose variations with a few number of parameters. In particular, SMPL encodes the body geometry in two vectors: $\beta \in \mathbb{R}^{10}$, containing the parameters modelling the shape, and $\theta \in \mathbb{R}^{72}$, which encode the information regarding the person pose. Using previous definitions, SMPL also implements a function:

$$M(\beta, \theta) : \theta \times \beta \mapsto \mathbf{V} \in \mathbb{R}^{3B}, \quad (1)$$

which estimates a set of B vertices \mathbf{V} of the body mesh from the pose and shape parameters. Additionally, the model also provides a linear mapping \mathcal{J} so that $\mathbf{J} = \mathcal{J}(\mathbf{V})$ which allows to estimate the position of K body joints $\mathbf{J} \in \mathbb{R}^{3K}$ from the mesh vertices \mathbf{V} . In the following sections, we will refer to the i -th joint of a body mesh with parameters θ and β as $\mathbf{J}(\beta, \theta)_i$.

Along with the body mesh, a weak-perspective camera $\Pi = [\sigma, t_x, t_y] \in \mathbb{R}^3$ is parametrized to project the mesh into image coordinates. To position the humans in the 3D space, the weak-perspective camera can be converted to the perspective camera projection by assuming a fixed focal length f for all images, where the depth of the person is determined by the relation $d = f/\sigma$ and the absolute translation is given by:

$$\mathbf{t} = [t_x, t_y, d]. \quad (2)$$

Given a set of 2D body joints $\tilde{\mathbf{J}} \in \mathbb{R}^{2K}$ in pixel coordinates and a confidence value c_i associated to each of them,

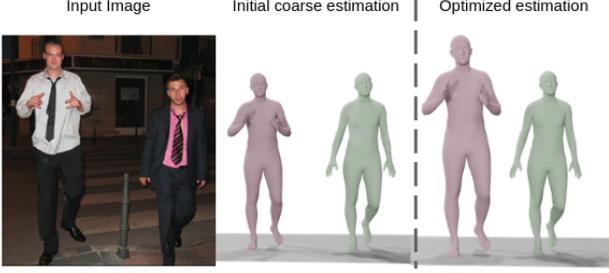


Figure 4: **An illustration of the difference in body sizes.** The reconstructed 3D bodies are placed at the same depth and on the same plane. Our method scales body sizes accordingly. In contrast, other state-of-the-art methods struggle capturing the variation of body sizes.

the pose and shape parameters defining the body geometry can be estimated by solving the following optimization problem:

$$\min_{\theta, \mathbf{R}, \mathbf{t}, s} \sum_i c_i \|\tilde{\mathbf{J}}_i - \Pi(\mathbf{R}\mathbf{J}(\beta, \theta)_i + \mathbf{t})\| + L_p \quad (3)$$

where \mathbf{R} is a global rotation matrix and L_p a regularizer. Intuitively, the defined problem attempts to find the pose and shape parameters which minimize the distance between the estimated 2D joints and the projection of the 3D mesh in image coordinates.

3.2. Problem Formulation

Consider an input RGB image $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$ with N persons and $(\theta^{1:N}, \beta^{1:N})$ to be the set of SMPL parameters obtained by solving Eq. (3) for each individual. Our goal is to obtain the set of translation and scale parameters $(\mathbf{t}^{1:N}, s^{1:N})$ for each person, which best explain the scene and produces feasible results with correct depths and sizes.

In order to obtain a set of translation and scale parameters leading to realistic results, we propose to solve a similar minimization problem to the one introduced in Eq. (3). In our case, however, we introduce two important modifications. Firstly, we jointly optimize the individual parameters for each person by defining the loss as:

$$L(\mathbf{t}^{1:N}, s^{1:N}) = \sum_{i,n} c_i^n \|\tilde{\mathbf{J}}_i^n - \Pi(s^n \mathbf{R}^n \mathbf{J}(\beta^n, \theta^n)_i + \mathbf{t}^n)\|, \quad (4)$$

where \mathbf{R}^n is the estimated rotation matrix for the person n and $(\tilde{\mathbf{J}}_i^n, c_i^n)$ is the i -th 2D joint in image coordinates and its associated confidence, respectively. Note that here, we introduce an additional scale parameter s that modifies the overall size of the body (including height) without the need to change β .

Secondly, we introduce a regularizer conditioned on the scale and translation parameters that aims to minimize the feet-to-ground distance and avoid body configurations flying on the air. We will detail this loss in the following subsection.

Finally, we jointly find the optimal translation and scale for each person by minimizing:

$$\min_{\mathbf{t}^{1:N}, s^{1:N}} L(\mathbf{t}^{1:N}, s^{1:N}) + \lambda L_p(\mathbf{t}^{1:N}, s^{1:N}), \quad (5)$$

where λ is a hyper-parameter controlling the trade-off between the minimized projection error and the regularization term.

3.3. Scale and Translation aware Feet-to-Ground Constraint

The inclusion of the scale and translation aware regularizer in Eq. (5) is one of the key contributions of our paper. This is motivated by the observation that in most real scenes, the depicted individuals are typically standing on a common ground. While this observation has already been exploited in [58], the ground-to-plane distance in this work is conditioned uniquely on the 3D joints of the feet. In contrast, we condition this distance using also the body size scale and translation vector. By doing this we make it possible to address the depth-size ambiguity problem. To illustrate this, consider Figure 3 (b). Given an estimated normal vector of a plane representing the scene ground, the estimation of the translation and scale parameters obtained with Eq. (3) produces unrealistic results where one of the persons is "flying". However, this ambiguity can be addressed by our optimization scheme, in which we translate each of the individuals and change the size of their bodies so as to get close to the ground plane, while still minimizing the 2D reprojection error.

Regularizer definition: Assume that, for a given image, we have access to the vector \mathbf{n} defining the normal of the ground plane. To enforce all people in the scene to touch the ground, we use:

$$L_p = \sum_n |(\mathbf{x}_l^n(\mathbf{t}^n, s^n) - \mathbf{p}) \cdot \mathbf{n}| + |(\mathbf{x}_r^n(\mathbf{t}^n, s^n) - \mathbf{p}) \cdot \mathbf{n}|, \quad (6)$$

where \mathbf{x}_l^n and \mathbf{x}_r^n are the left and right ankle 3D joints of the n -th person obtained after translating and scaling its mesh using the estimated parameters \mathbf{t}^n and s^n . Being k the index of these joints in $\mathbf{J}(\beta^n, \theta^n)$, they are computed as:

$$\mathbf{x}_k = s^n \mathbf{R}\mathbf{J}(\beta^n, \theta^n)_k + \mathbf{t}^n. \quad (7)$$

Additionally, \mathbf{p} is a reference point fixed in the plane which is chosen as the ankle of a reference person, see Figure 3(c). Notice that we purposefully translate the plane to \mathbf{p} . We select the reference person as the one with lower initial reprojection error.

Ground Normal Estimation: Restricting and determining a person's movement in the 3D space can be possible with

a correct ground plane normal \mathbf{n} . If, on top of this, we force the reprojected body joints to match a set of reference 2D keypoints, we can easily estimate correct translation and scale parameters. Unfortunately, we typically do not know the direction of \mathbf{n} in real scenarios and thus, it must be estimated. For this purpose, we first obtain a depth map $\hat{\mathbf{d}} \in \mathbb{R}^{\mathbf{H} \times \mathbf{W}}$ by applying an off-the-shelf depth estimator to the input image [45]. Then, we use panoptic segmentation [19] to identify the pixels corresponding to the ground regions in the image. Finally, we take the depth values from these ground regions and, using their corresponding x, y pixel coordinates, obtain a set of 3D points. With these points we fit a plane, acquiring its normal \mathbf{n} . To handle outliers we use the RANSAC [11].

To enforce the condition for Eq. (6) that \mathbf{p} must be a fixed point in the plane, we assume that \mathbf{p} is a reference point with known depth. So we take \mathbf{p} as a point in the 3D space and translate the entire plane to this point. All subsequent translation estimations are, therefore, relative to this point.

3.4. Optimization

The loss and regularization terms minimized in Eq. (5) are fully-differentiable w.r.t the optimized variables. Therefore, to solve our optimization problem we use a gradient descent approach. In particular, we initialize $\mathbf{t}^{1:N}$ with the estimations obtained using Eq. (3) and set all $s^{1:N}$ to one (mean scale). We iteratively minimize the objective by computing the gradients of the translation and scale parameters.

4. Experiments

4.1. Datasets

In our experiments, we validate our method conducting exhaustive experiments over two standard benchmarks for human pose/shape estimation.

MuPoTS-3D [31]: It is a multi-person dataset providing 3D ground truth for people in the scene. We use the same test sequences as in [18].

3DPW [52]: It is a multi-person in-the-wild dataset, which features diverse motions and scenes. It contains 60 video sequences with 3D joints and translations annotations. We use this dataset to especially test our models capability of generality of applicability to challenging in-the-wild scenarios. We use only the test sequences that present the ground truth for more than one person as we focus on multi-person reconstruction.

4.2. Metrics

In order to evaluate the quality of the multi-person 3D pose and shape estimations, we use three different metrics: (i) pairwise person’s depth order (d_{ord}), (ii) pairwise normalized distances between persons in the scene (d_{norm}) and

(iii) pairwise person’s height discrete comparison (h_{ord}). The first metric was introduced in previous works [18, 61] and quantifies the percentage of correctly estimated ordinal depth relations between all pairs of people in the image. The other two metrics are proposed by us and allow to conduct a more thorough evaluation.

Pairwise normalized distances d_{norm} : This metric complements the depth order by measuring how accurate the distance relations are between pairs of people in the image. Thus, measuring the quality of the estimated translations. This metric is motivated by the observation that an estimation can have correct depth order for all people but have unrealistic or disproportional distances between them. If d_{ord} is correct but each person is very far from each other, in comparison to what is perceived in the image, then the estimation cannot be considered accurate. The opposite also holds true, if d_{norm} has a low score, it does not necessarily mean that the depth order is correct. Formally, this metric is defined as follows:

$$d = \sum_i \sum_j \|(\mathbf{t}_i - \mathbf{t}_j)\| \quad (8)$$

$$d_{norm} = \frac{\sum_{f=0}^m (d_f / d_f^{max} - \hat{d}_f / \hat{d}_f^{max})}{m}, \quad (9)$$

where m is the number of frames and i and j are indices for the i -th and j -th persons in the frame f . Also, d_f corresponds to the pairwise distances of the ground truth values and \hat{d}_f , to estimated ones, which are both calculated with Eq. 8. Here, \mathbf{t} is the estimated translation of each person in the image. For each frame we normalize all distances by the maximum pairwise distance (d_f^{max}).

Discrete height comparison h_{ord} : To determine the quality of the estimated body scales, we use people’s height as a proxy measure. One can reason that if a model is able to correctly estimate the height of each person, providing a good reprojected, then a correct spatial arrangement is more likely to occur. This metric is the percentage of correctly estimated ordinal height relations between all pairs of people in the image, i.e., smaller to bigger or vices-versa. This is very similar to d_{ord} but applied to heights. This metric helps us determine that the scale has correctly increased or decreased a person’s size in order to provide a better estimation. The heights are calculated by the euclidean distance across the limbs connecting the head and feet joints.

Overall, these three metrics attempt to measure the aspects of multi-person reconstruction that our method focuses on improving: translation and scale.

4.3. Implementation details

We jointly optimize for both translation and scale using the ADAM optimizer [22] with learning rate 1e-2 for a total of 600 iterations. As reference 2D keypoints, we use the

Method	MuPoTS-3D			3DPW		
	$d_{ord} \uparrow$	$d_{norm} \downarrow$	$h_{ord} \uparrow$	$d_{ord} \uparrow$	$d_{norm} \downarrow$	$h_{ord} \uparrow$
FrankMocap [20]	85.56	0.492	52.79	81.61	0.463	50.53
Ours (w/ FrankMocap)	88.03	0.367	49.65	82.69	0.557	51.60
CRMH [18]	92.20	0.351	52.43	76.13	0.737	51.32
Ours (w/ CRMH)	95.58	0.243	57.75	85.00	0.516	55.80
BMP* [61]	94.50	-	-	-	-	-

Table 1: **Results on representative datasets.** We present the results on three state-of-the-art methods and use two of them as initialization for our approach. Evaluation is done on MuPoTS-3D and 3DPW datasets with the metrics described in Sec. 4.2. *The values were copied directly from the paper as there is no publicly code nor data available.

Method	MuPoTS-3D		
	$d_{ord} \uparrow$	$d_{norm} \downarrow$	$h_{ord} \uparrow$
Only Reprojection	92.51	0.415	55.48
Only Plane	95.10	0.293	51.14
Ours w/ GT-J2D	94.62	0.318	59.28
Ours	95.58	0.243	57.75

Table 2: **Ablative study for proposed losses.** We present the effect of each individual loss (reprojection, plane) and the effect of using ground truth 2D keypoints as inputs to the model instead of detected ones by and off-the-shelf detector.

3D reprojected joints obtained from the initial SMPL estimations from either CRMH [18] or FrankMocap[20]. We present experiments with these two initializations in Sec. 4, however, given the nature of our method, any other initialization can be used. We use semantic segmentation from Detectron2 [54] for our baseline and use MiDasv3. [45] as depth estimator. For all cases we use a fixed focal length of $f = 1000$.

4.4. Comparison with simple baseline

In this first experiment, we evaluate the effectiveness of the proposed ground-plane regularization. For this purpose, we compare our method with a baseline that incorporates depth information but does not explicitly model the scene ground. Concretely, our baseline estimates the depth of each person in the scene. This is done by averaging the depth values of all the pixels belonging to a person’s mask (obtained with from semantic segmentation). Then, we optimize for s and t only over the x and y axis while maintaining z equal to the estimated depth using Eq. 4 as objective function.

Comparing ourselves to a method that directly uses people’s depth information is important as one part of our goal is to improve depth order. We can see in Table 3 that our method outperforms this baseline by a significant margin.

Method	MuPoTS-3D		
	$d_{ord} \uparrow$	$d_{norm} \downarrow$	$h_{ord} \uparrow$
Baseline	87.34	0.635	55.09
Ours	95.58	0.243	57.75

Method	3DPW		
	$d_{ord} \uparrow$	$d_{norm} \downarrow$	$h_{ord} \uparrow$
Baseline	82.24	0.508	54.27
Ours	85.00	0.516	55.80

Table 3: **Comparison with our baseline.** As a sanity check we compare our method with the baseline described in Sec. 4.4.

This justifies the choice of using the ground plain as a constraint. It is more reliable to estimate the ground plane’s normal and use it as a constraint than to recover individual person’s depth from an estimated depth map. The baseline relies on segmentation masks for computing a person’s depth. In many cases these segmentation masks are not able to be recovered, especially under the presence of occlusions. This results in incorrect depth estimations. Thus, the poor performance of the baseline. In contrast, segmenting the ground is easier which facilitates the depth estimation of the ground plane and consequently the normal \mathbf{n} . For this reason, our method is able to make a better use of depth information from a given depth map to improve 3D human pose-shape recovery.

4.5. Comparison with the state-of-the-art

We compare our method to other multi-person approaches with regards to spatial arrangement, specially for correct depth order and body size and leave out any human pose comparison as we do not change the pose from the initial estimations. We measure the quality of the translation and scale estimations as they are crucial for facing the depth-size ambiguity.



Figure 5: **Qualitative results.** We visualize the results of our method from different viewpoints: front (green background), top (blue background), and side (red background) views. We compare the results with CRMH [18] and improve depth order and overall scene coherence. Note that CRMH does not estimate a plane. We add the same plane estimated by our method for the visualization purposes.

We compare ourselves with two state-of-the-art methods: CRMH [18] and BMP [61], and an additional baseline for which we apply FrankMocap [20] to the multi-person setting. The results in both datasets (MuPoTS-3D, 3DPW) are presented in Table 1.

Improving depth order. We first evaluate our proposed method for depth order and correct translation of body meshes in the 3D space, as they are both related. The results are shown in the first two columns of each dataset in Table 1. It can be seen that our approach is able to improve the spatial arrangement of the initial methods used ([20] and [18]). Concretely, it can be seen that our method improves

CRMH by 3.38% (95.58% vs. 92.20%) and FrankMocap by 2.47% (88.03% vs. 85.56%) in the d_{ord} metric for the MuPoTS-3D dataset and CRMH by 8.87% (85.00% vs. 76.13%) and FrankMocap by 1.08% (82.69% vs. 81.61%) in 3DPW dataset, which contains images in-the-wild. We also improve d_{norm} score by a good margin. For example, in MuPoTS-3D we improve CRMH’s score by 0.108 (0.243 vs. 0.351) and FrankMocap by 0.125 (0.367 vs. 0.492). A similar improvement is presented in 3DPW dataset when using CRMH as initial estimation. Our method also outperforms the current state-of-the-art [61] by 1.08% in the depth order metric. Please note that BMP [61] results are

taken from the original paper and not recomputed by us as we do for the other methods as their code and data is not yet publicly available.

Relative body sizes. After stabilising that we can estimate more accurate translation parameters, we now focus on the people’s body sizes, which also play an important role when trying to correctly perceive the 3D information of the scene. In the same fashion as before, we measure the correct height order (h_{ord}), meaning the scale relations (bigger/smaller) between people in a scene. Here we find that our method improves CRMH by 5.32% (57.75% vs. 52.43%) in MuPoTS-3D, CRMH by 4.48% (55.80% vs. 51.32%) in 3DPW and FrankMocap by 1.07% (51.60% vs. 50.53%) in 3DPW. A lower score on this metric shows that there are estimations where people that should be bigger in comparison to others in the scene, are actually estimated as smaller and vice-versa.

4.6. Ablation studies

In Table 2 we analyse the effect of using only one of the terms of our objective function. First, we use only the reprojection term for the optimization and, later, we use only the feet-ground constraint. We do this to validate that both of the terms contribute to the overall performance of the model. Additionally, we use ”perfect” 2D keypoints as input, taken from the ground truth, to see if there is an upper-bound limit to our method’s performance and how this affects the results. Note that a discussion on the effects and limitations of our feet-ground constraint is presented in the Supplementary Material.

It can be seen that using only reprojection helps significantly to estimate better height results as h_{ord} is relatively high in comparison to using only the ground plane constraint. However, this term has no contribution with the depth order nor correct proportional distance between people, having the worst scores on d_{ord} and d_{norm} , just as expected. It can be seen that using only the feet-ground constraint alone leads to better depth ordering and better d_{norm} score than when using only reprojection. This is expected as we first move the people towards the plane in the camera direction during optimization, meaning that reprojection is partially conserved even without the reprojection term. As we can see in the results, heights estimations here are the worst as h_{ord} scores lower, which is also expected.

We also perform an ablation study that uses ground truth 2D keypoints annotations used by the reprojection loss. As it turns out, our method surprisingly yields to better results in the d_{ord} and d_{norm} metrics than by using ground truth 2D keypoint annotations. However, this setting yields to the best results in height order (h_{ord}) showing that better reference keypoints results in better height estimates but not necessarily into better depth ordering.

4.7. Robustness to initialization

To show that our framework can work along with any other initial method, we present results in both datasets using FrankMocap [20] estimations as initial point. Note that this method is originally trained for single-person estimation. The results presented in Table 1 show a good improvement over [20] used directly as a multi-person approach. Our method is also robust to 2D keypoint initialization as it has been established in the previous section.

4.8. Qualitative Results

We present our qualitative results in Figure 5. As expected, our method produces results with more realistic spatial arrangement while capturing the diversity of body sizes. This yields improved depth ordering and better proportional distances among people in the scene when compared to method [18]. In all cases shown, we use CMRH for initial estimations. The image in the last row of Figure 5 presents a failure case. Here, one person is jumping and the ground constraint is not completely met. Although, in this case, the depth order is still improved by our method, the distance between each person may not be accurate. More qualitative results, including failure cases, and discussion of the effect of the ground plane constraint can be found in the Supplementary Material.

5. Conclusions

In this paper, we have presented a novel optimization scheme in order to address the inherent depth/size ambiguity in multi-person 3D body pose estimation. In particular, our method jointly computes the optimal translation and scale parameters of all the individuals in the scene by imposing a constraint forcing them to stand into a shared ground-plane. In order to estimate this information, we use off-the-shelf depth-map and semantic segmentation methods to automatically extract the ground-plane normal from the image. Comparing our approach to state-of-the-art methods over benchmark datasets, we show that the proposed pipeline provides more coherent solutions in terms of the relative scale and translation of the estimated persons. More concretely, we are able to consistently improve the performance by a considerable margin in multiple metrics evaluating the accuracy of the pose parameters. Finally, it is worth mentioning that our method is model agnostic and only requires a set of initial estimations for the 3D joints. Therefore, it can be combined with any previous or future method for human mesh reconstruction from still images.

Acknowledgements: This work is partly supported by the Spanish government with the project MoHuCo PID2020-120049RB-I00, the ERA-Net Chistera project IPALM PCI2019-103386 and María de Maeztu Seal of Excellence MDM-2016-0656. Adria Ruiz acknowledges financial support from MICINN (Spain) through the program Juan de la Cierva.

References

- [1] A. Agudo and F. Moreno-Noguer. DUST: dual union of spatio-temporal subspaces for monocular multiple object 3D reconstruction. In *CVPR*, 2017. [4321](#)
- [2] A. Agudo and F. Moreno-Noguer. Robust spatio-temporal clustering and reconstruction of multiple deformable bodies. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 41(4):971–984, 2019. [4321](#)
- [3] T. Alldieck, M. Magnor, B. L. Bhatnagar, C. Theobalt, and G. Pons-Moll. Learning to reconstruct people in clothing from a single RGB camera. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, jun 2019. [4321](#)
- [4] T. Alldieck, G. Pons-Moll, C. Theobalt, and M. Magnor. Tex2shape: Detailed full human body geometry from a single image. In *IEEE International Conference on Computer Vision (ICCV)*, 2019. [4321](#)
- [5] R. Alp Guler and I. Kokkinos. Holopose: Holistic 3d human reconstruction in-the-wild. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [4322](#)
- [6] F. Bogo, A. Kanazawa, C. Lassner, P. Gehler, J. Romero, and M. J. Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *ECCV*. Springer, 2016. [4321](#)
- [7] F. Bogo, A. Kanazawa, C. Lassner, P. Gehler, J. Romero, and M. J. Black. Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image. In *Computer Vision – ECCV 2016*, Lecture Notes in Computer Science. Springer International Publishing, Oct. 2016. [4322](#)
- [8] E. Corona, A. Pumarola, G. Alenyà, G. Pons-Moll, and F. Moreno-Noguer. Smplicit: Topology-aware generative model for clothed people. 2021. [4321](#)
- [9] Q. Fang, Q. Shuai, J. Dong, H. Bao, and X. Zhou. Reconstructing 3d human pose by watching humans in the mirror. In *CVPR*, 2021. [4322](#)
- [10] M. Fieraru, M. Zanfir, E. Oneata, A.-I. Popa, V. Olaru, and C. Sminchisescu. Three-dimensional reconstruction of human interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [4321](#), [4323](#)
- [11] M. A. Fischler and R. C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *24(6)*:381–395, June 1981. [4325](#)
- [12] G. Georgaki, R. Li, S. Karanam, T. Chen, J. Seck, and Z. Wu1. Hierarchical kinematic human mesh recovery. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. [4322](#)
- [13] P. Guan, A. Weiss, A. O. Balan, and M. J. Black. Estimating human shape and pose from a single image. In *ICCV*. IEEE, 2009. [4321](#)
- [14] W. Guo, E. Corona, F. Moreno-Noguer, and X. Alamedd-Pineda. Pi-net: Pose interacting network for multi-person monocular 3d pose estimation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2796–2806, 2021. [4321](#)
- [15] M. Hassan, V. Choutas, D. Tzionas, and M. J. Black. Resolving 3d human pose ambiguities with 3d scene constraints. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. [4323](#)
- [16] M. Hassan, P. Ghosh, J. Tesch, D. Tzionas, and M. J. Black. Populating 3D scenes by learning human-scene interaction. In *Proceedings IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2021. [4323](#)
- [17] 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*, 36(7):1325–1339, July 2014. [4322](#)
- [18] W. Jiang, N. Kolotouros, G. Pavlakos, X. Zhou, and K. Daniilidis. Coherent reconstruction of multiple humans from a single image. In *CVPR*, 2020. [4321](#), [4322](#), [4325](#), [4326](#), [4327](#), [4328](#)
- [19] H. Joo, H. Liu, L. Tan, L. Gui, B. Nabbe, I. Matthews, T. Kanade, S. Nobuhara, and Y. Sheikh. Panoptic studio: A massively multiview system for social motion capture. In *ICCV*, 2015. [4322](#), [4325](#)
- [20] H. Joo, T. Simon, and Y. Sheikh. Total capture: A 3d deformation model for tracking faces, hands, and bodies. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. [4321](#), [4322](#), [4326](#), [4327](#), [4328](#)
- [21] A. Kanazawa, M. J. Black, D. W. Jacobs, and J. Malik. End-to-end recovery of human shape and pose. In *Computer Vision and Pattern Recognition (CVPR)*, 2018. [4321](#), [4322](#)
- [22] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In Y. Bengio and Y. LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. [4325](#)
- [23] M. Kocabas, N. Athanasiou, and M. J. Black. Vibe: Video inference for human body pose and shape estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [4322](#)
- [24] N. Kolotouros, G. Pavlakos, M. J. Black, and K. Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *ICCV*, 2019. [4321](#), [4322](#)
- [25] N. Kolotouros, G. Pavlakos, and K. Daniilidis. Convolutional mesh regression for single-image human shape reconstruction. In *CVPR*, 2019. [4321](#), [4322](#)
- [26] L. Kumarapu and P. Mukherjee. Animepose: Multi-person 3d pose estimation and animation. *Pattern Recognition Letters*, 147:16–24, 2021. [4322](#)
- [27] C. Lassner, J. Romero, M. Kiefel, F. Bogo, M. J. Black, and P. V. Gehler. Unite the people: Closing the loop between 3d and 2d human representations. In *CVPR*, 2017. [4321](#)
- [28] S. Li and A. B. Chan. 3d human pose estimation from monocular images with deep convolutional neural network. In *Asian Conference on Computer Vision*, pages 332–347. Springer, 2014. [4322](#)
- [29] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black. SMPL: a skinned multi-person linear model. *ACM Transactions on Graphics*, 34(6):1–16, Oct. 2015. [4322](#)

- [30] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black. Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6):1–16, 2015. 4323
- [31] 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)*, pages 120–130. IEEE, 2018. 4322, 4325
- [32] 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 *3D Vision (3DV), 2018 Sixth International Conference on*. IEEE, sep 2018. 4322
- [33] D. Mehta, S. Sridhar, O. Sotnychenko, H. Rhodin, M. Shafiei, H.-P. Seidel, W. Xu, D. Casas, and C. Theobalt. Vnect: Real-time 3d human pose estimation with a single rgb camera. *ACM Transactions on Graphics (TOG)*, 36(4):1–14, 2017. 4322
- [34] G. Moon, J. Chang, and K. M. Lee. Camera distance-aware top-down approach for 3d multi-person pose estimation from a single rgb image. In *The IEEE Conference on International Conference on Computer Vision (ICCV)*, 2019. 4321
- [35] 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*, pages 10133–10142, 2019. 4322
- [36] F. Moreno-Noguer. 3d human pose estimation from a single image via distance matrix regression. In *CVPR*, 2017. 4322
- [37] M. Omran, C. Lassner, G. Pons-Moll, P. Gehler, and B. Schiele. Neural body fitting: Unifying deep learning and model based human pose and shape estimation. In *3DV*. IEEE, 2018. 4321
- [38] M. Omran, C. Lassner, G. Pons-Moll, P. V. Gehler, and B. Schiele. Neural body fitting: Unifying deep learning and model-based human pose and shape estimation. 2018. 4322
- [39] H. Onizuka, Z. Hayirci, D. Thomas, A. Sugimoto, H. Uchiyama, and R.-i. Taniguchi. Tetratsdf: 3d human reconstruction from a single image with a tetrahedral outer shell. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6011–6020, 2020. 4321
- [40] G. Pavlakos, V. Choutas, N. Ghorbani, T. Bolkart, A. A. A. Osman, D. Tzionas, and M. J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 10975–10985, 2019. 4321, 4322
- [41] 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*, pages 7025–7034, 2017. 4322
- [42] G. Pavlakos, L. Zhu, X. Zhou, and K. Daniilidis. Learning to estimate 3d human pose and shape from a single color image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 4322
- [43] M. Ponzo. Intorno ad alcune illusioni nel campo delle sensazioni tattili, sull’illusione di aristotele e fenomeni analoghi. *Archiv für die gesamte Psychologie*, 16:307–345, 1910. 4322
- [44] A. Pumarola, J. Sanchez, G. Choi, A. Sanfeliu, and F. Moreno-Noguer. 3DPeople: Modeling the Geometry of Dressed Humans. In *International Conference in Computer Vision (ICCV)*, 2019. 4322
- [45] R. Ranftl, A. Bochkovskiy, and V. Koltun. Vision transformers for dense prediction, 2021. 4322, 4325, 4326
- [46] 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*, pages 3433–3441, 2017. 4322
- [47] G. Rogez, P. Weinzaepfel, and C. Schmid. Lcr-net++: Multi-person 2d and 3d pose detection in natural images. *IEEE transactions on pattern analysis and machine intelligence*, 2019. 4322
- [48] S. Saito, T. Simon, J. Saragih, and H. Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, June 2020. 4321
- [49] E. Simo-Serra, C. Torras, and F. Moreno-Noguer. 3d human pose tracking priors using geodesic mixture models. *International Journal of Computer Vision (IJCV)*, 122(2):388–408, 2017. 4322
- [50] Y. Sun, Q. Bao, W. Liu, Y. Fu, B. Michael J., and T. Mei. Monocular, one-stage, regression of multiple 3d people. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021. 4322
- [51] B. Tekin, I. Katircioglu, M. Salzmann, V. Lepetit, and P. Fua. Structured prediction of 3d human pose with deep neural networks. *BMVC*, 2016. 4322
- [52] 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)*, September 2018. 4322, 4325
- [53] C. Wang, J. Li, W. Liu, C. Qian, and C. Lu. Hmor: Hierarchical multi-person ordinal relations for monocular multi-person 3d pose estimation. In A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, editors, *Computer Vision – ECCV 2020*, pages 242–259, Cham, 2020. Springer International Publishing. 4321
- [54] Y. Wu, A. Kirillov, F. Massa, W.-Y. Lo, and R. Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. 4322, 4326
- [55] H. Xu, E. G. Bazavan, A. Zanfir, W. T. Freeman, R. Suktanar, and C. Sminchisescu. Ghum ghuml: Generative 3d human shape and articulated pose models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 4322
- [56] X. Xu, H. Chen, F. Moreno-Noguer, L. A. Jeni, and F. D. la Torre. 3d human pose, shape and texture from low-resolution images and videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2021. 4321

- [57] 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*, pages 5255–5264, 2018. [4322](#)
- [58] A. Zanfir, E. Marinoiu, and C. Sminchisescu. Monocular 3d pose and shape estimation of multiple people in natural scenes - the importance of multiple scene constraints. 2018. [4321](#), [4322](#), [4324](#)
- [59] A. Zanfir, E. Marinoiu, 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 31*, pages 8410–8419. Curran Associates, Inc., 2018. [4321](#), [4322](#)
- [60] W. Zeng, W. Ouyang, P. Luo, W. Liu, and X. Wang. 3d human mesh regression with dense correspondence. In *CVPR*, 2020. [4322](#)
- [61] J. Zhang, D. Yu, J. H. Liew, X. Nie, and J. Feng. Body meshes as points. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 546–556, June 2021. [4321](#), [4322](#), [4325](#), [4326](#), [4327](#)
- [62] J. Y. Zhang, S. PePOSE, H. Joo, D. Ramanan, J. Malik, and A. Kanazawa. Perceiving 3d human-object spatial arrangements from a single image in the wild. In *European Conference on Computer Vision (ECCV)*, 2020. [4323](#)
- [63] T. Zhang, B. Huang, and Y. Wang. Object-occluded human shape and pose estimation from a single color image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [4322](#)
- [64] L. Zhao, X. Peng, Y. Tian, M. Kapadia, and D. N. Metaxas. Semantic graph convolutional networks for 3d human pose regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [4322](#)
- [65] X. Zhou, Q. Huang, X. Sun, X. Xue, and Y. Wei. Towards 3d human pose estimation in the wild: A weakly-supervised approach. In *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. [4322](#)