

Probabilistic Monocular 3D Human Pose Estimation with Normalizing Flows

Tom Wehrbein¹Marco Rudolph¹Bodo Rosenhahn¹Bastian Wandt²¹Leibniz University Hannover, ²University of British Columbia

wehrbein@tnt.uni-hannover.de

Abstract

3D human pose estimation from monocular images is a highly ill-posed problem due to depth ambiguities and occlusions. Nonetheless, most existing works ignore these ambiguities and only estimate a single solution. In contrast, we generate a diverse set of hypotheses that represents the full posterior distribution of feasible 3D poses. To this end, we propose a normalizing flow based method that exploits the deterministic 3D-to-2D mapping to solve the ambiguous inverse 2D-to-3D problem. Additionally, uncertain detections and occlusions are effectively modeled by incorporating uncertainty information of the 2D detector as condition. Further keys to success are a learned 3D pose prior and a generalization of the best-of-M loss. We evaluate our approach on the two benchmark datasets Human3.6M and MPI-INF-3DHP, outperforming all comparable methods in most metrics. The implementation is available on GitHub¹.

1. Introduction

Estimating the 3D pose of a human from a single monocular image is an active research field in computer vision. It has many applications *e.g.* in human computer interaction, animation, medicine and surveillance. A common approach is to decouple the problem into two stages. In the first stage, a 2D pose detector is used to estimate 2D keypoints which are then lifted to 3D joint locations in the second stage. By utilizing a 2D pose detector pretrained on diverse and richly annotated data, the 3D pose estimator becomes invariant to different scenes varying in lighting, background and clothing. However, reconstructing the correct 3D pose from 2D joint detections is a highly ill-posed problem because of depth ambiguities and occluded body parts. While some ambiguities can be resolved by utilizing information from the image (*e.g.* difference in shading due to depth disparity) or by exploiting known proportions of the human body, such as joint angle and bone length constraints, there

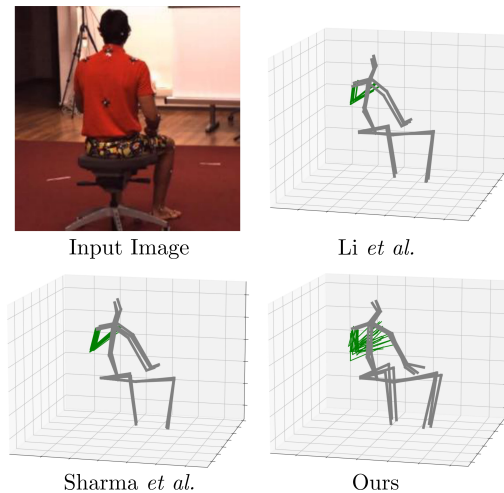


Figure 1. Our model generates diverse 3D pose hypotheses that are consistent with the input image. Compared to [28, 40] we achieve a higher diversity mainly where 2D detections are uncertain, in this case for the occluded left arm. For visualization purposes, more than three hypotheses are shown only for the highly ambiguous left arm.

still remain scenarios where multiple plausible 3D poses are consistent with the same image. Fig. 1 shows such a situation where the left arm is occluded by the upper body and therefore its position cannot be determined unambiguously. Nevertheless, most existing works ignore the ambiguities by assuming that only a single solution exists. In contrast, we model monocular 3D human pose estimation as an ambiguous inverse problem with multiple feasible solutions. Thus, in this work, we propose to estimate the full posterior distribution of plausible 3D poses conditioned on a monocular image.

Recently, few methods [20, 28, 29, 35, 40] have been proposed that follow the line of research to explicitly generate multiple 3D pose hypotheses from the 2D input. However, they only consider 2D joint coordinates and ignore the uncertainty of the 2D detector. While it is reasonable to infer depth ambiguities based on 2D coordinates only, directly modeling occlusions and uncertain detections is not meaningful. Fortunately, most 2D human joint detectors

¹<https://github.com/twehrbein/Probabilistic-Monocular-3D-Human-Pose-Estimation-with-Normalizing-Flows>

encode valuable information about uncertainties of the location of human joints in the predicted heatmaps. Instead of discarding this information, we propose to explicitly extract and utilize the uncertainties of the 2D detector from the estimated heatmaps. As shown in Fig. 1, this enables us to effectively model the uncertainties of the 2D detector together with the inherent depth ambiguities.

In this work, we propose a normalizing flow based method inspired by the framework for solving ambiguous inverse problems from Ardizzone *et al.* [3]. A normalizing flow [38, 46, 47] is a sequence of bijective transformations which allows evaluation in both directions. We propose to view 3D human pose estimation from a single image as an ambiguous inverse problem, since it is a deterministic forward process (*i.e.* projection of the 3D pose to 2D) with multiple different inverse mappings. Constructing a bijection between a 3D pose and the combination of a 2D pose with a latent vector allows to utilize the 3D-to-2D mapping (forward process) during training. Intuitively, depth information that otherwise gets lost in the forward process is encoded in the latent vector. Repeatedly sampling the latent vector and computing the inverse path of the normalizing flow generates arbitrary many 3D pose hypotheses that approximate the true posterior distribution. To incorporate the uncertainty information from the heatmaps, we employ a conditional variant of normalizing flows [4, 52]. We extract the uncertainty information by fitting 2D Gaussians to the heatmaps which are then used to form a conditioning vector. We optimize the model in both directions. The forward path learns the 3D-to-2D mapping and to produce latent vectors following a predefined distribution. For the inverse path, we utilize the 3D pose discriminator of [49] to penalize anthropometrically unfeasible poses. Additionally, we apply a loss enforcing the 3D pose hypotheses to reflect the uncertainties of the 2D detector. Motivated by common practice in particle filters, we further propose a generalization of the *best-of-M* loss [16] that minimizes the distance between the mean of the k best hypotheses and the corresponding ground truth.

We evaluate our approach on the two benchmark datasets Human3.6M [19] and MPI-INF-3DHP [33] and outperform all comparable methods in most metrics. Given the focus on ambiguous examples, we further evaluate on a subset of Human3.6M containing only samples with a high degree of 2D detector uncertainty. On this subset, our method outperforms the competitors by a large margin. To summarize, our contributions are:

- To the best of our knowledge, we are the first to employ a normalizing flow based method for modeling the posterior distribution of 3D poses given a single image.
- Uncertainty information from the predicted heatmaps

of the 2D detector is incorporated into our method, enabling to effectively model occlusions and uncertain detections.

- We propose a generalization of the *best-of-M* loss that noticeably improves prediction performance.

2. Related Work

In this section, we first give an overview of recent work in 3D human pose estimation focused on two-stage approaches. Afterwards, existing methods for multi-hypotheses 3D pose generation are discussed, followed by an overview of relevant work on normalizing flows. While there has been recent interest in estimating the 3D human body shape from monocular images [5, 21, 25, 26, 30, 37, 53, 55], this work focuses on predicting the 3D locations of a set of predefined joints.

Lifting 2D to 3D: Our approach belongs to the vast body of work that estimate 3D poses from the output of a 2D pose detector [8, 9, 12, 17, 18, 31, 41, 49, 50, 51, 54]. These two-stage approaches decouple the difficult problem of 3D depth estimation from the easier 2D pose localization. Furthermore, it allows to use both indoor and in-the-wild data for training the 2D detector, which effectively reduces the bias towards sterile indoor scenes. Akhter and Black [1] learn a pose-conditioned joint angle limit prior to restrict invalid 3D pose reconstructions. They perform 3D pose estimation using an over-complete dictionary of poses. Moreno-Noguer [34] casts the problem as a regression between 2D and 3D poses represented as distance matrices. Lifting 2D to 3D joints was further sparked by Martinez *et al.* [32], who employ a simple fully-connected network to lift 2D detections to 3D poses, surprisingly outperforming past approaches. Due to its simplicity and strong performance, it serves as a popular baseline for many following works.

Unlike the above mentioned approaches that assume a unimodal posterior distribution and only predict a single 3D pose for each input, we are able to generate a diverse set of plausible 3D poses. Additionally, anatomical constraints are learned implicitly by utilizing a strong 3D pose discriminator. In contrast to previous works that integrate uncertainty information of the 2D detector (*e.g.* [7, 50, 54]), we fit a 2D Gaussian to each heatmap instead of using only the maximum value of each heatmap as confidence score, thus better capturing the uncertainty distribution.

Multi-Hypotheses 3D Human Pose Estimation: There are early works [27, 42, 43, 44] that extensively analyze and discuss the ambiguities of monocular 3D human pose estimation and sample multiple 3D poses via heuristics. More recently, Jahangiri and Yuille [20] propose to generate multiple hypotheses from a predicted seed 3D pose by uniformly sampling from learned occupancy matrices [1]. Furthermore, they impose bone length constraints and re-

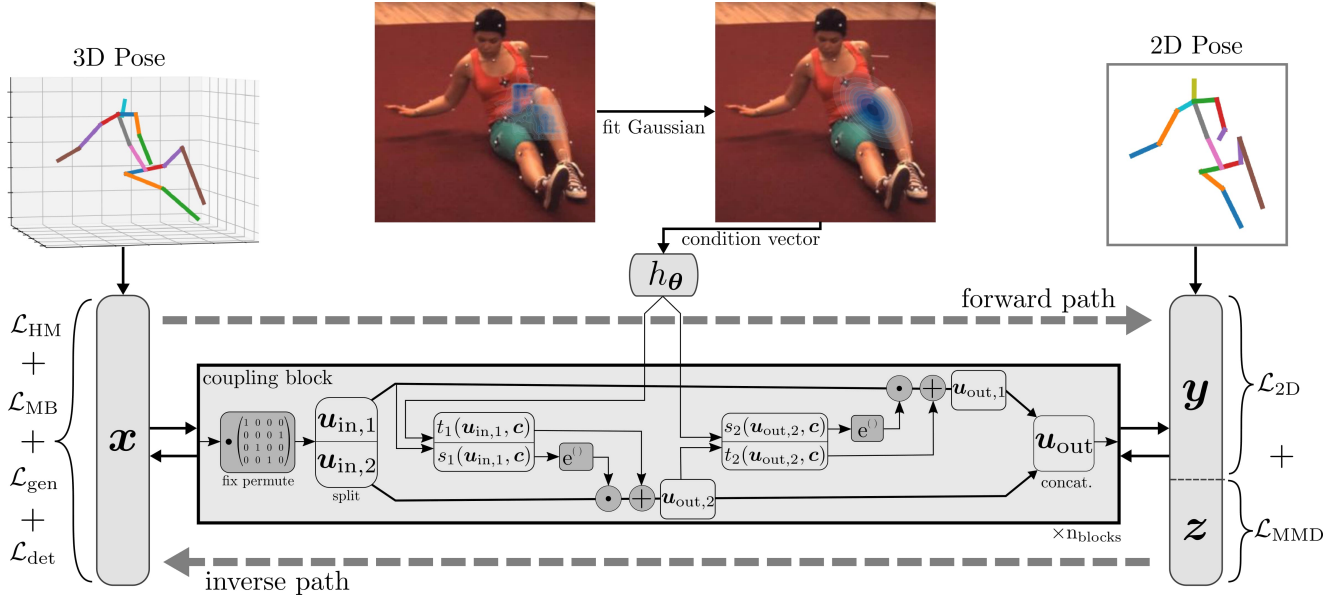


Figure 2. An overview of our proposed method. We employ a normalizing flow consisting of affine coupling blocks [11] for generating multiple 3D pose hypotheses. By constructing a bijection between a 3D pose and the concatenation of a 2D pose with a latent vector, we can exploit the 3D-to-2D mapping (forward path) during training. The model is optimized in both directions, whereas at inference, only the path from 2D to 3D (inverse path) is computed. Arbitrary many 3D pose hypotheses can be generated by repeatedly sampling the latent vector from a known distribution and computing the inverse path. Uncertainty information of the 2D detector in form of fitted Gaussians is incorporated by conditioning the coupling blocks. The architecture of a single coupling block is visualized in the gray box. For visualization purposes, only the forward computation of the coupling block is shown.

ject hypotheses with 2D reprojection error larger than some threshold. Li and Lee [28] employ a mixture density network (MDN) [6] to learn the multimodal posterior distribution. The conditional mean of each Gaussian kernel then denotes one 3D pose hypothesis. Oikarinen *et al.* [35] utilize the semantic graph neural network of [56] to improve upon the MDN approach of [28]. Contrary to our normalizing flow based approach, the number of generated hypotheses needs to be specified a priori and is fixed for every input. Furthermore, when increasing the number of generated hypotheses, significantly more computational resources are required. Sharma *et al.* [40] employ a conditional variational autoencoder to synthesize diverse 3D pose hypotheses conditioned on a 2D pose detection. They also propose to derive joint-ordinal depth relations from the image to rank the 3D pose samples. In contrast to [20], our normalizing flow based approach does not need to incorporate computationally heavy rejection sampling or requires to define the number of generated 3D pose hypotheses a priori. Our method is more flexible and is able to model any posterior distribution without requiring explicit hard constraints. Moreover, we are the only ones to incorporate the uncertainty information of the 2D detector, enabling us to significantly improve on highly ambiguous cases and to inherently handle an arbitrary number of occluded joints.

Normalizing Flows: A normalizing flow [38, 46, 47] is a sequence of bijective transformations that transforms

a simple tractable distribution into a complex target data distribution. Because of the bijectivity, evaluation in both directions is possible. Namely, sampling data from the modeled distribution as well as exact density estimation (*i.e.* assigning a likelihood to each data point). Most common state-of-the-art flow architectures are based on autoregressive models that utilize the Bayesian chain rule to decompose the density [10, 11, 13, 23, 36, 39]. For a more comprehensive introduction, we refer the reader to [24].

Ardizzone *et al.* [3] extend the real-valued non-volume preserving (Real-NVP) transformations from Dinh *et al.* [11] to the task of computing posteriors for ambiguous inverse problems. Given such an ambiguous inverse problem, they propose to learn the well-understood forward process in a supervised manner and encode otherwise lost information in additional latent variables. Thus, they learn a bijective mapping between the target data distribution and the joint distribution of latent variables and forward process solutions. Due to invertibility, the inverse is implicitly learned. By repeatedly sampling the latent variables from a simple tractable distribution, they can approximate the full posterior. Inspired by their work, we adopt and extend their framework for modeling the full posterior distribution of plausible 3D poses conditioned on a monocular image. We introduce a conditioning vector, a learnable prior and two additional loss functions.

To the best of our knowledge, the only previous works

in human pose estimation utilizing normalizing flows are [5, 53, 55]. However, they employ normalizing flows as 3D pose prior and not for directly modeling the posterior distribution of 3D poses conditioned on an image.

3. Method

Our aim is to learn the full posterior distribution of plausible 3D poses conditioned on a monocular image. We follow the popular two-stage approach by first applying a state-of-the-art 2D joint detector [45] and subsequently using its output to estimate corresponding 3D pose hypotheses. The core idea is that instead of conditioning the posterior distribution only on the 2D detections, we additionally utilize uncertainty information extracted from the predicted heatmaps in a novel way. This enables to effectively model the uncertainties of the 2D detector together with the inherent depth ambiguities.

An overview of the proposed method is shown in Fig. 2. To learn the posterior distribution, we employ a normalizing flow to construct a bijective mapping between a 3D pose $\mathbf{x} \in \mathbb{R}^{3J}$ and the concatenation of a 2D pose $\mathbf{y} \in \mathbb{R}^{2J}$ with a latent vector $\mathbf{z} \in \mathbb{R}^J$, where J is the number of joints in one pose. The introduction of the latent vector \mathbf{z} allows to utilize the well-defined forward process of projecting a 3D pose to its 2D observation during training. Intuitively, \mathbf{z} captures depth information that is otherwise lost in the mapping from 3D to 2D. Instead of simply using the argmax of the heatmaps, we incorporate the uncertainty information of the 2D detector by conditioning the normalizing flow on Gaussians fitted to the heatmaps. At inference, the full posterior is approximated by repeatedly sampling \mathbf{z} from the distribution of latent variables and computing the inverse path. If the forward process is simulated successfully, all generated hypotheses reproject to the corresponding 2D pose observation.

3.1. Conditional Normalizing Flow

As normalizing flow we adopt the Real-NVP [11] affine coupling block architecture. This architecture can straightforwardly be extended to incorporate a conditional input [4, 52]. A single coupling block is shown in the gray box in Fig. 2. The input \mathbf{u}_{in} is split into two parts $\mathbf{u}_{\text{in},1}$ and $\mathbf{u}_{\text{in},2}$. Subsequently, $\mathbf{u}_{\text{in},1}$ and $\mathbf{u}_{\text{in},2}$ undergo a scale and translation transformation parameterized by the functions s_i and t_i ($i \in \{1, 2\}$) on two separate paths. The outputs $\mathbf{u}_{\text{out},1}$ and $\mathbf{u}_{\text{out},2}$ are concatenated to form the overall output of the coupling block. Given the heatmap condition $\hat{\mathbf{c}}$, further encoded into the conditioning vector $\mathbf{c} = h_\theta(\hat{\mathbf{c}})$, the forward path of a coupling block is defined as

$$\begin{aligned} \mathbf{u}_{\text{out},2} &= \mathbf{u}_{\text{in},2} \odot e^{s_1(\mathbf{u}_{\text{in},1}, \mathbf{c})} + t_1(\mathbf{u}_{\text{in},1}, \mathbf{c}) \\ \mathbf{u}_{\text{out},1} &= \mathbf{u}_{\text{in},1} \odot e^{s_2(\mathbf{u}_{\text{out},2}, \mathbf{c})} + t_2(\mathbf{u}_{\text{out},2}, \mathbf{c}), \end{aligned} \quad (1)$$

where \odot denotes the element-wise multiplication. The exponential function is used to prevent multiplication by zero, which ensures the invertibility of the block. Note that s_i and t_i represent functions that do not need to be invertible. The only restriction is that their produced output matches the dimensions of the data on the corresponding path in the coupling block. Instead of regressing the scale and translation coefficients separately, we employ a fully-connected network that jointly predicts them by splitting its output. By construction, the coupling block can be trivially inverted without any computational overhead. The overall network consists of multiple chained blocks, each followed by a pre-defined random permutation which shuffles the path assignment of the variables. The output of the last block is split to form the 2D pose \mathbf{y} and the latent vector \mathbf{z} .

Following Ardizzone *et al.* [4], we adopt a parameterized soft clamping mechanism to prevent instabilities caused by the exponential function in the coupling block. The soft clamping is defined as

$$\sigma_\alpha(r) = \frac{2\alpha}{\pi} \arctan \frac{r}{\alpha}, \quad (2)$$

and is applied as the last layer of s_1 and s_2 . It prevents scaling components of exploding magnitude by restricting the output to the interval $(-\alpha, \alpha)$.

3.2. Heatmap Condition

Recent 2D detectors are optimized by applying a supervised loss between the predicted heatmap and a ground-truth heatmap consisting of a 2D Gaussian centered at the joint location. This leads to the predicted heatmaps being a valuable source of uncertainty of the 2D detector. Instead of estimating 3D poses solely based on 2D joint coordinates, we incorporate the uncertainties of the 2D detector encoded in the estimated heatmaps. Specifically, we fit a 2D Gaussian to each predicted heatmap to best capture the uncertainty distribution. The fitting process is done using non-linear least squares. As initial parameters, we set the amplitude to 1, the mean of each Gaussian to the corresponding regressed 2D joint location and the covariance matrix to a diagonal matrix with $\sigma_x^2 = \sigma_y^2 = \sigma_{\text{gt}}^2$, where σ_{gt}^2 is the ground-truth variance used for training the 2D detector. For each image, the fitted coefficients are stacked to form a single vector. We discard the Gaussian coefficients for the hip joints, since the typical alignment of the root joint of the 3D poses heavily reduces the possible variances in these joints. Thus, the heatmap conditioning vector is denoted as $\hat{\mathbf{c}} \in \mathbb{R}^{6(J-3)}$. We employ a fully-connected network as encoding network h_θ that further encodes $\hat{\mathbf{c}}$ into $\mathbf{c} = h_\theta(\hat{\mathbf{c}})$. For the 3D pose hypotheses to best reflect the uncertainties of the 2D detector, we explicitly optimize the network to match the 2D Gaussian distributions in the x - and y -direction of the 3D hypotheses for each joint. Let

$\hat{\Sigma} \in \mathbb{R}^{2 \times 2}$ be the covariance matrix of a single joint estimated from the positions of that joint in the L produced hypotheses. Defining $\Sigma \in \mathbb{R}^{2 \times 2}$ as the covariance matrix of the fitted 2D Gaussian of the corresponding heatmap, we minimize a masked lower bound Root Mean Square Error (RMSE) between both covariance matrices:

$$\mathcal{L}_{\text{HM}}(\Sigma, \hat{\Sigma}) = m \cdot \left(\max(0, \Sigma_{1,1} - \hat{\Sigma}_{1,1})^2 + \max(0, \Sigma_{2,2} - \hat{\Sigma}_{2,2})^2 + (\Sigma_{1,2} - \hat{\Sigma}_{1,2})^2 \right)^{\frac{1}{2}}, \quad (3)$$

where the masking scalar m is defined as

$$m = \begin{cases} 1 & \sqrt{\Sigma_{1,1}} > \sigma_t \vee \sqrt{\Sigma_{2,2}} > \sigma_t \\ 0 & \text{otherwise} \end{cases}. \quad (4)$$

Thus, the loss has no influence if the 2D detector is certain about the location of the specific joint, indicated by a fitted Gaussian with standard deviations smaller than the threshold σ_t . To not unnecessarily restrict the network, we only penalize the diagonal entries of the covariance matrices if they are smaller than the corresponding ground-truth values.

3.3. Optimization

The core idea of the optimization procedure is to train the 3D-to-2D mapping (forward path) in a supervised manner, while the highly ambiguous 2D-to-3D mapping (inverse path) is learned implicitly due to the invertibility of the normalizing flow and supported by additional supervision with the inverse process. Each training iteration consists of first calculating the forward path, followed by L computations of the inverse path and two additional, one for the discriminator and one for the deterministic 3D reconstruction. The gradients from both directions are accumulated before performing a parameter update. Note that due to the Real-NVP coupling block architecture, both directions can be computed efficiently.

Forward Path: In the forward process, the network predicts the corresponding 2D joint detections given a 3D pose. This is optimized using the L_1 distance:

$$\mathcal{L}_{2\text{D}} = \|\mathbf{y} - \hat{\mathbf{y}}\|_1, \quad (5)$$

where \mathbf{y} is the ground-truth and $\hat{\mathbf{y}}$ the estimated 2D observation. The estimated latent variables are optimized to follow a zero-mean isotropic Gaussian $p_{\mathbf{Z}} = \mathcal{N}(0, I)$ and to be independent from the distribution of 2D observations $p_{\mathbf{Y}}$. Both properties are enforced by minimizing the Maximum Mean Discrepancy (MMD) [14] between the joint distribution of network outputs $q(\mathbf{y}, \mathbf{z})$ and the product of marginal distributions $p_{\mathbf{Y}}$ and $p_{\mathbf{Z}}$. Given samples $\hat{\mathbf{V}} = \{\hat{\mathbf{v}}_i\}_{i=1}^n$ drawn i.i.d. from $q(\mathbf{y}, \mathbf{z})$ and $\mathbf{V} = \{\mathbf{v}_i\}_{i=1}^n$ with $\mathbf{v}_i = [\mathbf{y}, \mathbf{z}]$

and $\mathbf{y} \sim p_{\mathbf{Y}}, \mathbf{z} \sim p_{\mathbf{Z}}$, the unbiased estimator of the squared MMD with kernel φ is

$$\mathcal{L}_{\text{MMD}} = \text{MMD}_u^2(\mathbf{V}, \hat{\mathbf{V}}) = \frac{1}{n(n-1)} \sum_{i \neq j}^n \varphi(\mathbf{v}_i, \mathbf{v}_j) + \frac{1}{n(n-1)} \sum_{i \neq j}^n \varphi(\hat{\mathbf{v}}_i, \hat{\mathbf{v}}_j) - \frac{2}{n^2} \sum_{i,j=1}^n \varphi(\mathbf{v}_i, \hat{\mathbf{v}}_j). \quad (6)$$

Following [3], we block the gradients of \mathcal{L}_{MMD} with respect to \mathbf{y} to prevent the predictions of \mathbf{y} from deteriorating.

Inverse Path: Given a 2D pose \mathbf{y} , a latent vector \mathbf{z} is drawn from the base distribution $p_{\mathbf{Z}}$ and concatenated to form the input $[\mathbf{y}, \mathbf{z}]$ of the inverse path. By repeatedly sampling $\mathbf{z} \sim p_{\mathbf{Z}}$, arbitrary many 3D pose hypotheses can be created. Although $\mathcal{L}_{2\text{D}}$ and \mathcal{L}_{MMD} are in theory sufficient to best approximate the true posterior distribution [3], we apply additional losses to the inverse path to improve convergence. To penalize geometrically unfeasible 3D pose hypotheses, we introduce a discriminator network and adopt the Improved Wasserstein GAN training procedure of [15]. The inverse path acts as the generator by producing 3D poses that minimize the negated output of the discriminator. This loss is denoted as \mathcal{L}_{gen} . The architecture of the discriminator is taken from [49], including a Kinematic Chain Space layer [48] encoding bone lengths and angular representations. Additionally, we generate a 3D pose for each 2D input by using the corresponding latent vector \mathbf{z}_{det} produced in the forward path. Note that contrary to sampling a latent vector $\mathbf{z} \sim p_{\mathbf{Z}}$, when using the estimated latent vector \mathbf{z}_{det} , it is reasonable to apply a supervised loss \mathcal{L}_{det} linking a 2D input to a single 3D pose, since the combination of a predicted latent vector and matching 2D pose detection should correspond to the single exact solution of the ambiguous inverse problem. We minimize the L_1 distance between the ground-truth 3D pose \mathbf{x} and the estimated 3D pose $\hat{\mathbf{x}}_{\text{det}}$

$$\mathcal{L}_{\text{det}} = \|\mathbf{x} - \hat{\mathbf{x}}_{\text{det}}\|_1. \quad (7)$$

To further guide the optimization process, we propose a generalization of the *best-of-M* loss [16]. Given a set of 3D pose hypotheses $\mathbf{H} = \{\hat{\mathbf{x}}_i\}_{i=1}^L$ generated from the same 2D input, we select the subset $\mathbf{H}_{\text{topk}} \subseteq \mathbf{H}$ consisting of the k pose hypotheses with the lowest Mean Per Joint Position Error (MPJPE) to the corresponding ground-truth pose \mathbf{x} . We then minimize the L_1 distance between the ground-truth pose \mathbf{x} and the mean of the k best hypotheses:

$$\mathcal{L}_{\text{MB}} = \left\| \mathbf{x} - \frac{\sum_{\hat{\mathbf{x}} \in \mathbf{H}_{\text{topk}}} \hat{\mathbf{x}}}{k} \right\|_1. \quad (8)$$

Overall: In total, the objective function of our normalizing flow is

$$\mathcal{L}_{\text{NF}} = \mathcal{L}_{2\text{D}} + \mathcal{L}_{\text{gen}} + \lambda_{\text{MMD}} \mathcal{L}_{\text{MMD}} + \lambda_{\text{det}} \mathcal{L}_{\text{det}} + \lambda_{\text{MB}} \mathcal{L}_{\text{MB}} + \lambda_{\text{HM}} \mathcal{L}_{\text{HM}}, \quad (9)$$

where λ_{MMD} , λ_{det} , λ_{MB} and λ_{HM} represent the weights of the corresponding losses. The discriminator network is optimized to distinguish between the 3D poses produced by the normalizing flow and 3D poses from the training set by minimizing the WGAN-GP objective function [15]. The encoding network h_{θ} is jointly optimized with the normalizing flow by propagating gradients from \mathcal{L}_{NF} through h_{θ} .

4. Experiments

4.1. Datasets and Evaluation Metrics

Human3.6M [19] is the largest video pose dataset for 3D human pose estimation. It features 7 professional actors performing 15 different activities, such as *Sitting*, *Walking* and *Smoking*. For each frame, accurate 2D and 3D joint locations and camera parameters are provided. We follow the standard protocols and evaluate on every 64th frame of subjects 9 and 11. Protocol 1 computes the Mean Per Joint Position Error (MPJPE) between the reconstructed and ground-truth 3D joint coordinates directly, whereas Protocol 2 first applies a rigid alignment between the poses (PMPJPE). We additionally show results for the Correct Poses Score (CPS) metric proposed by [50]. It considers a pose as correct if and only if all joints have a Euclidean distance to the ground-truth below a threshold value ϑ . CPS is then defined as the area under curve for $\vartheta \in [0 \text{ mm}, 300 \text{ mm}]$. Instead of evaluating the reconstruction joint by joint, the CPS considers the whole pose. Compared to other common metrics, it is better suited to detect wrongly estimated poses that could negatively influence downstream tasks.

Human3.6M Ambiguous (H36MA): To focus evaluation on highly ambiguous examples, we select a subset² of the Human3.6M test split according to the uncertainties of the 2D detector. This subset only contains samples for which at least one fitted Gaussian has a standard deviation larger than 5 px, which holds true for 6.4% of all samples in the test split. These samples are extremely challenging since the joint detector gives inaccurate or wrong results.

MPI-INF-3DHP (3DHP) [33] is a 3D human pose dataset containing annotated images recorded in three different settings: studio with green screen, studio without green screen and outdoors. We evaluate on the test split without utilizing the training data to assess the generalization capability of our network. Following previous works, the Percentage of Correct Keypoints (PCK) under 150 mm is adopted as the metric for 3DHP.

4.2. Implementation Details

2D Detector: We use the publicly available HRNet [45] pretrained on MPII [2] as our 2D joint detector and finetune

²Information about the exact composition of the subset can be found in the official GitHub repository.

it on Human3.6M. Target ground-truth heatmaps are created with $\sigma_{\text{gt}} = 2 \text{ px}$.

Data Preprocessing: We center each 2D pose to its mean and divide it by its standard deviation. The 3D poses are processed in metres and also mean centered individually. Before evaluation, 3D poses are zero-centered around the hip joint to follow the standard protocols.

Network Details: The normalizing flow consists of 8 coupling blocks with fully-connected networks, denoted as *subnetworks*, as scale and translation functions. Each subnetwork upscales its input to 1024 dimensions with a fully-connected layer. This is followed by a ReLU and a second fully-connected layer with dimension 48. The condition encoding network h_{θ} follows the same design with 256 and 56 as output dimensions of the fully-connected layers. We set the clamping parameter inside the coupling blocks to $\alpha = 2.0$. For \mathcal{L}_{MMD} , we follow [3] and employ a mixture of inverse multiquadratics kernels

$$\varphi_{\mathcal{S}}^{\text{im}}(\mathbf{v}, \hat{\mathbf{v}}) = \sum_{b \in \mathcal{S}} \frac{b}{b + \|\mathbf{v} - \hat{\mathbf{v}}\|^2} \quad (10)$$

with bandwidth parameters $\mathcal{S} = \{0.0025, 0.04, 0.81\}$.

Training: The overall network is trained for 155 epochs using Adam [22] with an initial learning rate of $1 \cdot 10^{-4}$ and momentum values $\beta_1 = 0.5$ and $\beta_2 = 0.9$. The learning rate is halved after 150 epochs, and a batch size of 64 is used. To improve optimization stability, we clip the gradients in the range $[-15, 15]$. During training, the covariance matrices are computed from $L = 200$ 3D pose hypotheses and the standard deviation threshold for the masking of \mathcal{L}_{HM} (Eq. 4) is set to $\sigma_t = 1.05 \cdot \sigma_{\text{gt}} = 2.1$. The weights of the different losses are set to $\lambda_{\text{MMD}} = 10$, $\lambda_{\text{det}} = \lambda_{\text{MB}} = 4$ and $\lambda_{\text{HM}} = 750$, and the number of best hypotheses selected in \mathcal{L}_{MB} to $k = 5$. Since we estimate 3D poses in metric scale, there needs to be a conversion factor defined to relate between covariance matrices from pose hypotheses and from heatmaps. We empirically found a good conversion factor to be $1 \text{ px} \hat{=} 10 \text{ mm}$.

4.3. Evaluation on Human3.6M

We follow previous works and report metrics for the best 3D pose hypothesis generated by our network. This is especially reasonable for ambiguous examples, where multiple diverse 3D poses form a correct solution for the 3D pose reconstruction. Therefore, instead of validating whether predictions are equal to a specific solution, we evaluate if that specific solution is contained in the set of predictions. Additionally, we show results for 3D poses generated with an all-zero latent vector z_0 . Since we sample z from $\mathcal{N}(0, I)$ during training, such poses are approximately the highest likelihood solutions. Following [40], we produce $M = 200$ hypotheses for each 2D input. The results of our approach and other state-of-the-art methods are shown in Table 1. We

Protocol 1 (MPJPE)	Direct.	Disc.	Eat	Greet	Phone	Photo	Pose	Purch.	Sit	SitD	Smoke	Wait	WalkD	Walk	WalkT	Avg.
Martinez <i>et al.</i> [32] ($M = 1$)	51.8	56.2	58.1	59.0	69.5	78.4	55.2	58.1	74.0	94.6	62.3	59.1	65.1	49.5	52.4	62.9
Li <i>et al.</i> [29] ($M = 10$)	62.0	69.7	64.3	73.6	75.1	84.8	68.7	75.0	81.2	104.3	70.2	72.0	75.0	67.0	69.0	73.9
Li <i>et al.</i> [28] ($M = 5$)	43.8	48.6	49.1	49.8	57.6	61.5	45.9	48.3	62.0	73.4	54.8	50.6	56.0	43.4	45.5	52.7
Oikarinen <i>et al.</i> [35] ($M = 200$)	40.0	43.2	41.0	43.4	50.0	53.6	40.1	41.4	52.6	67.3	48.1	44.2	44.9	39.5	40.2	46.2
Sharma <i>et al.</i> [40] ($M = 200$)	37.8	<u>43.2</u>	43.0	44.3	51.1	57.0	39.7	43.0	56.3	<u>64.0</u>	<u>48.1</u>	<u>45.4</u>	50.4	<u>37.9</u>	39.9	46.8
Ours (z_0) ($M = 1$)	52.4	60.2	57.8	57.4	65.7	74.1	56.2	59.1	69.3	78.0	61.2	63.7	67.0	50.0	54.9	61.8
Ours ($M = 200$)	38.5	42.5	39.9	41.7	46.5	51.6	39.9	40.8	49.5	56.8	45.3	46.4	<u>46.8</u>	37.8	40.4	44.3
Protocol 2 (PMPJPE)	Direct.	Disc.	Eat	Greet	Phone	Photo	Pose	Purch.	Sit	SitD	Smoke	Wait	WalkD	Walk	WalkT	Avg.
Martinez <i>et al.</i> [32] ($M = 1$)	39.5	43.2	46.4	47.0	51.0	56.0	41.4	40.6	56.5	69.4	49.2	45.0	49.5	38.0	43.1	47.7
Li <i>et al.</i> [29] ($M = 10$)	38.5	41.7	39.6	45.2	45.8	46.5	37.8	42.7	52.4	62.9	45.3	40.9	45.3	38.6	38.4	44.3
Li <i>et al.</i> [28] ($M = 5$)	35.5	39.8	41.3	42.3	46.0	48.9	36.9	37.3	51.0	60.6	44.9	40.2	44.1	33.1	36.9	42.6
Oikarinen <i>et al.</i> [35] ($M = 200$)	30.8	34.7	<u>33.6</u>	<u>34.2</u>	<u>39.6</u>	<u>42.2</u>	<u>31.0</u>	31.9	<u>42.9</u>	53.5	<u>38.1</u>	<u>34.1</u>	<u>38.0</u>	<u>29.6</u>	<u>31.1</u>	<u>36.3</u>
*Sharma <i>et al.</i> [40] ($M = 200$)	<u>30.6</u>	<u>34.6</u>	35.7	36.4	41.2	43.6	31.8	<u>31.5</u>	46.2	<u>49.7</u>	39.7	35.8	39.6	29.7	32.8	37.3
Ours (z_0) ($M = 1$)	37.8	41.7	42.1	41.8	46.5	50.2	38.0	39.2	51.7	61.8	45.4	42.6	45.7	33.7	38.5	43.8
Ours ($M = 200$)	27.9	31.4	29.7	30.2	34.9	37.1	27.3	28.2	39.0	46.1	34.2	32.3	33.6	26.1	27.5	32.4

Table 1. Detailed results of MPJPE in millimetres on Human3.6M under Protocol 1 (no rigid alignment) and Protocol 2 (rigid alignment). Our model achieves state-of-the-art results, outperforming all other methods in nearly every activity. All scores are taken from the referenced papers, except the row marked with * which is computed using the publicly available official code and model from [40]. The number of samples estimated by the respective approaches is denoted as M .

Method	MPJPE↓	PMPJPE↓	PCK↑	CPS↑
Li <i>et al.</i> [28]	81.1	66.0	85.7	119.9
Sharma <i>et al.</i> [40]	<u>78.3</u>	<u>61.1</u>	<u>88.5</u>	<u>136.4</u>
Ours	71.0	54.2	93.4	171.0

Table 2. Evaluation results on the subset H36MA containing highly ambiguous examples. For each metric, the best hypothesis score is reported.

outperform every competitor and achieve a clear improvement of 4.1% and 10.7% over the previous best scores under Protocol 1 and Protocol 2. Note that Li *et al.* [28] only show detailed results for $M = 5$, but state that their model performance does not significantly improve when increasing M . We generated the numbers for [40] under Protocol 2 (row marked with *) using their publicly available model, code and data, because they only report scores for the PMPJPE on subject 11. Outperforming the single prediction baseline of [32] with z_0 generated poses (*i.e.* $M = 1$) shows that our model is additionally able to give strong single predictions.

To evaluate the performance on highly ambiguous examples, we compute results for the challenging subset H36MA. We use the publicly available code from [40] and [28] to compare with their approaches. As is shown in Table 2, we outperform both competitors significantly and by a larger margin than on the whole test set. This emphasizes the ability of our model to generate diverse hypotheses for highly ambiguous examples. We argue that the CPS is especially meaningful in this setting, since high individual joint errors that often occur for challenging poses cannot be averaged out as in *e.g.* MPJPE or PCK.

4.4. Transfer to MPI-INF-3DHP

To assess the generalization capability of our model, we evaluate on MPI-INF-3DHP. Note that neither the 2D detector nor the normalizing flow is trained on this dataset. The results are shown in Table 3. Even though [28] use the ground-truth 2D joints provided by the dataset, we clearly outperform them in all three settings. We also achieve com-

Method	Studio GS	Studio no GS	Outdoor	All PCK
Li <i>et al.</i> [29]	86.9	86.6	<u>79.3</u>	85.0
Li <i>et al.</i> [28]	70.1	68.2	66.6	67.9
Ours	<u>86.6</u>	<u>82.8</u>	82.5	<u>84.3</u>

Table 3. Quantitative results on MPI-INF-3DHP. We outperform the approach from [28] by a large margin which even uses ground truth 2D joint positions. Note that [29] is trained weakly supervised and therefore specifically built for transfer learning. However, we still achieve on par results and even outperform them in the challenging outdoor sequences.

petitive performance compared to the weakly supervised approach from [29] that focuses on transfer learning. Our strong results for outdoor scenes further emphasize the generalization capability to different settings.

4.5. Sample Diversity

Heatmap Variance: We visually inspect the distribution of generated joint locations and compare them with the corresponding fitted Gaussians in Fig. 3. For visualization purposes, only three hypotheses are shown for all joints except the one with the highest uncertainty. As can be seen, the uncertainties of the 2D detector are reflected in the 3D hypotheses.

Depth Ambiguities: Even though variance in the depth direction is not explicitly optimized, our model learns to generate feasible hypotheses with varying depth. In fact, the standard deviation of the hypotheses averaged over all joints in the test set of Human3.6M is highest in the depth direction with 42.4 mm, compared to 18.3 mm and 17.3 mm in the x - and y -directions. Ankle, elbow and wrist joints account for the highest amount of variance. A visual example of meaningful depth diversity is given in Fig. 4.

Sample Set Size and Noise Baseline: In Fig. 5, we plot the MPJPE on the subset H36MA with increasing number of samples. The best hypothesis performance of our model continues to improve significantly, further enlarging the gap to [40]. To validate that our approach is superior to directly




Figure 3. The uncertainties of the 2D detector are successfully reflected in the 3D pose hypotheses. For visualization purposes, we only show the fitted Gaussian and a high number of hypotheses for the joint with highest uncertainty.

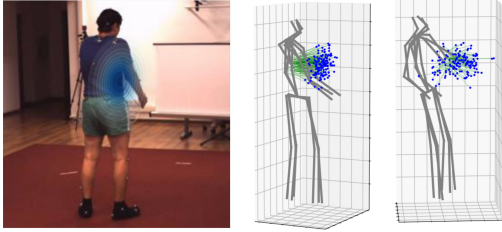


Figure 4. Depth ambiguities can be modeled together with the uncertainties of the 2D detector.

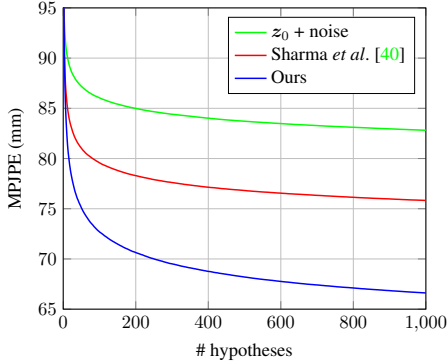


Figure 5. Evaluation results on the subset H36MA for an increasing number of hypotheses. Our model further improves and enlarges the gap to [40] and to a noise baseline.

sampling from the fitted Gaussians, we also plot results for a sampling baseline. The baseline is constructed by adding noise sampled from the fitted Gaussians to each joint of the z_0 predictions. A constant Gaussian for the depth dimension is assumed. Evidently, the performance of this baseline saturates earlier and at a higher error.

4.6. Ablation Studies

To quantify the influence of our proposed components and loss functions, we remove them individually and show the results in Table 4. As can be seen, the removal of each component leads to a degradation in performance. When removing the heatmap condition, a large drop in the CPS can be observed. This shows that some individual joints cannot be reconstructed without the uncertainty information of the 2D detector. Providing the condition alone already leads to a significant improvement of the CPS, indicating that the network can automatically leverage the informa-

Method	MPJPE↓	PMPJPE↓	PCK↑	CPS↑
w/o condition	71.7	57.2	91.2	137.6
w/o \mathcal{L}_{HM}	72.4	56.2	92.2	157.3
w/o \mathcal{L}_{gen}	73.6	58.2	92.5	165.5
w/o \mathcal{L}_{MB}	76.0	58.5	91.6	161.4
$\mathcal{L}_{\text{MB}}(k=1)$	71.8	54.9	92.6	167.4
$\mathcal{L}_{\text{MB}}(k=10)$	70.7	54.9	93.3	168.3
$\mathcal{L}_{\text{MB}}(k=50)$	71.0	55.2	93.1	168.0
Ours (Full)	<u>71.0</u>	54.2	93.4	171.0

Table 4. Ablation studies on the subset H36MA.

tion to model ambiguities. Adding \mathcal{L}_{HM} further improves all metrics. The importance of the discriminator becomes especially evident when considering the worst hypothesis error instead of the best. For example, Protocol 2 computed for the worst hypothesis deteriorates from 86.8 mm to 284.1 mm without the discriminator. Thus, the adversarial training procedure ensures the feasibility of the generated poses. Table 4 also shows the influence of the number of best hypotheses k selected for computing \mathcal{L}_{MB} . Note that \mathcal{L}_{MB} with $k=1$ is equivalent to the typical *best-of-M* loss.

5. Conclusion

This paper presents a normalizing flow based method for the ambiguous inverse problem of 3D human pose estimation from 2D inputs. We exploit the bijectivity of the normalizing flow by utilizing the known 3D to 2D projection during training. By incorporating uncertainty information from the heatmaps of a 2D pose detector, valuable information is maintained which is discarded by previous approaches. As demonstrated, the generated hypotheses reflect these uncertainties and additionally show meaningful diversity along the ambiguous depth of the joints. Furthermore, the introduction of a 3D pose discriminator ensures the geometrical feasibility of the poses and a proposed generalization of the *best-of-M* loss improves the performance. Experimental results show that our method outperforms all previous multi-hypotheses approaches in most metrics, especially on a challenging subset of Human3.6M containing highly ambiguous examples.

Acknowledgements. This work has been supported by the Federal Ministry of Education and Research (BMBF), Germany, under the project LeibnizKILabor (grant no. 01DD20003), the Center for Digital Innovations (ZDIN) and the Deutsche Forschungsgemeinschaft (DFG) under Germany’s Excellence Strategy within the Cluster of Excellence PhoenixD (EXC 2122).

References

- [1] Ijaz Akhter and Michael J. Black. Pose-conditioned joint angle limits for 3d human pose reconstruction. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 2
- [2] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 6
- [3] Lynton Ardizzone, Jakob Kruse, Sebastian J. Wirkert, Daniel Rahner, Eric W. Pellegrini, Ralf S. Klessen, Lena Maier-Hein, Carsten Rother, and Ullrich Köthe. Analyzing inverse problems with invertible neural networks. *International Conference on Learning Representations (ICLR)*, 2019. 2, 3, 5, 6
- [4] Lynton Ardizzone, Carsten Lüth, Jakob Kruse, Carsten Rother, and Ullrich Köthe. Guided image generation with conditional invertible neural networks. *arXiv preprint arXiv:1907.02392*, 2019. 2, 4
- [5] Benjamin Biggs, Sébastien Ehrhardt, Hanbyul Joo, Benjamin Graham, Andrea Vedaldi, and David Novotny. 3d multi-bodies: Fitting sets of plausible 3d human models to ambiguous image data. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2, 4
- [6] Christopher M. Bishop. Mixture density networks. Technical report, Aston University, 1994. 3
- [7] 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 (ECCV)*, 2016. 2
- [8] Ching-Hang Chen and Deva Ramanan. 3d human pose estimation = 2d pose estimation + matching. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [9] Hai Ci, Chunyu Wang, Xiaoxuan Ma, and Yizhou Wang. Optimizing network structure for 3d human pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019. 2
- [10] Laurent Dinh, David Krueger, and Yoshua Bengio. NICE: non-linear independent components estimation. In *International Conference on Learning Representations (ICLR)*, 2015. 3
- [11] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real NVP. In *International Conference on Learning Representations (ICLR)*, 2017. 3, 4
- [12] Haoshu 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*, 2018. 2
- [13] Mathieu Germain, Karol Gregor, Iain Murray, and Hugo Larochelle. Made: Masked autoencoder for distribution estimation. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2015. 3
- [14] Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research (JMLR)*, 13(25), 2012. 5
- [15] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 5, 6
- [16] Abner Guzmán-rivera, Dhruv Batra, and Pushmeet Kohli. Multiple choice learning: Learning to produce multiple structured outputs. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2012. 2, 5
- [17] 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 *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [18] Mir Rayat Imtiaz Hossain and James J. Little. Exploiting temporal information for 3d pose estimation. *European Conference on Computer Vision (ECCV)*, 2018. 2
- [19] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 36(7), 2014. 2, 6
- [20] Ehsan Jahangiri and Alan L. Yuille. Generating multiple diverse hypotheses for human 3d pose consistent with 2d joint detections. *International Conference on Computer Vision Workshops (ICCVW)*, 2017. 1, 2, 3
- [21] Angjoo Kanazawa, Michael J. Black, David W. Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [22] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015. 6
- [23] Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved variational inference with inverse autoregressive flow. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2016. 3
- [24] Ivan Kobyzev, Simon Prince, and Marcus Brubaker. Normalizing flows: An introduction and review of current methods. *TPAMI*, 2020. 3
- [25] Muhammed Kocabas, Nikos Athanasiou, and Michael J. Black. Vibe: Video inference for human body pose and shape estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 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. 2
- [27] Mun Wai Lee and Isaac Cohen. Proposal maps driven mcmc for estimating human body pose in static images. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2004. 2
- [28] 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)*, 2019. 1, 3, 7

- [29] Chen Li and Gim Hee Lee. Weakly supervised generative network for multiple 3d human pose hypotheses. *British Machine Vision Conference (BMVC)*, 2020. 1, 7
- [30] Jiefeng Li, Chao Xu, Zhicun Chen, Siyuan Bian, Lixin Yang, and Cewu Lu. Hybrik: A hybrid analytical-neural inverse kinematics solution for 3d human pose and shape estimation. In *CVPR*, 2021. 2
- [31] Shichao Li, Lei Ke, Kevin Pratama, Yu-Wing Tai, Chi-Keung Tang, and Kwang-Ting Cheng. Cascaded deep monocular 3d human pose estimation with evolutionary training data. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [32] Julieta Martinez, Rayat Hossain, Javier Romero, and James J. Little. A simple yet effective baseline for 3d human pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. 2, 7
- [33] 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 *International Conference on 3D Vision (3DV)*, 2017. 2, 6
- [34] Francesc Moreno-Noguer. 3d human pose estimation from a single image via distance matrix regression. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [35] Tuomas P. Oikarinen, Daniel C. Hannah, and Sohrob Kazerounian. Graphmdn: Leveraging graph structure and deep learning to solve inverse problems. *arXiv preprint arXiv:2010.13668*, 2020. 1, 3, 7
- [36] George Papamakarios, Theo Pavlakou, and Iain Murray. Masked autoregressive flow for density estimation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 3
- [37] Georgios Pavlakos, Luyang Zhu, Xiaowei Zhou, and Kostas Daniilidis. Vibe: Video inference for human body pose and shape estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2
- [38] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. *Proceedings of Machine Learning Research (PMLR)*, 2015. 2, 3
- [39] Marco Rudolph, Bastian Wandt, and Bodo Rosenhahn. Same same but different: Semi-supervised defect detection with normalizing flows. In *Winter Conference on Applications of Computer Vision (WACV)*, 2021. 3
- [40] Saurabh Sharma, Pavan Teja Varigonda, Prashast Bindal, Abhishek Sharma, and Arjun Jain. Monocular 3d human pose estimation by generation and ordinal ranking. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019. 1, 3, 6, 7, 8
- [41] Soshi Shimada, Vladislav Golyanik, Weipeng Xu, and Christian Theobalt. Physcap: Physically plausible monocular 3d motion capture in real time. *ACM Transactions on Graphics*, 39(6), 2020. 2
- [42] E. Simo-Serra, A. Ramisa, G. Alenyà, C. Torras, and F. Moreno-Noguer. Single image 3d human pose estimation from noisy observations. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012. 2
- [43] Cristian Sminchisescu and Bill Triggs. Covariance scaled sampling for monocular 3d body tracking. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2001. 2
- [44] Cristian Sminchisescu and Bill Triggs. Kinematic jump processes for monocular 3d human tracking. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2003. 2
- [45] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 4, 6
- [46] Esteban G. Tabak and Cristina V. Turner. A family of non-parametric density estimation algorithms. *Communications on Pure and Applied Mathematics*, 66(2), 2013. 2, 3
- [47] Esteban G. Tabak and Eric Vanden-Eijnden. Density estimation by dual ascent of the log-likelihood. *Communications in Mathematical Sciences*, 8(1), 2010. 2, 3
- [48] Bastian Wandt, Hanno Ackermann, and Bodo Rosenhahn. A kinematic chain space for monocular motion capture. In *European Conference on Computer Vision Workshops (ECCVW)*, 2018. 5
- [49] Bastian Wandt and Bodo Rosenhahn. Repnet: Weakly supervised training of an adversarial reprojection network for 3d human pose estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 5
- [50] Bastian Wandt, Marco Rudolph, Petriッサ Zell, Helge Rhodin, and Bodo Rosenhahn. Canonpose: Self-supervised monocular 3d human pose estimation in the wild. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 6
- [51] Jue Wang, Shaoli Huang, Xinchao Wang, and Dacheng Tao. Not all parts are created equal: 3d pose estimation by modelling bi-directional dependencies of body parts. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019. 2
- [52] Christina Winkler, Daniel Worrall, Emiel Hoogeboom, and Max Welling. Learning likelihoods with conditional normalizing flows. *arXiv preprint arXiv:1912.00042*, 2019. 2, 4
- [53] 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 *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 4
- [54] Jingwei Xu, Zhenbo Yu, Bingbing Ni, Jiancheng Yang, Xiaokang Yang, and Wenjun Zhang. Deep kinematics analysis for monocular 3d human pose estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [55] Andrei Zanfir, Eduard Gabriel Bazavan, Hongyi Xu, Bill Freeman, Rahul Sukthankar, and Cristian Sminchisescu. Weakly supervised 3d human pose and shape reconstruction with normalizing flows. *European Conference on Computer Vision (ECCV)*, 2020. 2, 4
- [56] Long Zhao, Xi Peng, Yu Tian, Mubbasir Kapadia, and Dimitris N. Metaxas. Semantic graph convolutional networks for 3d human pose regression. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3