

Utilizing Uncertainty in 2D Pose Detectors for Probabilistic 3D Human Mesh Recovery

Tom Wehrbein¹

Marco Rudolph¹

Bodo Rosenhahn¹

Bastian Wandt²

¹Leibniz University Hannover, ²Linköping University

wehrbein@tnt.uni-hannover.de

Abstract

Monocular 3D human pose and shape estimation is an inherently ill-posed problem due to depth ambiguities, occlusions, and truncations. Recent probabilistic approaches learn a distribution over plausible 3D human meshes by maximizing the likelihood of the ground-truth pose given an image. We show that this objective function alone is not sufficient to best capture the full distributions. Instead, we propose to additionally supervise the learned distributions by minimizing the distance to distributions encoded in heatmaps of a 2D pose detector. Moreover, we reveal that current methods often generate incorrect hypotheses for invisible joints which is not detected by the evaluation protocols. We demonstrate that person segmentation masks can be utilized during training to significantly decrease the number of invalid samples and introduce two metrics to evaluate it. Our normalizing flow-based approach predicts plausible 3D human mesh hypotheses that are consistent with the image evidence while maintaining high diversity for ambiguous body parts. Experiments on 3DPW and EMDB show that we outperform other state-of-the-art probabilistic methods. Code is available for research purposes at <https://github.com/twehrbein/humr>.

1. Introduction

Reconstructing 3D human pose and shape from monocular images is a long-standing computer vision problem with vast applications in e.g. robotics, medicine, sports, AR/VR and animation. It is a fundamentally ill-posed problem, because multiple 3D bodies can explain a given 2D image. Apart from the inherent depth ambiguity, body parts are often occluded or truncated and thus prevent the existence of a single unique solution. Consequently, there has been an increasing interest in either estimating a fixed-size set of solutions [3, 26, 39, 43, 47] or, more recently, modeling the full 3D pose distribution conditioned on the 2D input [7, 25, 35, 53–56, 68, 71]. In addition to the theoretic

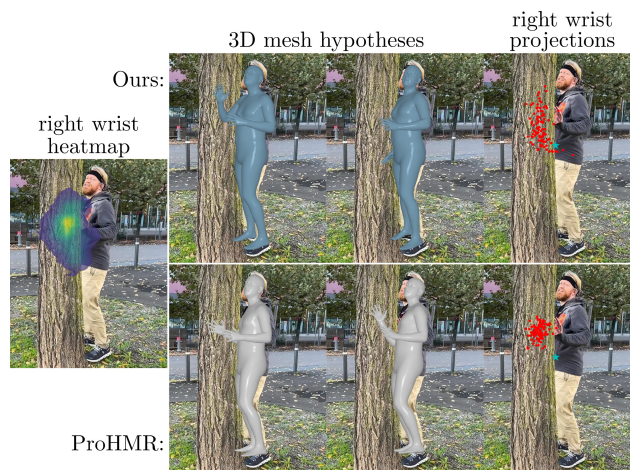


Figure 1. Our method models the full posterior distribution of plausible 3D human meshes given an RGB image. By utilizing heatmaps of a 2D pose detector, the learned distributions have more meaningful diversity and are more accurate than distributions predicted by ProHMR [35]. Two mesh hypotheses and the projection of 100 right wrist samples are shown. \star is the ground-truth 2D position of the right wrist.

cal advantages, several downstream tasks like autonomous driving and video surveillance can benefit from having access to multiple plausible solutions and their prediction uncertainties. Furthermore, methods estimating a distribution over body parameters can act as an effective conditional prior for parametric model fitting [35, 55]. In this work, we propose a method that learns the full posterior distribution of plausible 3D human body parameters which is thus well suited for all aforementioned applications.

Recent probabilistic 3D human mesh estimation approaches [7, 35, 53, 55, 71] employ a generative model trained by maximizing the likelihood of the ground-truth pose given an RGB image. To further constrain the learned distribution and encourage consistency with the input, pose hypotheses are typically projected to the image and supervised with ground-truth 2D keypoints. When only utiliz-

ing these two objective functions, learning solely depends on the availability of images with corresponding 3D human mesh annotations. Ideally, one would have highly similar images with significantly different plausible 3D bodies, especially for examples with occlusions. However, 3D annotations are still rather scarce and, as we argue, on its own not sufficient to best capture the conditional 3D human mesh distribution. On the other hand, there exists an abundant amount of images with ground-truth 2D poses (e.g. [1, 40, 42, 45, 70]). Wehrbein *et al.* [68] were the first to show that 2D detectors [63, 72] trained on such large scale datasets encode joint occurrence probabilities in the predicted heatmaps that can be leveraged for probabilistic 3D human pose estimation. They simplify each heatmap as a 2D Gaussian which is then used as condition and for supervision during training. DiffPose [25] computes a condition based on samples drawn from the heatmaps, but does not utilize them for supervising the generated distributions. In contrast, our method directly uses the distributions encoded in heatmaps as supervision target without having to oversimplify them. Furthermore, we go beyond modeling only the 3D positions of a small number of predefined joints.

In this work, we aim to train a probabilistic 3D human mesh estimation method that fully utilizes an off-the-shelf 2D pose detector [72] to best capture the full posterior distribution of plausible 3D body poses given an RGB image. We employ normalizing flows [31, 52] to model the conditional distributions. The condition consists of an image feature vector together with the maximum-likelihood pose of a 2D detector. We show that although the image is used as input, heatmaps provide important additional supervision signals (see Fig. 1). Our key contribution is to directly supervise the learned distributions during training by minimizing the distance to the distributions encoded in the heatmaps. This is done by generating multiple 3D hypotheses, projecting selected keypoints to the image, and comparing them joint-wise with samples drawn from the heatmaps. As a distance measure, we use the Maximum Mean Discrepancy (MMD) [18], which only requires samples from the distributions and no explicit density estimation. Hence, arbitrary complex distributions can be reproduced, avoiding the need of oversimplifying the heatmaps as in [68]. By distilling the uncertainty information of a 2D detector into our 3D model, we circumvent the problem of having insufficient 3D annotations to learn the full 3D human mesh distribution.

Our second major contribution reveals a problem of current multi-hypothesis human pose estimation approaches. As shown in Fig. 2, we observe that incorrect samples are often generated, where joints are visible that should be invisible. Such failure cases are not captured in current evaluation protocols and are likely caused by the common practice of only penalizing visible joints during training. To fill this evaluation gap, we propose two simple metrics utiliz-

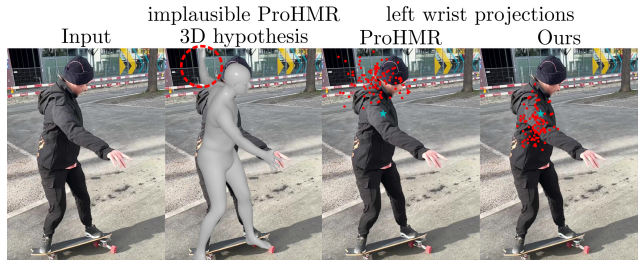


Figure 2. Multi-hypothesis 3D human pose estimation methods (e.g. ProHMR [35]) often generate implausible hypotheses, with joints visible that should be invisible. We significantly reduce the number of incorrect hypotheses by utilizing segmentation masks during training. \star is the ground-truth 2D position of the left wrist.

ing person segmentation masks. The first metric measures the percentage of generated joint samples that lie inside the masks, and the second metric computes the minimum distance to the masks. The idea is that while it is not feasible to describe all possible locations an invisible joint is allowed to be, it is clear that all joints outside the person mask are invalid. We additionally show that by utilizing the person masks during training, we can decrease the number of invalid hypotheses generated by our model.

We evaluate our approach on 3DPW [65] and EMDB [28], outperforming all competitors. Our main contributions are summarized as follows:

- We show that probabilistic 3D human mesh estimation benefits from predicted heatmaps of a 2D detector.
- We fully utilize the heatmaps by directly using them for supervising the learned distributions of our model in a sample-wise manner.
- We demonstrate that person masks can be used to decrease the number of incorrect hypotheses and introduce two new metrics to evaluate it.

2. Related Work

This section first briefly introduces deterministic methods for monocular 3D human pose and shape estimation. Subsequently, probabilistic approaches for reconstructing 3D human meshes and related work on probabilistic 3D human pose estimation from 2D keypoints are discussed.

Deterministic 3D pose and shape estimation. Most methods output a single deterministic estimate given a monocular image. Pioneering work [6, 19, 21, 36, 57] in this area follow an optimization-based approach by fitting parameters of a body model such as SMPL [46, 50] to 2D image observations. However, the optimization process tends to be slow and is sensitive to the initialization and the quality of the given image cues. Starting with HMR [27], learning-based approaches that directly regress the 3D body param-

ters from an image using a deep network became the leading paradigm [12, 13, 20, 33, 34, 41, 44, 62, 66, 67, 73]. They typically consist of an image backbone such as ResNet [23] or HRNet [63] followed by a regression head.

Probabilistic estimation. The ambiguity of reconstructing 3D human pose from a monocular image was already extensively analyzed by early work [8, 37, 38, 58–61]. They specify posterior distributions of 3D poses given 2D observations together with sampling strategies to generate multiple plausible solutions. More recently, several approaches are introduced that use a deep generative model to directly learn the posterior distribution. Li and Lee [39] and Oikarinen *et al.* [47] employ a Mixture Density Network [4] and generate a fixed number of 3D pose hypotheses defined by the mean of the Gaussian kernels. Sharma *et al.* [56] use a variational autoencoder conditioned on 2D pose detections which is more flexible and can produce an unlimited number of hypotheses. Since these approaches only take 2D keypoints as input, they cannot correctly model occlusions. Addressing this problem, Wehrbein *et al.* [68] propose to utilize Gaussians fitted to heatmaps of a 2D pose detector. They use them to condition a normalizing flow and for supervising the learned distributions during training. Instead of oversimplifying heatmaps as Gaussians, DiffPose [25] computes an embedding based on samples drawn from the heatmaps which is used to condition a diffusion model [24]. In contrast, we do not use an embedding of heatmap samples as condition, but use the samples directly to supervise the learned distributions.

Other methods extend beyond 3D joints and aim to learn the posterior of full 3D human bodies given an image. Biggs *et al.* [3] modify HMR [27] to produce a fixed-size set of hypotheses using multiple regression heads, while Sengupta *et al.* [54] predict a Gaussian distribution over SMPL body pose and shape parameters. HierProbHumans [54] and ProPose [14] output a matrix Fisher distribution over body pose rotations. To learn more expressive distributions over SMPL parameters, ProHMR [35] uses a normalizing flow based on the Glow architecture [30]. HuManiFlow [55] extends ProHMR to factorize full body pose into per-body-part pose distributions in an autoregressive manner. Other normalizing flow-based approaches are MHEntropy [7] and HuProSO3 [11]. Diffusion models [24] are also employed to tackle the ill-posed nature of 3D human mesh recovery from monocular images [15, 71] and from egocentric view [74]. However, they cannot assign likelihoods to each sample and cannot be used as conditional prior for parametric model fitting. None of the methods try to utilize uncertainty information encoded in the heatmaps of a 2D pose detector, nor do they address the issue of generating incorrect hypotheses for occluded body parts.

3. Preliminaries

Human body representation. We use SMPL [46] to represent the 3D human body. SMPL is a differentiable parametric model that given pose and shape parameters outputs a 3D mesh $\mathcal{M}(\theta, \beta) \in \mathbb{R}^{N \times 3}$ with $N = 6890$ vertices. The body pose $\theta \in \mathbb{R}^{24 \times 3}$ consists of 23 relative joint rotations plus a global orientation, whereas the shape parameters $\beta \in \mathbb{R}^{10}$ are coefficients of a PCA shape space. 3D joint locations can be expressed as a linear combination of mesh vertices $\mathcal{J}_{3D} = W\mathcal{M} \in \mathbb{R}^{J \times 3}$ using a regressor W .

Normalizing flows. A normalizing flow (NF) is a generative model used to learn arbitrarily complex data distributions through a series of invertible mappings of a simple base distribution. Let $Z \in \mathbb{R}^D$ be a random variable from a base distribution $p_Z(\mathbf{z})$, typically specified as $\mathcal{N}(\mathbf{0}, \mathbf{I})$, and $\mathbf{f} : \mathbb{R}^D \rightarrow \mathbb{R}^D$ an invertible and differentiable function. The normalizing flow then transforms Z into a target distribution $X = \mathbf{f}(Z)$ using the change-of-variables formula

$$p_X(\mathbf{x}) = p_Z(\mathbf{z}) \left| \det \frac{\partial \mathbf{f}(\mathbf{z})}{\partial \mathbf{z}} \right|^{-1}. \quad (1)$$

Typically, $\mathbf{f} = \mathbf{f}_L \circ \dots \circ \mathbf{f}_1$ is the composition of multiple functions, each implemented using a deep neural network. To ensure that all \mathbf{f}_i are invertible, specialized flow architectures have been introduced [9, 10, 16, 30, 31, 48]. The log-probability density of the target data distribution X is then defined as

$$\log p_X(\mathbf{x}) = \log p_Z(\mathbf{z}) - \sum_{l=1}^L \log \left| \det \frac{\partial \mathbf{f}_l(\mathbf{z}_{l-1})}{\partial \mathbf{z}_{l-1}} \right|, \quad (2)$$

which can be directly optimized during training. By introducing a condition vector $\mathbf{c} \in \mathbb{R}^{D_c}$ and using transformations $\mathbf{x} = \mathbf{f}(\mathbf{z}; \mathbf{c})$, $\mathbf{f} : \mathbb{R}^D \times \mathbb{R}^{D_c} \rightarrow \mathbb{R}^D$, NFs can model conditional probability distributions $p_{X|C}(\mathbf{x}|\mathbf{c})$ [69].

4. Method

Our goal is to learn the full posterior distribution of plausible 3D human meshes given an image. Following ProHMR [35], we model the distribution of SMPL pose parameters using a normalizing flow, while predicting deterministic estimates for shape and camera parameters (Sec. 4.1). Unlike all previous work, we *fully* utilize predicted heatmaps of a state-of-the-art 2D pose detector [72], since they encode useful information about joint occurrence probabilities. Our model not only uses the highest-likelihood 2D pose as condition, but additionally directly supervises the learned distribution with samples from the heatmaps (Sec. 4.2). Furthermore, we exploit human segmentation masks to penalize incorrect hypotheses for ambiguous body parts (Sec. 4.3). An overview of the proposed method is shown in Fig. 3.

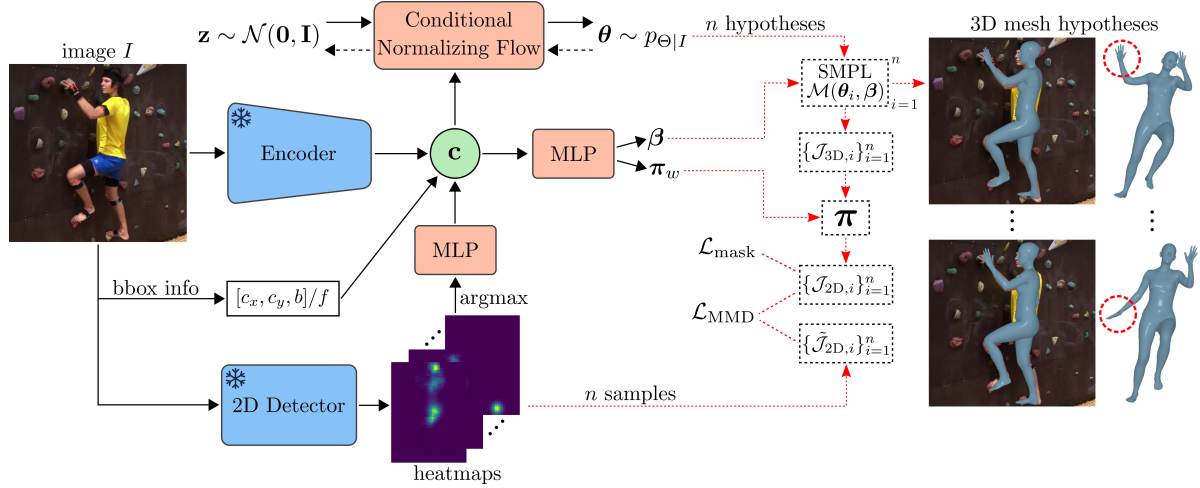


Figure 3. Overview of our approach. Given an image I , we model the full posterior distribution of plausible 3D human meshes using a normalizing flow. In addition to maximizing the likelihood of the ground-truth pose, we supervise the learned distributions by minimizing the Maximum Mean Discrepancy between heatmap samples and projections of 3D mesh hypotheses generated by our NF. Furthermore, a segmentation mask loss is used to penalize invalid hypotheses. Body shape β and camera parameters π_w are estimated deterministically.

4.1. Model Design

Given an image I cropped around a person, we first use HRNet [63] as an image encoder to compute its context vector $\mathbf{c}_I \in \mathbb{R}^{720}$. Additionally, I is fed to the 2D pose detector ViTPose [72] which outputs one heatmap per joint, encoding the probability of a joint’s presence at each location. The highest-likelihood 2D pose with its corresponding confidence scores is then used as input to a small MLP to produce a feature vector $\mathbf{c}_P \in \mathbb{R}^{256}$. Since human mesh estimation methods process images cropped around a person, the location of the person in the original full-frame image is lost, introducing extra ambiguity for global rotation estimation. To retain this information, we follow CLIFF [44] and use the center c_x, c_y and scale b of the person’s bounding box normalized by the focal length f as feature vector:

$$\mathbf{c}_B = [c_x, c_y, b]/f. \quad (3)$$

Finally, our condition is constructed by concatenating the three feature vectors: $\mathbf{c} = [\mathbf{c}_I, \mathbf{c}_P, \mathbf{c}_B] \in \mathbb{R}^{979}$.

We employ a conditional normalizing flow to model the posterior distribution of SMPL joint rotations $\log p_{\Theta|I}(\theta|\mathbf{c})$. The architecture is based on RealNVP [10] which consists of multiple stacked affine coupling layers, each predicting element-wise scale and shift coefficients for an affine transformation. To stabilize training, we apply a soft-clamping mechanism [2], preventing scale components from diverging. We found that it is important to use expressive non-volume preserving flows instead of simpler volume-preserving ones [9] as used in ProHMR [35].

The full condition vector \mathbf{c} is also utilized to regress deterministic estimates for the body shape β and for a weak-perspective camera model $\pi_w = [s, t_x, t_y]$, employing the

same MLP architecture as in HMR [27, 35]. Using either the ground-truth or an approximated focal length [32], we transform the weak-perspective projection to perspective projection parameters π_p of the original image [44].

4.2. Sampling from Heatmaps

Most 2D pose estimators are optimized to regress one heatmap per joint, with ground-truth heatmaps consisting of a 2D Gaussian centered at the joint location. As a by-product of training, heatmaps of a strong detector represent the confidence (*i.e.* a value between 0 and 1) of the corresponding keypoint being at any location. A heatmap is therefore a finite set of non-negative values, and can hence be interpreted as a multinomial distribution with 48×64 possible outcomes (the output dimensions of ViTPose [72]). We independently draw n samples per joint from the constructed distributions. Instead of using an embedding of the samples as condition [25], we directly use them for supervising the learned distributions of our normalizing flow. To this end, we draw n SMPL pose hypotheses from $p_{\Theta|I}(\theta|\mathbf{c})$ and project the target keypoints to the image using the regressed camera π_p . For joint k , given heatmap samples $\hat{\mathcal{S}}_k = \{\hat{s}_{k,i}\}_{i=1}^n$ and projected NF samples $\mathcal{S}_k = \{s_{k,i}\}_{i=1}^n$ both normalized to the range $[-1, 1]$, we compute the Maximum Mean Discrepancy with kernel φ as

$$\begin{aligned} \mathcal{L}_{\text{MMD}}(\mathcal{S}_k, \hat{\mathcal{S}}_k) &= \frac{1}{n(n-1)} \sum_{i \neq j}^n \varphi(\mathbf{s}_{k,i}, \mathbf{s}_{k,j}) \\ &+ \frac{1}{n(n-1)} \sum_{i \neq j}^n \varphi(\hat{\mathbf{s}}_{k,i}, \hat{\mathbf{s}}_{k,j}) - \frac{2}{n^2} \sum_{i,j=1}^n \varphi(\mathbf{s}_{k,i}, \hat{\mathbf{s}}_{k,j}). \end{aligned} \quad (4)$$

Note that we calculate the MMD joint-wise and not pose-wise, because the heatmaps encode the occurrence probabilities of joints independently. Since the detector tends to overestimate the diversity of clearly visible keypoints, we only use heatmap samples for uncertain joints, and n times duplicated ground-truth positions for the certain ones.

4.3. Penalizing Incorrect Hypotheses

We aim to produce diverse 3D hypotheses, especially for occluded body parts. However, we notice that the diversity produced for invisible joints is sometimes too large, as evidenced by samples generated that are visible (see Fig. 2). To penalize these incorrect hypotheses, we propose to utilize person segmentation masks. Intuitively, 3D hypotheses that project outside the person mask are misaligned with the image evidence and thus incorrect. As visualized in Fig. 4, we compute an l_1 loss, $\mathcal{L}_{\text{mask}}$, between projected 2D joints of NF samples that lie outside the person mask and the closest corresponding heatmap samples. Note that this is more efficient than optimizing the distance to the closest mask pixel and, as we observe, leads to more robust training. We remove outside of mask heatmap samples before computing the loss and find this more effective than removing these samples for MMD computation (Eq. 4). Unlike a typical 2D reprojection loss which minimizes the distance to the corresponding ground-truth keypoint, our mask loss $\mathcal{L}_{\text{mask}}$ does not falsely restrict the learned distributions and is hence suitable for occluded joints.

4.4. Optimization

Given an image with corresponding ground-truth SMPL parameters, we minimize the negative log-likelihood of the pose $\hat{\theta}$ and apply an MSE loss for the shape $\hat{\beta}$:

$$\begin{aligned}\mathcal{L}_{\text{NLL}} &= -\log p_{\Theta|I}(\hat{\theta}|\mathbf{c}), \\ \mathcal{L}_{\beta} &= \|\beta - \hat{\beta}\|_2^2.\end{aligned}\quad (5)$$

Additionally, we compute an l_1 loss between 2D ground-truth keypoints, $\hat{\mathcal{J}}_{2D}$, and projected keypoints of the approximated mode of the output distribution θ^* , corresponding to the NF sample with all-zeros latent vector $\mathbf{z} = \mathbf{0}$:

$$\mathcal{L}_{2D} = \|\pi_p(W\mathcal{M}(\theta^*, \beta)) - \hat{\mathcal{J}}_{2D}\|_1. \quad (6)$$

Note that we apply this loss only to the approximated mode prediction and not to random samples, as this helps producing better single-hypothesis estimates while not unnecessarily constraining the modeled distribution. Finally, since we use the 6D rotation representation of [75] in our NF, we regularize generated samples to be orthonormal with $\mathcal{L}_{\text{orth}}$ following [35, 74]. With the proposed MMD (Sec. 4.2) and mask loss (Sec. 4.3), the overall loss is then defined as:

$$\begin{aligned}\mathcal{L} &= \lambda_{\beta}\mathcal{L}_{\beta} + \lambda_{2D}\mathcal{L}_{2D} + \lambda_{\text{NLL}}\mathcal{L}_{\text{NLL}} \\ &\quad + \lambda_{\text{orth}}\mathcal{L}_{\text{orth}} + \lambda_{\text{MMD}}\mathcal{L}_{\text{MMD}} + \lambda_{\text{mask}}\mathcal{L}_{\text{mask}},\end{aligned}\quad (7)$$



Figure 4. Visualization of the person mask loss $\mathcal{L}_{\text{mask}}$ for the left wrist of the person in the center. We explicitly penalize hypotheses for invisible joints that lie outside the person masks by minimizing the l_1 distance to the closest corresponding heatmap samples. Plausible hypotheses not penalized by $\mathcal{L}_{\text{mask}}$ are shown as green dots, implausible ones as green triangles and heatmap samples as blue squares. Best viewed with zoom and in color.

where λ are the corresponding weight coefficients.

4.5. Implementation Details

2D detector and heatmap sampling. We employ the state-of-the-art model ViTPose-H [72] with the publicly available checkpoint trained on multiple human datasets. To reduce training time, we precompute ViTPose predictions with the ground-truth image crop and save 150 heatmap samples per joint using the 17 keypoints COCO skeleton [45] together with their minimum distance to the person mask. Samples with confidence smaller than 0.05 are discarded to suppress outliers. Since in 3D human mesh estimation, the root joint (pelvis) is typically centered in the crop and relative rotations are estimated, the possible variance of joints close to the root of the kinematic tree is limited. We found that ViTPose often generates distributions with too much diversity for these joints, as well as for the facial landmarks. Therefore, we only use heatmap samples for joints that are highly articulated (*i.e.* elbows, wrists, knees and ankles) and uncertain. We define a keypoint as being uncertain if its corresponding ViTPose heatmap has a maximum confidence value below 0.7, which typically happens due to *e.g.* occlusion, motion blur or unusual clothing. Furthermore, we use the maximum confidence values as proxy for the visibility of joints, and set the threshold to 0.5. For non-highly articulated joints, ground-truth positions are used and a prerequisite for applying \mathcal{L}_{MMD} is that the joint is visible. Finally, in order to not falsely restrict the learned distributions, \mathcal{L}_{MMD} is only applied if the ground-truth joint is inside the image crop. We randomly draw $n = 25$ samples and use a mixture of inverse multiquadratics kernels [2, 64]

$$\varphi_B^{im}(\mathbf{s}, \hat{\mathbf{s}}) = \sum_{a \in B} \frac{a^2}{a^2 + \|\mathbf{s} - \hat{\mathbf{s}}\|^2} \quad (8)$$

with bandwidth parameters $B = \{0.05, 0.20, 0.90\}$ [2].

Person segmentation masks can be either obtained from strong segmentation models (e.g. [17, 22, 29, 51]) or are already provided in synthetic human pose datasets [5, 49]. To penalize and later on evaluate incorrect hypotheses, it is crucial that the person mask contains at least every pixel of the image where the target person could be. However, this is not the case if the person is occluded by another person or by an object, because the mask only contains visible person pixels. Thus, as shown in Fig. 4, we modify the segmentations by taking the union of all person masks in an image. We detect object occlusions by projecting the ground-truth body to the image and measuring the overlap with the corresponding mask. The mask loss $\mathcal{L}_{\text{mask}}$ is then not applied to these examples, since it would be impractical to generate segmentations for every possible objects occluding a person. We define a ground-truth joint to be invisible if its location is outside the image crop or its corresponding maximum confidence score is below 0.5.

Network and training details. As image backbone, we use an HRNet-W48 [63] with pretrained checkpoint from BEDLAM-CLIFF [5]. We freeze the weights as this significantly speeds up training while resulting in similar performance. Since the training data contains images with heavily overlapping people, ViTPose sometimes computes estimates for a person other than the target. We filter out these examples by thresholding the Euclidean distance between ground-truth and ViTPose 2D poses. Further implementation and training details are provided in the supplementary.

5. Experiments

5.1. Datasets and Evaluation Metrics

Following BEDLAM-CLIFF [5], we train our model on BEDLAM [5], AGORA [49] and 3DPW [65]. Since BEDLAM and AGORA are synthetic datasets, they provide high quality person segmentation masks which we utilize during training. We do not apply our mask loss to examples from 3DPW, as computing accurate masks for the crowded scenes would incur a large overhead, while 3DPW only accounts for a small fraction of the total training data.

We evaluate on the challenging in-the-wild datasets 3DPW and EMDB (subset 1) [28], and follow HuManiFlow [55] to assess the *accuracy*, *input consistency* and *diversity* of the estimated 3D human mesh distributions. The distribution accuracy is measured by calculating the Mean Per Joint Position Error (MPJPE), the MPJPE after Procrustes Alignment (PA-MPJPE) and the Per Vertex Error (PVE), all in *mm*. We report these metrics for the best 3D hypothesis generated by the predicted distribution, evaluating whether the ground-truth is contained in the distribution. To validate that all hypotheses are consistent with the input image, we compute the Euclidean distance between *visible* ground-truth 2D keypoints (2DKP) and the corresponding

reprojected predicted samples, averaged over 100 hypotheses per image. Following [55], the 17 keypoints COCO skeleton [45] is used and the 2DKP error is calculated in 256×256 pixel space. ViTPose confidence values with a threshold of 0.5 are used as proxy for the visibility of keypoints. The distribution diversity is measured by generating 100 3D keypoint (3DKP) hypotheses, and computing the average Euclidean distance to the mean for each keypoint in *mm*. This is split into visible and invisible keypoints, expecting higher diversity for invisible ones. Since for invisible joints, it is not meaningful to compute the 2DKP error and a high diversity is even desired, hypotheses misaligned with the image evidence are currently not captured by the evaluation protocols. Thus, we introduce two metrics utilizing person segmentation masks which measure the *plausibility* of hypotheses generated for invisible joints. We define an invisible joint as plausible if it is inside the person mask, and compute the percentage of hypotheses inside the mask (PercIn) and the minimum Euclidean distance to the mask (MinDist). Mask-RCNN [17] is employed to generate the segmentations. To ensure that all masks are of high quality, we evaluate the mask metrics only on a subset of EMDB. This subset is constructed by first selecting all images with at least two invisible keypoints and then manually filtering for quality, resulting in 1760 person masks. EMDB is used for evaluation instead of 3DPW since it contains less cluttered scenes leading to more accurate mask estimates.

5.2. Quantitative Evaluation

Since models are often trained on different data and with different backbones, we establish a baseline by training ProHMR [35] on the same three datasets using the same image backbone and denote this competitor as ProHMR[†]. Unfortunately, we could not successfully reproduce the training of HuManiFlow [55] and thus use the provided checkpoint for evaluation (marked by *). All other numbers are taken from the respective papers. To evaluate the distribution accuracy of the probabilistic 3D human mesh estimation methods, we calculate the minimum MPJPE, PA-MPJPE and PVE of 100 hypotheses. Additionally, we report results for a single deterministic estimate which we generate by using an all-zeros latent vector for the normalizing flow, corresponding to the approximated mode of the distribution. The results of our method and the state-of-the-art competitors are shown in Table 1. We significantly outperform all competitors in all metrics on both 3DPW and EMDB. Our method additionally yields better performance for the conventional single prediction regression task than BEDLAM-CLIFF [5] which is also trained on the same data with the same backbone.

Table 2 compares the input consistency, sample diversity and sample plausibility of recent probabilistic methods. We achieve significantly better input consistency for the

		3DPW (14)			EMDB (24)		
Models		MPJPE ↓	PA-MPJPE ↓	PVE ↓	MPJPE ↓	PA-MPJPE ↓	PVE ↓
Best of 1	Biggs <i>et al.</i> [3] NeurIPS’20	93.8	59.9	-	-	-	-
	Sengupta <i>et al.</i> [54] CVPR’21	97.1	61.1	-	-	-	-
	ProHMR [35] ICCV’21	97.0	59.8	-	-	-	-
	HierProbHuman [53] ICCV’21	84.9	53.6	-	-	-	-
	HuManiFlow* [55] CVPR’23	83.1	53.9	98.6	113.9	76.4	133.0
	BEDLAM-CLIFF [5] CVPR’23	66.9	43.0	78.5	98.0	60.6	111.6
	ScoreHypo [71] CVPR’24	72.4	44.5	84.6	112.4*	77.9*	131.5*
	ProHMR [†] [35]	71.7	43.5	84.7	98.2	60.4	114.5
Ours		62.2	40.9	73.9	84.3	56.1	97.7
Best of 100	Biggs <i>et al.</i> [3] NeurIPS’20	74.6	48.3	-	-	-	-
	Sengupta <i>et al.</i> [54] CVPR’21	84.4	52.1	-	-	-	-
	ProHMR [35] ICCV’21	81.5	48.2	-	-	-	-
	HierProbHuman [53] ICCV’21	70.9	43.8	-	-	-	-
	HuManiFlow* [55] CVPR’23	64.5	40.0	75.5	88.7	56.5	100.6
	ScoreHypo [71] CVPR’24	63.0	37.6	73.4	87.4*	58.5*	99.6*
	ProHMR [†] [35]	55.0	34.7	65.7	76.7	47.1	87.3
Ours		46.2	29.8	54.4	63.6	40.9	72.0

Table 1. Comparison of the distribution accuracy of probabilistic 3D human mesh estimation methods on 3DPW [65] and EMDB [28]. Our model achieves state-of-the-art performance for both single hypothesis and best of 100 hypotheses 3D pose and mesh metrics. Rows or numbers marked with * are computed using publicly available checkpoints, and † denotes a retrained baseline. Parenthesis denotes the number of body joints used to calculate MPJPE and PA-MPJPE.




Figure 5. Qualitative results for challenging in-the-wild images with significant occlusions or truncations of body parts. Four samples from the learned 3D human mesh distribution are shown together with the 2D projections of 100 hypotheses for a highly ambiguous joint.

mode and for random hypotheses than the competitors. Furthermore, our generated 3D diversity is lowest for visible joints while being reasonable high for invisible ones. The ProHMR baseline fails to produce meaningful diversity, evidenced by a relatively small difference in 3D spread for visible and invisible keypoints. While HuManiFlow achieves reasonable diversity, its generated hypotheses are often in-

correct as evaluated by the introduced mask metrics. By employing our proposed mask loss $\mathcal{L}_{\text{mask}}$, the percentage of plausible hypotheses increases, and the minimum mask distance decreases significantly. Additionally, $\mathcal{L}_{\text{mask}}$ leads to a lower 3D spread for invisible joints, while the distribution accuracy does not deteriorate (see Table 3), showing that only unnecessary diversity is removed. Overall, our

Models	Consistency	Diversity	Plausibility	
	2DKP Error ↓	3DKP Spread	Mask Metrics	
	Mode / Samples	Vis. / Invis.	PercIn↑ / MinDist↓	
3DPW	HF* [55]	5.73 / 7.26	42.0 / 91.7	-
	PHMR† [35]	6.91 / 9.22	57.0 / 74.0	-
	Ours w/o $\mathcal{L}_{\text{mask}}$	4.79 / 5.96	35.3 / 80.0	-
	Ours w/ $\mathcal{L}_{\text{mask}}$	4.77 / 5.94	34.8 / 75.1	-
	HF* [55]	7.54 / 8.96	47.6 / 88.3	85.1 / 1.576
EMDB	PHMR† [35]	8.59 / 10.94	62.5 / 85.9	86.0 / 1.334
	Ours w/o $\mathcal{L}_{\text{mask}}$	5.03 / 6.37	41.0 / 94.9	88.8 / 1.049
	Ours w/ $\mathcal{L}_{\text{mask}}$	5.02 / 6.34	40.8 / 87.8	91.4 / 0.616

Table 2. Comparison of the input consistency, sample diversity and sample plausibility of recent probabilistic 3D human mesh estimation methods. By employing our mask loss, we significantly improve the sample plausibility.

learned 3D mesh distributions have by far the best characteristics *w.r.t.* accuracy, consistency, diversity and plausibility. Note that the HuManiFlow numbers for 3DPW in Table 2 slightly differ from their reported numbers since they use wider bounding boxes for 2DKP error computation and a different 2D pose detector to determine if a joint is visible.

Qualitative results are shown in Fig. 5. Our method generates diverse and plausible configurations for ambiguous body parts, with all hypotheses consistent with the image evidence. A wider variety of results and comparisons with competitors can be found in the supplementary material.

5.3. Ablation Study

To investigate the influence of our proposed design choices and loss functions, we start with the baseline ProHMR[†] and successively add all components. The distribution accuracy metrics for 100 hypotheses per image evaluated on 3DPW are presented in Table 3. We extend the findings of CLIFF [44] to the probabilistic domain, showing that bounding box information of the person crop is not only beneficial for deterministic regression, but also when used as condition for probabilistic human mesh modeling. Leveraging the most likely 2D pose from ViTPose as additional condition for our NF leads to further improvements after filtering out incorrect predictions. The non-volume preserving flow RealNVP is able to better model the distributions than the simpler volume-preserving flow NICE [9] used by ProHMR. We find that NICE especially struggles with predicting sharp distributions for unambiguous joints, while estimating diverse hypotheses for invisible body parts. By directly supervising (\mathcal{L}_{MMD}) the learned distributions with distributions encoded in heatmaps of ViTPose, the model learns to generate more meaningful diversity. Finally, while the proposed mask loss $\mathcal{L}_{\text{mask}}$ has little influence on the 3D pose and mesh metrics, it leads to a significant decrease in the number of incorrect hypotheses and thus to higher plausibility of the estimated distributions. This is qualitatively shown in Fig. 6 and quantitatively evaluated in Table 2.

Models	3DPW (14)		
	MPJPE ↓	PA-MPJPE ↓	PVE ↓
ProHMR [†] [35]	55.0	34.7	65.7
+ bbox info [44]	52.3	34.5	62.6
+ 2D pose condition	49.5	32.7	58.6
+ RealNVP	48.1	31.8	56.9
+ \mathcal{L}_{MMD}	46.5	29.7	54.8
+ $\mathcal{L}_{\text{mask}}$ (Ours full)	46.2	29.8	54.4

Table 3. Ablation study analyzing our proposed design choices and loss functions. Components are added successively, and the minimum errors out of 100 hypotheses are reported.




Figure 6. By using our mask loss $\mathcal{L}_{\text{mask}}$, significantly fewer incorrect hypotheses are generated while meaningful diversity is maintained. The 2D projections of 100 hypotheses for the left wrist are shown, together with the ground-truth position (*).

6. Conclusion

We present a probabilistic approach to the highly ill-posed problem of monocular 3D human mesh reconstruction. Our normalizing flow-based method utilizes joint occurrence probabilities encoded in heatmaps of a 2D pose detector as supervision target for the learned distributions. This is accomplished by minimizing the Maximum Mean Discrepancy between samples drawn from the heatmaps and hypotheses generated by our model, and leads to more meaningfully diverse distributions. Additionally, we reveal that current methods suffer from generating incorrect solutions for invisible joints and propose a simple loss based on segmentation masks that improves plausibility of the hypotheses. Given a monocular image, our method produces plausible 3D human mesh hypotheses that are consistent with the image evidence while maintaining high diversity for ambiguous body parts. Future work could extend our approach to the temporal domain to reconstruct temporally consistent motions for sequences with strong occlusions.

Acknowledgements. This work was supported by the Federal Ministry of Education and Research (BMBF), Germany, under the AI service center KISSKI (grant no. 01IS22093C), the Deutsche Forschungsgemeinschaft (DFG) under Germany’s Excellence Strategy within the Cluster of Excellence PhoenixD (EXC 2122), the Wallenberg Artificial Intelligence, Autonomous Systems and Software Program (WASP), funded by Knut and Alice Wallenberg Foundation, and the European Union within the Horizon Europe research and innovation programme under grant agreement no. 101136006 – XTREME.

References

- [1] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *CVPR*, 2014. 2
- [2] Lynton Ardizzone, Jakob Kruse, Sebastian J. Wirkert, Daniel Rahner, Eric W. Pellegrini, Ralf S. Klessen, Lena Maier-Hein, Carsten Rother, and Ullrich Köthe. Analyzing inverse problems with invertible neural networks. In *ICLR*, 2019. 4, 5
- [3] Benjamin Biggs, Sébastien Ehrhardt, Hanbyul Joo, Benjamin Graham, Andrea Vedaldi, and David Novotny. 3d multi-bodies: Fitting sets of plausible 3d human models to ambiguous image data. In *NeurIPS*, 2020. 1, 3, 7
- [4] Christopher M. Bishop. Mixture density networks. Technical report, Aston University, 1994. 3
- [5] Michael J. Black, Priyanka Patel, Joachim Tesch, and Jinlong Yang. BEDLAM: A synthetic dataset of bodies exhibiting detailed lifelike animated motion. In *CVPR*, 2023. 6, 7
- [6] 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. 2
- [7] Rongyu Chen, Linlin Yang, and Angela Yao. Mhentropy: Entropy meets multiple hypotheses for pose and shape recovery. In *ICCV*, 2023. 1, 3
- [8] Kiam Choo and D.J. Fleet. People tracking using hybrid monte carlo filtering. In *ICCV*, 2001. 3
- [9] Laurent Dinh, David Krueger, and Yoshua Bengio. NICE: non-linear independent components estimation. In *ICLR*, 2015. 3, 4, 8
- [10] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real NVP. In *ICLR*, 2017. 3, 4
- [11] Olaf Dünkel, Tim Salzmann, and Florian Pfaff. Normalizing flows on the product space of $so(3)$ manifolds for probabilistic human pose modeling. In *CVPR*, 2024. 3
- [12] Sai Kumar Dwivedi, Cordelia Schmid, Hongwei Yi, Michael J. Black, and Dimitrios Tzionas. POCO: 3D pose and shape estimation using confidence. In *3DV*, 2024. 3
- [13] Sai Kumar Dwivedi, Yu Sun, Priyanka Patel, Yao Feng, and Michael J. Black. TokenHMR: Advancing human mesh recovery with a tokenized pose representation. In *CVPR*, 2024. 3
- [14] Qi Fang, Kang Chen, Yinghui Fan, Qing Shuai, Jiefeng Li, and Weidong Zhang. Learning analytical posterior probability for human mesh recovery. In *CVPR*, 2023. 3
- [15] Lin Geng Foo, Jia Gong, Hossein Rahmani, and Jun Liu. Distribution-aligned diffusion for human mesh recovery. In *ICCV*, 2023. 3
- [16] Mathieu Germain, Karol Gregor, Iain Murray, and Hugo Larochelle. Made: Masked autoencoder for distribution estimation. In *ICML*, 2015. 3
- [17] Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin D. Cubuk, Quoc V. Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. In *CVPR*, 2021. 6
- [18] Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(25):723–773, 2012. 2
- [19] Peng Guan, Alexander Weiss, Alexandru O. Bălan, and Michael J. Black. Estimating human shape and pose from a single image. In *ICCV*, 2009. 2
- [20] Mi-Gyeong Gwon, Gi-Mun Um, Won-Sik Cheong, and Wonjun Kim. Instance-aware contrastive learning for occluded human mesh reconstruction. In *CVPR*, 2024. 3
- [21] Nils Hasler, Hanno Ackermann, Bodo Rosenhahn, Thorsten Thormählen, and Hans-Peter Seidel. Multilinear pose and body shape estimation of dressed subjects from image sets. In *CVPR*, 2010. 2
- [22] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask r-cnn. In *ICCV*, 2017. 6
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *CVPR*, 2016. 3
- [24] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 3
- [25] Karl Holmquist and Bastian Wandt. Diffpose: Multi-hypothesis human pose estimation using diffusion models. In *ICCV*, 2023. 1, 2, 3, 4
- [26] Ehsan Jahangiri and Alan L. Yuille. Generating multiple diverse hypotheses for human 3d pose consistent with 2d joint detections. In *ICCV Workshops*, 2017. 1
- [27] Angjoo Kanazawa, Michael J. Black, David W. Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *CVPR*, 2018. 2, 3, 4
- [28] Manuel Kaufmann, Jie Song, Chen Guo, Kaiyue Shen, Tianjian Jiang, Chengcheng Tang, Juan José Zárate, and Otmar Hilliges. EMDB: The Electromagnetic Database of Global 3D Human Pose and Shape in the Wild. In *ICCV*, 2023. 2, 6, 7
- [29] Rawal Khrodgar, Timur Bagautdinov, Julieta Martinez, Su Zhaoen, Austin James, Peter Selednik, Stuart Anderson, and Shunsuke Saito. Sapiens: Foundation for human vision models. In *ECCV*, 2024. 6
- [30] Durk P. Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1×1 convolutions. In *NeurIPS*, 2018. 3
- [31] Durk P. Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved variational inference with inverse autoregressive flow. In *NeurIPS*, 2016. 2, 3
- [32] Imry Kissos, Lior Fritz, Matan Goldman, Omer Meir, Eduard Oks, and Mark Kliger. Beyond weak perspective for monocular 3d human pose estimation. In *ECCV Workshops*, 2020. 4
- [33] Muhammed Kocabas, Chun-Hao P. Huang, Otmar Hilliges, and Michael J. Black. PARE: Part attention regressor for 3D human body estimation. In *ICCV*, 2021. 3
- [34] 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. 3
- [35] Nikos Kolotouros, Georgios Pavlakos, Dinesh Jayaraman, and Kostas Daniilidis. Probabilistic modeling for human mesh recovery. In *ICCV*, 2021. 1, 2, 3, 4, 5, 6, 7, 8

- [36] 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. 2
- [37] Hsi-Jian Lee and Zen Chen. Determination of 3d human body postures from a single view. *Computer Vision, Graphics, and Image Processing*, 30(2):148–168, 1985. 3
- [38] Mun Wai Lee and Isaac Cohen. Proposal maps driven mcmc for estimating human body pose in static images. In *CVPR*, 2004. 3
- [39] Chen Li and Gim Hee Lee. Generating multiple hypotheses for 3d human pose estimation with mixture density network. In *CVPR*, 2019. 1, 3
- [40] Jiefeng Li, Can Wang, Hao Zhu, Yihuan Mao, Hao-Shu Fang, and Cewu Lu. Crowdpose: Efficient crowded scenes pose estimation and a new benchmark. In *CVPR*, 2019. 2
- [41] Jiefeng Li, Chao Xu, Zhicun Chen, Siyuan Bian, Lixin Yang, and Cewu Lu. Hybrik: A hybrid analytical-neural inverse kinematics solution for 3d human pose and shape estimation. In *CVPR*, 2021. 3
- [42] Ruilong Li, Xin Dong, Zixi Cai, Dingcheng Yang, Haozhi Huang, Song-Hai Zhang, Paul L. Rosin, and Shi-Min Hu. Pose2seg: Human instance segmentation without detection. In *CVPR*, 2019. 2
- [43] Wenhao Li, Hong Liu, Hao Tang, Pichao Wang, and Luc Van Gool. Mhformer: Multi-hypothesis transformer for 3d human pose estimation. In *CVPR*, 2022. 1
- [44] Zhihao Li, Jianzhuang Liu, Zhensong Zhang, Songcen Xu, and Youliang Yan. Cliff: Carrying location information in full frames into human pose and shape estimation. In *ECCV*, 2022. 3, 4, 8
- [45] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, Lubomir D. Bourdev, Ross B. Girshick, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In *ECCV*, 2014. 2, 5, 6
- [46] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. In *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*. ACM, 2015. 2, 3
- [47] Tuomas P. Oikarinen, Daniel C. Hannah, and Sohrob Kazerounian. Graphmdn: Leveraging graph structure and deep learning to solve inverse problems. In *IJCNN*, 2021. 1, 3
- [48] George Papamakarios, Theo Pavlakou, and Iain Murray. Masked autoregressive flow for density estimation. In *NeurIPS*, 2017. 3
- [49] Priyanka Patel, Chun-Hao P. Huang, Joachim Tesch, David T. Hoffmann, Shashank Tripathi, and Michael J. Black. AGORA: Avatars in geography optimized for regression analysis. In *CVPR*, 2021. 6
- [50] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3d hands, face, and body from a single image. In *CVPR*, 2019. 2
- [51] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollár, and Christoph Feichtenhofer. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. 6
- [52] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *Proceedings of Machine Learning Research (PMLR)*, 2015. 2
- [53] Akash Sengupta, Ignas Budvytis, and Roberto Cipolla. Hierarchical kinematic probability distributions for 3d human shape and pose estimation from images in the wild. In *ICCV*, 2021. 1, 7
- [54] Akash Sengupta, Ignas Budvytis, and Roberto Cipolla. Probabilistic 3d human shape and pose estimation from multiple unconstrained images in the wild. In *CVPR*, 2021. 1, 3, 7
- [55] Akash Sengupta, Ignas Budvytis, and Roberto Cipolla. Humaniflow: Ancestor-conditioned normalising flows on so(3) manifolds for human pose and shape distribution estimation. In *CVPR*, 2023. 1, 3, 6, 7, 8
- [56] Saurabh Sharma, Pavan Teja Varigonda, Prashast Bindal, Abhishek Sharma, and Arjun Jain. Monocular 3d human pose estimation by generation and ordinal ranking. In *ICCV*, 2019. 1, 3
- [57] Leonid Sigal, Alexandru Balan, and Michael Black. Combined discriminative and generative articulated pose and non-rigid shape estimation. In *NeurIPS*, 2007. 2
- [58] E. Simo-Serra, A. Ramisa, G. Alenyà, C. Torras, and F. Moreno-Noguer. Single image 3d human pose estimation from noisy observations. In *CVPR*, 2012. 3
- [59] Cristian Sminchisescu and Bill Triggs. Covariance scaled sampling for monocular 3d body tracking. In *CVPR*, 2001. 3
- [60] Cristian Sminchisescu and Bill Triggs. Hyperdynamics importance sampling. In *ECCV*, 2002. 3
- [61] Cristian Sminchisescu and Bill Triggs. Kinematic jump processes for monocular 3d human tracking. In *CVPR*, 2003. 3
- [62] Yu-Pei Song, Xiao Wu, Zhaoquan Yuan, Jian-Jun Qiao, and Qiang Peng. Posturehmr: Posture transformation for 3d human mesh recovery. In *CVPR*, 2024. 3
- [63] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *CVPR*, 2019. 2, 3, 4, 6
- [64] Ilya Tolstikhin, Olivier Bousquet, Sylvain Gelly, and Bernhard Schölkopf. Wasserstein auto-encoders. In *ICLR*, 2018. 5
- [65] Timo von Marcard, Roberto Henschel, Michael J. Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3D human pose in the wild using IMUs and a moving camera. In *ECCV*, 2018. 2, 6, 7
- [66] Yufu Wang and Kostas Daniilidis. Refit: Recurrent fitting network for 3d human recovery. In *ICCV*, 2023. 3
- [67] Tom Wehrbein, Bodo Rosenhahn, Iain Matthews, and Carsten Stoll. Personalized 3d human pose and shape refinement. In *ICCV Workshops*, 2023. 3
- [68] Tom Wehrbein, Marco Rudolph, Bodo Rosenhahn, and Bastian Wandt. Probabilistic monocular 3d human pose estimation with normalizing flows. In *ICCV*, 2021. 1, 2, 3

- [69] Christina Winkler, Daniel Worrall, Emiel Hoogeboom, and Max Welling. Learning likelihoods with conditional normalizing flows. *arXiv preprint arXiv:1912.00042*, 2019. 3
- [70] Jiahong Wu, He Zheng, Bo Zhao, Yixin Li, Baoming Yan, Rui Liang, Wenjia Wang, Shipei Zhou, Guosen Lin, Yanwei Fu, Yizhou Wang, and Yonggang Wang. Large-scale datasets for going deeper in image understanding. In *2019 IEEE International Conference on Multimedia and Expo (ICME)*, 2019. 2
- [71] Yuan Xu, Xiaoxuan Ma, Jiajun Su, Wentao Zhu, Yu Qiao, and Yizhou Wang. Scorehypo: Probabilistic human mesh estimation with hypothesis scoring. In *CVPR*, 2024. 1, 3, 7
- [72] Yufei Xu, Jing Zhang, Qiming Zhang, and Dacheng Tao. ViTPose: Simple vision transformer baselines for human pose estimation. In *NeurIPS*, 2022. 2, 3, 4, 5
- [73] Hongwen Zhang, Yating Tian, Xinchu Zhou, Wanli Ouyang, Yebin Liu, Limin Wang, and Zhenan Sun. Pymaf: 3d human pose and shape regression with pyramidal mesh alignment feedback loop. In *ICCV*, 2021. 3
- [74] Siwei Zhang, Qianli Ma, Yan Zhang, Sadegh Aliakbarian, Darren Cosker, and Siyu Tang. Probabilistic human mesh recovery in 3d scenes from egocentric views. In *ICCV*, 2023. 3, 5
- [75] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *CVPR*, 2019. 5