

Body Knowledge and Uncertainty Modeling for Monocular 3D Human Body Reconstruction

Yufei Zhang¹, Hanjing Wang¹, Jeffrey O. Kephart², Qiang Ji¹

¹Rensselaer Polytechnic Institute, ²IBM Research

{zhangy76, wangh36, jiq}@rpi.edu, kephart@us.ibm.com

Abstract

While 3D body reconstruction methods have made remarkable progress recently, it remains difficult to acquire the sufficiently accurate and numerous 3D supervisions required for training. In this paper, we propose **KNOWN**, a framework that effectively utilizes body **KNOW**ledge and **uN**certainty modeling to compensate for insufficient 3D supervisions. **KNOWN** exploits a comprehensive set of generic body constraints derived from well-established body knowledge. These generic constraints precisely and explicitly characterize the reconstruction plausibility and enable 3D reconstruction models to be trained without any 3D data. Moreover, existing methods typically use images from multiple datasets during training, which can result in data noise (e.g., inconsistent joint annotation) and data imbalance (e.g., minority images representing unusual poses or captured from challenging camera views). **KNOWN** solves these problems through a novel probabilistic framework that models both aleatoric and epistemic uncertainty. Aleatoric uncertainty is encoded in a robust Negative Log-Likelihood (NLL) training loss, while epistemic uncertainty is used to guide model refinement. Experiments demonstrate that **KNOWN**'s body reconstruction outperforms prior weakly-supervised approaches, particularly on the challenging minority images.

1. Introduction

Recovering 3D human body configurations from monocular RGB images has broad applications in robotics, human-computer interaction, and human behaviour analysis. This is a challenging task, as the inherent depth ambiguity under-constrains the reconstruction problem. To alleviate the high dimensionality of the reconstruction space, deformable body models such as SMPL [54] or GHUM [82] have been used widely to represent a 3D human body in terms of body pose and shape parameters. Building upon these parametric models with the aid of deep learning,

model-based methods have made promising progress on existing benchmarks [60, 38, 44, 35, 15, 49, 51, 45, 43]. However, obtaining 3D annotations for training deep models often requires a Motion Capture (MoCap) system [73], which is cumbersome, expensive, and limited to specific environments and subjects. Moreover, labeling noise can be introduced when obtaining the 3D annotations by fitting parametric models to sparse 3D MoCap markers [53, 55] or generating pseudo-labels using existing reconstruction models [37, 57, 50]. Thus it is important to reduce the dependency on 3D annotations to the greatest possible extent.

One approach is to employ weak supervisions, such as utilizing 2D landmark annotations combined with prior knowledge of constraints on human body pose and shape that are extracted from data [9, 38, 59, 29, 90, 94, 17, 27, 85]. While this method avoids the demand for 3D annotations, this problem is replaced by another: it is challenging to collect a sufficiently homogeneous and dense data set from which to extract a precise prior [32]. Moreover, such data-driven priors do not explicitly capture reconstruction plausibility, as they are only modeled via a black-box distribution. In this work, we avoid data-driven priors by exploiting generic body constraints derived from well-established studies on human body structure and movements [28, 58, 80] to effectively utilize 2D annotations for 3D body reconstruction.

Another problematic aspect of existing approaches is that they typically combine images from multiple datasets for training, neglecting the subsequent problem of data noise and data imbalance. In particular, data noise can stem from inconsistent joint definitions across different datasets [62] and the presence of low quality images [43, 84]. Data imbalance occurs because 3D datasets are rich in images from videos but suffer from diversity in subjects and actions, whereas 2D datasets are diverse in subjects, poses, and backgrounds but usually suffer from having few images per subject-pose-background combination. **KNOWN** deals with the problems of data noise and data imbalance by modeling uncertainty. Specifically, **KNOWN** employs a novel probabilistic framework that captures both aleatoric uncer-

tainty (also called data uncertainty), which measures the inherent data noise, and epistemic uncertainty (also called model uncertainty), which reflects the lack of knowledge due to limited data [40, 47, 41]. KNOWN is trained using NLL, which achieves robustness to data noise by adaptively assigning weights based on the captured uncertainty. Because previous methods treat all the data samples equally, their performance suffers on the minority images (those that are distinct from the training data and have low data density). KNOWN uses an uncertainty-guided refinement strategy to improve model performance, particularly on these minorities.

In summary, our main contributions include:

- A systematic study of body knowledge and its encoding as generic constraints that allow training 3D body reconstruction models without requiring any 3D data;
- The first 3D body reconstruction model that accounts for both aleatoric and epistemic uncertainty, thereby supporting both data noise and data density characterization and model performance improvement;
- Extensive experiments that demonstrates improved reconstruction accuracy over existing works. In particular, KNOWN outperforms fully-supervised 3D reconstruction models on challenging minority test images.

2. Related Work

2.1. 3D Body Reconstruction with Prior Knowledge

Prior knowledge can be extracted from data (data-driven prior) or derived from established body knowledge (generic body prior). Here we review how the existing works encode and leverage these two types of priors.

Data-driven priors are typically learned from MoCap data [1, 2, 55]. Some optimization-based methods encode pose feasibility via nonlinear inequality constraints [2], a Gaussian Mixture Model [9, 4], or a Variational Auto-Encoder [59]. They incorporate these learned priors as constraints or penalties to avoid infeasible joint angle estimation. Additionally, average bone length of training data is exploited to ensure body scale validity [5, 95, 61, 75]. Among learning-based methods, adversarial framework [34, 85, 74, 11, 46, 93, 23, 86] or normalizing flow [88, 6, 77, 45, 89, 91] is employed to encode body pose feasibility. These models are utilized in building prior body models for constrained parameter estimation or to regularize model training. In particular, HMR [38] learns a pose and shape prior from MoCap data, which facilitates subsequent training on images that possess only 2D keypoint annotations. These data-driven priors, however, assume that sufficient 3D data be available. In contrast, KNOWN does not require any 3D data or annotations.

Generic body priors recognize that human body structure and pose comply with basic functionality principles that generalize across different people and activities [28, 58, 80]. Some works design a neural network architecture to encode body knowledge [48, 18, 16, 25, 15, 14, 92, 83, 87, 49], such as predicting joint angles from the torso outward based on body kinematics [22]. These models still require either 3D supervisions or pose priors for training. Moreover, some authors [78, 17] introduce bone symmetry loss to reduce the dependency on 3D annotations. We supplement this body anatomy constraint by adding statistics of different bones from an anthropometric study [28] and by introducing special geometry characteristics of human body joints. Other authors [12, 26, 67] utilize body biomechanics by imposing constraints on individual joint rotations. We supplement these constraints by adding inter-joint dependencies based on body functional anatomy [28]. Additionally, Proxy Geometries [9, 90] and Signed Distance Field (SDF) [59, 39, 70] are used in optimization-based frameworks to guard against a key body physics constraint: body parts should not inter-penetrate. We introduce a SDF-based loss [59] into our learning framework to avoid collisions. To summarize, existing works consider body knowledge partially, while we combine constraints comprehensively based upon body anatomy, biomechanics, and physics and properly encode them into a probabilistic model that produces valid 3D predictions without requiring any 3D data.

2.2. 3D Body Reconstruction with Uncertainty

Although uncertainty modeling has shown benefits in various computer vision tasks [41, 76, 40], such as robust 3D face model fitting via incorporating 2D landmark prediction uncertainty [81], it has not been well explored in 3D body reconstruction. For 3D body reconstruction, some methods output a probability distribution to generate multiple hypotheses to account for the depth ambiguity [45, 77, 6]. However, they do not explicitly quantify uncertainty, nor do they resolve it into aleatoric and epistemic components. Other methods quantify aleatoric uncertainty and relate it with image occlusion [65, 64]. We are unaware of any existing works that quantify epistemic uncertainty. Furthermore, there is a noticeable lack of research on effectively leveraging uncertainty for model improvements.

To our knowledge, KNOWN is the first 3D body reconstruction model that captures both aleatoric and epistemic uncertainties. It does so using a single two-stage probabilistic neural network, the uncertainty quantification efficiency of which compares favorably with standard Bayesian [21, 13, 30, 7] and non-Bayesian [47, 79, 20, 19, 72] methods. Moreover, KNOWN is unique in that it leverages epistemic uncertainty to improve reconstruction accuracy and (unlike existing probabilistic 3D body reconstruction models) it does not require 3D data.

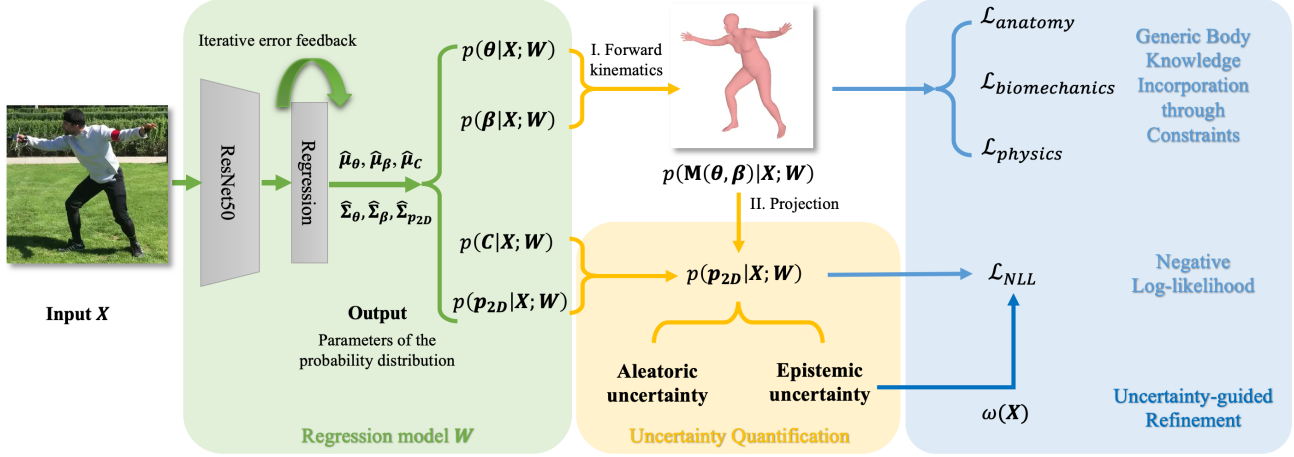


Figure 1. **Method overview.** KNOWN employs a two-stage probabilistic neural network for 3D body reconstruction from 2D images. Given an input image, the first stage models the conditional probability distribution of the body pose, shape, and camera parameters, while the second stage captures the conditional distribution of the corresponding image projection.

3. Method

An overview of KNOWN is illustrated in Fig. 1. As will be elaborated in Sec. 3.1, KNOWN is a two-stage probabilistic neural network that captures both aleatoric and epistemic uncertainty. It is trained by combining the generic body constraints (Sec. 3.2) with a new NLL loss (Sec. 3.3). KNOWN’s final step (uncertainty-guided refinement) and the total training loss are described in Sec. 3.4.

To represent a 3D human body, we use SMPL [54], which includes the pose parameters $\theta \in \mathbb{R}^{23 \times 6}$ and the shape parameters $\beta \in \mathbb{R}^{10}$ to characterize the rotation of 23 body joints and body shape variation, respectively. To establish 3D-2D consistency, we employ a weak-perspective projection model with parameters $\mathbf{C} = [s, \mathbf{R}, \mathbf{t}]$, where $\mathbf{t} \in \mathbb{R}^2$, $s \in \mathbb{R}$, and $\mathbf{R} \in \mathbb{R}^6$ denotes the global translation, scale factor, and camera rotation, respectively. The projection of 3D body keypoints is $\mathbf{p}_{2D} = Proj(\mathbf{P}_{3D}, \beta, \mathbf{C})$, where $Proj(\cdot)$ denotes the weak-perspective projection function. We use a 6D representation for the body pose and camera rotation to avoid the singularity problem [96].

3.1. Two-stage Probabilistic Regression Model

As shown in Fig. 1, Stage I models the distributions of the body model and camera parameters given an input image, while stage II models the distributions of the corresponding 2D projections of body keypoints. As Gaussian distributions have been widely used to efficiently and effectively model continuous data, we assume Gaussian parameters. To specify the modeled probability distributions, we employ a ResNet50 [31] as a backbone to extract image features plus a regression network with iterative error feedback [10] to predict the distribution parameters. Details of

the modeled distributions are introduced below.

Stage I: The distributions of the body model and camera parameters dependent upon an input image \mathbf{X} and neural network weights \mathbf{W} are:

$$p(\theta|\mathbf{X}; \mathbf{W}) = \mathcal{N}(\mu_\theta(\mathbf{X}; \mathbf{W}), \Sigma_\theta(\mathbf{X}; \mathbf{W})), \quad (1)$$

$$p(\beta|\mathbf{X}; \mathbf{W}) = \mathcal{N}(\mu_\beta(\mathbf{X}; \mathbf{W}), \Sigma_\beta(\mathbf{X}; \mathbf{W})), \quad (2)$$

$$p(\mathbf{C}|\mathbf{X}; \mathbf{W}) = \mathcal{N}(\mu_C(\mathbf{X}; \mathbf{W}), \Sigma_C(\mathbf{X}; \mathbf{W})), \quad (3)$$

where $\mu_{\theta, \beta, C}$ and $\Sigma_{\theta, \beta, C}$ are mean and covariance matrices. A nonlinear function $\mathbf{M}(\theta, \beta) \in \mathbb{R}^{6890 \times 3}$ that represents a forward kinematic process can be applied to obtain the 3D vertex positions \mathbf{M} . Then the 3D body joint positions are computed as a linear combination of the vertex positions via $\mathbf{P}_{3D}(\theta, \beta) = \mathbf{H}\mathbf{M}(\theta, \beta)$, where J denotes the number of body joints, and $\mathbf{H} \in \mathbb{R}^{J \times 6890}$ is a predefined joint regressor.

Stage II: The distribution of the corresponding projection of 3D body keypoints is assumed to be

$$p(\mathbf{p}_{2D}|\mathbf{Y}, \mathbf{C}; \mathbf{X}, \mathbf{W}) = \mathcal{N}(\mu_{\mathbf{p}_{2D}}(\mathbf{Y}, \mathbf{C}), \Sigma_{\mathbf{p}_{2D}}(\mathbf{X}; \mathbf{W})), \quad (4)$$

where $\mathbf{Y} = [\theta, \beta]$, $\mu_{\mathbf{p}_{2D}}(\mathbf{Y}, \mathbf{C})$ is the mean that is computed from 3D body keypoints $\mathbf{P}_{3D}(\mathbf{Y})$ and camera parameters \mathbf{C} using the projection function, and $\Sigma_{\mathbf{p}_{2D}}(\mathbf{X}; \mathbf{W})$ is the covariance matrix that is directly estimated by the neural network. Given the definition above, we obtain the conditional probability distribution of \mathbf{p}_{2D} as

$$p(\mathbf{p}_{2D}|\mathbf{X}; \mathbf{W}) = \iint p(\mathbf{p}_{2D}|\mathbf{Y}, \mathbf{C}; \mathbf{X}, \mathbf{W})p(\mathbf{Y}, \mathbf{C}|\mathbf{X}; \mathbf{W})d\mathbf{Y}d\mathbf{C}. \quad (5)$$

Assuming that the body pose θ , shape β , and camera parameters \mathbf{C} are independent, $p(\mathbf{Y}, \mathbf{C}|\mathbf{X}; \mathbf{W}) = p(\theta|\mathbf{X}; \mathbf{W})p(\beta|\mathbf{X}; \mathbf{W})p(\mathbf{C}|\mathbf{X}; \mathbf{W})$.

Inference and uncertainty quantification: During inference, the final 3D human body is recovered via the estimated mean pose and shape, which are the mode of the respective distributions (Eq. (1-2)). Then, following the typical uncertainty quantification strategy [40, 71], we compute the covariance matrix of the keypoint projection to quantify the aleatoric and epistemic uncertainties:

$$\underbrace{\text{Cov}_{p(\mathbf{p}_{2D}|\mathbf{X};\mathbf{W})}[\mathbf{p}_{2D}]}_{\text{Total uncertainty}} = \underbrace{\text{E}_{p(\mathbf{Y},\mathbf{C}|\mathbf{X};\mathbf{W})}[\text{Cov}_{p(\mathbf{p}_{2D}|\mathbf{Y},\mathbf{C};\mathbf{X},\mathbf{W})}[\mathbf{p}_{2D}]]}_{\text{Aleatoric uncertainty}} + \underbrace{\text{Cov}_{p(\mathbf{Y},\mathbf{C}|\mathbf{X};\mathbf{W})}[\text{E}_{p(\mathbf{p}_{2D}|\mathbf{Y},\mathbf{C};\mathbf{X},\mathbf{W})}[\mathbf{p}_{2D}]]}_{\text{Epistemic uncertainty}}. \quad (6)$$

Based on Eq. (4), the aleatoric uncertainty is

$$\begin{aligned} & \underbrace{\text{E}_{p(\mathbf{Y},\mathbf{C}|\mathbf{X};\mathbf{W})}[\text{Cov}_{p(\mathbf{p}_{2D}|\mathbf{Y},\mathbf{C};\mathbf{X},\mathbf{W})}[\mathbf{p}_{2D}]]}_{\text{Aleatoric uncertainty}} \\ &= \text{E}_{p(\mathbf{Y},\mathbf{C}|\mathbf{X};\mathbf{W})}[\boldsymbol{\Sigma}_{\mathbf{p}_{2D}}(\mathbf{X};\mathbf{W})] \\ &= \boldsymbol{\Sigma}_{\mathbf{p}_{2D}}(\mathbf{X};\mathbf{W}). \end{aligned} \quad (7)$$

The epistemic uncertainty is

$$\begin{aligned} & \underbrace{\text{Cov}_{p(\mathbf{Y},\mathbf{C}|\mathbf{X};\mathbf{W})}[\text{E}_{p(\mathbf{p}_{2D}|\mathbf{Y},\mathbf{C};\mathbf{X},\mathbf{W})}[\mathbf{p}_{2D}]]}_{\text{Epistemic uncertainty}} \\ &= \text{Cov}_{p(\mathbf{Y},\mathbf{C}|\mathbf{X};\mathbf{W})}[\boldsymbol{\mu}_{\mathbf{p}_{2D}}(\mathbf{Y},\mathbf{C})]. \end{aligned} \quad (8)$$

The aleatoric uncertainty equals to the predicted variance as derived in Eq. (7), while computing the epistemic uncertainty in Eq. (8) is difficult because $\boldsymbol{\mu}_{\mathbf{p}_{2D}}(\mathbf{Y},\mathbf{C})$ is a nonlinear function of \mathbf{Y} and \mathbf{C} (the forward kinematic and projection functions are nonlinear). We hence use samples to approximate the covariance matrix. Denote $\{\mathbf{Y}^s\}_{s=1}^S$ and $\{\mathbf{C}^s\}_{s=1}^S$ as the samples drawn from $p(\mathbf{Y}|\mathbf{X};\mathbf{W})$ and $p(\mathbf{C}|\mathbf{X};\mathbf{W})$, respectively. We compute the keypoint projection of each sample and obtain $\{\boldsymbol{\mu}_{\mathbf{p}_{2D}}^s\}_{s=1}^S$. The epistemic uncertainty is computed as

$$\begin{aligned} & \underbrace{\text{Cov}_{p(\mathbf{Y},\mathbf{C}|\mathbf{X};\mathbf{W})}[\text{E}_{p(\mathbf{p}_{2D}|\mathbf{Y},\mathbf{C})}[\mathbf{p}_{2D}]]}_{\text{Epistemic uncertainty}} \\ & \approx \text{Cov}[\{\boldsymbol{\mu}_{\mathbf{p}_{2D}}^s\}_{s=1}^S], \end{aligned} \quad (9)$$

which is sample covariance matrix. To obtain a scalar quantification of uncertainty, we use the trace of the covariance matrix for both the aleatoric and epistemic uncertainty.

3.2. Generic Body Constraints

Based on body anatomy, biomechanics, and physics, we introduce hard and soft generic body constraints on the body pose and shape parameters to characterize 3D reconstruction feasibility and encourage realistic prediction.

Body anatomy studies the structure of the human body. While body symmetry and the connection of body parts by body joints are already encoded via SMPL shape bases,

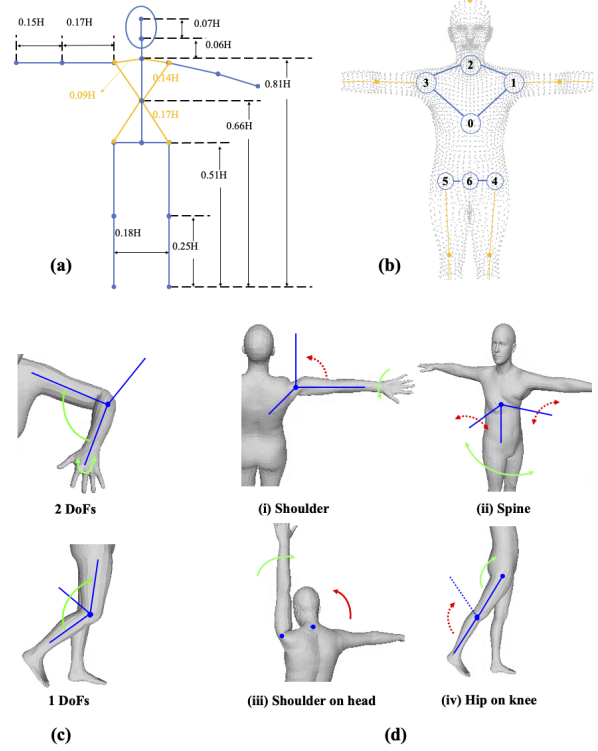


Figure 2. **Illustration of anatomical and biomechanical constraints.** (a) Consistent human body proportions across individuals [80]. (b) Body joints that are collinear and coplanar. (c) Degrees of freedom (DoFs) and movement ranges of body joints. (d) Examples of inter-joint dependencies.

there are several other anatomical principles that can be exploited. Specifically, anthropometry [80] provides a set of 20 constraints on relative body proportions (Fig. 2a), from which we formulate an anthropometric loss term:

$$\mathcal{L}_{anthropometry} = \frac{1}{20} \sum_{i=1}^{20} \left(\frac{\hat{L}_i - L_i}{L_i} \right)^2, \quad (10)$$

where \hat{L} and L are the predicted and anthropometric bone lengths, respectively. Moreover, as shown in Fig. 2b, the body torso is rigid under different actions, the shoulders, neck, and spine joints are coplanar, and the hips and pelvis joints are collinear. We find that these geometry constraints can be satisfied by imposing the anthropometric loss. As is explicated more fully in Appx. A, the geometry constraints (1) ensure realistic 3D reconstruction, especially if the body structure is not constrained; and (2) meaningfully solve the inherent depth ambiguity. Additionally, to explicitly constrain shape estimation, we add an L2 regularization term $\mathcal{L}_{beta-reg}$ on the shape parameters [54], resulting in a total anatomy loss term:

$$\mathcal{L}_{anatomy} = \mathcal{L}_{anthropometry} + \lambda_{beta} \mathcal{L}_{beta-reg}. \quad (11)$$

Body biomechanics, the study of body movement mechanisms, indicates that the degrees of freedom (DoFs) and ranges of body joint are restricted [58]. For example, as illustrated in Fig. 2c, the knee is a 1 DoF joint and the flexion can not exceed 146 degrees. The elbow is a joint with 2 DoFs that exhibits flexion and pronation/supination. The joint angle limits are naturally defined using three Euler angles, while we employ the 6D rotation representation. To impose the constraints, we recover the Euler angles from the output rotation matrix, using octant lookup based on the identified joint ranges of motion to determine the rotation order and thereby ensure a *unique* solution. The loss is formulated as:

$$\mathcal{L}_{biomechanics} = \sum_{i=1}^{69} (\max\{\hat{\varphi}_i - \varphi_{i,max}, \varphi_{i,min} - \hat{\varphi}_i, 0\})^2, \quad (12)$$

where $\{\hat{\varphi}_i\}_{i=1}^{69}$ are the recovered Euler angles of the 23 body joints, $\varphi_{i,max}$ and $\varphi_{i,min}$ represent the joint angle limits obtained from literature [58], and $\max\{\cdot\}$ selects out the value that violates the bound.

From functional anatomy [28], which studies how anatomical structure restricts joint movement, we derive additional biomechanical constraints on inter-joint dependencies. They fall into two classes: (1) those among pose parameters of an individual joint, and (2) those among pose parameters of different joints. We exemplify (1) and (2) in Fig. 2(d,i-ii) and Fig. 2(d,iii-iv), respectively. Specifically, (i) For shoulder joints, rotation along the arm (green arrow) limits upward movement (red arrow). (ii) For the spine joint, rotation along the torso restricts lateral and anterior-posterior rotation. (iii) Raising an arm is accompanied by an anterior translation of the head. (iv) Knee flexion is limited when the thigh is flexed. To impose the inter-joint dependency constraints, we cannot use a constant joint rotation range. Instead, we dynamically update the bounds given the current rotation and use the updated bounds to compute the loss defined in Eq. (12). For the case of Fig. 2(d,i), denote the maximal joint angle for the shoulder joint to move up as $\varphi_{shoulder,up,max}$ and the predicted rotation angle as $\hat{\varphi}_{shoulder}$. The bound is updated according to

$$\begin{aligned} \hat{\varphi}_{shoulder,up,max} \\ = \varphi_{shoulder,up,max} - \alpha_0 \hat{\varphi}_{shoulder}, \end{aligned} \quad (13)$$

where α_0 is determined from the functional anatomy literature [28]. Other inter-joint dependencies are imposed similarly with different values of α_0 given different underlying dependencies. The biomechanical constraints are hard constraints. Large weights are employed to encourage them to be better satisfied.

Body physics stipulates that different body parts can not penetrate each other. To impose the non-penetration constraints, we follow the approach introduced by [39, 70]. We

first detect colliding triangle pairs [68] and measure the collision through predefined signed distance fields Ψ for each pair of colliding triangles. The loss to avoid penetration is then formulated as

$$\begin{aligned} \mathcal{L}_{physics} = \sum_{(f_s, f_t) \in \mathcal{C}} \left\{ \sum_{\mathbf{v}_s \in f_s} \|\Psi_{f_t}(\mathbf{v}_s) \cdot \mathbf{n}_s\|^2 \right. \\ \left. + \sum_{\mathbf{v}_t \in f_t} \|\Psi_{f_s}(\mathbf{v}_t) \cdot \mathbf{n}_t\|^2 \right\}, \end{aligned} \quad (14)$$

where \mathcal{C} is the set of colliding triangles, f_s and f_t represent intruding and receiving triangles with the corresponding norm \mathbf{n}_s and \mathbf{n}_t , respectively.

The **total generic loss** obtained via assembling the anatomical, biomechanical, and physics loss functions is:

$$\begin{aligned} \mathcal{L}_{generic} = \lambda_1 \mathcal{L}_{anatomy} + \lambda_2 \mathcal{L}_{biomechanics} \\ + \lambda_3 \mathcal{L}_{physics}. \end{aligned} \quad (15)$$

We highlight that these generic constraints (1) apply to *all* subjects and activities, not just those represented in the training set; (2) do not require a laborious 3D data collection process; and (3) are not affected by data noise.

3.3. Negative Log Likelihood for 3D-2D Consistency

To further ensure that 3D prediction is consistent with the input image, we consider a novel training loss based on Negative Log Likelihood (NLL). Specifically, we formulate NLL to align 2D keypoint projections with input images \mathbf{X}_i with 2D keypoint labels $\bar{\mathbf{p}}_{2D,i} \in \mathbb{R}^{J \times 2}$:

$$\begin{aligned} \mathcal{L}_{NLL,i} &= -\log p(\bar{\mathbf{p}}_{2D,i} | \mathbf{X}_i; \mathbf{W}) \\ &= -\log \iint p(\bar{\mathbf{p}}_{2D,i} | \mathbf{Y}, \mathbf{C}) p(\mathbf{Y}, \mathbf{C} | \mathbf{X}_i; \mathbf{W}) d\mathbf{Y} d\mathbf{C} \\ &\approx -\log \frac{1}{S} \sum_{s=1}^S p(\bar{\mathbf{p}}_{2D,i} | \mathbf{Y}_s, \mathbf{C}_s), \end{aligned} \quad (16)$$

where $\{\mathbf{Y}_s, \mathbf{C}_s\}_{s=1}^S$ are the samples drawn from the corresponding distribution. Directly solving the integration in Eq. (16) is intractable, so we approximate the integration through sampling. Meanwhile, we employ a reparameterization trick [42] to efficiently select samples in a differentiable way as

$$\mathbf{Y} = \boldsymbol{\mu}_{\mathbf{Y}}(\mathbf{X}; \mathbf{W}) + \mathbf{L}(\mathbf{X}; \mathbf{W})\boldsymbol{\sigma}, \quad \boldsymbol{\sigma} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (17)$$

where $\boldsymbol{\mu}_{\mathbf{Y}}(\mathbf{X})$ is the mean and \mathbf{I} is an identity matrix with the same dimension of $\text{Cov}[\mathbf{Y}]$, $\mathbf{L}(\mathbf{X}; \mathbf{W})$ is the Cholesky decomposition of $\text{Cov}[\mathbf{Y}]$ ($\text{Cov}[\mathbf{Y}] = \mathbf{L}(\mathbf{X}; \mathbf{W})\mathbf{L}^T(\mathbf{X}; \mathbf{W})$). In this work, we consider a diagonal covariance matrix of the body model parameters and zero covariance of the camera parameters. The mean of $p(\bar{\mathbf{p}}_{2D} | \mathbf{X}; \mathbf{W})$ is therefore computed using $\boldsymbol{\mu}_{\mathbf{C}}$ and the body model parameters $\mathbf{Y} = [\boldsymbol{\theta}, \boldsymbol{\beta}]$ sampled using Eq. (17).

3.4. Uncertainty-guided Model Refinement Loss

Given that the training data span multiple datasets that vary in diversity and size (e.g. some contain limited data on in-the-wild scenarios with challenging poses), KNOWN augments the initial training with a final uncertainty-guided training step based on the epistemic uncertainty quantified via Eq. (9). Note that the epistemic uncertainty generated by the initial training is negatively correlated with training data density, and thus explains where the model may fail due to lack of data [40]. We compute a refinement weight $w(\mathbf{X}_i)$ based on the quantified epistemic uncertainty $\mathcal{U}_e(\mathbf{X}_i)$ for the i th training data over a minibatch of size B :

$$w(\mathbf{X}_i) = 1 + \frac{\exp(\mathcal{U}_e(\mathbf{X}_i))}{\sum_{b=1}^B \exp(\mathcal{U}_e(\mathbf{X}_b))}. \quad (18)$$

The overall training loss is then given by:

$$\mathcal{L} = \sum_{i=1}^N w_i \mathcal{L}_{NLL,i} + \mathcal{L}_{generic,i}, \quad (19)$$

where $w_i = \{w(\mathbf{X}_i), 1\}$ with $w_i = 1$ when training an initial model while $w_i = w(\mathbf{X}_i)$ during the refinement, and N is the total number of training images.

4. Experiment

We first perform ablation studies to show the effectiveness of KNOWN’s body knowledge and uncertainty modeling. Then, we quantitatively compare KNOWN to the weakly-supervised state-of-the-art methods (SOTAs), *i.e.* those that do not require 3D supervisions paired with input images during training. We highlight the advantages of KNOWN on the minority testing samples. Last, we qualitatively evaluate KNOWN’s 3D reconstruction performance and illustrate how uncertainty modeling can lead to model improvements. This is followed by a discussion on KNOWN’s generalization ability.

Datasets and implementation. Following the typical strategy [38, 44] for model training, we consider both 3D and 2D datasets. 3D datasets include Human3.6M (H36M) [33] and MPI-INF-3DHP (MPI-3D) [56], which are collected in constrained environments using MoCap system. 2D datasets include COCO [52], LSP and LSP-Extended [36], and MPII [3], which are diverse in poses, subjects, and backgrounds. The input images are augmented with random scaling and flipping. ResNet-50 [31] is pretrained on the ImageNet classification task. Body images are scaled to 224×224 pixels with the aspect ratio preserved. Images from different datasets are fed into one minibatch. For Eq. (15), we use the predicted mean of the body pose and shape parameters. For Eq. (16), only one sample is drawn to efficiently approximate the integration [8]. Appx. B provides detailed implementation settings.

Loss functions					P-MPE (H36M, P2)	
\mathcal{L}_{ana}	\mathcal{L}_{bio}	\mathcal{L}_{phy}	\mathcal{L}_{NLL}	U-refine	All	Minority
	✓	✓	✓		146.0	160.9
✓		✓	✓		107.4	148.5
✓	✓		✓		60.7	81.7
✓	✓	✓			63.3	78.8
✓	✓	✓	✓		58.1	81.3
✓	✓	✓	✓	✓	55.9	70.3

Table 1. **Ablation study of different loss terms.** \mathcal{L}_{ana} , \mathcal{L}_{bio} , and \mathcal{L}_{phy} stand for $\mathcal{L}_{anatomy}$, $\mathcal{L}_{biomechanics}$, and $\mathcal{L}_{physics}$, respectively. “U-refine” indicates whether utilizing the uncertainty-guided refinement. When not using NLL, we use Mean Square Error (MSE) to compute the 2D keypoint reprojection error. We report evaluation on all the testing images (All) and those testing images having high epistemic uncertainty (Minority).

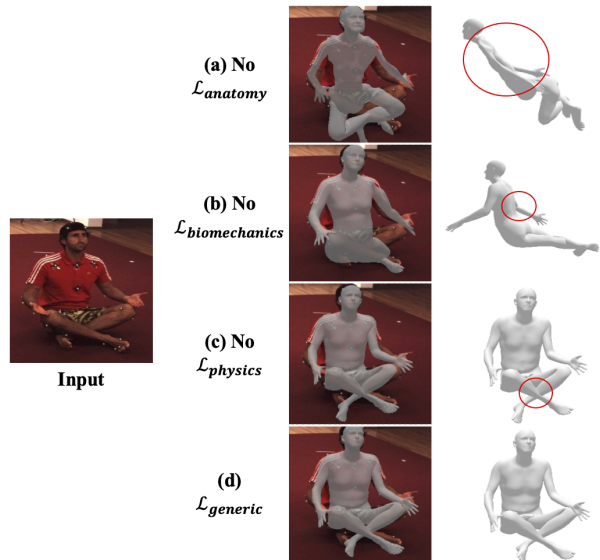


Figure 3. **Different generic constraints in reconstructing realistic 3D human body.** Distinct violations of the constraints are marked by red circles.

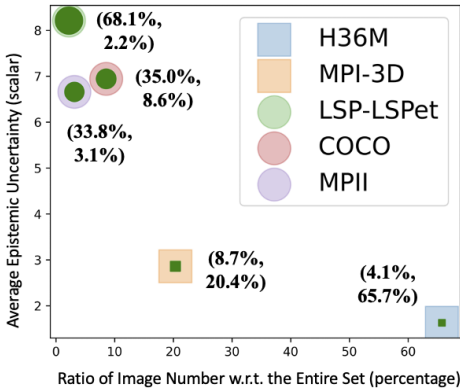
Evaluation metrics. We evaluate the 3D body pose estimation task using Mean Per-Joint Position Error (MPE) and MPE after rigid alignment (P-MPE) in units of millimeters. MPE and P-MPE measure the average distance between the predicted joint positions and the ground truth; thus smaller values are better. For the evaluation on H36M, the MPE and P-MPE computations follow two typical protocols: P1 (all images) and P2 (just frontal camera images).

4.1. Ablation Study

Generic constraints. Tab. 1 (Rows 1-4) and Fig. 3 summarize and illustrate the impact of the various generic constraints used by KNOWN. With all three constraints in



(i) Human3.6M (ii) MPI-3D (iii) LSP-LSPet (iv) COCO (v) MPII
(a) Sample images with high aleatoric uncertainty



(b) Statistics of epistemic uncertainty

Figure 4. **Aleatoric and epistemic uncertainty.** (a) Aleatoric uncertainty (Eq. (7)). Example training images and the corresponding label of left (green) and right (blue) side of 2D body skeleton are shown. (b) Epistemic uncertainty (Eq. (8)). The percentage of images within a dataset that are minorities (top) and the percentage of images in a dataset relative to the entire training set (bottom) are shown alongside each dataset.

place, the average P-MPE reconstruction error is 58.1mm. If the body anatomy constraints are removed, the predicted shape is very unreasonable on a relatively challenging image containing occlusion (Fig. 3(a)), with large depth estimation error leading to a P-MPE of 146.0mm. Removing just the biomechanics constraint also degrades the reconstruction significantly: the P-MPE is 107.4mm and Fig. 3(b) is a clearly invalid pose. While removing the physics constraint has a relatively small impact on P-MPE (60.7mm), it can still lead to physically implausible poses because of body part penetration like that illustrated in Fig. 3(c). Thus the three types of generic constraints are evidently synergistic, providing good, plausible 3D reconstructions from 2D body landmarks. The impact of the NLL constraint is to reduce the P-MPE somewhat, but the three generic constraints appear to be enough to ensure realistic 3D body reconstruc-

tion, as seen in Fig. 3(d).

Uncertainty modeling. KNOWN effectively captures both aleatoric and epistemic uncertainty. By leveraging aleatoric uncertainty, KNOWN successfully identifies the images containing large data noise due to self-occlusion (Fig. 4(a), columns 1-4), poor image quality (Fig. 4(a), column 3), or annotation errors (Fig. 4(a), columns 2,5). For example, images in MPI-3D are collected by a markerless MoCap system. They can have severe labeling errors, including inaccurate body landmark positions or mislabeling of left and right body sides (Fig. 4(a), column 2). Moreover, leveraging epistemic uncertainty, KNOWN characterizes data imbalance as shown in Fig. 4(b). Specifically, we use KNOWN (before refinement) to compute epistemic uncertainty of the training images in different datasets and report statistics related to the image number and epistemic uncertainty. To better illustrate the data imbalance problem, we regard the images whose epistemic uncertainty are larger than 90% of the training data as *minorities*. As shown, compared to 3D datasets (H36M, MPI-3D), 2D datasets (LSP, COCO, MPII) have smaller dataset size but larger average epistemic uncertainty and minority ratio due to their greater diversity in backgrounds, poses, and subjects. For example, images in LSP-LSPet are only 2.2% of the total training images but these images have high epistemic uncertainty and 68.1% of them are minorities. The success of KNOWN in exploiting uncertainty modeling to capture data noise and data imbalance allows further model improvements through the new NLL training loss and the novel uncertainty-guided refinement training strategy as illustrated below.

NLL and uncertainty-guided refinement. The NLL loss term adaptively assigns smaller weights to noisy inputs based on their larger aleatoric uncertainty (which relates to the predicted variances as derived in Eq. (7)). Utilizing the NLL loss results in a training loss function that is robust to data noise and leads to better model performance. As seen in Tab. 1 (Rows 4-5), when added to the generic loss term, NLL reduces the error from 63.6mm to 58.1mm. The final step in the KNOWN pipeline is the uncertainty-guided refinement, which exploits the well-captured epistemic uncertainty as guidance to assign larger weights on the minority images exhibiting high epistemic uncertainty. Utilizing the uncertainty-guided refinement effectively handles the data imbalance problem and further improves the model performance, especially for the challenging minorities. As shown in Tab. 1 (Rows 5-6), it reduces reconstruction error from 58.1mm to 55.9mm, with particular advantage on the minorities (from 81.3mm to 70.3mm).

4.2. Comparison to SOTAs

Tab. 2 compares KNOWN with prior works. In order to ensure a fair comparison, we exclude methods that rely on additional 2D features as input [63, 65, 64, 24]. KNOWN

Method	Source of the Prior	H36M		MPI-3D
		MPE	P-MPE	P-MPE
HMR [38]	3D	106.8	66.5	113.2
SPIN [44]	3D	-	62.0	80.4
Song <i>et al.</i> [66]	3D	-	56.4	-
Kundu <i>et al.</i> [46]	3D	86.4	58.2	-
HUND [89]	3D	91.8	-	-
THUNDER [91]	3D	87.0	59.7	-
PoseNet [69]	2D	-	59.4	102.4
Yu <i>et al.</i> [87]	2D	87.1	-	87.4
Ours	Knowledge	79.2	55.9	79.3

Table 2. **Comparisons to SOTAs.** Existing works rely on data-drive priors learned from 3D MoCap data (3D) or 2D body pose data (2D), while we derive generic prior from well-established body knowledge (Knowledge). Results of other works are from their paper. Missing values present due to their absence in the original paper.

achieves the best performance on both H36M and MPI-3D. In particular, HMR adopts a data-driven prior learned from 3D MoCap data (including H36M but not MPI-3D) as a regularization for training with 2D body landmarks only. HMR fails to perform well on MPI-3D, where the poses are different from the learned prior. In contrast, KNOWN leverages the generic constraints to achieve consistent performance on both of the datasets. Building upon HMR, SPIN further incorporates an optimization procedure into training to iteratively refine the estimation. Instead of treating all the samples equally and refining the model via multiple rounds as SPIN, KNOWN effectually targets the challenging minority images by employing its uncertainty-guided refinement strategy, outperforming SPIN substantially on H36M and slightly on MPI-3D. Moreover, Kundu *et al.* depend on a 3D pose prior and additional self-supervision from appearance cues for training. HUND and THUNDER build prior models from MoCap data to constrain the reconstruction space. Yu *et al.* and PoseNet extract a 2D projection prior from 2D pose data to encourage prediction feasibility, while Yu *et al.* relies on an extra inverse kinematic mapping module for incorporating body knowledge. Compared to them, KNOWN achieves better performance by (1) leveraging the generic body constraints, which are systematic, physically meaningful, and readily incorporated into different frameworks, and (2) using uncertainty modeling, thereby achieving more efficient and effective training.

To further demonstrate the advantages of KNOWN, we highlight the evaluation on the minorities. Tab. 3 reports the evaluation of KNOWN on the minority testing images of H36M compared to two recent works. For both HMR and SPIN, performance drops severely on the minorities even with the usage of 3D supervisions (3D body model param-

Method	All	Minority
HMR [38] (w.o./w. \mathcal{P})	65.2/57.2	108.5/107.3
SPIN [44] (w.o./w. \mathcal{P})	62.0/41.1	-/71.0
Ours	55.9	70.3

Table 3. **Qualitative evaluation on the minorities.** The numbers are P-MPE evaluated on the testing images of H36M (P2). \mathcal{P} denotes the usage of the 3D annotations paired with the input images.




Figure 5. **Qualitative evaluation with uncertainty visualization.** Images are the testing set from: H36M (top row), and MPI-3D (bottom row). In columns 2-3, the size of the red (aleatoric) and yellow (epistemic) circles indicate the level of quantified uncertainty. Larger sizes indicate higher uncertainties. For the 3D reconstruction results before (columns 4) and after the refinement (columns 5), the color of each vertex represents the level of epistemic uncertainty on each vertex with lighter colors representing higher epistemic uncertainty.

eter annotations paired with the input images). In contrast, KNOWN outperforms their fully supervised setting without requiring any 3D information. For example, we achieve 70.3mm, significantly better than HMR’s 107.3mm. The evaluation on minorities further establishes the importance of uncertainty modeling and generic knowledge and demonstrates the advantages of KNOWN.

4.3. Qualitative Evaluation

3D reconstruction with uncertainty visualization. As shown in Fig. 5, KNOWN generates accurate 3D reconstruction results. Beyond 3D reconstruction, KNOWN also provides both aleatoric and epistemic uncertainty of the prediction. The estimated uncertainty captures the hard regions, such as large aleatoric and epistemic uncertainty are exhibited at the occluded regions or at the leaf joints (re-

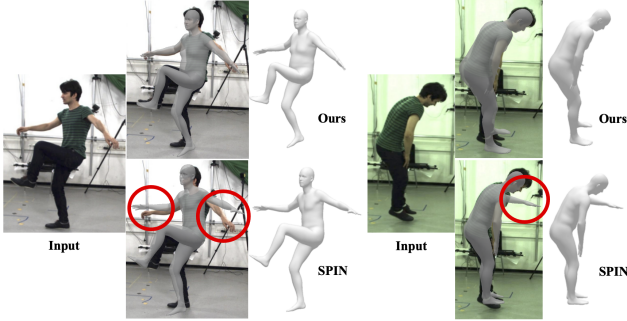


Figure 6. **Qualitative evaluation with comparison of using the generic prior (KNOWN) and the data-driven prior (SPIN).**

gions circled by green in Fig. 5 Columns 2-3). By further incorporating uncertainty-guided refinement, KNOWN improves accuracy for these challenging and low confidence cases (region circled by blue over region circled by red in Fig. 5 Columns 4-5). Moreover, besides estimating the uncertainty in 2D keypoint projections, KNOWN captures the epistemic uncertainty of each 3D vertex prediction (columns 4-5 of Fig. 5), providing different insights into the model’s behaviour. We discuss them and the way of quantifying 3D vertex prediction uncertainty in Appx. C.

Generalization. Fig. 6 compares KNOWN with the data-driven SPIN method on testing images from MPI-3D. As illustrated in these examples, when the poses in images are distinct to the training data, SPIN can fail by generating results that, while plausible in 3D pose, are not well-aligned with the input images (regions circled by red in Fig. 6). In contrast, KNOWN generalizes better to different data sets because it utilizes generic body knowledge from literature rather than data.

5. Discussion

Thus far we have focused primarily on the benefits of KNOWN considering the challenging scenario where 3D annotations are unavailable. As described more fully in Appx. D), KNOWN can also utilize any additional available annotations, such as paired 3D annotations. Here we compare KNOWN with prior works that use uncertainty modeling in conjunction with 3D annotations.

Tab. 4 shows that KNOWN compares favorably to recent methods on H36M and 3DPW [73] (a dataset captured outdoors with more diverse backgrounds). Compared to the deterministic data-driven baselines (HMR and SPIN) and SOTAs with uncertainty modeling (Biggs *et al.* and ProHMR), KNOWN achieves better reconstruction performance on both datasets without using additional 3D Mo-Cap data that are required by these methods. Moreover, existing approaches with uncertainty modeling only capture aleatoric uncertainty, which is used solely during testing to

Methods	\mathcal{U}	\mathcal{M}	\mathcal{S}	P-MPE	
				H36M	3DPW
HMR [38] (w. \mathcal{P})	N/A	103.1	10.2	56.8	81.3
SPIN [44] (w. \mathcal{P})	N/A	103.1	10.2	41.1	59.1
Biggs <i>et al.</i> [6]	\mathcal{U}_a	566.6	-	41.6	59.9
ProHMR [45]	\mathcal{U}_a	223.0	16.4	41.2	59.8
Ours				55.9	70.4
Ours (w. \mathcal{P})	$\mathcal{U}_a+\mathcal{U}_e$	103.9	11.6	40.5	58.1

Table 4. **Evaluation of using paired 3D annotations.** \mathcal{U} indicates the ability to capture aleatoric (\mathcal{U}_a) or epistemic uncertainty (\mathcal{U}_e). The unit of model memory (\mathcal{M}) and running speed (\mathcal{S}) is MB and ms, respectively. The P-MPE computation on H36M follows P2. Note that KNOWN not only captures uncertainty, but leverages it effectively to improve the model.

generate plausible hypotheses. In contrast, KNOWN captures both aleatoric and epistemic uncertainty and improves the training process by utilizing the well-captured uncertainty. Tab. 4 illustrates that KNOWN achieves these advantages while sacrificing very little in terms of memory footprint and speed, consuming just 103.9MB in model size with a near-real-time running speed of 11.6ms. Compared to existing uncertainty modeling approaches, KNOWN is much more efficient. This suggests that KNOWN would be quite practical in a variety of real-world applications, such as in safety-critical scenarios that also require accurate uncertainty quantification.

6. Conclusion

We have introduced KNOWN, a knowledge-encoded probabilistic model that tackles the data insufficiency problem in 3D body reconstruction. KNOWN introduces a set of constraints derived from a systematic study of body knowledge available from literature. These constraints are generalizable, explainable, and easy to adapt to different frameworks. Moreover, KNOWN is a new probabilistic framework that efficiently and effectively captures both aleatoric and epistemic uncertainty. The captured uncertainty handles the data noise and data imbalance through training with a robust NLL loss and a novel uncertainty-guided refinement strategy. KNOWN achieves remarkable performance without relying on any 3D data, and is efficient in memory footprint and computation speed, suggesting that it would be useful in a wide variety of applications — particularly those for which acquiring 3D data is virtually impossible, such as 3D animal body pose reconstruction.

Acknowledgement This work is supported in part by the Cognitive Immersive Systems Laboratory (CISL), a collaboration between IBM and RPI, and also a center in IBM’s AI Horizon Network.

References

- [1] NSF Grant 0196217. Cmu graphics lab motion capture database. [2](#)
- [2] Ijaz Akhter and Michael J Black. Pose-conditioned joint angle limits for 3d human pose reconstruction. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1446–1455, 2015. [2](#)
- [3] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3686–3693, 2014. [6](#)
- [4] Anurag Arnab, Carl Doersch, and Andrew Zisserman. Exploiting temporal context for 3d human pose estimation in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3395–3404, 2019. [2](#)
- [5] Chiraz BenAbdelkader and Yaser Yacoob. Statistical estimation of human anthropometry from a single uncalibrated image. *Computational Forensics*, pages 200–20, 2008. [2](#)
- [6] Benjamin Biggs, Sébastien Ehrhardt, Hanbyul Joo, Benjamin Graham, Andrea Vedaldi, and David Novotny. 3d multi-bodies: Fitting sets of plausible 3d human models to ambiguous image data. *arXiv preprint arXiv:2011.00980*, 2020. [2](#), [9](#)
- [7] David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American statistical Association*, 112(518):859–877, 2017. [2](#)
- [8] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *International conference on machine learning*, pages 1613–1622. PMLR, 2015. [6](#)
- [9] 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 *European Conference on Computer Vision*, pages 561–578. Springer, 2016. [1](#), [2](#)
- [10] Joao Carreira, Pulkit Agrawal, Katerina Fragkiadaki, and Jitendra Malik. Human pose estimation with iterative error feedback. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4733–4742, 2016. [3](#)
- [11] Ching-Hang Chen, Amrbrish Tyagi, Amit Agrawal, Dylan Drover, Stefan Stojanov, and James M Rehg. Unsupervised 3d pose estimation with geometric self-supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5714–5724, 2019. [2](#)
- [12] Jixu Chen, Siqi Nie, and Qiang Ji. Data-free prior model for upper body pose estimation and tracking. *IEEE transactions on image processing*, 22(12):4627–4639, 2013. [2](#)
- [13] Tianqi Chen, Emily B. Fox, and Carlos Guestrin. Stochastic gradient hamiltonian monte carlo. In *Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32, ICML'14*, page II–1683–II–1691. JMLR.org, 2014. [2](#)
- [14] Zerui Chen, Yan Huang, Hongyuan Yu, Bin Xue, Ke Han, Yiru Guo, and Liang Wang. Towards part-aware monocular 3d human pose estimation: An architecture search approach. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 715–732. Springer, 2020. [2](#)
- [15] 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 *European Conference on Computer Vision*, pages 769–787. Springer, 2020. [1](#), [2](#)
- [16] Hai Ci, Chunyu Wang, Xiaoxuan Ma, and Yizhou Wang. Optimizing network structure for 3d human pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2262–2271, 2019. [2](#)
- [17] Rishabh Dabral, Anurag Mundhada, Uday Kusupati, Safer Afaque, Abhishek Sharma, and Arjun Jain. Learning 3d human pose from structure and motion. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 668–683, 2018. [1](#), [2](#)
- [18] Hao-Shu Fang, Yuanlu Xu, Wenguan Wang, Xiaobai Liu, and Song-Chun Zhu. Learning pose grammar to encode human body configuration for 3d pose estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. [2](#)
- [19] Loic Le Folgoc, Vasileios Baltatzis, Sujal Desai, Anand Devaraj, Sam Ellis, Octavio E Martinez Manzanera, Arjun Nair, Huaqi Qiu, Julia Schnabel, and Ben Glocker. Is mc dropout bayesian? *arXiv preprint arXiv:2110.04286*, 2021. [2](#)
- [20] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR, 2016. [2](#)
- [21] Stuart Geman and Donald Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on pattern analysis and machine intelligence*, (6):721–741, 1984. [2](#)
- [22] Georgios Georgakis, Ren Li, Srikrishna Karanam, Terrence Chen, Jana Košecká, and Ziyang Wu. Hierarchical kinematic human mesh recovery. In *European Conference on Computer Vision*, pages 768–784. Springer, 2020. [2](#)
- [23] Kehong Gong, Jianfeng Zhang, and Jiashi Feng. Poseaug: A differentiable pose augmentation framework for 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8575–8584, 2021. [2](#)
- [24] Xuan Gong, Meng Zheng, Benjamin Planche, Srikrishna Karanam, Terrence Chen, David Doermann, and Ziyang Wu. Self-supervised human mesh recovery with cross-representation alignment. *arXiv preprint arXiv:2209.04596*, 2022. [7](#)
- [25] 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*, pages 10884–10894, 2019. [2](#)
- [26] Marc Habermann, Weipeng Xu, Michael Zollhofer, Gerard Pons-Moll, and Christian Theobalt. Deepcap: Monocular

- human performance capture using weak supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5052–5063, 2020. 2
- [27] Ikhsanul Habibie, Weipeng Xu, Dushyant Mehta, Gerard Pons-Moll, and Christian Theobalt. In the wild human pose estimation using explicit 2d features and intermediate 3d representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10905–10914, 2019. 1
- [28] Joseph Hamill and Kathleen M Knutzen. *Biomechanical basis of human movement*. Lippincott Williams & Wilkins, 2006. 1, 2, 5
- [29] Mohamed Hassan, Vasileios Choutas, Dimitrios Tzionas, and Michael J Black. Resolving 3d human pose ambiguities with 3d scene constraints. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2282–2292, 2019. 1
- [30] W Keith Hastings. Monte carlo sampling methods using markov chains and their applications. 1970. 2
- [31] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016. 3, 6
- [32] Lorna Herda, Raquel Urtasun, and Pascal Fua. Hierarchical implicit surface joint limits for human body tracking. *Computer Vision and Image Understanding*, 99(2):189–209, 2005. 1
- [33] 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*, 36(7):1325–1339, 2013. 6, 14
- [34] Dominic Jack, Frederic Maire, Anders Eriksson, and Sareh Shirazi. Adversarially parameterized optimization for 3d human pose estimation. In *2017 International Conference on 3D Vision (3DV)*, pages 145–154. IEEE, 2017. 2
- [35] Wen Jiang, Nikos Kolotouros, Georgios Pavlakos, Xiaowei Zhou, and Kostas Daniilidis. Coherent reconstruction of multiple humans from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5579–5588, 2020. 1
- [36] Sam Johnson and Mark Everingham. Clustered pose and nonlinear appearance models for human pose estimation. In *bmvc*, volume 2, page 5. Citeseer, 2010. 6
- [37] Hanbyul Joo, Natalia Neverova, and Andrea Vedaldi. Exemplar fine-tuning for 3d human model fitting towards in-the-wild 3d human pose estimation. In *2021 International Conference on 3D Vision (3DV)*, pages 42–52. IEEE, 2021. 1
- [38] 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*, pages 7122–7131, 2018. 1, 2, 6, 8, 9, 17
- [39] Tero Karras. Maximizing parallelism in the construction of bvhs, octrees, and k-d trees. In *Proceedings of the Fourth ACM SIGGRAPH / Eurographics Conference on High-Performance Graphics*, pages 33–37. Eurographics Association, 2012. 2, 5
- [40] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? *arXiv preprint arXiv:1703.04977*, 2017. 2, 4, 6
- [41] S. Khan, M. Hayat, S. W. Zamir, J. Shen, and L. Shao. Striking the right balance with uncertainty. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 103–112, 2019. 2
- [42] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 5
- [43] Muhammed Kocabas, Chun-Hao P Huang, Otmar Hilliges, and Michael J Black. Pare: Part attention regressor for 3d human body estimation. *arXiv preprint arXiv:2104.08527*, 2021. 1
- [44] 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/CVF International Conference on Computer Vision*, pages 2252–2261, 2019. 1, 6, 8, 9
- [45] Nikos Kolotouros, Georgios Pavlakos, Dinesh Jayaraman, and Kostas Daniilidis. Probabilistic modeling for human mesh recovery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11605–11614, 2021. 1, 2, 9
- [46] Jogendra Nath Kundu, Mugalodi Rakesh, Varun Jampani, Rahul Mysore Venkatesh, and R Venkatesh Babu. Appearance consensus driven self-supervised human mesh recovery. In *European Conference on Computer Vision*, pages 794–812. Springer, 2020. 2, 8
- [47] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 6402–6413. Curran Associates, Inc., 2017. 2
- [48] Kyoungoh Lee, Inwoong Lee, and Sanghoon Lee. Propagating lstm: 3d pose estimation based on joint interdependency. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 119–135, 2018. 2
- [49] 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 3383–3393, 2021. 1, 2
- [50] Zhihao Li, Jianzhuang Liu, Zhensong Zhang, Songcen Xu, and Youliang Yan. Cliff: Carrying location information in full frames into human pose and shape estimation. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part V*, pages 590–606. Springer, 2022. 1
- [51] Kevin Lin, Lijuan Wang, and Zicheng Liu. End-to-end human pose and mesh reconstruction with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1954–1963, 2021. 1

- [52] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. [6](#)
- [53] Matthew Loper, Naureen Mahmood, and Michael J Black. Mosh: Motion and shape capture from sparse markers. *ACM Transactions on Graphics (TOG)*, 33(6):1–13, 2014. [1](#), [17](#)
- [54] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6):1–16, 2015. [1](#), [3](#), [4](#)
- [55] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5442–5451, 2019. [1](#), [2](#)
- [56] Dushyant Mehta, Helge Rhodin, Dan Casas, Pascal Fua, Oleksandr Sotnychenko, Weipeng Xu, and Christian Theobalt. Monocular 3d human pose estimation in the wild using improved cnn supervision. In *2017 international conference on 3D vision (3DV)*, pages 506–516. IEEE, 2017. [6](#)
- [57] Gyeongsik Moon, Hongsuk Choi, and Kyoung Mu Lee. Neuralannot: Neural annotator for 3d human mesh training sets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2299–2307, 2022. [1](#)
- [58] NASA NASA. Std-3000. man systems integration standards. *National Aeronautics and Space Administration: Houston, USA*, 1995. [1](#), [2](#), [5](#)
- [59] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. Expressive body capture: 3d hands, face, and body from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10975–10985, 2019. [1](#), [2](#)
- [60] 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*, pages 459–468, 2018. [1](#)
- [61] Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Reconstructing 3d human pose from 2d image landmarks. In *European conference on computer vision*, pages 573–586. Springer, 2012. [2](#)
- [62] István Sárándi, Alexander Hermans, and Bastian Leibe. Learning 3d human pose estimation from dozens of datasets using a geometry-aware autoencoder to bridge between skeleton formats. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2956–2966, 2023. [1](#)
- [63] Akash Sengupta, Ignas Budvytis, and Roberto Cipolla. Synthetic training for accurate 3d human pose and shape estimation in the wild. *arXiv preprint arXiv:2009.10013*, 2020. [7](#)
- [64] 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 IEEE/CVF International Conference on Computer Vision*, pages 11219–11229, 2021. [2](#), [7](#)
- [65] 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 16094–16104, 2021. [2](#), [7](#)
- [66] Jie Song, Xu Chen, and Otmar Hilliges. Human body model fitting by learned gradient descent. In *European Conference on Computer Vision*, pages 744–760. Springer, 2020. [8](#)
- [67] Adrian Spurr, Umar Iqbal, Pavlo Molchanov, Otmar Hilliges, and Jan Kautz. Weakly supervised 3d hand pose estimation via biomechanical constraints. *arXiv preprint arXiv:2003.09282*, 2020. [2](#)
- [68] Matthias Teschner, Stefan Kimmerle, Bruno Heidelberger, Gabriel Zachmann, Laks Raghupathi, Arnulph Fuhrmann, M-P Cani, François Faure, Nadia Magnenat-Thalmann, Wolfgang Strasser, et al. Collision detection for deformable objects. In *Computer graphics forum*, volume 24, pages 61–81. Wiley Online Library, 2005. [5](#)
- [69] Shashank Tripathi, Siddhant Ranade, Amrith Tyagi, and Amit Agrawal. Posenet3d: Learning temporally consistent 3d human pose via knowledge distillation. In *2020 International Conference on 3D Vision (3DV)*, pages 311–321. IEEE, 2020. [8](#)
- [70] Dimitrios Tzionas, Luca Ballan, Abhilash Srikantha, Pablo Aponte, Marc Pollefeys, and Juergen Gall. Capturing hands in action using discriminative salient points and physics simulation. *International Journal of Computer Vision (IJCV)*, 118(2):172–193, June 2016. [2](#), [5](#)
- [71] Matias Valdenegro-Toro and Daniel Saromo Mori. A deeper look into aleatoric and epistemic uncertainty disentanglement. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1508–1516. IEEE, 2022. [4](#)
- [72] Joost Van Amersfoort, Lewis Smith, Yee Whye Teh, and Yarin Gal. Uncertainty estimation using a single deep deterministic neural network. In *International conference on machine learning*, pages 9690–9700. PMLR, 2020. [2](#)
- [73] 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 *European Conference on Computer Vision (ECCV)*, sep 2018. [1](#), [9](#), [17](#)
- [74] Bastian Wandt and Bodo Rosenhahn. Repnet: Weakly supervised training of an adversarial reprojection network for 3d human pose estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7782–7791, 2019. [2](#)
- [75] Chunyu Wang, Yizhou Wang, Zhouchen Lin, Alan L Yuille, and Wen Gao. Robust estimation of 3d human poses from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2361–2368, 2014. [2](#)
- [76] Zhenyu Wang, Yali Li, Ye Guo, Lu Fang, and Shengjin Wang. Data-uncertainty guided multi-phase learning for

- semi-supervised object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4568–4577, 2021. 2
- [77] Tom Wehrbein, Marco Rudolph, Bodo Rosenhahn, and Bastian Wandt. Probabilistic monocular 3d human pose estimation with normalizing flows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11199–11208, 2021. 2
- [78] Xiaolin K Wei and Jinxiang Chai. Modeling 3d human poses from uncalibrated monocular images. In *2009 IEEE 12th International Conference on Computer Vision*, pages 1873–1880. IEEE, 2009. 2
- [79] Yeming Wen, Dustin Tran, and Jimmy Ba. Batchensemble: an alternative approach to efficient ensemble and lifelong learning. *arXiv preprint arXiv:2002.06715*, 2020. 2
- [80] David A Winter. *Biomechanics and motor control of human movement*. John Wiley & Sons, 2009. 1, 2, 4
- [81] Erroll Wood, Tadas Baltrušaitis, Charlie Hewitt, Matthew Johnson, Jingjing Shen, Nikola Milosavljević, Daniel Wilde, Stephan Garbin, Toby Sharp, Ivan Stojiljković, et al. 3d face reconstruction with dense landmarks. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XIII*, pages 160–177. Springer, 2022. 2
- [82] 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. 1
- [83] Jingwei Xu, Zhenbo Yu, Bingbing Ni, Jiancheng Yang, Xiaokang Yang, and Wenjun Zhang. Deep kinematics analysis for monocular 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 899–908, 2020. 2
- [84] Xiangyu Xu, Hao Chen, Francesc Moreno-Noguer, Laszlo A Jeni, and Fernando De la Torre. 3d human pose, shape and texture from low-resolution images and videos. *IEEE transactions on pattern analysis and machine intelligence*, 44(9):4490–4504, 2021. 1
- [85] Wei Yang, Wanli Ouyang, Xiaolong Wang, Jimmy Ren, Hongsheng Li, and Xiaogang Wang. 3d human pose estimation in the wild by adversarial learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5255–5264, 2018. 1, 2
- [86] Zhenbo Yu, Bingbing Ni, Jingwei Xu, Junjie Wang, Chenglong Zhao, and Wenjun Zhang. Towards alleviating the modeling ambiguity of unsupervised monocular 3d human pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8651–8660, 2021. 2
- [87] Zhenbo Yu, Junjie Wang, Jingwei Xu, Bingbing Ni, Chenglong Zhao, Minsi Wang, and Wenjun Zhang. Skeleton2mesh: Kinematics prior injected unsupervised human mesh recovery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8619–8629, 2021. 2, 8
- [88] Andrei Zanfir, Eduard Gabriel Bazavan, Hongyi Xu, William T Freeman, Rahul Sukthankar, and Cristian Sminchisescu. Weakly supervised 3d human pose and shape reconstruction with normalizing flows. In *European Conference on Computer Vision*, pages 465–481. Springer, 2020. 2
- [89] 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 14484–14493, 2021. 2, 8
- [90] 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*, pages 2148–2157, 2018. 1, 2
- [91] Mihai Zanfir, Andrei Zanfir, Eduard Gabriel Bazavan, William T Freeman, Rahul Sukthankar, and Cristian Sminchisescu. Thundr: Transformer-based 3d human reconstruction with markers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12971–12980, 2021. 2, 8
- [92] Ailing Zeng, Xiao Sun, Fuyang Huang, Minhao Liu, Qiang Xu, and Stephen Lin. Srnet: Improving generalization in 3d human pose estimation with a split-and-recombine approach. In *European Conference on Computer Vision*, pages 507–523. Springer, 2020. 2
- [93] Jianfeng Zhang, Xuecheng Nie, and Jiashi Feng. Inference stage optimization for cross-scenario 3d human pose estimation. *arXiv preprint arXiv:2007.02054*, 2020. 2
- [94] Xingyi Zhou, Qixing Huang, Xiao Sun, Xiangyang Xue, and Yichen Wei. Towards 3d human pose estimation in the wild: a weakly-supervised approach. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 398–407, 2017. 1
- [95] Xiaowei Zhou, Spyridon Leonardos, Xiaoyan Hu, and Kostas Daniilidis. 3d shape estimation from 2d landmarks: A convex relaxation approach. In *proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4447–4455, 2015. 2
- [96] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5745–5753, 2019. 3