

Hierarchical Kinematic Probability Distributions for 3D Human Shape and Pose Estimation from Images in the Wild

Akash Sengupta
University of Cambridge
as2562@cam.ac.uk

Ignas Budvytis
University of Cambridge
ib255@cam.ac.uk

Roberto Cipolla
University of Cambridge
rc10001@cam.ac.uk

Abstract

This paper addresses the problem of 3D human body shape and pose estimation from an RGB image. This is often an ill-posed problem, since multiple plausible 3D bodies may match the visual evidence present in the input - particularly when the subject is occluded. Thus, it is desirable to estimate a distribution over 3D body shape and pose conditioned on the input image instead of a single 3D reconstruction. We train a deep neural network to estimate a hierarchical matrix-Fisher distribution over relative 3D joint rotation matrices (i.e. body pose), which exploits the human body's kinematic tree structure, as well as a Gaussian distribution over SMPL body shape parameters. To further ensure that the predicted shape and pose distributions match the visual evidence in the input image, we implement a differentiable rejection sampler to impose a reprojection loss between ground-truth 2D joint coordinates and samples from the predicted distributions, projected onto the image plane. We show that our method is competitive with the state-of-the-art in terms of 3D shape and pose metrics on the SSP-3D and 3DPW datasets, while also yielding a structured probability distribution over 3D body shape and pose, with which we can meaningfully quantify prediction uncertainty and sample multiple plausible 3D reconstructions to explain a given input image.

1. Introduction

3D human body shape and pose estimation from an RGB image is a challenging computer vision problem, partly due to its under-constrained nature wherein multiple 3D human bodies may explain a given 2D image, especially when the subject is significantly occluded, as is common for in-the-wild images. Several recent works [53, 20, 27, 26, 46, 63, 12, 37, 14, 40, 39, 54, 35, 52] use deep neural networks to regress a single body shape and pose solution, which can result in impressive 3D body reconstructions given sufficient visual evidence in the input image. However, when visual evidence of the subject's shape and pose is obscured, e.g. due to occluding objects

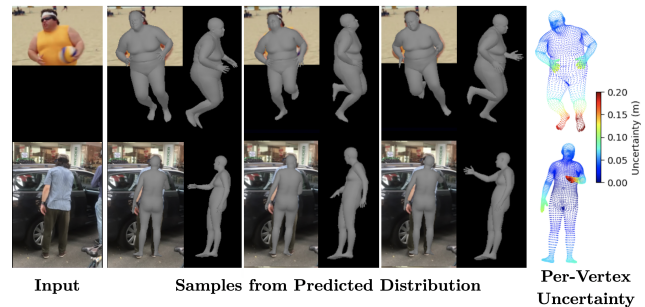


Figure 1. 3D reconstruction samples and per-vertex uncertainty corresponding to the predicted hierarchical shape and pose distributions computed from the given input images.

or self-occlusions, a single solution does not fully describe the space of plausible 3D reconstructions. In contrast, we aim to estimate a *structured probability distribution* over 3D body shape and pose, conditioned on the input image, thereby allowing us to sample any number of plausible 3D reconstructions and quantify prediction uncertainty over the 3D body surface, as shown in Figure 1.

We use the SMPL body model [32] to represent human shape and pose. Identity-dependent body shape is parameterised by coefficients of a PCA basis - hence, a simple multivariate Gaussian distribution over the shape parameters is suitable. Body pose is parameterised by relative 3D joint rotations along the SMPL kinematic tree, which may be represented using rotation matrices. Regressing rotation matrices using neural networks is non-trivial, since they lie in $SO(3)$, a non-linear 3D manifold with a different topology to $\mathbb{R}^{3 \times 3}$ or \mathbb{R}^9 , the space in which unconstrained neural network outputs lie. However, one can define probability density functions over the Lie group $SO(3)$, such as the matrix-Fisher distribution [33, 11, 22], the parameter of which is an element of $\mathbb{R}^{3 \times 3}$ and may be easily regressed with a neural network [34]. We propose a hierarchical probability distribution over relative 3D joint rotations along the SMPL kinematic tree, wherein the probability density function of each joint's relative rotation matrix is a matrix-Fisher distribution conditioned on the parents of that joint in the

kinematic tree. We train a deep neural network to predict the parameters of such a distribution over body pose, alongside a Gaussian distribution over SMPL shape.

Moreover, to ensure that 3D bodies sampled from the predicted distributions match the 2D input image, we implement a reprojection loss between predicted samples and ground-truth visible 2D joint annotations. To allow for the backpropagation of gradients through the sampling operation, we present a differentiable rejection sampler for matrix-Fisher distributions over relative 3D joint rotations.

Finally, a key obstacle for SMPL body shape regression from in-the-wild images is the lack of training datasets with *accurate* and *diverse* body shape labels [46]. To overcome this, we follow [46, 52, 40, 47] and utilise synthetic data, randomly generated on-the-fly during training. Inspired by [7], we use convolutional edge filters to close the large synthetic-to-real gap and show that using edge-based inputs yields better performance than commonly-used silhouette-based inputs [46, 52, 47, 40], due to improved robustness and capacity to retain visual shape information.

In summary, our main contributions are as follows:

- Given an input image, we predict a novel hierarchical matrix-Fisher distribution over relative 3D joint rotation matrices, whose structure is explicitly informed by the SMPL kinematic tree, alongside a Gaussian distribution over SMPL shape parameters.
- We present a differentiable rejection sampler to sample any number of plausible 3D reconstructions and quantify prediction uncertainty over the body surface. This enables a reprojection loss between predicted samples and ground-truth coordinates of visible 2D joints, further ensuring that the predicted distributions are consistent with the input image.
- We use simple convolutional edge filters to improve the random synthetic training framework used by [46, 47]. Edge filtering is a computationally-cheap and robust method for closing the domain gap between synthetic RGB training data and real RGB test data.

2. Related Work

This section reviews approaches to monocular 3D human body shape and pose estimation, as well as deep-learning-based methods for probabilistic rotation estimation.

Monocular 3D shape and pose estimation methods can be classified as optimisation-based or learning-based. Optimisation-based approaches fit a parametric 3D body model [32, 1, 38, 19] to 2D observations, such as 2D keypoints [5, 29], silhouettes [29] or body part segmentations [61], by optimising a suitable cost function. These methods do not require expensive 3D-labelled training data, but are sensitive to poor initialisations and noisy observations.

Learning-based approaches can be further split into model-free or model-based. Model-free methods use deep networks to directly output human body vertex meshes [27, 35, 63, 62, 8], voxel grids [54] or implicit surfaces [44, 45] from an input image. In contrast, model-based methods [20, 46, 37, 12, 53, 14, 40, 39, 59] regress 3D body model parameters [38, 32, 19, 1], which give a low-dimensional representation of a 3D human body. To overcome the lack of in-the-wild 3D-labelled training data, several methods [20, 59, 27, 12, 14] use diverse 2D-labelled data as a source of weak supervision. [26] extends this approach by incorporating optimisation into their model training loop, lifting 2D labels to self-improving 3D labels. These approaches often result in impressive 3D pose predictions, but struggle to accurately predict a diverse range of body shapes, since 2D keypoint supervision only provides a sparse shape signal. Shape prediction accuracy may be improved using synthetic training data [46, 52, 40, 47] consisting of synthetic input proxy representations (PRs) paired with ground-truth body shape and pose. PRs commonly consist of silhouettes and 2D joint heatmaps [46, 40, 47], necessitating accurate silhouette segmentations [25, 15] at test-time, which is not guaranteed for challenging in-the-wild inputs. Other methods [54] pre-train on synthetic RGB inputs [55] and then fine-tune on the scarce and limited-shape-diversity real 3D training data available [17, 56], to avoid over-fitting to artefacts in low-fidelity synthetic data. In contrast, we utilise edge-based PRs, hence dropping the reliance on accurate segmentation networks without requiring fine-tuning on real data or high-fidelity synthetic data.

3D human shape and pose distribution estimation. Early optimisation-based 3D pose estimators [49, 50, 51, 9, 10] specified a cost function corresponding to the posterior probability of 3D pose given 2D observations and analysed its multi-modal structure due to ill-posedness. Strategies to sample multiple 3D poses with high posterior probability included cost-covariance-scaled [49] and inverse-kinematics-based [51] global search and local refinement, as well as cost-function-modifying MCMC [50]. Recently, several learning-based methods [48, 31, 18, 57, 36] predict multi-modal distributions over 3D joint locations conditioned on 2D inputs, using Bayesian mixture of experts [48], mixture density networks [31, 4, 36] or normalising flows [57, 43]. Our method extends beyond 3D joints and predicts distributions over human pose *and shape*. This has been addressed by Biggs *et al.* [3], who predict a categorical distribution over a set of SMPL [32] parameter hypotheses. Sengupta *et al.* [47] estimate an independent Gaussian distribution over both SMPL shape and joint rotation vectors. In contrast, we note that 3D rotations lie in $SO(3)$, motivating our hierarchical matrix-Fisher distribution.

Rotation distribution estimation via deep learning. Prokudin *et al.* [41] use biternion networks to predict a

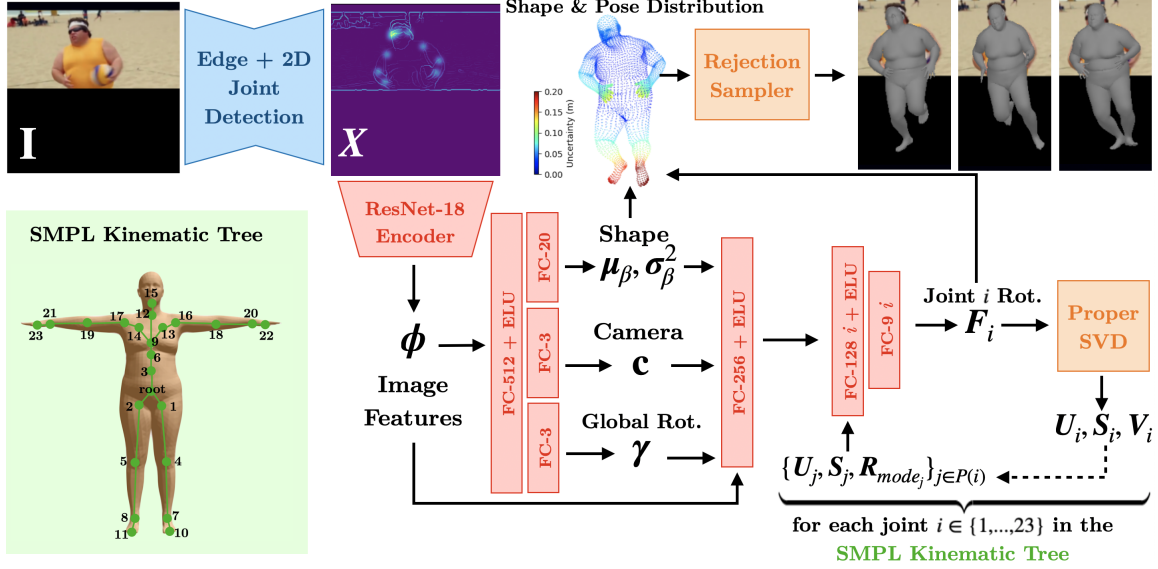


Figure 2. Network architecture of our hierarchical SMPL shape and pose distribution predictor. The given input image is converted into an edge-and-joint-heatmap proxy representation, which is passed through the prediction network to produce distributions over shape parameters and relative 3D joint rotation matrices. Rejection sampling is used to sample 3D reconstructions from the predicted distributions.

mixture-of-von-Mises distribution over object pose angle. Gilitschenski *et al.* [13] use a Bingham distribution over unit quaternions to represent orientation uncertainty. However, these works have to enforce constraints on the parameters of their predicted distributions (e.g. positive semi-definiteness). To overcome this, Mohlin *et al.* [34] train a deep network to regress a matrix-Fisher distribution [33, 11, 22] over 3D rotation matrices. We adapt this approach to define our hierarchical matrix-Fisher distribution over relative 3D joint rotation matrices.

3. Method

This section provides an overview of SMPL [32] and the matrix-Fisher distribution [11, 22, 33], presents our structured, hierarchical pose and shape distribution estimation architecture and discusses the loss functions used to train it.

3.1. SMPL model

SMPL [32] is a parametric 3D human body model. Identity-dependent body shape is represented by shape parameters $\beta \in \mathbb{R}^{10}$, which are coefficients of a PCA body shape basis. Body pose is defined by the relative 3D rotations of the bones formed by the 23 body (i.e. non-root) joints in the SMPL kinematic tree. The rotations may be represented using rotation matrices $\{\mathbf{R}_i\}_{i=1}^{23}$, where $\mathbf{R}_i \in SO(3)$. We parameterise the global rotation (i.e. rotation of the root joint) in axis-angle form by $\gamma \in \mathbb{R}^3$. A differentiable function $\mathcal{S}(\{\mathbf{R}_i\}_{i=1}^{23}, \beta, \gamma)$ maps the input pose and shape parameters to an output vertex mesh $\mathbf{V} \in \mathbb{R}^{6890 \times 3}$. 3D joint locations, for L joints of interest, are obtained as

$\mathcal{J}^{3D} = \mathcal{J}\mathbf{V}$ where $\mathcal{J} \in \mathbb{R}^{L \times 6890}$ is a linear vertex-to-joint regression matrix.

3.2. Matrix-Fisher distribution over $SO(3)$

The 3D special orthogonal group 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 matrix-Fisher distribution [11, 22, 33] defines a probability density function over $SO(3)$, given by

$$p(\mathbf{R}|\mathbf{F}) = \frac{1}{c(\mathbf{F})} \exp(\text{tr}(\mathbf{F}^T \mathbf{R})) = \mathcal{M}(\mathbf{R}; \mathbf{F}) \quad (1)$$

where $\mathbf{F} \in \mathbb{R}^{3 \times 3}$ is the matrix parameter of the distribution, $c(\mathbf{F})$ is the normalising constant and $\mathbf{R} \in SO(3)$. We present some key properties of the matrix-Fisher distribution below, but refer the reader to [30, 34] for further details, visualisations and a method for approximating the intractable normalising constant and its gradient w.r.t. \mathbf{F} .

The properties of $\mathcal{M}(\mathbf{R}; \mathbf{F})$ can be described in terms of the singular value decomposition (SVD) of \mathbf{F} , denoted by $\mathbf{F} = \mathbf{U}' \mathbf{S}' \mathbf{V}'^T$, with $\mathbf{S}' = \text{diag}(s'_1, s'_2, s'_3)$. \mathbf{U}' and \mathbf{V}' are orthonormal matrices, but they may have a determinant of -1 and thus are not necessarily elements of $SO(3)$. Therefore, a *proper* SVD [30] $\mathbf{F} = \mathbf{U} \mathbf{S} \mathbf{V}^T$ is used, where

$$\begin{aligned} \mathbf{U} &= \mathbf{U}' \text{diag}(1, 1, \det(\mathbf{U}')) \\ \mathbf{V} &= \mathbf{V}' \text{diag}(1, 1, \det(\mathbf{V}')) \\ \mathbf{S} &= \text{diag}(s_1, s_2, s_3) = \text{diag}(s'_1, s'_2, \det(\mathbf{U}' \mathbf{V}') s'_3) \end{aligned} \quad (2)$$

which ensures that $\mathbf{U}, \mathbf{V} \in SO(3)$. Then, the *mode* of the distribution is given by [30]

$$\mathbf{R}_{\text{mode}} = \arg \max_{\mathbf{R} \in SO(3)} p(\mathbf{R}|\mathbf{F}) = \mathbf{U} \mathbf{V}^T. \quad (3)$$

The columns of \mathbf{U} define the distribution’s *principal axes* of rotation (analogous to the principal axes of a multivariate Gaussian distribution), while the proper singular values in \mathbf{S} give the *concentration* of the distribution for rotations about the principal axes [30]. Specifically, the concentration along rotations of \mathbf{R}_{mode} about the i -th principal axis (i -th column of \mathbf{U}) is given by $s_j + s_k$ for $(i, j, k) \in \{(1, 2, 3), (2, 3, 1), (3, 1, 2)\}$. The concentration of the distribution may be different about each principal axis, allowing for axis-dependent rotation uncertainty modelling.

3.3. Proxy representation computation

Given an input RGB image \mathbf{I} , we first compute a proxy representation \mathbf{X} (see Figure 2), consisting of an edge-image concatenated with joint heatmaps. Comparisons with silhouette- and RGB-based representations are given in Section 5.1. Edge-images are obtained with Canny edge detection [6]. 2D joint heatmaps are computed using Detectron2 [16, 58], and joint predictions with low confidence scores are thresholded out. The edge-image and joint heatmaps are stacked along the channel dimension to produce $\mathbf{X} \in \mathbb{R}^{H \times W \times (L+1)}$. Proxy representations [46, 40] are used to close the domain gap between synthetic training images and real test-time RGB images, since synthetic proxy representations are more similar to their real counterparts than synthetic RGB images are to real RGB images.

3.4. Body shape and pose distribution prediction

Our goal is to predict a probability distribution over relative 3D joint rotations $\{\mathbf{R}_i\}_{i=1}^{23}$ and SMPL shape parameters β conditioned upon a given input proxy representation \mathbf{X} . We also predict deterministic estimates of the global body rotation γ and weak-perspective camera parameters $\mathbf{c} = [s, t_x, t_y]$, representing scale and xy translation.

Since β represents the linear coefficients of a PCA shape-space, a Gaussian distribution with a diagonal covariance matrix is suitable [47],

$$p(\beta|\mathbf{X}) = \mathcal{N}(\beta; \mu_\beta(\mathbf{X}), \text{diag}(\sigma_\beta^2(\mathbf{X}))) \quad (4)$$

where the mean μ_β and variances σ_β^2 are functions of \mathbf{X} .

The matrix-Fisher distribution (Equation 1) may be naively used to define a distribution over 3D joint rotations

$$p(\mathbf{R}_i|\mathbf{X}) = \mathcal{M}(\mathbf{R}_i; \mathbf{F}_i(\mathbf{X})) \quad (5)$$

for $i \in \{1, 2, \dots, 23\}$. Here, each joint is modelled *independently* of all the other joints. Thus, the matrix parameter of the i -th joint, \mathbf{F}_i , is a function of the input \mathbf{X} only.

To predict the parameters of this naive, independent distribution over 3D joint rotations, in addition to the shape distribution parameters, global body rotation and weak-perspective camera, we learn a function f_{indep} mapping the input \mathbf{X} to the set of desired outputs $\mathbf{Y} =$

$\{\{\mathbf{F}_i\}_{i=1}^{23}, \mu_\beta, \sigma_\beta^2, \gamma, \mathbf{c}\}$, where f_{indep} is represented by a deep neural network with weights $\mathbf{W}_{\text{indep}}$.

However, the independent matrix-Fisher distribution in Equation 5 does not model SMPL 3D joint rotations faithfully, since the rotation of each part/bone is defined *relative* to its parent joint in the SMPL kinematic tree. Hence, a distribution over the i -th rotation matrix \mathbf{R}_i conditioned on the input \mathbf{X} should be informed by the distributions over all its parent joints $P(i)$, as well as the global body rotation γ , to enable the distribution to match the 2D visual pose evidence present in \mathbf{X} . Furthermore, 3D joints in the SMPL rest-pose skeleton are dependent upon the shape parameters β , while the mapping from 3D to the 2D image plane is given by the camera model. Hence, a distribution over \mathbf{R}_i given \mathbf{X} should also consider the predicted shape mean μ_β and variance σ_β^2 , as well as the predicted camera \mathbf{c} . This is similar to the rationale behind the deterministic iterative/hierarchical predictors in [20, 12], except we model these relationships in a probabilistic sense, by defining

$$p(\mathbf{R}_i|\mathbf{X}, \{\mathbf{F}_j\}_{j \in P(i)}, \gamma, \mu_\beta, \sigma_\beta^2, \mathbf{c}) = \mathcal{M}(\mathbf{R}_i; \mathbf{F}_i) \quad (6)$$

$$\mathbf{F}_i = f_i(\mathbf{X}, \{\mathbf{U}_j, \mathbf{S}_j, \mathbf{R}_{\text{mode}_j}\}_{j \in P(i)}, \gamma, \mu_\beta, \sigma_\beta^2, \mathbf{c})$$

for $i \in \{1, 2, \dots, 23\}$. Now, the matrix parameter of the i -th joint is a function of all its parent distributions, represented by the principal axes \mathbf{U}_j , singular values \mathbf{S}_j and modes $\mathbf{R}_{\text{mode}_j} = \mathbf{U}_j \mathbf{V}_j^T$ for $j \in P(i)$, as well as the shape distribution $\{\mu_\beta, \sigma_\beta^2\}$, global rotation γ , camera parameters \mathbf{c} and the input \mathbf{X} . Note that the parent distributions are themselves functions of *their* respective parent joints, while $\gamma, \mu_\beta, \sigma_\beta^2$ and \mathbf{c} are all functions of \mathbf{X} .

To predict the parameters of the hierarchical matrix-Fisher distribution in Equation 6, we propose a hierarchical neural network architecture f_{hier} , with weights \mathbf{W}_{hier} (Figure 2). When considered as a black-box, f_{hier} yields the same set of outputs \mathbf{Y} as f_{indep} . However, f_{hier} utilises the iterative hierarchical architecture presented in Figure 2, which amounts to multiple streams of fully-connected layers, each following one “limb” of the kinematic tree. In contrast, f_{indep} predicts pose similarly to shape, camera and global rotation parameters, using a single stream of fully-connected layers. We compare the naive independent formulation with the hierarchical formulation in Section 5.1.

3.5. Loss functions

Distribution prediction networks are trained with a synthetic dataset $\{\mathbf{X}^n, (\{\mathbf{R}_i^n\}_{i=1}^{23}, \beta^n, \gamma^n)\}_{n=1}^N$ (Section 4). **Negative log-likelihood (NLL) loss on distribution parameters.** The NLL corresponding to the Gaussian body shape distribution (Equation 4) is given by:

$$\mathcal{L}_{\beta\text{-NLL}} = - \sum_{n=1}^N \log \mathcal{N}(\beta^n; \mu_\beta(\mathbf{X}^n), \text{diag}(\sigma_\beta^2(\mathbf{X}^n))). \quad (7)$$

Algorithm 1: Differentiable Rejection Sampler

Input: $U, S = \text{diag}(s_1, s_2, s_3), V, b$
Output: $\hat{R} \in SO(3)$ s.t. $\hat{R} \sim \mathcal{M}(R; USV^T)$
1 $A = \text{diag}(0, 2(s_2 + s_3), 2(s_1 + s_3), 2(s_1 + s_2))$
2 $\Omega = I_4 + \frac{2}{b}A$
3 $M = \exp\left(\frac{b-A}{2}\right)\left(\frac{4}{b}\right)^2$
4 **repeat**
5 Sample $\epsilon \sim \mathcal{N}(\mathbf{0}_4, I_4)$
6 $\mathbf{y} = (\Omega^{-1})^{\frac{1}{2}}\epsilon$
7 Propose $\mathbf{x} = \frac{\mathbf{y}}{\|\mathbf{y}\|}$ s.t. $\mathbf{x} \in S^3$
8 Sample $w \sim \text{Unif}[0, 1]$
9 **until** $w < \frac{\exp(-\mathbf{x}^T A \mathbf{x})}{M(\mathbf{x}^T \Omega \mathbf{x})^{-2}}$;
10 $\hat{Q} = \text{quaternion_to_matrix}(\mathbf{x})$ s.t. $\hat{Q} \in SO(3)$
11 **return** $\hat{R} = U\hat{Q}V^T$

The NLL corresponding to the matrix-Fisher distribution over relative 3D joint rotations is defined as [34]:

$$\begin{aligned} \mathcal{L}_{R\text{-NLL}} &= -\sum_{n=1}^N \log \mathcal{M}(\mathbf{R}_i^n; \mathbf{F}_i^n) \\ &= \sum_{n=1}^N \log c(\mathbf{F}_i^n) - \text{tr}(\mathbf{F}_i^{nT} \mathbf{R}_i^n) \end{aligned} \quad (8)$$

for $i \in \{1, 2, \dots, 23\}$, where \mathbf{F}_i^n may be obtained via the independent or hierarchical matrix-Fisher models presented above. Intuitively, the trace term pushes the predicted distribution mode $\mathbf{R}_{\text{mode}_i}^n$ (Equation 3) towards the target \mathbf{R}_i^n , while the log normalising constant acts as a regulariser, preventing the singular values of \mathbf{F}_i^n from getting too large [34]. All predicted distribution parameters are dependent on the model weights, $\mathbf{W}_{\text{indep}}$ or \mathbf{W}_{hier} , which are learnt in a maximum likelihood framework aiming to minimise the joint shape and pose NLL: $\mathcal{L}_{\text{NLL}} = \mathcal{L}_{\beta\text{-NLL}} + \mathcal{L}_{R\text{-NLL}}$.

Loss on global body rotation. We predict deterministic estimates of the global body rotation vectors $\hat{\gamma}^n$, which are supervised using ground-truth global rotations γ^n , with loss $\mathcal{L}_{\text{global}} = \sum_{n=1}^N \|\mathbf{R}(\gamma^n) - \mathbf{R}(\hat{\gamma}^n)\|_F^2$. $\mathbf{R}(\gamma) \in SO(3)$ is the rotation matrix corresponding to γ .

2D joints loss on samples. Applying \mathcal{L}_{NLL} alone results in overly uncertain predicted 3D shape and pose distributions (see Section 5.1). To ensure that the predicted distributions match the visual evidence in the input \mathbf{X}^n , we impose a reprojection loss between ground-truth 2D joint coordinates (in the image plane) and predicted 2D joint samples, which are obtained by *differentiably sampling* 3D bodies from the predicted distributions and projecting to 2D using the predicted camera $\mathbf{c}^n = [s^n, t_x^n, t_y^n]$. Ground-truth 2D joints \mathbf{J}_{2D}^n are computed from $\{\{\mathbf{R}_i^n\}_{i=1}^{23}, \beta^n, \gamma^n\}$ during synthetic training data generation (see Section 4).

We adapt the rejection sampler presented in [21] to sample from a matrix-Fisher distribution $\mathcal{M}(R; F)$, modifying it to allow for backpropagation of gradients through the proposal sampling step (lines 5-7 in Algorithm 1). We refer the reader to [21] for further details about the rejection sampler. In short, to simulate a matrix-Fisher distribution with parameter $F = USV^T$, we sample unit quaternions from a Bingham distribution [33] over the unit 3-sphere S^3 , with Bingham parameter A computed from S , and then convert the sampled quaternions into rotation matrices [21, 33] with the desired matrix-Fisher distribution. Rejection sampling is used to sample from the Bingham distribution, which has pdf $p_{\text{Bing}}(\mathbf{x}) \propto \exp(-\mathbf{x}^T A \mathbf{x})$ for $\mathbf{x} \in S^3$. The proposal distribution for the rejection sampler is an angular central Gaussian (ACG) distribution, with pdf $p_{\text{ACG}}(\mathbf{x}) \propto (\mathbf{x}^T \Omega \mathbf{x})^{-2}$. The ACG distribution is easily simulated [21] by sampling from a zero-mean Gaussian distribution with covariance matrix Ω^{-1} and normalising to unit-length (lines 5-7 in Algorithm 1). The re-parameterisation trick [24] is used to differentially sample from this zero-mean Gaussian, thus allowing for backpropagation of gradients through the rejection sampler.

Algorithm 1 samples K sets of relative 3D joint rotation matrices $\{\{\hat{\mathbf{R}}_{i,k}^n\}_{i=1}^{23}\}_{k=1}^K$ from the corresponding distributions $\{\mathcal{M}(\mathbf{R}_i^n; \mathbf{F}_i^n)\}_{i=1}^{23}$. Furthermore, we differentially sample K SMPL shape vectors from the predicted Gaussian distribution $\{\hat{\beta}_k^n \sim \mathcal{N}(\beta; \mu_\beta(\mathbf{X}^n), \text{diag}(\sigma_\beta^2(\mathbf{X}^n)))\}_{k=1}^K$, again using the re-parameterisation trick [24].

The body shape and 3D joint rotation samples are converted into 2D joint samples using the SMPL model and weak-perspective camera parameters

$$\hat{\mathbf{J}}_{2D_k}^n = s^n \Pi(\mathcal{J}\mathcal{S}(\{\hat{\mathbf{R}}_{i,k}^n\}_{i=1}^{23}, \hat{\beta}_k^n, \hat{\gamma}^n)) + [t_x^n, t_y^n] \quad (9)$$

where $\Pi()$ is an orthographic projection. The reprojection loss applied between the predicted 2D joint samples and the visible target 2D joint coordinates is given by

$$\mathcal{L}_{2D \text{ Samples}} = \sum_{n=1}^N \sum_{k=1}^K \|\omega^n (\mathbf{J}_{2D}^n - \hat{\mathbf{J}}_{2D_k}^n)\|_2^2 \quad (10)$$

where the visibilities of the target joints are denoted by $\omega^n \in \{0, 1\}^L$ (1 if visible, 0 otherwise).

4. Implementation Details

Synthetic training data. To train our 3D body shape and pose distribution prediction networks, we require a training dataset $\{\mathbf{X}^n, (\{\mathbf{R}_i^n\}_{i=1}^{23}, \beta^n, \gamma^n)\}_{n=1}^N$. We extend the synthetic training frameworks presented in [46, 47], which involve generating inputs and corresponding SMPL body shape and pose (i.e. 3D joint rotation) labels randomly and on-the-fly during training. In brief, for every training iteration, SMPL shapes β^n are randomly sampled from a

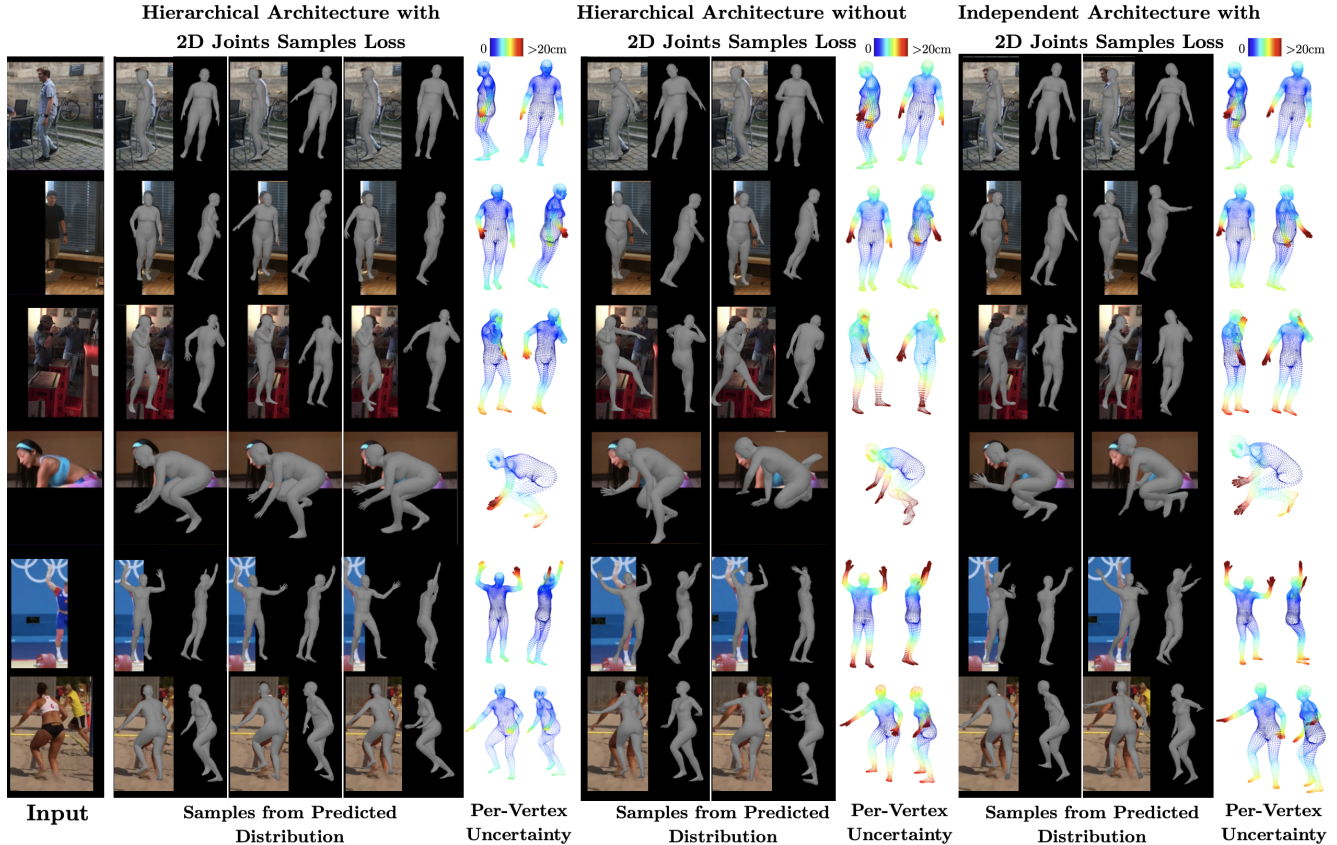


Figure 3. 3D reconstruction samples and per-vertex uncertainty corresponding to shape and pose distributions predicted using the hierarchical architecture with 2D samples loss (left), hierarchical architecture without 2D samples loss (centre) and independent architecture with 2D samples loss (right). Per-vertex uncertainty (in cm) is estimated by sampling 100 SMPL meshes from the predicted distributions and determining the average Euclidean distance from the sample mean for each vertex. Both the hierarchical architecture *and* the sample reprojection loss are required for predicted distributions to match the inputs, while demonstrating greater uncertainty for ambiguous parts.

prior Gaussian distribution while relative 3D joint rotations $\{\mathbf{R}_i^n\}_{i=1}^{23}$ and global rotation γ^n are chosen from the training sets of UP-3D [29], 3DPW [56] or Human3.6M [17]. These are converted into training inputs \mathbf{X}^n and ground-truth 2D joint coordinates \mathbf{J}^n using the SMPL model and a light-weight renderer [42]. Cropping, occlusion and noise augmentations are then applied to the synthetic inputs.

Previous synthetic training frameworks [46, 47, 52] often use silhouette-based training inputs. This necessitates accurate human silhouette segmentation at test-time, which may be challenging to do robustly. In contrast, our input representations consist of edge-images concatenated with 2D joint heatmaps. To generate edge-images, we first create synthetic RGB images by rendering textured SMPL meshes. For each training mesh, clothing textures are randomly chosen from [55, 2]. The textured SMPL mesh is rendered onto a background image (randomly chosen from LSUN [60]), using randomly-sampled lighting and camera parameters. Canny edge detection [6] is used to compute edge-images from the synthetic RGB images. We show in Section 5.1 that, despite the lack of photorealism in the

synthetic RGB images, edge-filtering bridges the synthetic-to-real domain gap at test-time - and performs better than either silhouette-based or synthetic-RGB-based training inputs in our experiments. Examples of synthetic training samples are given in the supplementary material.

Training details. We use Adam [23] with a learning rate of 0.0001, batch size of 80 and train for 150 epochs. For stability, the 2D joints reprojection loss is only applied on the mode pose and shape (projected to 2D) in the first 50 epochs and not on the samples, which are supervised in the next 100 epochs. To boost 3D pose metrics, an MSE loss on the mode 3D joint locations is applied in the final 50 epochs.

Evaluation datasets. 3DPW [56] is used to evaluate 3D pose prediction accuracy. We report mean-per-joint-position-error after scale correction (MPJPE-SC) [46] and after Procrustes analysis (MPJPE-PA), both in mm. Both metrics are computed using the mode 3D joint coordinates of the predicted shape and pose distributions.

SSP-3D is primarily used to evaluate 3D body shape prediction accuracy, using per-vertex Euclidean error in a T-pose after scale-correction (PVE-T-SC) [46] in mm, com-

| Input Type | Architecture | 2D Samples Loss | Synthetic Test Data | | | SSP-3D | | 3DPW |
|-----------------|--------------|-----------------|---------------------|-------------|------------------|--------------|------------------|-------------|
| | | | MPJPE-SC | PVE-T-SC | 2D Joint Err. | PVE-T-SC | 2D Joint Err. | MPJPE-SC |
| | | | Mode/Samples | | | Mode/Samples | | |
| Silh. + J2DHmap | Independent | No | 84.9 | 12.8 | 7.2 / 11.6 | 14.3 | 6.0 / 11.9 | 93.0 |
| RGB + J2DHmap | Independent | No | 79.9 | 11.3 | 7.1 / 11.7 | 14.0 | 5.9 / 12.0 | 92.8 |
| Edge + J2DHmap | Independent | No | 85.8 | 12.9 | 7.5 / 12.0 | 13.7 | 5.9 / 11.8 | 88.4 |
| Edge + J2DHmap | Independent | Yes | 86.3 | 13.2 | 7.6 / 8.9 | 13.9 | 6.2 / 9.6 | 91.3 |
| Edge + J2DHmap | Hierarchical | No | 84.4 | 12.8 | 7.3 / 10.4 | 13.6 | 5.3 / 11.2 | 87.7 |
| Edge + J2DHmap | Hierarchical | Yes | 79.1 | 12.6 | 6.7 / 6.9 | 13.6 | 4.8 / 6.9 | 84.7 |

Table 1. Experiments investigating different input representations, hierarchical versus independent distribution prediction networks and the 2D samples reprojection loss, evaluated in terms of shape and pose prediction metrics on synthetic data, SSP-3D [46] and 3DPW [56].

puted with the mode 3D body shape from the predicted shape distribution. We also evaluate 2D joint prediction error (2D Joint Err. Mode/Samples) in pixels, computed using both the mode 3D body and 10 3D bodies randomly sampled from the predicted shape and pose distributions, projected onto the image plane using the camera prediction. 2D joint error is evaluated on *visible* target 2D joints only.

Finally, we use a synthetic test dataset for our ablation studies investigating different input representations. It consists of 1000 synthetic input-label pairs, generated in the same way as the synthetic training data, with poses sampled from the test set of Human3.6M. [17].

5. Experimental Results

This section investigates different input representations and the benefits of the 2D joints samples loss, compares independent and hierarchical distribution predictors and benchmarks our method against the state-of-the-art.

5.1. Ablation studies

Input proxy representation. Rows 1-3 in Table 1 compare different choices of input proxy representation: binary silhouettes, RGB images and edge-filtered images (each additionally concatenated with 2D joint heatmaps). The independent network architecture is used for all three input types. To investigate the synthetic-to-real domain gap, metrics are presented for synthetic test data, as well as real test images from SSP-3D and 3DPW. For the latter, silhouette segmentation is carried out with DensePose [15]. Using RGB-based input representations (row 2) results in the best 3D shape and pose metrics on synthetic data, which is reasonable since RGB contains more information than both silhouettes and edge-filtered images. However, metrics are significantly worse on real datasets, suggesting that the network has over-fitted to unrealistic artefacts present in synthetic RGB. Silhouette-based input representations (row 1) also demonstrate a deterioration of 3D metrics on real test data compared to synthetic data, since they are heavily reliant upon accurate silhouettes, which are difficult to robustly segment in test images containing challenging poses or severe occlusions. Inaccurate silhouette segmentations

critically impair the network’s ability to predict 3D body pose and shape. In contrast, edge-filtering is a simpler and more robust operation than segmentation, but is still able to retain important shape information from the RGB image. Thus, edge-images (concatenated with 2D joint heatmaps) can better bridge the synthetic-to-real domain gap, resulting in improved metrics on real test inputs (row 3).

Hierarchical architecture and reprojection loss on 2D joints samples. Figure 3 and rows 3-6 in Table 1 compare the independent and hierarchical distribution prediction architectures (f_{indep} and f_{hier}) presented in Section 3.4, both with and without the reprojection loss on sampled 2D joints ($\mathcal{L}_{2D \text{ Samples}}$) from Section 3.5. When $\mathcal{L}_{2D \text{ Samples}}$ is not applied, the shape and pose distributions predicted by both the independent and hierarchical network architectures do not consistently match the the input image, as evidenced by the significant gap between the visible 2D joint error computed using the distributions’ modes versus samples drawn from the distributions (in rows 3 and 5 of Table 1) on both synthetic test data and SSP-3D [46]. This implies that the predicted distributions are overly uncertain about parts of the subject’s body that are visible and unambiguous in the input image. The visualisations corresponding to the hierarchical architecture trained without $\mathcal{L}_{2D \text{ Samples}}$ in Figure 3 (centre) further demonstrate that the predicted samples often do not match the input image, particularly at the extreme ends of the body. This results in significant undesirable per-vertex uncertainty over unambiguous body parts.

Applying $\mathcal{L}_{2D \text{ Samples}}$ to the independent network f_{indep} partially alleviates the mismatch between inputs and predicted samples, as shown by Figure 3 (right) and row 4 in Table 1, where the mode versus sample 2D joint error gap has reduced. However, training with $\mathcal{L}_{2D \text{ Samples}}$ deteriorates the independent architecture’s mode pose prediction metrics (MPJPE-SC and 2D Joint Err. Mode in row 3 vs 4 of Table 1) on both synthetic and real test data. This is because f_{indep} naively models each joint’s relative rotation independently of its parents’ rotations (Equation 5); however, to predict realistic human pose samples that match the visible input, each joint’s rotation distribution *must* be informed by its parents. $\mathcal{L}_{2D \text{ Samples}}$ attempts to force predicted samples to match the input despite this logical inconsistency, which

| Method | 3DPW | |
|-----------------------------|-------------|-------------|
| | MPJPE-SC | MPJPE-PA |
| HMR [20] | 102.8 | 71.5 |
| GraphCMR [27] | 102.0 | 70.2 |
| SPIN [26] | 89.4 | 59.0 |
| I2L-MeshNet [35] | 80.7 | 57.7 |
| Biggs <i>et al.</i> [3] | - | 55.6 |
| DaNet [63] | 82.4 | 54.8 |
| HMR (unpaired) [20] | 126.3 | 92.0 |
| Kundu <i>et al.</i> [28] | - | 89.8 |
| STRAPS [46] | 99.0 | 66.8 |
| Sengupta <i>et al.</i> [47] | 90.9 | 61.0 |
| Ours | 84.7 | 59.2 |

Table 2. Comparison with SOTA in terms of MPJPE-SC and MPJPE-PA (both mm) on 3DPW [56]. Methods in the top half require training images paired with 3D ground-truth, methods in the bottom half do not.

causes a trade-off between mode and sample pose prediction metrics, particularly worsening MPJPE-SC.

In contrast, applying $\mathcal{L}_{2D \text{ Samples}}$ to the hierarchical network f_{hier} improves metrics corresponding to both mode and sample predictions, as shown by row 6 in Table 1. Now, each SMPL joint’s relative rotation distribution is conditioned on all its parents’ distributions (Equation 6). Thus, $\mathcal{L}_{2D \text{ Samples}}$ and \mathcal{L}_{NLL} work in conjunction in enabling predicted hierarchical distributions (and samples) to match the visible input, while yielding improved 3D metrics. Figure 3 (left) exhibits such visually-consistent samples and demonstrates greater prediction uncertainty for ambiguous parts. Note that uncertainty can arise even without occlusion in a monocular setting, e.g. due to depth ambiguities [49, 51] as shown by the left arm samples in the last row of Figure 3. Further visual results are in the supplementary material.

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

Shape prediction. Table 3 evaluates 3D body shape metrics on SSP-3D [46] for single image inputs and multi-image input sets, which we evaluate using both mean and probabilistic combination methods from [47]. Our network surpasses the state-of-the-art [47], mainly due to our use of an edge-based proxy representation, instead of the silhouette-based representations used in [46] and [47]. These methods rely on accurate human silhouettes, which may be difficult to compute at test-time, as discussed in Section 5.1, while our method does not have such dependencies. However, our method may result in erroneous shape predictions when the subject is wearing loose clothing which obscures body shape, in which case the shape prediction over-estimates the subject’s true proportions (see rows 1-2 in Figure 3).

Pose prediction. Table 2 evaluates 3D pose metrics on 3DPW [56]. Our method is competitive with the state-of-the-art and surpasses other methods that do not require 3D-labelled training images [46, 47, 28, 20]. Figure 4(a) shows

| Max. input set size | Method | SSP-3D |
|---------------------|---|-------------|
| | | PVE-T-SC |
| 1 | HMR [20] | 22.9 |
| | GraphCMR [27] | 19.5 |
| | SPIN [26] | 22.2 |
| | DaNet [63] | 22.1 |
| | STRAPS [46] | 15.9 |
| | Sengupta <i>et al.</i> [47] | 15.2 |
| | Ours | 13.6 |
| 5 | HMR [20] + Mean | 22.9 |
| | GraphCMR [27] + Mean | 19.3 |
| | SPIN [26] + Mean | 21.9 |
| | DaNet [63] + Mean | 22.1 |
| | STRAPS [46] + Mean | 14.4 |
| | Sengupta <i>et al.</i> [47] + Mean | 13.6 |
| | Sengupta <i>et al.</i> [47] + Prob. Comb. | 13.3 |
| | Ours + Mean | 12.2 |
| | Ours + Prob. Comb. | 12.0 |

Table 3. Comparison with SOTA in terms of PVE-T-SC (mm) on SSP-3D [46]. Top half: single-input, bottom half: multi-input.

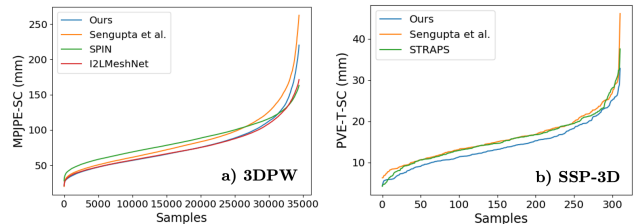


Figure 4. Comparison with SOTA using sorted per-sample distributions of a) MPJPE-SC on 3DPW and b) PVE-T-SC on SSP-3D.

that our method performs well for most test examples in 3DPW, even matching pose-focused approaches that do not attempt to accurately predict diverse body shapes [35, 26]. However, some images in 3DPW contain significant occlusion, which can lead to noisy 2D joint heatmaps in the proxy representations, resulting in poor 3D pose metrics as shown by the right end of the curve in Figure 4(a).

6. Conclusion

In this paper, we have proposed a probabilistic approach to the ill-posed problem of monocular 3D human shape and pose estimation, motivated by the fact that multiple 3D bodies may explain a given 2D image. Our method predicts a novel hierarchical matrix-Fisher distribution over relative 3D joint rotations and a Gaussian distribution over SMPL [32] shape parameters, from which we can sample any number of plausible 3D reconstructions. To ensure that the predicted distributions match the input image, we have implemented a differentiable rejection sampler to impose a loss between predicted 2D joint samples and ground-truth 2D joint coordinates. Our method is competitive with the state-of-the-art in terms of pose metrics on 3DPW, while surpassing the state-of-the-art for shape accuracy on SSP-3D.

References

- [1] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. SCAPE: Shape completion and animation of people. In *ACM Transactions on Graphics (TOG) - Proceedings of SIGGRAPH*, volume 24, pages 408–416, 2005. [2](#)
- [2] Bharat Lal Bhatnagar, Garvita Tiwari, Christian Theobalt, and Gerard Pons-Moll. Multi-garment net: Learning to dress 3D people from images. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2019. [6](#)
- [3] 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. [2](#), [8](#)
- [4] Christopher M. Bishop. Mixture density networks. Technical report, 1994. [2](#)
- [5] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J. Black. Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image. In *Proceedings of the European Conference on Computer Vision (ECCV)*, Oct. 2016. [2](#)
- [6] John F. Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 8(6):679–698, 1986. [4](#), [6](#)
- [7] 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. [2](#)
- [8] 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](#)
- [9] 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](#)
- [10] J. Deutscher, A. Blake, and I. Reid. Articulated body motion capture by annealed particle filtering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2000. [2](#)
- [11] Thomas D. Downs. Orientation statistics. *Biometrika*, 59(3):665–676, 12 1972. [1](#), [3](#)
- [12] 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. [1](#), [2](#), [4](#)
- [13] Igor Gilitschenski, Roshni Sahoo, Wilko Schwarting, Alexander Amini, Sertac Karaman, and Daniela Rus. Deep orientation uncertainty learning based on a bingham loss. In *International Conference on Learning Representations*, 2020. [3](#)
- [14] Riza Alp Güler 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. [1](#), [2](#)
- [15] Riza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [2](#), [7](#)
- [16] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. [4](#)
- [17] 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. [2](#), [6](#), [7](#)
- [18] Ehsan Jahangiri and Alan L. Yuille. Generating multiple diverse hypotheses for human 3D pose consistent with 2D joint detections. In *IEEE International Conference on Computer Vision (ICCV) Workshops (PeopleCap)*, 2017. [2](#)
- [19] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total capture: A 3D deformation model for tracking faces, hands, and bodies. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. [2](#)
- [20] 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](#), [4](#), [8](#)
- [21] John T. Kent, Asaad M. Ganeiber, and Kanti V. Mardia. A new method to simulate the Bingham and related distributions in directional data analysis with applications, 2013. [5](#)
- [22] C. G. Khatri and K. V. Mardia. The Von Mises-Fisher matrix distribution in orientation statistics. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):95–106, 1977. [1](#), [3](#)
- [23] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2014. [6](#)
- [24] Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2014. [5](#)
- [25] Alexander Kirillov, Yuxin Wu, Kaiming He, and Ross Girshick. PointRend: Image segmentation as rendering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [2](#)
- [26] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3D human pose and shape via model-fitting in the loop. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019. [1](#), [2](#), [8](#)
- [27] 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. [1](#), [2](#), [8](#)
- [28] Jogendra Nath Kundu, Mugalodi Rakesh, Varun Jampani, Rahul M Venkatesh, and R. Venkatesh Babu. Appearance consensus driven self-supervised human mesh recovery. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. [8](#)
- [29] Christoph Lassner, Javier Romero, Martin Kiefel, Federica Bogo, Michael J. Black, and Peter V. Gehler. Unite the Peo-

- ple: 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
- [30] T. Lee. Bayesian attitude estimation with the matrix fisher distribution on $so(3)$. *IEEE Transactions on Automatic Control*, 63(10):3377–3392, 2018. 3, 4
- [31] Chen Li and Gim Hee Lee. Generating multiple hypotheses for 3D human pose estimation with mixture density network. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [32] 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. 1, 2, 3, 8
- [33] K. V. Mardia and P. E. Jupp. *Directional statistics*. Wiley, 2000. 1, 3, 5
- [34] David Mohlin, Josephine Sullivan, and Gérald Bianchi. Probabilistic orientation estimation with matrix fisher distributions. In *Advances in Neural Information Processing Systems*, volume 33, 2020. 1, 3, 5
- [35] 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. 1, 2, 8
- [36] Tuomas P. Oikarinen, Daniel C. Hannah, and Sohrob Kazerounian. GraphMDN: Leveraging graph structure and deep learning to solve inverse problems. *CoRR*, abs/2010.13668, 2020. 2
- [37] Mohamed Omran, Christoph Lassner, Gerard Pons-Moll, Peter V. Gehler, and Bernt Schiele. Neural body fitting: Unifying deep learning and model-based human pose and shape estimation. In *Proceedings of the International Conference on 3D Vision (3DV)*, 2018. 1, 2
- [38] 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
- [39] Georgios Pavlakos, Nikos Kolotouros, and Kostas Daniilidis. Texturepose: Supervising human mesh estimation with texture consistency. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019. 1, 2
- [40] Georgios Pavlakos, Luyang Zhu, Xiaowei Zhou, and Kostas Daniilidis. Learning to estimate 3D human pose and shape from a single color image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2, 4
- [41] Sergey Prokudin, Peter Gehler, and Sebastian Nowozin. Deep directional statistics: Pose estimation with uncertainty quantification. In *Proceedings of the European Conference on Computer Vision (ECCV)*, Sept. 2018. 2
- [42] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3D deep learning with PyTorch3D. *arXiv:2007.08501*, 2020. 6
- [43] 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
- [44] 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
- [45] 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
- [46] 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. 1, 2, 4, 5, 6, 7, 8
- [47] 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. 2, 4, 5, 6, 8
- [48] C. Sminchisescu, A. Kanaujia, Zhiguo Li, and D. Metaxas. Discriminative density propagation for 3D human motion estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 390–397 vol. 1, 2005. 2
- [49] 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, 8
- [50] Cristian Sminchisescu and Bill Trigg. Hyperdynamics importance sampling. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2002. 2
- [51] 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. 2, 8
- [52] B. M. Smith, V. Chari, A. Agrawal, J. M. Rehg, and R. Sever. Towards accurate 3D human body reconstruction from silhouettes. In *Proceedings of the International Conference on 3D Vision (3DV)*, 2019. 1, 2, 6
- [53] 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. 1, 2
- [54] 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. 1, 2
- [55] Gül Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J. Black, Ivan Laptev, and Cordelia Schmid.

- Learning from synthetic humans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2, 6
- [56] 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
- [57] 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
- [58] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. 4
- [59] Yuanlu Xu, Song-Chun Zhu, and Tony Tung. DenseRaC: Joint 3D pose and shape estimation by dense render-and-compare. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019. 2
- [60] Fisher Yu, Yinda Zhang, Shuran Song, Ari Seff, and Jianxiong Xiao. LSUN: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015. 6
- [61] 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
- [62] Wang Zeng, Wanli Ouyang, Ping Luo, Wentao Liu, and Xiaogang Wang. 3d human mesh regression with dense correspondence. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2
- [63] 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. 1, 2, 8