

# HuManiFlow: Ancestor-Conditioned Normalising Flows on $SO(3)$ Manifolds for Human Pose and Shape Distribution Estimation

Akash Sengupta    Ignas Budvytis    Roberto Cipolla

University of Cambridge

{as2562, ib255, rc10001}@cam.ac.uk

## Abstract

Monocular 3D human pose and shape estimation is an ill-posed problem since multiple 3D solutions can explain a 2D image of a subject. Recent approaches predict a probability distribution over plausible 3D pose and shape parameters conditioned on the image. We show that these approaches exhibit a trade-off between three key properties: (i) **accuracy** - the likelihood of the ground-truth 3D solution under the predicted distribution, (ii) **sample-input consistency** - the extent to which 3D samples from the predicted distribution match the visible 2D image evidence, and (iii) **sample diversity** - the range of plausible 3D solutions modelled by the predicted distribution. Our method, HuManiFlow, predicts simultaneously accurate, consistent and diverse distributions. We use the human kinematic tree to factorise full body pose into ancestor-conditioned per-body-part pose distributions in an autoregressive manner. Per-body-part distributions are implemented using normalising flows that respect the manifold structure of  $SO(3)$ , the Lie group of per-body-part poses. We show that ill-posed, but ubiquitous, 3D point estimate losses reduce sample diversity, and employ only probabilistic training losses. HuManiFlow outperforms state-of-the-art probabilistic approaches on the 3DPW and SSP-3D datasets.

## 1. Introduction

Estimating 3D human pose and shape from a single RGB image is an inherently ill-posed [25, 47] computer vision task. Many 3D solutions can correspond to an input 2D observation, due to depth ambiguity, occlusion and truncation. Thus, several recent approaches [1, 2, 23, 42, 43] use deep neural networks to predict a probability *distribution* over 3D pose and shape, conditioned on the 2D input. In theory, this has a few advantages over deterministic single-solution predictors [18, 20, 21, 27, 55] - such as the quantification of prediction uncertainty, sampling of multiple plausible 3D solutions, and usage in downstream tasks such as multi-input fusion [43, 44] or as a prior in parametric model fitting [23].

To fully realise the advantages of probabilistic 3D pose




Figure 1. Comparison between pose and shape distributions from HuManiFlow and ProHMR [23]. 3D samples from HuManiFlow are *consistent* with the visible 2D evidence, while being more *diverse* than samples from ProHMR. Per-vertex sample variances along the x/y/z-axes highlight interpretable uncertainty due to occlusion/truncation (all axes), and depth ambiguity (z-axis-specific).

and shape estimation in practice, we suggest that predicted distributions should exhibit three properties: *accuracy*, *sample-input consistency* and *sample diversity*. Accuracy denotes the likelihood of the ground-truth (GT) 3D pose and shape under the distribution. Sample-input consistency measures the extent to which 3D samples from the distribution match the 2D input. In particular, after projection to the image plane, samples should agree with any pose and shape information visible in the image. Sample diversity refers to the range of 3D poses and shapes modelled by the distribution. The GT pose and shape is but one 3D solution - the predicted distribution should model several plausible solutions when ill-posedness arises due to occlusion, truncation and depth ambiguity in the 2D input. More diverse samples enable better estimates of prediction uncertainty.

We show that recent probabilistic approaches suffer from a trade-off between accuracy, consistency and diversity.

Several methods [23, 42–44] predict uni-modal pose and shape distributions with limited expressiveness, and use non-probabilistic loss functions such as L1/L2 losses between GT 3D keypoints and a 3D point estimate (usually the mode of the predicted distribution). These choices favour accuracy and consistency but harm diversity, as shown in Figure 1. Approaches that generate diverse samples [2], through the use of more expressive probability distributions, often yield samples that are not consistent with the 2D input image.

We aim to balance accuracy, consistency and diversity with our approach, **HuManiFlow**, which outputs a distribution over SMPL [29] pose and shape parameters conditioned on an input image. We use normalising flows [37] to construct expressive full body pose distributions, which are factorised into per-body-part distributions autoregressively conditioned on ancestors along the human kinematic tree. We account for the manifold structure of the Lie Group of per-body-part poses (or 3D rotations)  $SO(3)$  by predicting distributions on the corresponding Lie algebra  $\mathfrak{so}(3)$ , and “pushing forward” the algebra distributions onto the group via the exponential map [12]. We follow [42, 43] in predicting a Gaussian distribution over SMPL’s shape-space PCA coefficients. Notably, our method is trained without commonly-used point estimate losses on 3D keypoints. We demonstrate that such non-probabilistic losses reduce sample diversity while providing negligible accuracy improvements when expressive distribution estimation models are used.

In summary, our main contributions are as follows:

- We demonstrate that current probabilistic approaches to monocular 3D human pose and shape estimation suffer from a trade-off between distribution accuracy, sample-input consistency and sample diversity.
- We propose HuManiFlow, a normalising-flow-based method to predict distributions over SMPL pose and shape parameters that (i) considers the manifold structure of the 3D body-part rotation group  $SO(3)$ , (ii) exploits the human kinematic tree via autoregressive factorisation of full body pose into per-body-part rotation distributions, and (iii) is trained without any non-probabilistic point estimate losses on 3D keypoints (such as vertices or body joints).
- We show that HuManiFlow provides more accurate, input-consistent and diverse pose and shape distributions than current approaches, using the 3DPW [51] and SSP-3D [41] datasets (see Figure 1). Our method interpretably and intuitively models uncertainty due to occlusion, truncation and depth ambiguities.

## 2. Related Work

**Monocular 3D pose and shape estimation** approaches can be labelled as optimisation-based or learning-based.

Optimisation-based approaches involve iteratively updating the parameters of a 3D body model [29, 36, 54] (“model fitting”) to match 2D observations, such as 2D keypoints [4, 24], silhouettes [24], body-part masks [56] or dense correspondences [14]. These methods do not need expensive 3D-labelled training data, but require accurate 2D observations, good parameter initialisations and suitable 3D pose priors.

Learning-based approaches may be model-free or model-based. Model-free approaches directly regress a 3D human representation, such as a voxel grid [50], vertex mesh [7, 22, 33, 57] or implicit surface [39, 40]. Model-based methods [18, 20, 41, 49, 55] regress body model parameters [29, 36, 54]. Both typically use deep neural networks (DNNs).

Recently, several approaches have combined optimisation and learning. SPIN [21] initialises model fitting with a DNN regression, then supervises the DNN with the optimised parameters. EFT [17] optimises the weights of a pre-trained DNN at test-time, instead of body model parameters. HybriK [27] uses a DNN to regress the longitudinal (twist) rotation of 3D joints, and computes the in-plane (swing) rotation analytically using inverse kinematics.

**3D pose and shape distribution estimation.** Early optimisation-based approaches to 3D pose estimation [8, 45–47] specified multi-modal posterior probabilities of 3D pose given 2D observations, and provided methods to sample multiple plausible 3D poses from the posterior. Recent learning-based approaches predict distributions over 3D keypoint locations conditioned on 2D images, using mixture density networks [3, 26, 34] or normalising flows [52]. Other methods extend this to distributions over 3D pose represented by body-part rotations. 3D Multibodies [2] predicts a categorical distribution over SMPL [29] pose and shape parameters, while Sengupta *et al.* [43] estimate a Gaussian distribution. HierProbHumans [42] outputs a hierarchical matrix-Fisher distribution over body-part rotations, informed by the SMPL kinematic tree. ProHMR [23] uses additive-coupling normalising flows [9] to learn more-expressive distributions over SMPL pose parameters. These methods exhibit a trade-off between distribution accuracy, sample-input consistency and sample diversity, as we show in Section 5.

## 3. Method

This section provides preliminary overviews of normalising flows [37], the Lie group  $SO(3)$  and SMPL [29], then details our pose and shape distribution prediction method.

### 3.1. Normalising flows

Normalising flows [35, 37] are a method for constructing expressive probability distributions using the change-of-variables formula. Specifically, pushing a  $D$ -dimensional continuous random vector  $Z \sim p_Z(\mathbf{z})$  through a diffeomorphism  $f : \mathbb{R}^D \rightarrow \mathbb{R}^D$  (i.e. a bijective differentiable map with differentiable inverse  $f^{-1}$ ) induces a random vector

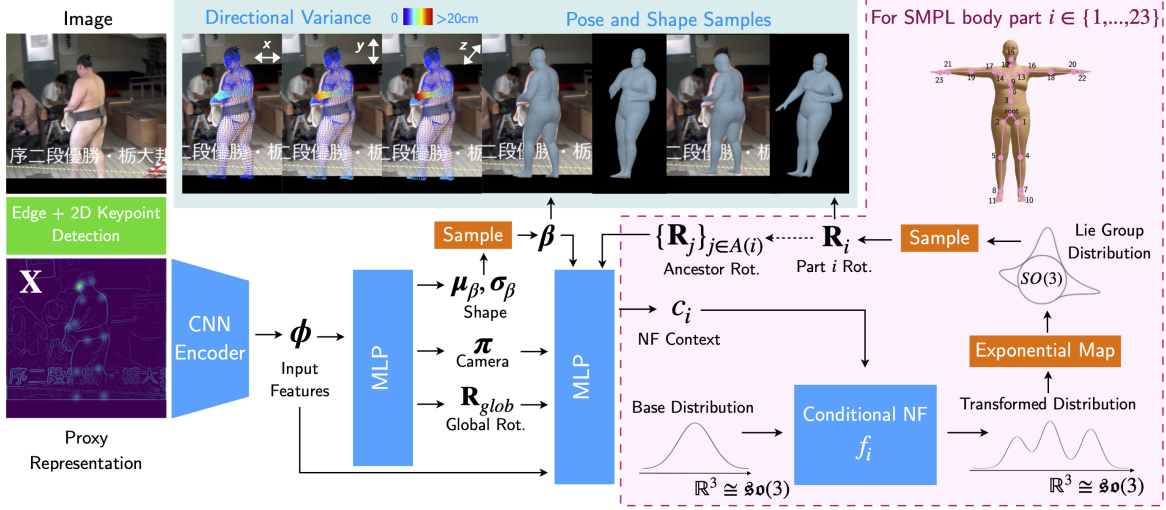


Figure 2. Overview of HuManiFlow, our method for image-conditioned pose and shape distribution prediction. We output a Gaussian distribution over SMPL [29] shape, and learn normalising flows [10, 37] over full body pose, factorised into ancestor-conditioned per-body-part distributions using the SMPL kinematic tree. Normalising flows are defined on  $SO(3)$ , the Lie group of per-body-part poses, by pushing distributions over the Lie algebra  $\mathfrak{so}(3) \cong \mathbb{R}^3$  through the exponential map [12]. 3D body samples are obtained using ancestral sampling.

$Y = f(Z)$  with density function given by

$$\begin{aligned} p_Y(\mathbf{y}) &= p_Z(f^{-1}(\mathbf{y})) |\det J_{f^{-1}}(\mathbf{y})| \\ &= p_Z(\mathbf{z}) |\det J_f(\mathbf{z})|^{-1} \end{aligned} \quad (1)$$

where  $\mathbf{y} \in \mathbb{R}^D$ ,  $\mathbf{z} = f^{-1}(\mathbf{y})$  and  $J_f(\mathbf{z}) \in \mathbb{R}^{D \times D}$  is the Jacobian matrix representing the differential of  $f$  at  $\mathbf{z}$ . Intuitively,  $|\det J_f(\mathbf{z})|$  gives the relative change of volume of an infinitesimal neighbourhood around  $\mathbf{z}$  due to  $f$ .

In a normalising flow model,  $f$  is the composition of multiple simple diffeomorphisms  $f = f_K \circ \dots \circ f_1$ . Each  $f_k$  is implemented using a deep neural network. Typically, the base distribution  $p_Z(\mathbf{z})$  is specified as  $\mathcal{N}(\mathbf{0}, \mathbf{I})$ .

This formulation may be extended to model conditional probability distributions [53]  $p_{Y|C}(\mathbf{y}|\mathbf{c})$ , where  $\mathbf{c} \in \mathbb{R}^C$  is a context vector, using a transformation  $f: \mathbb{R}^D \times \mathbb{R}^C \rightarrow \mathbb{R}^D$  such that  $\mathbf{y} = f(\mathbf{z}; \mathbf{c})$ .  $f$  is bijective in  $\mathbf{y}$  and  $\mathbf{z}$ .

### 3.2. Lie group structure of $SO(3)$

The Lie group of 3D rotations may be defined as  $SO(3) = \{\mathbf{R} \in \mathbb{R}^{3 \times 3} | \mathbf{R}^T \mathbf{R} = \mathbf{I}, \det \mathbf{R} = 1\}$ . The corresponding Lie algebra  $\mathfrak{so}(3)$  (i.e. tangent space at the identity  $\mathbf{I}$ ) consists of  $3 \times 3$  skew-symmetric matrices. Since  $\mathfrak{so}(3)$  is a real 3D vector space, an isomorphism from  $\mathbb{R}^3$  to  $\mathfrak{so}(3)$  may be defined by the hat operator  $\hat{\cdot}: \mathbb{R}^3 \rightarrow \mathfrak{so}(3)$  where

$$\hat{\mathbf{v}} = \begin{bmatrix} 0 & -v_3 & v_2 \\ v_3 & 0 & -v_1 \\ -v_2 & v_1 & 0 \end{bmatrix} \in \mathfrak{so}(3) \quad (2)$$

for  $\mathbf{v} = [v_1, v_2, v_3]^T \in \mathbb{R}^3$ .

$SO(3)$  is a matrix Lie group; thus, the exponential map  $\exp: \mathfrak{so}(3) \rightarrow SO(3)$  coincides with the matrix exponential

$\exp \hat{\mathbf{v}} = \sum_{k=0}^{\infty} \frac{\hat{\mathbf{v}}^k}{k!}$  for  $\hat{\mathbf{v}} \in \mathfrak{so}(3)$ . In practice, we use

$$\exp \hat{\mathbf{v}} = \mathbf{I} + (\sin \theta) \hat{\mathbf{u}} + (1 - \cos \theta) \hat{\mathbf{u}}^2 \quad (3)$$

i.e. the Rodrigues' rotation formula with  $\hat{\mathbf{v}} = \theta \hat{\mathbf{u}}$ . Here,  $\theta \in \mathbb{R}$  is the rotation angle and the unit vector  $\mathbf{u} \in \mathbb{R}^3$  is the axis. Thus, there is a correspondence between axis-angle vectors  $\mathbf{v} = \theta \mathbf{u} \in \mathbb{R}^3$  and Lie algebra elements  $\hat{\mathbf{v}} = \theta \hat{\mathbf{u}} \in \mathfrak{so}(3)$ .

The exponential map for  $SO(3)$  is surjective, as  $SO(3)$  is connected and compact, and smooth. However, it is not injective, and thus not a diffeomorphism. This is because

$$\exp(\theta \hat{\mathbf{u}}) = \exp((\theta + 2\pi k) \hat{\mathbf{u}}) = \exp(\theta_k \hat{\mathbf{u}}) \quad (4)$$

for any  $\theta \hat{\mathbf{u}} \in \mathfrak{so}(3)$ ,  $k \in \mathbb{Z}$  and  $\theta_k = \theta + 2\pi k$ . Given a rotation  $\mathbf{R} \in SO(3)$ , the log operator may be used to find the corresponding axis-angle vector with the smallest angle (or minimum 2-norm). Specifically,  $\theta \mathbf{u} = \log \mathbf{R}$  such that

$$\theta = \arccos\left(\frac{\text{tr} \mathbf{R} - 1}{2}\right), \quad \mathbf{u} = \frac{1}{2 \sin \theta} \begin{bmatrix} R_{32} - R_{23} \\ R_{13} - R_{31} \\ R_{21} - R_{12} \end{bmatrix} \quad (5)$$

where  $\theta \in [0, \pi]$ . The set of all equivalent axis-angle vectors, or  $\mathfrak{so}(3)$  elements, can be obtained with Eqn. 4 - thus providing a many-valued inverse function to  $\exp$ .

### 3.3. Constructing distributions on $SO(3)$

We ultimately aim to learn expressive probability distributions on the Lie group  $SO(3)$ . We can construct a distribution on the Lie algebra  $\mathfrak{so}(3)$  using a distribution over axis-angle vectors in  $\mathbb{R}^3$ , since  $\mathbb{R}^3$  is isomorphic to  $\mathfrak{so}(3)$  (or  $\mathbb{R}^3 \cong \mathfrak{so}(3)$ ). The axis-angle density function,  $p_{\mathbb{R}^3}(\mathbf{v})$  for



Figure 3. Comparison between HuManiFlow and recent probabilistic approaches to monocular pose and shape estimation. 3D Multibodies [2] generates diverse 3D samples, but these do not consistently match the input image. Samples from ProHMR [23] and HierProbHumans [42] are input-consistent but not diverse, and tend to cluster around the predicted distribution’s mode. Our method outputs consistent and diverse samples, and interpretably captures uncertainty due to depth ambiguity (z-axis variance), truncations and occlusions (all-axes variance).

$\mathbf{v} \in \mathbb{R}^3$ , can be modelled by, for example, normalising flows (NFs) [37] or mixture density networks (MDNs) [3].

$p_{\mathbb{R}^3}(\mathbf{v})$  is pushed from  $\mathbb{R}^3 \cong \mathfrak{so}(3)$  onto  $SO(3)$  via the exponential map, using a variant of the change-of-variables formula in Eqn. 1.  $\exp$  is not a diffeomorphism over all of  $\mathfrak{so}(3)$ , which obstructs change-of-variables. However, it is a *local* diffeomorphism (and surjective) within an open Euclidean ball  $B_r(\mathbf{0})$  of radius  $\pi < r < 2\pi$  about  $\mathbf{0} \in \mathbb{R}^3 \cong \mathfrak{so}(3)$ , as shown by [12], who then use this property to derive a change-of-variables formula for the  $\exp$  map:

$$p_{SO(3)}(\mathbf{R}) = \sum_{k \in \{0, \pm 1\}} p_{\mathbb{R}^3}(\theta_k \mathbf{u}) |\det J_{\exp}(\theta_k \mathbf{u})|^{-1} \quad (6)$$

where  $\mathbf{R} \in SO(3)$ ,  $\theta \mathbf{u} = \mathbf{v} = \log \mathbf{R}$  and  $\theta_k = \theta + 2\pi k$ .  $p_{\mathbb{R}^3}$  must be constructed with *compact support*, such that  $p_{\mathbb{R}^3}(\mathbf{v}) = 0$  for  $\mathbf{v} \notin B_r(\mathbf{0})$ , to allow for a finite sum over  $k \in \{0, \pm 1\}$  instead of  $k \in \mathbb{Z}$  (as in Eqn. 4).  $p_{SO(3)}$  is the density function on  $SO(3)$  induced by pushing  $p_{\mathbb{R}^3}$  through  $\exp$ . The determinant of  $J_{\exp}$  at  $\mathbf{v} = \theta \mathbf{u}$  is given by [12]

$$\det J_{\exp}(\mathbf{v}) = \det J_{\exp}(\theta \mathbf{u}) = \frac{2 - 2 \cos \theta}{\theta^2} \quad (7)$$

which, intuitively, gives the relative change of volume due to  $\exp$  of an infinitesimal neighbourhood around  $\mathbf{v} = \theta \mathbf{u}$ .

### 3.4. SMPL model

SMPL [29] is a parametric human body model. Body shape is parameterised by coefficients  $\beta \in \mathbb{R}^{10}$  corresponding to a PCA shape-space basis. Pose is given by the 3D rotation of each body-part relative to its parent joint in the kinematic tree, which consists of 23 body (i.e. non-root) joints. Full body pose can be represented as a set of per-body-part relative rotations  $\{\mathbf{R}_i\}_{i=1}^{23}$  where  $\mathbf{R}_i \in SO(3)$ . We denote the global rotation about the root joint as  $\mathbf{R}_{glob}$ . SMPL provides a function  $\mathcal{M}(\beta, \{\mathbf{R}_i\}_{i=1}^{23}, \mathbf{R}_{glob})$  that outputs a vertex mesh  $\mathbf{V} \in \mathbb{R}^{6890 \times 3}$ . 3D keypoints are given by  $\mathbf{J}^{3D} = \mathcal{J}\mathbf{V}$  where  $\mathcal{J}$  is a pre-trained linear regressor.

### 3.5. Pose and shape distribution prediction

Our method, HuManiFlow, predicts probability distributions over SMPL pose and shape parameters conditioned on an input image, as shown in Figure 2. It also outputs deterministic estimates of weak-perspective camera parameters  $\pi = [s, t_x, t_y]$ , denoting scale and  $xy$  translation, and  $\mathbf{R}_{glob}$ .

Body Pose Distribution On	Distribution Type	3DPW			3DPW Cropped		
		Accuracy	Consistency	Diversity	Accuracy	Consistency	Diversity
		MPJPE-PA Point / Sample Min.	2DKP Error Point / Samples	3DKP Spread Vis. / Invis.	MPJPE-PA Point / Sample Min.	2DKP Error Point / Samples	3DKP Spread Vis. / Invis.
<i>Full Body (concatenated)</i>	Gaussian	55.3 / 47.8 (13.6%)	5.9 / 8.2	48.2 / 118.6	78.8 / 62.8 (20.3%)	10.9 / 15.5	57.4 / 140.7
	MDN [3]	54.6 / 46.9 (14.1%)	5.9 / 8.1	50.5 / 122.3	78.4 / 62.2 (20.7%)	10.9 / 15.5	57.4 / 140.7
<i>Axis-Angles</i> $\Theta \in \mathbb{R}^{69}$	LRS-NF [10]	54.5 / 46.1 (15.4%)	5.8 / 8.1	51.3 / 124.0	78.6 / 61.9 (21.2%)	10.8 / 15.5	57.6 / 140.9
	Gaussian	55.0 / 43.6 (20.7%)	5.4 / 6.7	43.2 / 105.3	84.8 / 60.5 (28.7%)	9.8 / 12.5	43.8 / 136.4
<i>Ancestor-Conditioned</i> <i>Axis-Angles</i> $\{\mathbf{v}_i\}_{i=1}^{23}$ , $\mathbf{v}_i \in \mathbb{R}^3 \cong \mathfrak{so}(3)$	MDN [3]	54.9 / 41.7 (24.0%)	5.3 / 6.7	47.6 / 118.7	84.4 / 60.0 (28.9%)	<b>9.7</b> / 12.2	43.7 / 138.9
	LRS-NF [10]	<b>53.4</b> / 41.1 (23.0%)	<b>5.1</b> / 6.6	44.7 / 110.5	83.6 / 59.3 (29.1%)	9.8 / 11.9	38.9 / 130.2
<i>Ancestor-Conditioned</i> <i>Matrices</i> $\{\mathbf{R}_i\}_{i=1}^{23}$ , $\mathbf{R}_i \in SO(3)$ manifold	Matrix-Fisher	54.0 / 43.4 (19.7%)	5.1 / 6.8	51.4 / <b>131.7</b>	80.4 / 58.5 (27.2%)	9.9 / 11.6	49.6 / <b>142.7</b>
	MDN [3]	54.3 / 40.8 (24.9%)	5.2 / 6.7	47.2 / 119.2	80.3 / 57.8 (28.0%)	9.8 / 11.5	42.8 / 139.2
	LRS-NF [10]	<b>53.4</b> / <b>39.9</b> (25.3%)	<b>5.1</b> / <b>6.2</b>	42.8 / 116.0	<b>78.2</b> / <b>54.9</b> (29.8%)	9.8 / <b>11.3</b>	40.0 / 128.5

Table 1. Ablation study comparing pose distribution modelling choices in terms of accuracy, sample-input consistency and sample diversity metrics (see Section 4) on 3DPW. ‘‘Point’’ indicates point estimate metrics. MPJPE-PA and 3DKP spread are in mm, while 2DKP reprojection error is in pixels. Brackets contain % decreases from point estimate MPJPE-PA to the min. sample MPJPE-PA computed using 100 samples.

Given an input image, we first compute a proxy representation  $\mathbf{X} \in \mathbb{R}^{H \times W \times C}$  consisting of an edge map and 2D keypoint heatmaps [42], stacked along the channel dimension (see Figure 2). These are obtained using Canny edge detection [5] and HRNet-W48 [48] respectively. Proxy representations are often used to bridge the gap between synthetic training images and real test images [6, 41].

The proxy representation is passed through a CNN encoder [15] to give input features. The camera  $\pi$  and global rotation  $\mathbf{R}_{\text{glob}}$  are regressed from these features with an MLP.

Next, we predict a joint distribution over SMPL pose and shape parameters,  $p_{\text{joint}}(\{\mathbf{R}_i\}_{i=1}^{23}, \beta | \mathbf{X})$ , conditioned on the input proxy representation  $\mathbf{X}$ .  $p_{\text{joint}}$  is factorised into

$$p_{\text{joint}}(\{\mathbf{R}_i\}_{i=1}^{23}, \beta | \mathbf{X}) = p_{\text{shape}}(\beta | \mathbf{X}) p_{\text{pose}}(\{\mathbf{R}_i\}_{i=1}^{23} | \beta, \mathbf{X}). \quad (8)$$

We condition the full body pose distribution  $p_{\text{pose}}$  on  $\beta$  as this determines 3D body-part proportions, which affects the posed locations of mesh vertices. In practice, we also explicitly condition  $p_{\text{pose}}$  on  $\pi$  and  $\mathbf{R}_{\text{glob}}$  (see Figure 2). These are functions of  $\mathbf{X}$ , and are notationally omitted for simplicity.

Following [43], we predict a Gaussian shape distribution

$$p_{\text{shape}}(\beta | \mathbf{X}) = \mathcal{N}(\beta; \mu_{\beta}(\mathbf{X}), \text{diag}(\sigma_{\beta}^2(\mathbf{X}))) \quad (9)$$

where  $\mu_{\beta}$  and  $\sigma_{\beta}^2$  are obtained with an MLP.

In SMPL, each body-part’s pose is defined relative to its parent joint. The parent joint is rotated about its own parent, all the way up the kinematic tree. Thus, it is reasonable to inform the  $i$ -th body-part’s pose  $\mathbf{R}_i$  on  $\{\mathbf{R}_j\}_{j \in A(i)}$ , the rotations of all its kinematic ancestors  $A(i)$ , as noted by [13, 42]. This motivates an autoregressive factorisation of  $p_{\text{pose}}$  into ancestor-conditioned per-body-part rotation distributions:

$$\begin{aligned} p_{\text{pose}}(\{\mathbf{R}_i\}_{i=1}^{23} | \beta, \mathbf{X}) &= \prod_{i=1}^{23} p_{SO(3)}(\mathbf{R}_i | \{\mathbf{R}_j\}_{j \in A(i)}, \beta, \mathbf{X}) \\ &= \prod_{i=1}^{23} p_{SO(3)}(\mathbf{R}_i | \mathbf{c}_i) \end{aligned} \quad (10)$$

where  $\mathbf{c}_i$  is a context vector, which is computed as a function of  $\{\mathbf{R}_j\}_{j \in A(i)}$ ,  $\beta$  and  $\mathbf{X}$ , as shown in Figure 2. Autoregressive factorisation is similar to the hierarchical distribution proposed in [42]. However, we condition part rotations directly on ancestor *rotations*, instead of ancestor *distribution parameters* as in [42]. This enables more input-consistent distributions, since rotation samples give the *exact* 3D locations of ancestor joints, while distribution parameters only say what the ancestors’ rotations are *likely* to be.

We implement  $p_{SO(3)}(\mathbf{R}_i | \mathbf{c}_i)$ , the  $i$ -th body-part rotation distribution, by first defining a conditional NF over the axis-angle vector  $\mathbf{v}_i \in \mathbb{R}^3 \cong \mathfrak{so}(3)$ , with density function  $p_{\mathbb{R}^3}(\mathbf{v}_i | \mathbf{c}_i)$ . This is shown in Figure 2, where the  $i$ -th flow diffeomorphism is denoted as  $f_i : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ . Then, letting  $\mathbf{v}_i = \theta_i \mathbf{u}_i$ , Eqn. 6 is used to push  $p_{\mathbb{R}^3}(\mathbf{v}_i | \mathbf{c}_i)$  onto  $SO(3)$ , finally yielding  $p_{SO(3)}(\mathbf{R}_i | \mathbf{c}_i)$  via the exponential map.

Eqn. 6 requires  $p_{\mathbb{R}^3}(\mathbf{v}_i | \mathbf{c}_i)$  to have compact support within  $B_r(\mathbf{0})$  with  $\pi < r < 2\pi$ . Thus, we implement a bijective radial tanh transform [12]  $t : \mathbb{R}^3 \rightarrow B_r(\mathbf{0})$ , where

$$t(\mathbf{x}) = r \tanh\left(\frac{\|\mathbf{x}\|}{r}\right) \frac{\mathbf{x}}{\|\mathbf{x}\|}, \quad (11)$$

as the last layer of each  $f_i$ . We improve the transform proposed in [12], by ensuring that  $t(\mathbf{x}) \approx \mathbf{x}$  for small  $\|\mathbf{x}\|$ , which empirically aids training as shown in the supplement.

### 3.6. Pose and shape sampling and point estimation

SMPL pose and shape samples can be obtained from  $p_{\text{joint}}(\{\mathbf{R}_i\}_{i=1}^{23}, \beta | \mathbf{X})$  via ancestral sampling, as shown in Figure 2. Specifically, we first sample  $\beta \sim p_{\text{shape}}(\beta | \mathbf{X})$ , which will be used to condition  $p_{\text{pose}}$ . Then, each body-part’s rotation  $\mathbf{R}_i \sim p_{SO(3)}(\mathbf{R}_i | \mathbf{c}_i)$  is sampled following the corresponding ‘‘limb’’ of the kinematic tree, by first sampling all ancestor rotations  $\{\mathbf{R}_j\}_{j \in A(i)}$ , then obtaining the context vector  $\mathbf{c}_i$  from  $\{\mathbf{R}_j\}_{j \in A(i)}$ ,  $\beta$  and  $\mathbf{X}$  using an MLP. Pose and shape samples are converted into 3D vertex mesh samples with the SMPL function  $\mathcal{M}$ . The variance of each vertex along the x, y and z directions highlights the uncertainty cap-

Losses Used			3DPW			3DPW Cropped		
NLL	2DKP Samples	3DKP + 3D Vert. Point Estimate	Accuracy MPJPE-PA Point / Sample Min.	Consistency 2DKP Error Point / Samples	Diversity 3DKP Spread Vis. / Invis.	Accuracy MPJPE-PA Point / Sample Min.	Consistency 2DKP Error Point / Samples	Diversity 3DKP Spread Vis. / Invis.
✓			56.3 / 42.3 (24.9%)	5.6 / 8.1	55.7 / <b>124.0</b>	90.2 / 64.9 (28.0%)	11.0 / 15.1	66.3 / <b>141.0</b>
✓	✓		53.4 / <b>39.9</b> (25.3%)	<b>5.1</b> / <b>6.2</b>	42.8 / 116.0	<b>78.2</b> / <b>54.9</b> (29.8%)	<b>9.8</b> / <b>11.3</b>	40.0 / 128.5
✓	✓	✓	<b>53.3</b> / <b>39.9</b> (25.1%)	<b>5.1</b> / 6.3	39.3 / 109.6	83.4 / 61.1 (26.7%)	<b>9.8</b> / 11.4	39.7 / 115.1

Table 2. Ablation study comparing probabilistic losses - i.e. negative log-likelihood and visibility-masked 2DKP samples loss - and non-probabilistic point estimate losses on 3D keypoints and vertices. Distribution accuracy, sample-input consistency and sample diversity metrics are detailed in Section 4. “Point” indicates point estimate metrics. MPJPE-PA and 3DKP spread are in mm, while 2DKP reprojection error is in pixels. Brackets contain % decreases from point estimate MPJPE-PA to the min. sample MPJPE-PA computed using 100 samples.

tured by the predicted distribution  $p_{\text{joint}}$ , arising due to depth (i.e. z-axis) ambiguity, occlusion and truncation.

To obtain a point estimate of SMPL pose and shape given an input image, we would want to compute the mode  $(\{\mathbf{R}_i^*\}_{i=1}^{23}, \beta^*) = \operatorname{argmax} p_{\text{joint}}(\{\mathbf{R}_i\}_{i=1}^{23}, \beta | X)$ . This is challenging since  $p_{\text{joint}}$  is the product of multiple complex NF distributions. The mode of each per-body-part NF distribution is itself non-trivial. As an approximation, we use

$$\beta^* = \mu_\beta, \mathbf{R}_i^* = \exp(f_i(\mathbf{0}; \mathbf{c}_i)) \quad (12)$$

where the  $i$ -th body-part’s pose estimate  $\mathbf{R}_i^*$  is acquired by passing the base distribution mode ( $\mathbf{0}$ ) through the  $i$ -th flow transform and exp map. While the resulting  $(\{\mathbf{R}_i^*\}_{i=1}^{23}, \beta^*)$  is not, in general, the mode of  $p_{\text{joint}}$ , we show that it typically has high likelihood under  $p_{\text{joint}}$  in the supplement.

### 3.7. Loss functions

We train our model using a dataset of synthetic inputs paired with ground-truth pose, shape and global rotation labels  $\{\mathbf{X}^n, \{\bar{\mathbf{R}}_i^n\}_{i=1}^{23}, \bar{\beta}^n, \bar{\mathbf{R}}_{\text{glob}}^n\}_{n=1}^N$ , as discussed in Section 4. We apply a negative log-likelihood (NLL) loss

$$\mathcal{L}_{\text{NLL}} = - \sum_{n=1}^N \ln p_{\text{joint}}(\{\bar{\mathbf{R}}_i^n\}_{i=1}^{23}, \bar{\beta}^n | \mathbf{X}^n) \quad (13)$$

over pose and shape parameters.  $\mathbf{R}_{\text{glob}}$  is supervised using

$$\mathcal{L}_{\text{glob}} = \sum_{n=1}^N \|\mathbf{R}_{\text{glob}}(\mathbf{X}_n) - \bar{\mathbf{R}}_{\text{glob}}^n\|_F^2. \quad (14)$$

Following [23, 42], we apply a loss between 2D keypoint samples and *visible* GT 2D keypoints,  $\mathcal{L}_{2\text{D}}$ , to encourage sample-input consistency. 2D keypoint samples are obtained by sampling pose and shape from  $p_{\text{joint}}$ , computing the corresponding 3D keypoint samples with SMPL, and projecting these onto the image plane using the predicted camera  $\pi$ .

The overall training loss is given by  $\mathcal{L} = \lambda_{\text{NLL}} \mathcal{L}_{\text{NLL}} + \lambda_{\text{glob}} \mathcal{L}_{\text{glob}} + \lambda_{2\text{D}} \mathcal{L}_{2\text{D}}$  where the  $\lambda$ s are weights. We do not use ubiquitous, but non-probabilistic, point estimate losses on 3D keypoints, as justified by Section 5.1 and Table 2.

## 4. Implementation Details

**Model architecture.** We use a ResNet-18 [15] CNN encoder. Per-body-part axis-angle probability densities  $p_{\mathbb{R}^3}(\mathbf{v}_i | \mathbf{c}_i)$  are implemented with Linear Rational Spline normalising flows (LRS-NFs) [10]. Further architecture and hyperparameter details are provided in the supplementary material.

**Synthetic training data.** We adopt the same training data generation pipeline as [42], which renders synthetic proxy representation inputs  $\{\mathbf{X}^n\}_{n=1}^N$  from ground-truth (GT) poses, shapes and global rotations  $\{\{\bar{\mathbf{R}}_i^n\}_{i=1}^{23}, \bar{\beta}^n, \bar{\mathbf{R}}_{\text{glob}}^n\}_{n=1}^N$ . GT poses and global rotations are sampled from the training sets of UP-3D [24], 3DPW [51] and Human3.6M [16]. GT shapes are randomly sampled from a prior Gaussian distribution. Truncation, occlusion and noise augmentations bridge the synthetic-to-real gap.

**Training details.** We use Adam [19], with a learning rate of 1e-4 and batch size of 72, and train for 200 epochs.

**Evaluation datasets and metrics.** We use the 3DPW [51] and SSP-3D [41] datasets to evaluate the *accuracy*, sample-input *consistency* and sample *diversity* of predicted pose and shape distributions. Moreover, we generate cropped versions of 3DPW and SSP-3D, resulting in more ambiguous test data to evaluate sample diversity. Cropped dataset generation details are provided in the supplementary material.

*Distribution accuracy* refers to the likelihood of the GT pose and shape under the predicted distribution. We measure accuracy on 3DPW using MPJPE and MPJPE-PA computed with the *minimum* error sample out of  $N$  samples from the predicted distribution, where  $N$  is increased from 1 to 100. If the GT pose and shape have high likelihood under the predicted distribution, we expect the minimum MPJPE over  $N$  samples to improve significantly with increasing  $N$ . For  $N = 1$ , the sample is obtained as a point estimate from the predicted pose and shape distribution, using Eqn. 12. Similarly, shape-specific distribution accuracy is measured on SSP-3D, using PVE-T-SC (per-T-pose-vertex-error after scale correction) computed with the minimum error sample.

*Sample-input consistency* denotes the extent to which predicted 3D samples match the *visible* pose and shape evidence in the input 2D image. We measure consistency on 3DPW and SSP-3D using the average reprojection error be-

Method	3DPW - Accuracy			
	MPJPE (mm)		MPJPE-PA (mm)	
	Point	Sample Min.	Point	Sample Min.
HMR [18]	130.0	-	76.7	-
SPIN [21]	96.9	-	59.0	-
I2L-MeshNet [33]	93.2	-	57.7	-
DaNet [57]	85.5	-	54.8	-
HUND [55]	81.4	-	57.5	-
PARE [20]	74.5	-	46.5	-
HybriK [27]	<b>74.1</b>	-	<b>45.0</b>	-
3D Multibodies [2]	93.8	74.6 (20.5%)	59.9	48.3 (19.4%)
Sengupta <i>et al.</i> [43]	97.1	84.4 (13.1%)	61.1	52.1 (14.7%)
ProHMR [23]	97.0	81.5 (16.0%)	59.8	48.2 (19.4%)
HierProbHuman [42]	84.9	70.9 (16.5%)	53.6	43.8 (18.3%)
HuManiFlow	<b>83.9</b>	<b>65.1 (22.4%)</b>	<b>53.4</b>	<b>39.9 (25.3%)</b>

Table 3. Comparison of recent deterministic (top) and probabilistic (bottom) methods in terms of 3D point estimate and distribution accuracy on 3DPW [51].

tween predicted 2D keypoint (2DKP) samples projected onto the image plane and visible GT 2DKPs, averaged over 100 samples for each test image. The 17 COCO keypoint convention is used [28]. We also compute the reprojection error between 2DKP point estimates (Eqn. 12) and GT 2DKPs. An input-consistent distribution should have low average sample reprojection error, close to that of the point estimate.

*Sample diversity* refers to the range of 3D reconstructions modelled by the predicted distribution. We measure diversity by drawing 100 predicted samples, and computing the average 3D Euclidean distance from the mean for each 3D keypoint (3DKP), split into *visible* and *invisible* keypoints. A diverse distribution should exhibit significant spread in sample 3DKP locations - typically along the z-axis for visible 3DKPs (depth ambiguity) and all axes for invisible 3DKPs. This is a simplistic metric; defining a good diversity metric for high-dimensional, complex distributions is non-trivial, and a potential area for future research.

## 5. Experimental Results

This section first presents ablation studies on distribution predictors and losses, and then compares our method’s accuracy, consistency and diversity with the state-of-the-art.

### 5.1. Ablation studies

**Pose distribution modelling choices.** Table 1 compares several different pose distribution prediction models, in terms of accuracy, consistency and diversity on 3DPW. Rows 1-3 report metrics from a naive approach, where SMPL pose parameters are simply treated as a vector  $\Theta \in \mathbb{R}^{69}$  formed by concatenating per-body-part axis-angle vectors  $\mathbf{v}_i \in \mathbb{R}^3$ . The distribution over  $\Theta$  is modelled as a multivariate Gaussian, MDN [3] or LRS-NF [10]. Notably, all three models perform similarly, despite the greater theoretical expressiveness of MDNs and LRS-NFs. This suggests that it is challenging

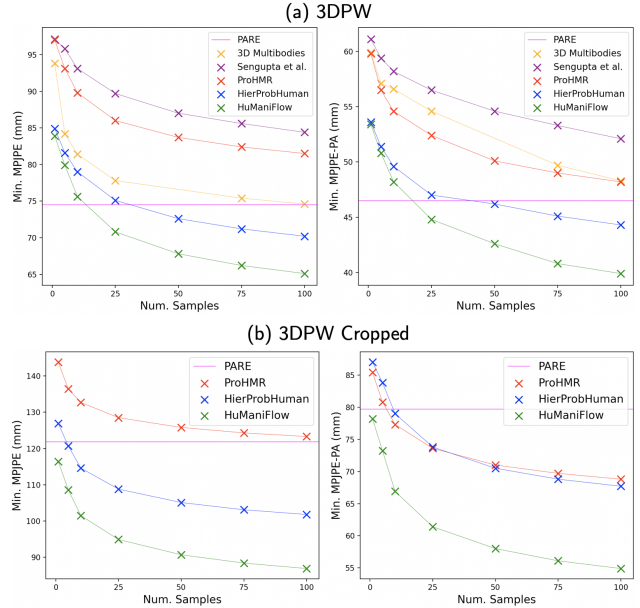


Figure 4. Rate of improvement in min. sample MPJPE(-PA) with increasing number of samples. A faster improvement rate indicates that the GT solution has higher likelihood under the predicted distribution - i.e. better distribution accuracy (see Section 4).

to directly predict complex probability distributions over high-dimensional full body pose. The shortcomings of high-dimensional MDNs in particular are well known [30, 38].

Rows 4-6 in Table 1 investigate an autoregressive approach, where the full body pose distribution is factorised into per-body-part distributions on axis-angle vectors  $\mathbf{v}_i$ , similar to Eqn. 10. The axis-angle distributions are modelled as a 3D Gaussian, MDN or LRS-NF. However, they are not pushed onto  $SO(3)$  using the change-of-variables formula for the exp map (Eqn. 6). Thus, the manifold structure of  $SO(3)$  and local “change of volume” due to exp are disregarded. These models improve consistency and accuracy over the naive approach (Rows 1-3), shown by better average sample 2DKP error and minimum sample MPJPE-PA. Despite the latter, they yield worse point estimate MPJPE-PA on the 3DPW Cropped split. This suggests that the point estimate may not be close to the GT, even for an accurate distribution, when faced with highly ambiguous inputs. The naive approach has greater sample diversity - but this is meaningless if the samples are not input-consistent.

Rows 7-9 in Table 1 investigate autoregressive per-body-part distributions over  $SO(3)$  (Eqn. 10). These are defined by pushing either MDNs or LRS-NFs over  $\mathbb{R}^3 \cong \mathfrak{so}(3)$  through exp onto  $SO(3)$ , or as a matrix-Fisher distribution [11, 31, 32]. Autoregressive LRS-NFs on  $SO(3)$  yield the best accuracy and consistency metrics, with only slightly reduced diversity; thus, we use this model in HuManiFlow.

**Distribution prediction losses.** Table 2 explores probabilistic losses - i.e. NLL and a loss on 2DKP *samples* (see Section

Method	SSP-3D - Accuracy (Shape)		
	PVE-T-SC (mm)		
	Point	Sample Min.	Multi-Input
SPIN [21]	22.2	-	21.9 (Mean)
PARE [20]	21.7	-	21.6 (Mean)
HybrIK [27]	22.9	-	22.8 (Mean)
STRAPS [41]	15.9	-	14.4 (Mean)
3D Multibodies [2]	22.3	19.2 (13.9%)	22.1 (Mean)
Sengupta <i>et al.</i> [43]	15.2	10.4 (31.6%)	13.3 (Prob. Comb. [43])
ProHMR [23]	22.2	-	21.9 (Mean)
HierProbHuman [42]	13.6	8.7 (36.0%)	12.0 (Prob. Comb.)
HuManiFlow	<b>13.5</b>	<b>8.3 (39.0%)</b>	<b>11.9</b> (Prob. Comb.)

Table 4. Comparison of recent deterministic (top) and probabilistic (bottom) methods in terms of shape accuracy on SSP-3D [41].

3.7) - and non-probabilistic losses (e.g. MSE) applied between GT 3DKPs and point estimates (Eqn. 12). Despite their ubiquity, we find that 3D point estimate losses do not improve distribution accuracy (i.e. min. Sample MPJPE-PA), or even point estimate MPJPE-PA, which is actually worsened on 3DPW Cropped. This is likely due to the ill-posedness of such losses for ambiguous inputs, where the GT 3DKPs represent but one of many plausible solutions. Point estimate losses also reduce sample diversity. Using NLL and a visibility-masked loss on 2DKP samples results in the best overall performance. Omitting the 2DKP samples loss degrades sample-input consistency, as expected.

## 5.2. Comparison with the state-of-the-art

**Distribution accuracy.** Tables 3 and 4 evaluate the accuracy of current pose and shape estimators. HuManiFlow predicts more accurate distributions (i.e. lower min. sample metrics) than other probabilistic approaches. This is corroborated by Figure 4, where HuManiFlow has the fastest decrease in min. sample MPJPE as the number of samples is increased. On 3DPW Cropped and SSP-3D, HuManiFlow’s point estimates have lower error than the deterministic SOTA approaches.

**Sample-input consistency and sample diversity.** Table 5 compares the consistency and diversity of current probabilistic methods. Samples from [23, 42] are input-consistent but not diverse, likely due to the use of unimodal pose distributions and 3D point estimate losses. [2] generates diverse samples but does not always match the input. Our method is generally the most input-consistent, with reasonably diverse samples. A qualitative demonstration is given in Figure 3.

**Model fitting with an image-conditioned prior.** Distributions with greater accuracy, consistency and diversity should be better for downstream tasks. An example task is model fitting [4] i.e. optimising SMPL point estimates to better fit observed 2DKPs. ProHMR [23] use their predicted distribution as an *image-conditioned prior* during fitting, outperforming generic pose priors [4, 36]. Table 6 shows that HuManiFlow surpasses [23] in this task. Note the difference between prior-

Dataset	Method	Consistency	Diversity
		2DKP Error Point / Samples	3DKP Spread Vis. / Invis.
3DPW	3D Multibodies [2]	5.2 / 7.8	80.1 / <b>126.9</b>
	Sengupta <i>et al.</i> [43]	5.6 / 8.1	48.3 / 98.8
	ProHMR [23]	6.8 / 7.5	35.1 / 60.8
	HierProbHuman [42]	<b>5.1</b> / 7.2	47.6 / 101.4
	HuManiFlow	<b>5.1</b> / <b>6.2</b>	42.8 / 116.0
3DPW	Sengupta <i>et al.</i> [43]	11.5 / 17.1	47.8 / 96.5
	ProHMR [23]	11.9 / 13.4	32.1 / 57.1
Cropped	HierProbHuman [42]	<b>9.7</b> / 12.8	38.5 / 100.2
	HuManiFlow	9.8 / <b>11.3</b>	40.0 / <b>128.5</b>
SSP-3D	3D Multibodies [2]	5.3 / 7.8	80.7 / -
	Sengupta <i>et al.</i> [43]	6.2 / 8.0	50.1 / -
	ProHMR [23]	6.9 / 7.6	36.6 / -
	HierProbHuman [42]	<b>4.8</b> / 6.9	48.5 / -
	HuManiFlow	<b>4.8</b> / <b>6.0</b>	47.3 / -
SSP-3D	Sengupta <i>et al.</i> [43]	13.3 / 18.9	60.2 / <b>139.6</b>
	ProHMR [23]	13.8 / 15.2	41.9 / 60.8
Cropped	HierProbHuman [42]	<b>10.6</b> / 14.1	58.7 / 105.5
	HuManiFlow	<b>10.6</b> / <b>13.0</b>	45.3 / 134.0

Table 5. Comparison of recent probabilistic methods in terms of sample-input consistency and sample diversity on 3DPW [51] and SSP-3D [41]. 3DKP spread is in mm and 2DKP error is in pixels.

Method	Prior	3DPW MPJPE-PA (mm)
SPIN [21] + Fit	GMM [4]	66.5
SPIN + Fit	VPoser [36]	70.9
SPIN + EFT [17]	-	56.6
ProHMR [23] + Fit	Image-conditioned	55.1
HuManiFlow + Fit	None	53.4
HuManiFlow + Fit	Image-conditioned	<b>51.2</b>

Table 6. Evaluation of model fitting methods with different SMPL parameter priors, including image-conditioned priors from distribution prediction methods ([23] and HuManiFlow). Fitting does not necessarily improve 3D point estimate accuracy, despite better model-image alignment, unless image-conditioned priors are used.

less optimisation of HuManiFlow’s point estimates (row 5) versus using the predicted distribution as a prior (row 6).

## 6. Conclusion

This work proposes a probabilistic approach to the ill-posed problem of monocular 3D human pose and shape estimation. We show that current methods suffer from a trade-off between distribution accuracy, sample-input consistency and sample diversity, which affects their utility in downstream tasks. Our method, HuManiFlow, uses a normalising-flow-based pose distribution which (i) accounts for the manifold structure of  $SO(3)$ , (ii) has an autoregressive factorisation informed by the human kinematic tree, and (iii) is trained without any ill-posed 3D point estimate losses. HuManiFlow yields more accurate, consistent and diverse distributions.



## References

- [1] Sadegh Aliakbarian, Pashmina Cameron, Federica Bogo, Andrew Fitzgibbon, and Thomas J. Cashman. Flag: Flow-based 3d avatar generation from sparse observations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [1](#)
- [2] Benjamin Biggs, Sébastien Erhardt, Hanbyul Joo, Benjamin Graham, Andrea Vedaldi, and David Novotny. 3D multibodies: Fitting sets of plausible 3D models to ambiguous image data. In *NeurIPS*, 2020. [1](#), [2](#), [4](#), [7](#), [8](#)
- [3] Christopher M. Bishop. Mixture density networks. Technical report, 1994. [2](#), [4](#), [5](#), [7](#)
- [4] 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 *Proceedings of the European Conference on Computer Vision (ECCV)*, Oct. 2016. [2](#), [8](#)
- [5] John F. Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 8(6):679–698, 1986. [5](#)
- [6] J. Charles, S. Bucciarelli, and R. Cipolla. Real-time screen reading: reducing domain shift for one-shot learning. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2020. [5](#)
- [7] Hongsuk Choi, Gyeongsik Moon, and Kyoung Mu Lee. Pose2mesh: Graph convolutional network for 3D human pose and mesh recovery from a 2D human pose. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. [2](#)
- [8] Kiam Choo and D.J. Fleet. People tracking using hybrid monte carlo filtering. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, volume 2, pages 321–328 vol.2, 2001. [2](#)
- [9] Laurent Dinh, David Krueger, and Yoshua Bengio. NICE: Non-linear independent components estimation. In *ICLR Workshop Track*, 2015. [2](#)
- [10] Hadi Mohaghegh Dolatabadi, Sarah Erfani, and Christopher Leckie. Invertible generative modeling using linear rational splines. In *The 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 4236–4246, 2020. [3](#), [5](#), [6](#), [7](#)
- [11] Thomas D. Downs. Orientation statistics. *Biometrika*, 59(3):665–676, 12 1972. [7](#)
- [12] L. Falorsi, P. de Haan, T.R. Davidson, and P. Forré. Reparameterizing distributions on lie groups. *22nd International Conference on Artificial Intelligence and Statistics (AISTATS-19)*, 2019. [2](#), [3](#), [4](#), [5](#)
- [13] Georgios Georgakis, Ren Li, Srikrishna Karanam, Terrence Chen, Jana Kosecka, and Ziyang Wu. Hierarchical kinematic human mesh recovery. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. [5](#)
- [14] Riza Alp Guler and Iasonas Kokkinos. Holopose: Holistic 3D human reconstruction in-the-wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [2](#)
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. [5](#), [6](#)
- [16] 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 Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 36(7):1325–1339, July 2014. [6](#)
- [17] Hanbyul Joo, Natalia Neverova, and Andrea Vedaldi. Exemplar fine-tuning for 3D human pose fitting towards in-the-wild 3D human pose estimation. *arXiv preprint arXiv:2004.03686*, 2020. [2](#), [8](#)
- [18] Angjoo Kanazawa, Michael J. Black, David W. Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [1](#), [2](#), [7](#)
- [19] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2014. [6](#)
- [20] Muhammed Kocabas, Chun-Hao P. Huang, Otmar Hilliges, and Michael J. Black. PARE: Part attention regressor for 3D human body estimation. In *Proceedings of the International Conference on Computer Vision (ICCV)*, pages 11127–11137, Oct. 2021. [1](#), [2](#), [7](#), [8](#)
- [21] 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 *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019. [1](#), [2](#), [7](#), [8](#)
- [22] Nikos Kolotouros, Georgios Pavlakos, and Kostas Daniilidis. Convolutional mesh regression for single-image human shape reconstruction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [2](#)
- [23] Nikos Kolotouros, Georgios Pavlakos, Dinesh Jayaraman, and Kostas Daniilidis. Probabilistic modeling for human mesh recovery. In *ICCV*, 2021. [1](#), [2](#), [4](#), [6](#), [7](#), [8](#)
- [24] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. [2](#), [6](#)
- [25] 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. [1](#)
- [26] Chen Li and Gim Hee Lee. Generating multiple hypotheses for 3D human pose estimation with mixture density network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [2](#)
- [27] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. [1](#), [2](#), [7](#), [8](#)
- [28] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2014. [7](#)

- [29] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. In *ACM Transactions on Graphics (TOG) - Proceedings of ACM SIGGRAPH Asia*, volume 34, pages 248:1–248:16. ACM, 2015. 2, 3, 4
- [30] O. Makansi et al. Overcoming limitations of mixture density networks: A sampling and fitting framework for multimodal future prediction. In *CVPR*, 2019. 7
- [31] K. V. Mardia and P. E. Jupp. *Directional statistics*. Wiley, 2000. 7
- [32] David Mohlin, Josephine Sullivan, and Gérald Bianchi. Probabilistic orientation estimation with matrix fisher distributions. In *Advances in Neural Information Processing Systems*, volume 33, 2020. 7
- [33] Gyeongsik Moon and Kyoung Mu Lee. I2L-MeshNet: Image-to-lixel prediction network for accurate 3D human pose and mesh estimation from a single rgb image. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2, 7
- [34] Tuomas P. Oikarinen, Daniel C. Hannah, and Sohrob Kazerooni. GraphMDN: Leveraging graph structure and deep learning to solve inverse problems. *CoRR*, abs/2010.13668, 2020. 2
- [35] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22(57):1–64, 2021. 2
- [36] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 8
- [37] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In Francis Bach and David Blei, editors, *Proceedings of the International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 1530–1538, Lille, France, 07–09 Jul 2015. PMLR. 2, 3, 4
- [38] C. Rupprecht et al. Learning in an uncertain world: Representing ambiguity through multiple hypotheses. In *CVPR*, 2017. 7
- [39] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, October 2019. 2
- [40] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3D human digitization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2
- [41] Akash Sengupta, Ignas Budvytis, and Roberto Cipolla. Synthetic training for accurate 3D human pose and shape estimation in the wild. In *Proceedings of the British Machine Vision Conference (BMVC)*, September 2020. 2, 5, 6, 8
- [42] Akash Sengupta, Ignas Budvytis, and Roberto Cipolla. Hierarchical Kinematic Probability Distributions for 3D Human Shape and Pose Estimation from Images in the Wild. In *Proceedings of the International Conference on Computer Vision*, October 2021. 1, 2, 4, 5, 6, 7, 8
- [43] Akash Sengupta, Ignas Budvytis, and Roberto Cipolla. Probabilistic 3D human shape and pose estimation from multiple unconstrained images in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2, 5, 7, 8
- [44] Akash Sengupta, Ignas Budvytis, and Roberto Cipolla. Probabilistic estimation of 3D human shape and pose with a semantic local parametric model. In *Proceedings of the British Machine Vision Conference (BMVC)*, September 2021. 1, 2
- [45] Cristian Sminchisescu and Bill Trigg. Covariance scaled sampling for monocular 3D body tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2001. 2
- [46] Cristian Sminchisescu and Bill Trigg. Hyperdynamics importance sampling. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2002. 2
- [47] Cristian Sminchisescu and Bill Trigg. Kinematic jump processes for monocular 3D human tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2003. 1, 2
- [48] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 5
- [49] Vince J. K. Tan, Ignas Budvytis, and Roberto Cipolla. Indirect deep structured learning for 3D human shape and pose prediction. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2017. 2
- [50] Gül Varol, Duygu Ceylan, Bryan Russell, Jimei Yang, Ersin Yumer, Ivan Laptev, and Cordelia Schmid. BodyNet: Volumetric inference of 3D human body shapes. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 2
- [51] 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 *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 2, 6, 7, 8
- [52] Tom Wehrbein, Marco Rudolph, Bodo Rosenhahn, and Bastian Wandt. Probabilistic monocular 3D human pose estimation with normalizing flows. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 2
- [53] Christina Winkler, Daniel Worrall, Emiel Hooeboom, and Max Welling. Learning likelihoods with conditional normalizing flows, 2019. 3
- [54] Hongyi Xu, Eduard Gabriel Bazavan, Andrei Zanfir, William T Freeman, Rahul Sukthankar, and Cristian Sminchisescu. Ghum & ghuml: Generative 3d human shape and articulated pose models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6184–6193, 2020. 2
- [55] Andrei Zanfir, Eduard Gabriel Bazavan, Mihai Zanfir, William T. Freeman, Rahul Sukthankar, and Cristian Sminchisescu. Neural descent for visual 3D human pose and shape.

In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. [1](#), [2](#), [7](#)

- [56] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [2](#)
- [57] Hongwen Zhang, Jie Cao, Guo Lu, Wanli Ouyang, and Zhenan Sun. Danet: Decompose-and-aggregate network for 3D human shape and pose estimation. In *Proceedings of the 27th ACM International Conference on Multimedia*, pages 935–944, 2019. [2](#), [7](#)