

Body Meshes as Points

Jianfeng Zhang¹ Dongdong Yu² Jun Hao Liew¹ Xuecheng Nie³ Jiashi Feng¹

¹National University of Singapore ²ByteDance AI Lab ³Yitu Technology

{zhangjianfeng, liewjunhao}@u.nus.edu yudongdong@bytedance.com elefjia@nus.edu.sg

Abstract

We consider the challenging multi-person 3D body mesh estimation task in this work. Existing methods are mostly two-stage based—one stage for person localization and the other stage for individual body mesh estimation, leading to redundant pipelines with high computation cost and degraded performance for complex scenes (e.g., occluded person instances). In this work, we present a single-stage model, *Body Meshes as Points (BMP)*, to simplify the pipeline and lift both efficiency and performance. In particular, *BMP* adopts a new method that represents multiple person instances as points in the spatial-depth space where each point is associated with one body mesh. Hinging on such representations, *BMP* can directly predict body meshes for multiple persons in a single stage by concurrently localizing person instance points and estimating the corresponding body meshes. To better reason about depth ordering of all the persons within the same scene, *BMP* designs a simple yet effective inter-instance ordinal depth loss to obtain depth-coherent body mesh estimation. *BMP* also introduces a novel keypoint-aware augmentation to enhance model robustness to occluded person instances. Comprehensive experiments on benchmarks *Panoptic*, *MuPoTS-3D* and *3DPW* clearly demonstrate the state-of-the-art efficiency of *BMP* for multi-person body mesh estimation, together with outstanding accuracy. Code can be found at: <https://github.com/jfzhang95/BMP>.

1. Introduction

3D human body mesh recovery aims to reconstruct the 3D full-body mesh of the person instance from images or videos. As a fundamental yet challenging task, it has been widely applied for action recognition [63], virtual try-on [41], motion retargeting [35], etc. With the recent notable progress in single-person based full-body mesh recovery [24, 27, 3, 25], a more realistic and challenging setting has attracted increasing attention, i.e. to estimate body meshes for multiple persons from a single image.

Existing methods for multi-person mesh recovery are




Figure 1. **Our single-stage solution.** The proposed model represents each person instance as the center point of its body. Instance localization and body mesh recovery are then directly predicted from the center point features, enabling simultaneous reconstruction of multiple persons in a single stage. Best viewed in color.

mainly two-stage solutions, including top-down [20] and bottom-up [69] approaches. The former first localizes person instances via a person detector, based on which it then recovers the 3D meshes individually; the bottom-up approach estimates person keypoints at first, and then jointly reconstructs multiple 3D human bodies in the image via constrained optimization [69]. Though with notable accuracy, the above paradigms are inefficient with computational redundancy. For instance, the former one estimates body mesh for each person separately, and consequently the total computation cost linearly grows with the number of persons in the image, while the latter requires grouping the keypoints into corresponding persons and inferring the body meshes iteratively, leading to high computational cost.

Targeted at a more efficient and compact pipeline, we consider exploring a single-stage solution. Despite the recent popularity and promising performance of single-stage methods on 2D keypoints estimation [43] and object detection tasks [75, 59], a single-stage pipeline for multi-person mesh recovery is barely explored as it remains unclear how to effectively integrate both person localization and mesh recovery steps within a single stage. In this work, we propose a new instance representation for multi-person body mesh recovery that represents multiple person instances as points in the spatial-depth space where each point is associated with one body mesh. Such an representation allows effective parallelism of person localization and body mesh recovery. Based on it, we develop a new model architecture that exploits shareable features for both localization and mesh recovery and thus achieve a single-stage solution.

In particular, the model has two parallel branches, one

for instance localization and the other for body mesh recovery. In the localization branch, we model each person instance as a single point in a 3-dimensional space, *i.e.* spatial (2D) and depth (1D), where each localized point (detected person) is associated with a body mesh in the body mesh branch represented by the SMPL parametric model [36]. This in turn converts the multi-person mesh recovery into a single-shot regression problem (Fig. 1). Specifically, the spatial location is represented by discrete coordinates w.r.t. regular grids over the input image. Similarly, we discretize depth into several levels to obtain the depth representation. To learn better feature representation to differentiate instances at different depth, motivated by the phenomenon that a person closer to the camera tends to seem larger in the image, we adopt the feature pyramid network (FPN) [30] to extract multi-scale features and use features from the lower scales to represent the closer (and larger) instances. In this way, each instance is represented as one point, whose associated features (extracted from its corresponding spatial location and FPN scale) are used to effectively estimate its body mesh. We name this **Body Meshes as Points (BMP)**.

Applying the BMP model for estimating multi-person body mesh simultaneously faces two challenges in realistic scenarios: how to coherently reconstruct instances with correct depth ordering, and how to handle the common occlusion issue (*e.g.*, overlapping instances and partial observations). For the first challenge, we consider explicitly using the ordinal relations among all the persons in the scene to supervise the model to learn to output body meshes with correct depth order. However, obtaining such ordinal relations is non-trivial for the scenes captured in the wild, since there is no 3D annotation available. Inspired by the recent success of depth estimation for human body joints [42, 73], we propose to take the depth of each person (center point) predicted by a model pre-trained on 3D datasets with depth annotations as the pseudo ordinal relation for model training on the in-the-wild data, which is experimentally proved beneficial to depth-coherent body mesh reconstruction.

Also, to tackle the common occlusion issue, we propose a novel keypoint-aware occlusion augmentation strategy to improve the model robustness to occluded person instances. Different from the previous method [55] that randomly simulates occlusion in images, we generate synthetic occlusion based on the position of skeleton keypoints. Such keypoint-aware occlusion explicitly forces the model to focus on body structure, making it more robust to occlusion.

Comprehensive experiments on 3D pose benchmarks Panoptic [23], MuPoTS-3D [38] and 3DPW [64] evidently demonstrate the high efficiency of the proposed model. Moreover, it achieves new state-of-the-art on Panoptic and MuPoTS-3D datasets, and competitive performance on 3DPW dataset. Our contributions are summarized as follows: 1) To our best knowledge, we are among the first

to explore the *single-stage* solution to multi-person mesh recovery. We introduce a new person instance representation that enables simultaneous person localization and body mesh recovery for all person instances in an image within a single stage, and design a novel model architecture accordingly. 2) We propose a simple yet effective inter-instance ordinal relation supervision to encourage depth-coherent reconstruction. 3) We propose a keypoint-aware occlusion augmentation strategy that takes body structure into consideration, to improve model robustness to occlusion.

2. Related Work

Single-person 3D pose and shape Previous works estimate 3D poses in the form of body skeleton [37, 40, 60, 74, 49, 47, 58, 71, 15] or non-parametric 3D shape [13, 56, 62]. In this work, we use the 3D mesh to represent the full-body pose and shape, and adopt the SMPL parametric model [36] for body mesh recovery. In literature, Bogo *et al.* [5] proposed SMPLify, the first optimization-based method to fit SMPL on the detected 2D joints iteratively. Later works extend SMPLify by either using more dense reference points to replace sparse keypoints like silhouettes and voxel occupancy grids for SMPL fitting [28, 62], or fitting a more expressive model (*e.g.*, SMPL-X) than SMPL [46].

Some recent works directly regress the SMPL parameters from images via deep neural networks in a two-stage manner. They first estimate the intermediate representation (*e.g.*, keypoints, silhouettes, etc) from images and then map it to SMPL parameters [48, 44, 61, 27]. Some others directly estimate SMPL parameters from images, either using complex model training strategies [24, 16] or leveraging temporal information [3, 25]. Although high accuracy is achieved in single-person cases, it remains unclear how to extend them to the more general multi-person cases.

Multi-person 3D pose and shape For multi-person 3D pose estimation, most existing methods adopt a top-down paradigm [53, 9, 54]. They first detect each person instance and then regress the locations of the body joints. Follow-up improvements are made by estimating additional absolute depth [42], considering multi-person interaction [17, 29] or extending to whole-body pose estimation [66]. Alternatively, some approaches also explore the bottom-up paradigm. SSMP3D [39] and SMAP [73] estimate 3D poses from occlusion-aware pose maps and use Part Affinity Fields [7] to infer their association. LoCO [12] maps the image to the volumetric heatmaps and then estimates multi-person 3D poses from them by an encoder-decoder framework. PandaNet [4] is an anchor-based model where 3D poses are regressed for each anchor position.

In contrast to the prosperity of multi-person 3D pose estimation, there is a limited number of works denoted to body mesh recovery for multiple people. Zanfir *et al.* [69] first estimate 3D joints of persons in the image and then optimize

their 3D shapes jointly with multiple constraints. They also propose a two-stage regression-based scheme that first estimates 3D joints for all the persons and then regresses their 3D shapes based on these 3D joints [70]. Instead of regressing SMPL parameters from an intermediate representation (e.g., 3D joints), Jiang *et al.* [20] attach an SMPL head to the Faster R-CNN framework [51] for estimating SMPL parameters directly from the input image in a top-down manner. Despite the encouraging results, these methods are based on the indirect multi-stage framework and suffer low efficiency. Different from all previous methods that rely on a multi-stage pipeline with computation redundancy, our method unifies person localization and body mesh, and enables a box-free and (ad hoc) optimization-free single stage solution to multi-person body mesh recovery.

Point-based representation The point-based methods [10, 75, 59] represent instances by a single point at their center. This approach is regarded as a simple replacement of the anchor-based representation, which has been widely used in many tasks, including object detection [10, 75, 59], 2D keypoints estimation [43] and instance segmentation [65]. However, these methods cannot be directly applied to body mesh recovery. In this work, we extend the point-based representation to multi-person body mesh recovery. A concurrent work [68] adopts a similar solution to body mesh recovery. Our model differs from it in two significant aspects: 1) BMP aims at more coherent reconstruction of persons in the scenes. It handles challenging spatial arrangement and occlusion problems by exploiting the ordinal depth loss and the keypoint-aware augmentation strategy, which are not considered in [68]. 2) BMP adopts a novel 3D point-based representation to differentiate instances at different depths, thus is more robust to overlapped instances; whereas [68] uses only 2D representation, and would fail in such cases.

3. Body Meshes as Points

3.1. Proposed single-stage solution

Given an image I , multi-person body mesh recovery targets at recovering body meshes of all the person instances in I . Existing approaches [70, 69, 20] solve this task via sequentially localizing and estimating the body mesh in a multi-stage manner, leading to computation redundancy. Differently, this work aims to unify the instance localization and body mesh recovery into a single-stage solution to enable a more efficient and concise framework.

We represent each person instance as a single point (i, j, k) in a 3-dimensional space (spanned by 2D spatial and 1D depth dimensions). By dividing the input image uniformly into $G \times G$ grids, its spatial dimension can be easily represented within such a grid coordinate. If the body center of a person falls into grid cell (i, j) , it is assigned with spatial coordinate (i, j) . Similarly, for the depth dimension, we discretize the depth value to K levels and obtain the k value

for each instance according to its depth. Such discretized depth value is beneficial for handling occluding instances, especially when the body centers of multiple instances fall into the same spatial grid coordinate.

Given this representation, we reformulate multi-person mesh recovery as two simultaneous prediction tasks: 1) instance localization and 2) body mesh recovery.

Instance localization For the first task, we employ the *instance map* $\mathcal{C} = \{\mathcal{C}_1, \dots, \mathcal{C}_K\}$, where $\mathcal{C}_k \in \mathbb{R}^{G \times G \times 1}$, to locate each person instance in the image, where G denotes the number of grid cells along one side, while K refers to the number of total depth levels. For each depth level, the network is trained to regress a scalar indicating the probability of every grid cell containing a person.

To construct ground truth (GT) for training, we first determine the depth value k for each instance. We observe that a person tends to seem larger (smaller) in the image when standing closer to (away from) the camera. In other words, the depth of an instance is roughly inversely proportional to its scale. Inspired by it, we employ a Feature Pyramid Network (FPN) [30] with K pyramid levels to capture K different scales, each of which is used to represent the instance with the corresponding depth. More specifically, for each instance, we compute its scale $s = \sqrt{hw}$ where (h, w) denotes the GT body size, and associate it to the corresponding pyramid level k , according to Table 1.

Pyramid	P ₂	P ₃	P ₄	P ₅	P ₆
Stride	8	8	16	32	32
Grid number G	40	36	24	16	12
Instance scale s	<64	32~128	64~256	128~512	≥256

Table 1. We employ FPN with five pyramid levels. P_{k+1} is used to predict instance \mathcal{C}_k and body mesh maps \mathcal{P}_k , where $k = 1, \dots, 5$.

Next, we locate the grid cell (i, j) in \mathcal{C}_k where the *central region* of that person lies. Inspired by [75, 10], the central region is defined as follows: given the GT body center (x^c, y^c) , body size (h, w) of each person and a controllable scale factor ϵ , the position and size of the central region are defined as $(x^c, y^c, \epsilon w, \epsilon h)$. In this work, we set the position of pelvis as body center and $\epsilon = 0.2$. Once identified, the grid cell (i, j) of the k -th pyramid level, i.e., $\mathcal{C}_k(i, j)$ is labeled as positive (label 1). The above steps are repeated for all the instances in the image.

Body mesh representation In parallel with the instance localization, we use the *body mesh map* $\mathcal{P} = \{\mathcal{P}_1, \dots, \mathcal{P}_K\}$, for body mesh recovery, where $\mathcal{P}_k \in \mathbb{R}^{G \times G \times S}$ and S is the dimension of body mesh representation. Concretely, given a positive response in \mathcal{C} that indicates the presence of a person, we regress the body mesh representation using the features from the corresponding grid cell, as shown in Fig. 2. In this work, we use the SMPL parametric model [36] for body mesh representation, which renders a body mesh using the pose parameters $\theta \in \mathbb{R}^{72}$ and shape parameters $\beta \in \mathbb{R}^{10}$. To improve the training stability, we adopt the

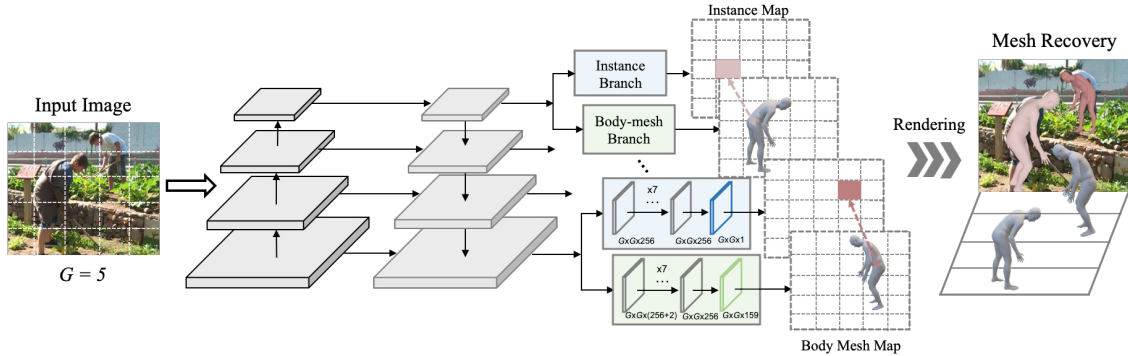


Figure 2. **Illustration of our BMP framework.** An input image is uniformly divided into $G \times G$ grids with $G = 5$ in this example. The model adopts an FPN with K levels ($K = 4$ here). Each person instance is thus represented by its residing grid cell and its associated FPN level (according to its depth). BMP uses the features from the grid cell and FPN level to localize the contained person (top) and estimate the body mesh (bottom) simultaneously.

6D rotation representation [76] for the pose parameters with $\theta \in \mathbb{R}^{144}$. The body mesh map also predicts a camera parameter $\pi = \{s, t_x, t_y\} \in \mathbb{R}^3$ for projecting body joints from 3D back to 2D space, which enables training on in-the-wild 2D pose datasets [21, 32, 2] to improve model generalization [24]. We further introduce a scalar confidence score c defined as the OKS [14] between the projected and GT 2D keypoints, to reflect the confidence level of the SMPL prediction; and we also propose an absolute depth variable d for the corresponding person instance that will be used for penalizing body mesh estimations with incoherent depth ordering (see Sec. 3.2 for details). Therefore, the total channel number of the body mesh map S is 159.

Network architecture We employ ResNet-50 [18] as our backbone. FPN is built on top of the backbone to extract a pyramid of feature maps (256-d). To perform body mesh recovery, we attach two task-specific heads to each level of the feature pyramid, one for instance localization and the other for the corresponding body mesh recovery, responsible for obtaining the instance map $\mathcal{C}_k \in \mathbb{R}^{G \times G \times 1}$ and the body mesh map $\mathcal{P}_k \in \mathbb{R}^{G \times G \times 159}$, respectively. As shown in Fig. 2, each head consists of 7 stacked 3×3 convolutions and one task-specific prediction layer. However, directly estimating the camera parameter from the whole image is non-trivial since it is sensitive to instance position. Inspired by CoordConv [34], we concatenate normalized pixel coordinates to the input feature map at the beginning of the mesh recovery head to encode position information into the network for better estimating camera parameter. Additionally, Group Normalization [67] is used in both prediction heads for facilitating model training. In order to match the features of size $H \times W$ to $G \times G$, we apply bilinear interpolation before the instance and body-mesh recovery branch.

3.2. Inter-instance ordinal depth supervision

Multi-person body mesh recovery is inherently ill-posed as multiple 3D predictions can correspond to the same 2D

projection. Therefore, the trained model would produce ambiguous body mesh estimations with incorrect depth order due to lack of priors. To alleviate such a problem, we use ordinal depth relations among all the persons in the input as supervision to guide reasoning about the depth ordering during the training process.

More concretely, given any two persons (p_m, p_n) in the image, we define the ordinal depth relation between them as $R(p_m, p_n)$, taking the value:

$$R(p_m, p_n) = \begin{cases} +1, & \text{if } d_n - d_m > T, \\ -1, & \text{if } d_m - d_n > T, \\ 0, & \text{if } |d_m - d_n| \leq T, \end{cases} \quad (1)$$

where $d_m \in \mathbb{R}^1$ denotes the depth of the person p_m and T is a pre-defined threshold to determine the ordinal relation. The ordinal relation $R(p_m, p_n) = 0$ means both instances are at roughly the same depth; otherwise one of them is closer to the camera than the other. With the ordinal relation of (p_m, p_n), we define the ordinal depth loss for this pair as

$$L_{(p_m, p_n)} = \begin{cases} \log(1 + \exp(z_m - z_n)), & \text{if } R = +1, \\ \log(1 + \exp(z_n - z_m)), & \text{if } R = -1, \\ (z_m - z_n)^2, & \text{if } R = 0, \end{cases} \quad (2)$$

where $z_m = \frac{2f}{s_z \alpha}$ denotes m -th person's body mesh depth calculated from the predicted camera parameter with focal length f , scale s_m and images's long edge width α . The ordinal depth loss enforces a large margin between z_m and z_n if $R(p_m, p_n) \neq 0$, *i.e.*, one of them measured as closer than the other, and otherwise enforces them to be equal.

However in practice, such ordinal depth relations are rarely available for the scenes captured in the wild due to lack of 3D annotations. To solve this issue, we propose to use pseudo ordinal relations for model training on the in-the-wild data. Specifically, we first train the model on 3D datasets [19, 38] with depth annotations to learn to estimate

the depth of each person in the images. We define the depth d of each person as the depth of body center (*i.e.*, pelvis joint). The model is trained by minimizing a depth loss L_{depth} , which is defined as the mean square errors (MSE) between the predicted and GT depths. After that, given an unlabeled data, we first leverage the pre-trained model to estimate the depth which is then used to obtain the pseudo ordinal relations for all the people in the image. Finally, given the pseudo ordinal relations, we adopt an OKS score-weighted ordinal depth loss to supervise the model training for images in the wild. The total loss for image I is computed as the average loss of all instances pairs:

$$L_{rank} = \frac{1}{N} \sum_{(p_m, p_n)} c_m c_n L_{(p_m, p_n)}, \quad (3)$$

where N denotes the number of paired instances in the image, c_m denotes OKS-score of the m -th person. Intuitively, training the model with such inter-instance ordinal depth supervision can help the model build a global understanding of the depth layout in the input scene and thus ensure more coherent reconstructions.

3.3. Keypoint-aware occlusion augmentation

SMPL-based body mesh recovery is highly sensitive to (partial) occlusion (*e.g.*, overlapping persons, truncation) [72, 52]. To improve model robustness to occlusion without requiring extra training data and annotations, we propose a keypoint-aware occlusion augmentation strategy during the training process. The proposed augmentation strategy aims to generate synthetic occlusion to synthesize real challenging cases such as partial observation for model training. Compared with previous work [55] that randomly simulates occlusion on the images, which may produce easy training samples that are less helpful for boosting model performance, our method directly generates synthetic occlusion based on the positions of skeleton keypoints, which can force the model to pay more attention to the body structure, leading to notable enhancement. More concretely, given a set of J keypoints $\{j_1, \dots, j_J\}$ of a person in the image, we first randomly choose a keypoint j_i . Then we randomly sample a non-human object from the PASCAL VOC [11] dataset and composite it at the location of the selected keypoint j_i . We randomly resize the sampled object to the range of $[0.1 * A, 0.2 * A]$ before compositing, where $A = wh$ denotes the area of that person. Additionally, we randomly shift the keypoint position by an offset δ to avoid over-fitting. During training, we set the probability of the occlusion augmentation as 0.5.

3.4. Training and inference

Training For training our proposed BMP model, we define the loss function L as follows:

$$L = L_{inst} + L_{mesh} + L_{depth} + 0.1 * L_{rank}, \quad (4)$$

where L_{inst} is a modified two-class Focal Loss [31] for instance localization; L_{depth} is the depth loss (Sec. 3.2); L_{mesh} is the loss for body mesh estimation. The training details of the body-mesh branch are similar to those in HMR [24]. Specifically, we formulate L_{mesh} as

$$L_{mesh} = L_{pose} + L_{vert} + \lambda_{3D} L_{3D} + \lambda_{2D} L_{2D} + \lambda_{shape} L_{shape} + \lambda_{conf} L_{conf} + \lambda_{adv} L_{adv}. \quad (5)$$

Here L_{pose} , L_{shape} , L_{3D} , L_{vert} denote MSE between the predicted and GT pose and shape parameters as well as 3D keypoints and vertices, respectively. L_{2D} is the 2D keypoints loss that minimizes the distance between the 2D projection from 3D keypoints and GT 2D keypoints. L_{conf} is the MSE of the predicted and GT confidences, where the GT confidence is computed as the OKS [14] between the projected and GT 2D keypoints. Moreover, we use a discriminator and apply an adversarial loss L_{adv} on the regressed pose and shape parameters, to encourage the outputs to lie on the distributions of real human bodies. $\lambda_{3D} = 4$, $\lambda_{2D} = 4$, $\lambda_{shape} = 0.01$, $\lambda_{conf} = 1$ and $\lambda_{adv} = 0.01$ are the weights of the corresponding loss terms. The loss L_{mesh} is applied independently to each positive grid cell. The ordinal depth loss L_{rank} illustrated in Eqn. (3) is adopted when the image contains more than one instance.

Inference The overall inference procedure for BMP is illustrated in Fig. 2. Given an image, BMP first obtains the instance map \mathcal{C} and the body mesh map \mathcal{P} from the prediction heads. Then it performs max pooling operation to find the local maximum on \mathcal{C} to obtain center point positions $\{(x_i^c, y_i^c, k_i^c)\}_{i=1}^{\hat{N}}$, where k_i^c and (x_i^c, y_i^c) denote the pyramid level and body center location for i -th person, respectively, and \hat{N} is the number of estimated persons. After that, BMP extracts the body mesh parameters of each person p_i via $\mathcal{P}_{k_i^c}(x_i^c, y_i^c)$. Finally, BMP outputs body mesh estimations by deforming the SMPL model using the predicted parameters. A keypoint-based NMS [14] is applied to remove the redundant predictions if they exist. We take the multiplication of the predicted OKS score and the probability score from the instance map as the confidence score for NMS.

3.5. Implementation details

We implement BMP with PyTorch [45] and mmdetection library [8] and utilize Rectified Adam [33] as the optimizer with an initial learning rate of $1e-4$. We resize all images to 832×512 while keeping the same aspect ratio following the original COCO training scheme [57, 65, 20]. During training, we augment the samples with horizontally flip and keypoints-aware occlusion (Sec. 3.3). Flip augmentation is conducted during testing. Moreover, since the BMP model directly extracts image-level features for estimations instead of features from cropped bounding boxes, it can take images with smaller resolution (512×512) as inputs. We denote such a setting as *BMP-Lite*. Other training and testing

settings are the same between BMP-Lite and BMP. Please refer to supplementary for more details.

4. Experiments

In this section, we aim to answer following questions. 1) Can BMP provide both efficient and accurate multi-person mesh recovery? 2) Is BMP able to give coherent meshes for multiple persons with correct depth ordering? 3) Is BMP robust to cases where person instances are occluded or partially observed? To this end, we conduct extensive experiments on several large-scale benchmarks.

4.1. Datasets

Human3.6M [19] is the most widely used single-person 3D pose benchmark collected in an indoor environment. It contains 3.6 million 3D poses and corresponding videos for 15 subjects. Due to its high-quality annotations, we use it following [20] for both training and testing.

Panoptic [23] is a large-scale dataset captured in the Panoptic studio, offering 3D pose annotations for multiple people engaging in diverse social activities. We use this dataset for evaluation with the same protocol as [69].

MuPoTS-3D [38] is a multi-person dataset with 3D pose annotations for both indoor and in-the-wild scenes. We follow [38] and use it for evaluation.

3DPW [64] is a multi-person in-the-wild dataset, which features diverse motions and scenes. It contains 60 video sequences (24 train, 24 test, 12 validation) with full-body mesh annotations. To verify generalizability of the proposed model to challenging in-the-wild scenarios, we use its test set for evaluation, following the same protocol as [25].

MPI-INF-3DHP [40] is a single-person multi-view 3D pose dataset. It contains 8 actors performing 8 activities, captured from 14 cameras. Mehta *et al.*, [38] generate a multi-person dataset called MuCo-3DHP, from MPI-INF-3DHP via mixing up segmented foreground human appearance. We use both datasets for training.

COCO [32], **LSP** [21], **LSP Extended** [22], **PoseTrack** [1], **MPII** [2] are in-the-wild datasets with annotations for 2D joints. We use them for training with the weakly-supervised training strategy [24] (Eqn. (5)).

4.2. Comparison with state-of-the-arts

Single-person setting We first evaluate our proposed BMP model on the single-person setting to validate the strategy of BMP on factorizing the instance localization and mesh recovery does not sacrifice on the performance. Concretely, we evaluate and compare the performance of BMP on the large-scale Human3.6M dataset with most competitive approaches [24, 20] sharing the similar regression target and learning strategy. The results are shown in Table 2. We can observe BMP outperforms all these methods.

Method	HMR [24]	CRMH [20]	BMP
PA-MPJPE	56.8	52.7	51.3

Table 2. **Results on Human3.6M.** We use mean per joint position errors in mm after Procrustes alignment (PA-MPJPE) as metric.

Multi-person settings Then we evaluate our BMP model for multi-person body mesh recovery. We first evaluate it on the multi-person dataset captured in the indoor Panoptic Studio [23] and compare with the most competitive approaches [69, 70, 20]. As shown in Table 3, our BMP model achieves the best performance in all scenarios. Overall, it improves upon the state-of-the-art top-down model CRMH [20] by 5.4% (135.4 mm vs. 143.2 mm in MPJPE), while offering a faster inference speed¹. Moreover, it significantly outperforms CRMH for Ultimatum and Pizza scenarios with crowded scenes and severe occlusion, verifying its robustness to occlusion cases. In addition, its lite version, BMP-Lite, is even faster, which only requires 0.038s to process an image, about 2× faster than CRMH while achieving comparable performance. These results demonstrate both effectiveness and efficiency of BMP for estimating body meshes of multiple persons in a single stage.

Method	Haggl.	Mafia	Ultim.	Pizza	Mean	Time[s]
Zanfir <i>et al.</i> [69]	140.0	165.9	150.7	156.0	153.4	-
MubyNet [70]	141.4	152.3	145.0	162.5	150.3	-
CRMH [20]	129.6	133.5	153.0	156.7	143.2	0.077
BMP-Lite	124.2	138.1	155.2	157.3	143.7	0.038
BMP	120.4	132.7	140.9	147.5	135.4	0.056

Table 3. **Results on the Panoptic.** We use MPJPE as evaluation metric. The lower the better. Best in **bold**.

Another popular 3D pose estimation benchmark is the MuPoTS-3D dataset [40]. We compare our method against two strong baselines, 1) the combination of OpenPose [6] with single-person mesh recovery methods (SMPLify-X [46] and HMR [24]), and 2) the state-of-the-art top-down approach CRMH [20]. We report the results in Table 4. As we can see, BMP outperforms significantly previous methods on both evaluation protocols.

Method	All	Matched	Time[s]
SMPLify-X [46]	62.84	68.04	6.4
HMR [24]	66.09	70.90	0.26
CRMH [20]	69.12	72.22	0.083
BMP-Lite	68.63	71.92	0.038
BMP	73.83	75.34	0.056

Table 4. **Results on MuPoTS-3D.** The numbers are 3DPCK. We report the overall accuracy (All), and the accuracy only for person annotations matched to a prediction (Matched). Best in **bold**.

Lastly, we compare our BMP model with state-of-the-art approaches on the challenging in-the-wild 3DPW dataset. Some approaches use the self-training strategy

¹We count per-image inference time in seconds. For all methods, the time is counted on GPU Tesla P100 and CPU Intel E5-2650 v2 @ 2.60GHz, without using test-time augmentation.

(*i.e.*, SPIN [26]) or temporal information (*i.e.*, VIBE [25]), and they rely on off-the-shelf person detectors [6, 50]. As shown in Table 5, our BMP outperforms CRMH [20] and SPIN [26] in terms of 3DPCK while maintaining an attractive efficiency, and achieves comparable results with VIBE [25] without relying on any temporal information. Additionally, BMP-Lite obtains roughly the same performance as the state-of-the-art CRMH model while achieving $2.1\times$ faster inference speed. These results further confirm the effectiveness of our single-stage solution over existing multi-stage strategies, with very competitive efficiency.

Method	PCK	AUC	MPJPE	PA-MPJPE	PVE	Time[s]
SPIN [26]	30.8	53.4	99.4	68.1	-	0.31
VIBE [25]	33.9	56.6	94.7	66.1	112.7	-
CRMH [20]	25.8	51.6	105.3	62.3	122.2	0.09
BMP-Lite	26.2	51.3	108.5	64.0	126.2	0.038
BMP	32.1	54.5	104.1	63.8	119.3	0.056

Table 5. **Results on 3DPW.** We use 3DPCK, AUC, MPJPE, PA-MPJPE and per-vertex error (PVE) as evaluation metrics.

Qualitative results We visualize some body mesh reconstructions of BMP on the challenging PoseTrack, MPII and COCO datasets, as shown in Fig. 3. It can be observed that BMP is robust to severe occlusion and crowded scenes and can reconstruct human bodies with correct depth ordering.

4.3. Ablative studies

We conduct ablation analysis on Panoptic, 3DPW and MuPoTS-3D datasets both qualitatively and quantitatively to justify our design choices. The qualitative analysis of the proposed method is illustrated in Fig. 4.

Person instance representation We first evaluate the proposed 3D point-based representation for person instances. The main difference between the proposed representation and previous 2D spatial representation [43, 75, 68] is that we use an additional depth dimension to differentiate person instances in the discretized depth space through FPN. We then compare BMP with a baseline model (*i.e.*, BMP using 2D spatial representation). For fair comparison, we aggregate features from all levels of FPN pyramid in the baseline model to obtain a single output for both instance localization and body mesh recovery. Specifically, we study three methods for the aggregation: we resize all feature pyramids to $1/8$ scale and then aggregate them by 1) element-wise addition (*Baseline-Add*), 2) concatenation (*Baseline-Concat*), or 3) adopting a convolutional layer after concatenating them (*Baseline-Conv*). Results are shown in Table 6. We can see our BMP model improves upon the baseline models by a large margin on all datasets, proving its efficacy for body mesh recovery. Additionally, from Fig. 4 (1st row), we observe BMP with the proposed representation is more robust in handling occluding instances, especially when the body centers of multiple instances fall at the same spatial grid coordinate, while the 2D based representation would

usually fail. Please refer to supplementary for the analysis on the pyramid level K .

Method	Panoptic (\downarrow)	3DPW (\downarrow)	MuPoTS-3D (\uparrow)
Baseline-Add	159.1	120.4	68.03
Baseline-Concat	150.3	114.6	68.52
Baseline-Conv	145.6	110.8	69.34
BMP	135.4	104.1	73.83

Table 6. **Ablation for person instance representation.** We report MPJPE for Panoptic and 3DPW, and 3DPCK for MuPoTS-3D.

Ordinal depth loss To investigate whether the ordinal depth loss L_{rank} can help produce more coherent results with correct depth ordering, we conduct experiments on the MuPoTS-3D dataset. Specifically, we evaluate the ordinal depth relations of all instance pairs in the scene and report the percentage of correctly estimated ordinal depth relations in Table 7. The model trained with L_{rank} significantly improves upon the baseline (BMP trained w/o L_{rank}) (from 91.42% to 94.50%). Such improvements can also be observed from Fig. 4 (2nd row). Additionally, by comparing our method with Moon *et al.* [42] and CRMH [20], we observe BMP achieves higher accuracy w.r.t. relative depth ordering than CRMH that only considers ordinal loss for overlapped pairs (94.50% *vs.* 93.68%). This demonstrates our full pair-wise ordinal loss can provide a more comprehensive supervision on the depth layout of the scene and thus train the model to give more coherent results.

Method	Moon [42]	CRMH [20]	BMP w/o L_{rank}	BMP
Accuracy	90.85%	93.68%	91.42%	94.50%

Table 7. **Ablation for ordinal depth loss.** Relative depth ordering results on MuPoTS-3D are shown. We evaluate the ordinal depth relations of all instance pairs in the scene and report the percentage of correctly estimated ordinal depth relations.

Keypoint-aware occlusion augmentation Finally, we study the impact of the proposed keypoint-aware occlusion augmentation strategy. We compare our BMP model with the models trained without occlusion augmentation (*BMP-NoAug*) and trained using randomly Synthetic Occlusion [55] (*BMP-RandOcc*) in Table 8. We can see BMP outperforms both of them by a large margin on all datasets. Notably, it respectively brings 9.1% and 17.3% improvements over BMP-NoAug on Panoptic and 3DPW datasets, which feature crowded scenes with severe overlap and partial observation. In contrast, the random augmentation hurts model performance on MuPoTS-3D (71.71 *vs.* 70.78). This verifies that our proposed occlusion augmentation can force the model to focus on body structure and thus improve its robustness to occlusion.

5. Conclusions

In this work, we present the first single-stage model, Body Meshes as Points (BMP), for multi-person body mesh



Figure 3. **Qualitative results.** We visualize the reconstructions of our approach on PoseTrack (1st row), MPII (2nd row) and COCO (3rd row) from different viewpoints: front (green background), top (blue background) and side (red background), respectively. Please refer to supplementary for more qualitative results.

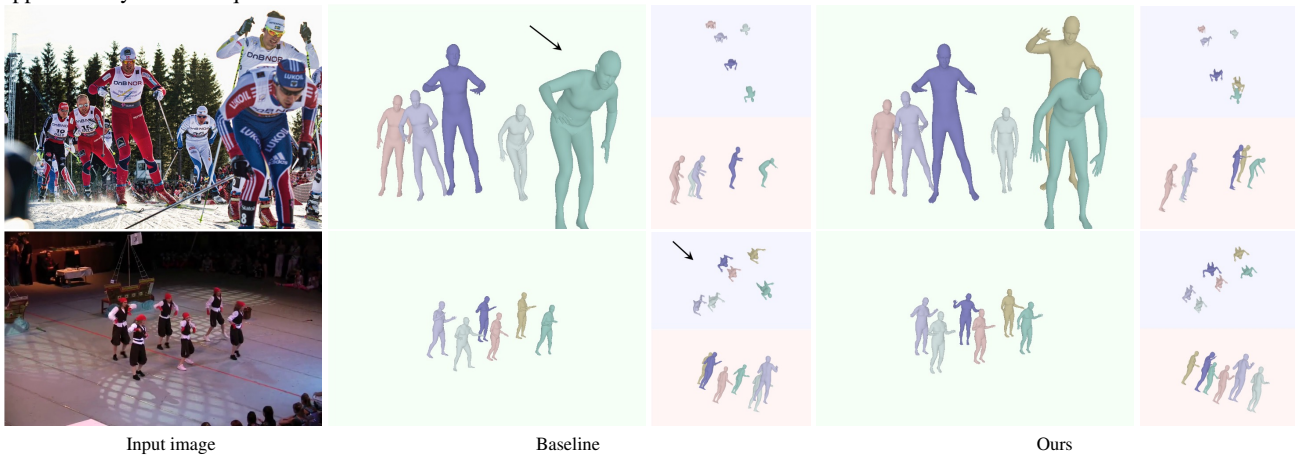


Figure 4. **Qualitative effect of proposed method.** Results of baseline 1 (BMP using 2D representation) (middle 1st row), baseline 2 (BMP trained w/o L_{rank}) (middle 2nd row) and BMP (right). Errors are highlighted by black arrows. As expected, the proposed methods take effect on producing better results (*i.e.*, robust to overlapping instances, more consistent depth ordering for estimated body meshes).

Method	Panoptic (\downarrow)	3DPW (\downarrow)	MuPoTS-3D (\uparrow)
BMP-NoAug	148.9	125.9	71.71
BMP-RandOcc	144.6	110.3	70.78
BMP	135.4	104.1	73.83

Table 8. **Ablation for occlusion augmentation.** We use MPJPE for the first two, and 3DPCK for the last one datasets as metrics.

recovery. BMP introduces a new representation method to enable such a compact pipeline: each person instance is represented as a point in the spatial-depth space which is associated with a parameterized body mesh. With such a representation, BMP can fully exploit shared features and perform person localization and body mesh recovery simul-

taneously. BMP significantly improves upon conventional two-stage paradigms, and offers outstanding efficiency and accuracy, as validated by extensive experiments on multiple benchmarks. Besides, BMP develops several new techniques to further improve the coherence and robustness of recovered body meshes, which are of broad interest for other applications like human pose estimation and detection. In future, we will explore how to make the model more compact and further improve its efficiency, as well as extend to inter-person interactions modeling.

Acknowledgement This research was partially supported by AISG-100E-2019-035, MOE2017-T2-2-151, NUS_ECRA_FY17_P08 and CRP20-2017-0006.

References

- [1] Mykhaylo Andriluka, Umar Iqbal, Eldar Insafutdinov, Leonid Pishchulin, Anton Milan, Juergen Gall, and Bernt Schiele. Posetrack: A benchmark for human pose estimation and tracking. In *CVPR*, 2018.
- [2] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *CVPR*, 2014.
- [3] Anurag Arnab, Carl Doersch, and Andrew Zisserman. Exploiting temporal context for 3d human pose estimation in the wild. In *CVPR*, 2019.
- [4] Abdallah Benzine, Florian Chabot, Bertrand Luvison, Quoc Cuong Pham, and Catherine Achard. Pandanet: Anchor-based single-shot multi-person 3d pose estimation. In *CVPR*, 2020.
- [5] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *ECCV*, 2016.
- [6] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Openpose: realtime multi-person 2d pose estimation using part affinity fields. *arXiv*, 2018.
- [7] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *CVPR*, 2017.
- [8] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv*, 2019.
- [9] Rishabh Dabral, Anurag Mundhada, Uday Kusupati, Safeer Afaq, Abhishek Sharma, and Arjun Jain. Learning 3d human pose from structure and motion. In *ECCV*, 2018.
- [10] Kaiwen Duan, Song Bai, Lingxi Xie, Honggang Qi, Qingming Huang, and Qi Tian. Centernet: Keypoint triplets for object detection. In *ICCV*, 2019.
- [11] Mark Everingham and John Winn. The pascal visual object classes challenge 2012 (voc2012) development kit. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 8, 2011.
- [12] Matteo Fabbri, Fabio Lanzi, Simone Calderara, Stefano Alletto, and Rita Cucchiara. Compressed volumetric heatmaps for multi-person 3d pose estimation. In *CVPR*, 2020.
- [13] Valentin Gabeur, Jean-Sébastien Franco, Xavier Martin, Cordelia Schmid, and Gregory Rogez. Moulding humans: Non-parametric 3d human shape estimation from single images. In *CVPR*, 2019.
- [14] Rohit Girdhar, Georgia Gkioxari, Lorenzo Torresani, Manohar Paluri, and Du Tran. Detect-and-track: Efficient pose estimation in videos. In *CVPR*, 2018.
- [15] Kehong Gong, Jianfeng Zhang, and Jiashi Feng. Poseaug: A differentiable pose augmentation framework for 3d human pose estimation in the wild. In *CVPR*, 2021.
- [16] Riza Alp Guler and Iasonas Kokkinos. Holopose: Holistic 3d human reconstruction in-the-wild. In *CVPR*, 2019.
- [17] Wen Guo, Enric Corona, Francesc Moreno-Noguer, and Xavier Alameda-Pineda. Pi-net: Pose interacting network for multi-person monocular 3d pose estimation. *arXiv*, 2020.
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [19] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, 2014.
- [20] Wen Jiang, Nikos Kolotouros, Georgios Pavlakos, Xiaowei Zhou, and Kostas Daniilidis. Coherent reconstruction of multiple humans from a single image. In *CVPR*, 2020.
- [21] Sam Johnson and Mark Everingham. Clustered pose and nonlinear appearance models for human pose estimation. In *BMVC*, 2010.
- [22] Sam Johnson and Mark Everingham. Learning effective human pose estimation from inaccurate annotation. In *CVPR*, 2011.
- [23] Hanbyul Joo, Hao Liu, Lei Tan, Lin Gui, Bart Nabbe, Iain Matthews, Takeo Kanade, Shohei Nobuhara, and Yaser Sheikh. Panoptic studio: A massively multiview system for social motion capture. In *ICCV*, 2015.
- [24] Angjoo Kanazawa, Michael J. Black, David W. Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *CVPR*, 2018.
- [25] Muhammed Kocabas, Nikos Athanasiou, and Michael J Black. Vibe: Video inference for human body pose and shape estimation. In *CVPR*, 2020.
- [26] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *ICCV*, 2019.
- [27] Nikos Kolotouros, Georgios Pavlakos, and Kostas Daniilidis. Convolutional mesh regression for single-image human shape reconstruction. In *CVPR*, 2019.
- [28] Christoph Lassner, Javier Romero, Martin Kiefel, Federica Bogo, Michael J Black, and Peter V Gehler. Unite the people: Closing the loop between 3d and 2d human representations. In *CVPR*, 2017.
- [29] Jiefeng Li, Can Wang, Wentao Liu, Chen Qian, and Cewu Lu. Hmor: Hierarchical multi-person ordinal relations for monocular multi-person 3d pose estimation. In *ECCV*, 2020.
- [30] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017.
- [31] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017.
- [32] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014.
- [33] Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. On the variance of the adaptive learning rate and beyond. *arXiv*, 2019.

- [34] Rosanne Liu, Joel Lehman, Piero Molino, Felipe Petroski Such, Eric Frank, Alex Sergeev, and Jason Yosinski. An intriguing failing of convolutional neural networks and the coordconv solution. In *NeurIPS*, 2018.
- [35] Wen Liu, Zhixin Piao, Jie Min, Wenhan Luo, Lin Ma, and Shenghua Gao. Liquid warping gan: A unified framework for human motion imitation, appearance transfer and novel view synthesis. In *ICCV*, 2019.
- [36] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6):1–16, 2015.
- [37] Julieta Martinez, Rayat Hossain, Javier Romero, and James J Little. A simple yet effective baseline for 3d human pose estimation. In *ICCV*, 2017.
- [38] Dushyant Mehta, Oleksandr Sotnychenko, Franziska Mueller, Weipeng Xu, Srinath Sridhar, Gerard Pons-Moll, and Christian Theobalt. Single-shot multi-person 3d pose estimation from monocular rgb. In *3DV*, 2018.
- [39] Dushyant Mehta, Oleksandr Sotnychenko, Franziska Mueller, Weipeng Xu, Srinath Sridhar, Gerard Pons-Moll, and Christian Theobalt. Single-shot multi-person 3d pose estimation from monocular rgb. In *3DV*, 2018.
- [40] Dushyant Mehta, Srinath Sridhar, Oleksandr Sotnychenko, Helge Rhodin, Mohammad Shafiei, Hans-Peter Seidel, Weipeng Xu, Dan Casas, and Christian Theobalt. Vnect: Real-time 3d human pose estimation with a single rgb camera. *ACM Trans. on Graphics*, 36(4):44, 2017.
- [41] Aymen Mir, Thiemo Alldieck, and Gerard Pons-Moll. Learning to transfer texture from clothing images to 3d humans. In *CVPR*, 2020.
- [42] Gyeongsik Moon, Ju Yong Chang, and Kyoung Mu Lee. Camera distance-aware top-down approach for 3d multi-person pose estimation from a single rgb image. In *ICCV*, 2019.
- [43] Xuecheng Nie, Jianfeng Zhang, Shuicheng Yan, and Jiashi Feng. Single-stage multi-person pose machines. In *ICCV*, 2019.
- [44] Mohamed Omran, Christoph Lassner, Gerard Pons-Moll, Peter Gehler, and Bernt Schiele. Neural body fitting: Unifying deep learning and model based human pose and shape estimation. In *3DV*, 2018.
- [45] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NeurIPS*, 2017.
- [46] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. Expressive body capture: 3d hands, face, and body from a single image. In *CVPR*, 2019.
- [47] Georgios Pavlakos, Xiaowei Zhou, and Kostas Daniilidis. Ordinal depth supervision for 3d human pose estimation. In *CVPR*, 2018.
- [48] Georgios Pavlakos, Luyang Zhu, Xiaowei Zhou, and Kostas Daniilidis. Learning to estimate 3d human pose and shape from a single color image. In *CVPR*, 2018.
- [49] Alin-Ionut Popa, Mihai Zanfir, and Cristian Sminchisescu. Deep multitask architecture for integrated 2d and 3d human sensing. In *CVPR*, 2017.
- [50] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv*, 2018.
- [51] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015.
- [52] Chris Rockwell and David F Fouhey. Full-body awareness from partial observations. In *ECCV*, 2020.
- [53] Gregory Rogez, Philippe Weinzaepfel, and Cordelia Schmid. Lcr-net: Localization-classification-regression for human pose. In *CVPR*, 2017.
- [54] Gregory Rogez, Philippe Weinzaepfel, and Cordelia Schmid. Lcr-net++: Multi-person 2d and 3d pose detection in natural images. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 42(5):1146–1161, 2019.
- [55] István Sárándi, Timm Linder, Kai O Arras, and Bastian Leibe. How robust is 3d human pose estimation to occlusion? In *ICRAw*, 2018.
- [56] David Smith, Matthew Loper, Xiaochen Hu, Paris Mavroidis, and Javier Romero. Facsimile: Fast and accurate scans from an image in less than a second. In *ICCV*, 2019.
- [57] Kai Su, Dongdong Yu, Zhenqi Xu, Xin Geng, and Changhu Wang. Multi-person pose estimation with enhanced channel-wise and spatial information. In *CVPR*, 2019.
- [58] Xiao Sun, Bin Xiao, Fangyin Wei, Shuang Liang, and Yichen Wei. Integral human pose regression. In *ECCV*, 2018.
- [59] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *ICCV*, 2019.
- [60] Denis Tome, Chris Russell, and Lourdes Agapito. Lifting from the deep: Convolutional 3d pose estimation from a single image. In *CVPR*, 2017.
- [61] Hsiao-Yu Tung, Hsiao-Wei Tung, Ersin Yumer, and Katerina Fragkiadaki. Self-supervised learning of motion capture. In *NeurIPS*, 2017.
- [62] Gul Varol, Duygu Ceylan, Bryan Russell, Jimei Yang, Ersin Yumer, Ivan Laptev, and Cordelia Schmid. Bodynet: Volumetric inference of 3d human body shapes. In *ECCV*, 2018.
- [63] Gül Varol, Ivan Laptev, Cordelia Schmid, and Andrew Zisserman. Synthetic humans for action recognition from unseen viewpoints. *arXiv*, 2019.
- [64] Timo von Marcard, Roberto Henschel, Michael Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3d human pose in the wild using imus and a moving camera. In *ECCV*, 2018.
- [65] Xinlong Wang, Tao Kong, Chunhua Shen, Yuning Jiang, and Lei Li. Solo: Segmenting objects by locations. In *ECCV*, 2019.
- [66] Philippe Weinzaepfel, Romain Brégier, Hadrien Combaluzier, Vincent Leroy, and Grégory Rogez. Dope: Distillation of part experts for whole-body 3d pose estimation in the wild. In *ECCV*, 2020.
- [67] Yuxin Wu and Kaiming He. Group normalization. In *ECCV*, 2018.

- [68] Sun Yu, Bao Qian, Liu Wu, Fu Yili, and Mei Tao. Centerhmr: a bottom-up single-shot method for multi-person 3d mesh recovery from a single image. *arXiv*, 2020.
- [69] Andrei Zanfir, Elisabeta Marinoiu, and Cristian Sminchisescu. Monocular 3d pose and shape estimation of multiple people in natural scenes-the importance of multiple scene constraints. In *CVPR*, 2018.
- [70] Andrei Zanfir, Elisabeta Marinoiu, Mihai Zanfir, Alin-Ionut Popa, and Cristian Sminchisescu. Deep network for the integrated 3d sensing of multiple people in natural images. In *NeurIPS*, 2018.
- [71] Jianfeng Zhang, Xuecheng Nie, and Jiashi Feng. Inference stage optimization for cross-scenario 3d human pose estimation. In *NeurIPS*, 2020.
- [72] Tianshu Zhang, Buzhen Huang, and Yangang Wang. Object-occluded human shape and pose estimation from a single color image. In *CVPR*, 2020.
- [73] Jianan Zhen, Qi Fang, Jiaming Sun, Wentao Liu, Wei Jiang, Hujun Bao, and Xiaowei Zhou. Smap: Single-shot multi-person absolute 3d pose estimation. In *ECCV*, 2020.
- [74] Xingyi Zhou, Qixing Huang, Xiao Sun, Xiangyang Xue, and Yichen Wei. Towards 3d human pose estimation in the wild: a weakly-supervised approach. In *ICCV*, 2017.
- [75] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv*, 2019.
- [76] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *CVPR*, 2019.