

AdaptPose: Cross-Dataset Adaptation for 3D Human Pose Estimation by Learnable Motion Generation

Mohsen Gholami, Bastian Wandt, Helge Rhodin, Rabab Ward, and Z. Jane Wang
 University of British Columbia

{mgholami, rababw, zjanew}@ece.ubc.ca {wandt, rhodin}@cs.ubc.ca

Abstract

This paper addresses the problem of cross-dataset generalization of 3D human pose estimation models. Testing a pre-trained 3D pose estimator on a new dataset results in a major performance drop. Previous methods have mainly addressed this problem by improving the diversity of the training data. We argue that diversity alone is not sufficient and that the characteristics of the training data need to be adapted to those of the new dataset such as camera viewpoint, position, human actions, and body size. To this end, we propose AdaptPose, an end-to-end framework that generates synthetic 3D human motions from a source dataset and uses them to fine-tune a 3D pose estimator. AdaptPose follows an adversarial training scheme. From a source 3D pose the generator generates a sequence of 3D poses and a camera orientation that is used to project the generated poses to a novel view. Without any 3D labels or camera information AdaptPose successfully learns to create synthetic 3D poses from the target dataset while only being trained on 2D poses. In experiments on the Human3.6M, MPI-INF-3DHP, 3DPW, and Ski-Pose datasets our method outperforms previous work in cross-dataset evaluations by 14% and previous semi-supervised learning methods that use partial 3D annotations by 16%.

1. Introduction

Monocular 3D human pose estimation aims to reconstruct the 3D skeleton of the human body from 2D images. Due to pose and depth ambiguities, it is well known to be an inherently ill-posed problem. However, deep learning models are able to learn 2D to 3D correspondences and achieve impressively accurate results when trained and tested on similar datasets [1, 3, 6, 14, 27, 31, 32].

An often disregarded aspect is that the distribution of features in a dataset e.g. camera orientation and body poses differ from one dataset to another. Therefore, a pre-trained network underperforms when applied to images captured




Figure 1. AdaptPose generates synthetic motions to improve the cross-dataset generalization. The source dataset has 3D labels and camera information, while the target dataset has only sample videos. The synthetic motions are generated to belong to the target dataset. Therefore fine-tuning the 3D pose estimator with synthetic motions improves the generalization of the model.

from a different viewpoint or when they contain an activity that is not present in the training dataset [42, 45]. As an example, Figure 1 shows images from the Human3.6M [15] dataset on the left and images from the Ski-Pose [33, 35] dataset on the right which we define as *source domain* and *target domain*, respectively. Camera viewpoint, position, human action, speed of motion, and body size significantly differ between the source and target domain. This large domain gap causes 3D pose estimation models trained on the source domain to make unreliable predictions for the target domain [42, 45, 46]. We address this problem by generating synthetic 3D data that lies within the distribution of the target domain and fine-tuning the pose estimation network by the generated synthetic data. Our method does not require 3D labels or camera information from the target domain and is only trained on sample videos from the target domain.

To the best of our knowledge, there are only two approaches that generate synthetic 2D-3D human poses for cross-dataset generalization of 3D human pose estimators [13, 23]. Li *et al.* [23] randomly generate new 2D-3D pairs of the source dataset by substituting parts of the human body in 3D space and projecting the new 3D pose to

2D. PoseAug [11] proposes a differential data augmentation framework that is trained along with a pose estimator. Both, [23] and [11], merely improve the diversity of the source domain without considering the distribution of the target domain. Moreover, these methods are based on single images and do not consider temporal information.

We formulate the data augmentation process as a domain adaptation problem. Figure 2 shows our training pipeline. Our goal is to generate plausible synthetic 2D-3D pairs that lie within the distribution of the target domain. Our framework, *AdaptPose*, introduces a human motion generator network that takes 3D samples from the source dataset and modifies them by a learned deformation to generate a sequence of new 3D samples. We project the generated 3D samples to 2D and feed them to a domain discriminator network. The domain discriminator is trained with real 2D samples from the target dataset and fake samples from the generator. We use the generated samples to fine-tune a pose estimation network. Therefore, our network adapts to any target using only images from the target dataset. 3D annotation from the target domain is not required. Unlike [13, 23], this enables our network to generate plausible 3D poses from the target domain. Another contribution is the extension of the camera viewpoint generation from a deterministic approach to a probabilistic approach. We assume that the camera viewpoint of the target domain comes from a specific well-defined, but unknown distribution. Therefore, we propose to learn a distribution of camera viewpoints instead of learning to generate a deterministic rotation matrix. Our network rotates the generated 3D poses into a random camera coordinate system within the learned distribution. The generated sample is a sequence of 2D-3D pose pairs that entails plausibility in the temporal domain. We believe that the application of the proposed motion generator is not limited to improving only cross-dataset performance of 3D pose estimation, but it could also be used in other tasks such as human action recognition.

Contributions. 1) we propose to close the domain gap between the training and test datasets by a kinematics-aware domain discriminator. The domain discriminator is trained along with a *human motion generator* (HMG) that uses a source training dataset to generate human motions close to those in the target dataset. 2) We show that learning the distribution of the camera viewpoint is more effective than learning to generate a deterministic camera matrix. 3) To the best of our knowledge, this is the first approach that proposes generating human motions specifically for cross-dataset generalization for 3D human pose estimation, unlike previous work that focuses on single-frame data augmentation.

2. Related Work

In the following, we discuss the related work with a focus on cross-dataset adaptation.

Weakly-supervised Learning. Weakly supervised learning has been proposed to diminish the dependency of networks on 3D annotations. These methods rely on unpaired 3D annotation [21, 39, 43], multi-view images [10, 16, 19, 33, 40], or cycle-consistency [4, 9]. Most related to our work is the adaptation of a network to the target domain via weakly supervised learning. Zhang *et al.* [46] propose an online adaptation to target test data based on the weakly supervised learning method of [4]. Yang *et al.* [43] use unpaired 3D annotation to further fine-tune a network on in-the-wild images. Kundu *et al.* [22] use a self-supervised learning method to improve the generalization of a pre-trained network on images with occlusion.

Cross-dataset Generalization. Cross-dataset adaptation of 3D pose estimators has recently gained attention. Guan *et al.* and Zhang *et al.* [13, 46] propose an online adaptation of the pose estimator during the inference stage over test data. Guan *et al.* [13] use a temporal consistency loss and a 2D projection loss on the streaming test data to adapt the network to the target test dataset. Zhang *et al.* [46] use a cycle consistency approach to optimize the network on every single test frame. Although the online-adaptation approach improves cross-dataset generalizability, it also increases the inference time, especially if the networks exploit temporal information. Wang *et al.* [42] argues that estimating the camera viewpoint beside the 3D keypoints improves cross-dataset generalization of the 3D pose estimator. However, the camera viewpoint is not the only criterion that differs between datasets. Split-and-Recombine [45] proposes to split the human skeleton into different body parts so that different body parts of a rare pose from the target dataset could have been seen in the source dataset.

Data Augmentation. Data augmentation is another way to diminish cross-dataset errors. Previous methods perform data augmentation on images [34], 3D mesh models [5, 36, 47], or 2D-3D pairs [7, 13, 23]. Most related to our work is augmenting 2D-3D pairs. Li *et al.* [23] generate synthetic 3D human samples by substituting body parts from a source training set. The evolutionary process of [23] is successful in generating new poses, however, the generation of natural camera viewpoints is overlooked. Instead, it randomly perturbs source camera poses. PoseAug [11] proposes an end-to-end data augmentation framework that trains along with a pose estimator network. Although it improves the diversity of the training data, there is no guarantee that the generated samples are in the distribution of the target dataset. Moreover, according to the ablation studies of PoseAug, the main improvement comes from generating camera viewpoints instead of generating new poses. This means that PoseAug has limited abilities to effectively im-




Figure 2. Overview of the proposed network. The input is a vector of 3D keypoints from the source dataset concatenated with Gaussian noise. The motion generator learns to generate a sequence of 3D keypoints X_{3D}^b and the mean and standard deviation of a normal distribution \mathcal{N} . A random rotation matrix is sampled from the learned normal distribution and X_{3D}^b is transformed to X_{3D}^r and projected to 2D. The domain discriminator is trained with X_{2D}^r and 2D keypoints from the target domain. The lifting network is a pretrained pose estimator that estimates 3D from 2D. It is used to evaluate X_{2D}^r, X_{3D}^r , provide feedback to the motion generator, and to select a subset of samples for fine-tuning the lifting network. The pipeline is trained end-to-end.

prove pose diversities in the training set. In contrast, we enforce the generated synthetic data to be in the distribution of the target data. Unlike PoseAug, we show that our motion generation network significantly improves cross-dataset results even without augmenting the camera-viewpoints.

3. Problem Formulation

Let $\mathbf{X}^{\text{src}} = (X_{2D}^{\text{src}}, X_{3D}^{\text{src}})$ be a pair of 2D and 3D poses from the source dataset and $\mathbf{X}^{\text{tar}} = X_{2D}^{\text{tar}}$ a 2D pose from the target dataset. The input to our model are sequences of frames with length n , $X_{2D}^{\text{src}} : [x_{2D}]_{t=0}^n$, $X_{3D}^{\text{src}} : [x_{3D}]_{t=0}^n$, and $X_{2D}^{\text{tar}} : [y_{2D}]_{t=0}^n$ where $x_{2D}, y_{2D} \in \mathbb{R}^{J \times 3}$. AdaptPose consists of a generator function

$$G(\mathbf{X}^{\text{src}}, \mathbf{z}; \theta_G) \rightarrow \mathbf{X}^{\text{fake}}, \quad (1)$$

with parameters θ_G , that maps source samples \mathbf{X}^{src} and a noise vector $\mathbf{z} \sim p_z$ to a fake 2D-3D pair $\mathbf{X}^{\text{fake}} = (X_{2D}^{\text{fake}}, X_{3D}^{\text{fake}})$. The fake samples $(X_{2D}^{\text{fake}}, X_{3D}^{\text{fake}})$ are a sequence of 2D-3D keypoints $X_{2D}^{\text{fake}} : [x_{2D}^{\text{fake}}]_{t=0}^n$, $X_{3D}^{\text{fake}} : [x_{3D}^{\text{fake}}]_{t=0}^n$. The generator G generates an adapted dataset $\mathbf{X}^{\text{fake}} = G(\mathbf{X}^{\text{src}}, \mathbf{z})$ of any desired size. In order to adapt the source to the target domain in the absence of 3D target poses we introduce a domain discriminator D_D and a 3D discriminator D_{3D} . The domain discriminator $D_D(\mathbf{x}; \theta_D)$ gives the likelihood d that the 2D input \mathbf{x} is sampled from the target domain X_{2D}^{tar} . The generator tries to generate fake samples X_{2D}^{fake} as close as possible to target samples X_{2D}^{tar} , while the discriminator tries to distinguish between them. Unlike a standard GAN network [12] where generator is conditioned only on a noise vector, our generator is conditioned on both a noise vector and a sample from the source dataset which was shown to be effective in generating synthetic images [2]. Additionally, the model is conditioned

on a 3D discriminator $D_{3D}(\mathbf{x}; \theta_D)$ that outputs the likelihood d' that the generated 3D, X_{3D}^{fake} , is sampled from the real 3D distribution. Ideally, we would like to condition on the target 3D dataset. Since 3D data from the target domain is not available we condition it on the source 3D dataset. However, conditioning the 3D discriminator D_{3D} directly on the source 3D poses restrains the motion generator to the source distribution. Instead, we condition the 3D discriminator D_{3D} on a perturbed version of data $X_{3D}^{\text{psrc}} = \mathbf{y} + X_{3D}^{\text{src}}$ where $\mathbf{y} \sim p_y$ is a small noise vector. The noise vector \mathbf{y} is selected such that X_{3D}^{psrc} is a valid pose from the source distribution. The goal of AdaptPose is to optimize the following objective function

$$\mathcal{L} = \min_{\theta_G} \max_{(\theta_{D_D}, \theta_{D_{3D}})} \alpha L(G, D_D) + \beta L(G, D_{3D}), \quad (2)$$

where α and β are the weights of the losses.

4. Human Motion Generator

We name the generator of our GAN network *Human Motion Generator* (HMG). The HMG consists of two main components. 1) A *bone generator* that rotates the bone vectors and changes the bone length ratios. The bone generation operation produces new 3D keypoints X_{3D}^b . 2) A *camera generator* that generates a new camera viewpoint $\{\mathbf{R}, \mathbf{T}\}$, where $\mathbf{R} \in \mathbb{R}^{3 \times 3}$ is a rotation matrix and \mathbf{T} is a translation vector. X_{3D}^b is transformed to the generated camera viewpoint by

$$X_{3D}^{\text{fake}} = \mathbf{R}X_{3D}^b + \mathbf{T}, \quad (3)$$

with the corresponding 2D keypoints

$$X_{2D}^{\text{fake}} = \Pi(X_{3D}^{\text{fake}}), \quad (4)$$

where Π is the perspective projection that uses the intrinsic parameters from the source dataset.

4.1. Bone Generation

In this section, we analyze different methods of bone vector generation in the temporal domain. The main challenge is to keep the bone changes plausible for every single frame and temporally consistent in the time domain. We propose and analyze the three different methods BG1, BG2, and BG3 shown in Figure 3.

BG1. The bone generation network accepts a sequence of 3D keypoints from the source dataset. The sequence of 3D keypoints is transformed into a bone vector representation $[\vec{B}_t^{\text{src}}]_{t=t_0}^{t_0+n}$ where $\vec{B}_t^{\text{src}} \in \mathbb{R}^{(J-1) \times 3}$ and J is the number of keypoints. BG1 generates a displacement vector $\Delta\vec{B} \in \mathbb{R}^{(J-1) \times 3}$ and a bone ratio $\lambda \in \mathbb{R}^{(J-1) \times 1}$. The new bone vector is $[\vec{B}_t^{\text{fake}}]_{t=t_0}^{t_0+n}$ where

$$\vec{B}_t^{\text{fake}} = \frac{\vec{B}_t^{\text{src}} + \Delta\vec{B}}{\|\vec{B}_t^{\text{src}} + \Delta\vec{B}\|} \|\vec{B}_t^{\text{src}}\| (1 + \lambda). \quad (5)$$

$\Delta\vec{B}$ may change the bone length instead of rotating to a new configuration as shown in Figure 3. To avoid this, we divide the generated bones by $\|\vec{B}_t^{\text{src}} + \Delta\vec{B}\|$ in Eq. 5.

BG2. The bone generation network accepts a single sample of 3D keypoints from the source dataset and converts it to a bone representation $\vec{B}_{t_0}^{\text{src}}$. BG2 generates $\Delta\vec{B}$ and λ . The new bone vector is $[\vec{B}_t^{\text{fake}}]_{t=t_0}^{t_0+n}$ where

$$\vec{B}_{t+j}^{\text{fake}} = \frac{\vec{B}_{t_0}^{\text{src}} + j\Delta\vec{B}/n}{\|\vec{B}_{t_0}^{\text{src}} + j\Delta\vec{B}/n\|} \|\vec{B}_{t_0}^{\text{src}}\| (1 + \lambda). \quad (6)$$

BG3. The bone generation network generates the vector $\vec{r} \in \mathbb{R}^{(J-1) \times 3}$ and the angle $\theta \in \mathbb{R}^{(J-1) \times 1}$. A sequence of rotation matrices $[R_t]_{t=0}^n$ is calculated by

$$R_{t+j} = \mathcal{H}\left(\frac{\vec{r}}{\|\vec{r}\|} \frac{j\theta}{n}\right), \quad (7)$$

where \mathcal{H} transforms axis-angle rotation of (θ, \vec{r}) to rotation matrix representation via quaternions $q = q_r + q_x \mathbf{i} + q_y \mathbf{j} + q_z \mathbf{k}$ by

$$q = \cos\left(\frac{\theta}{2}\right) + \frac{\vec{r}}{\|\vec{r}\|} \sin\left(\frac{\theta}{2}\right), \quad (8)$$

$$R = v \otimes v + q_r^2 \mathbf{I} + 2q_r [v]_{\times} + [v]_{\times}^2, \quad (9)$$

where \otimes is the outer product, \mathbf{I} is the identity matrix, and

$$[v]_{\times} = \begin{bmatrix} 0 & -v_3 & v_2 \\ v_3 & 0 & -v_1 \\ -v_2 & v_1 & 0 \end{bmatrix}. \quad (10)$$




Figure 3. Bone generation methods. Blue vectors indicate bone vectors before rotation and green vectors are bone vectors after rotation. $\Delta\vec{B}$ is rotating bone direction produced by the network. \vec{r} and θ are the axis and angle of the rotation, respectively.

4.2. Camera Generation

In this section, we introduce two different methods of camera generation: 1) *Deterministic*, which generates a single camera rotation matrix and translation and 2) *probabilistic*. The network learns a distribution of rotation matrices. A random rotation matrix is sampled from the learned distribution. Additionally, we explore three different rotation representations: axis-angle, Euler-angles, and quaternions. In the following, we will discuss each of the procedures for each of the rotation representations.

Deterministic Axis-angle. The network generates an axis \vec{r} and a translation T where the angle of rotation is $\|\vec{r}\|$. The rotation matrix $R \in \mathbb{R}^{3 \times 3}$ is produced by $R = \mathcal{H}(\vec{r})$ where \mathcal{H} is explained in the equation 8.

Probabilistic Axis-angle. The network learns three separate normal distributions $\mathcal{N}_1(\mu_1, \sigma_1)$, $\mathcal{N}_2(\mu_2, \sigma_2)$, and $\mathcal{N}_3(\mu_3, \sigma_3)$, an angle θ , and a translation T . The axis $r = \{r_1, r_2, r_3\}$ is sampled from the learned normal distributions and converted to a rotation matrix by

$$R = \mathcal{H}\left(\frac{\vec{r}}{\|\vec{r}\|} \theta\right). \quad (11)$$

Probabilistic Euler-angles. The network learns three Gaussian distributions $\mathcal{N}_1, \mathcal{N}_2$, and \mathcal{N}_3 to sample the Euler-angles (α, β, γ) from the specified distributions. The rotation matrix is obtained as follows:

$$R = R_z(\alpha)R_y(\beta)R_x(\lambda), \quad (12)$$

where $R_z(\alpha)$, $R_y(\beta)$, and $R_x(\lambda)$ are rotations of $(\alpha, \beta$, and $\gamma)$ degrees around z, y, x axis, respectively.

Probabilistic Quaternion. A quaternion represents a rotation around axis $\vec{u} = (u_x, u_y, u_z)$ with angle θ as

$$q = \cos\left(\frac{\theta}{2}\right) + \vec{u} \sin\left(\frac{\theta}{2}\right). \quad (13)$$

Therefore, q can be represented by four elements. Our network learns four distributions $\mathcal{N}_{1, \dots, 4}$ and randomly samples elements of q from the distributions. The quaternion q

is then converted to a rotation matrix representation as explained in section 4.1.

4.3. Domain and 3D Discriminators

We adopt the kinematic chain space (KCS) [38, 39] in 2D space to generate a matrix of joint angles and limb lengths in the image plane. The domain discriminator has two branches that accept 2D keypoints and the KCS matrix, respectively. The diagonal of the KCS matrix contains the limb lengths in the image space. Other components of the KCS matrix represent angular relationships of the 2D pose. It is important to mention that we do not normalize input 2D keypoints relative to the root joint as it causes perspective ambiguities [44]. Therefore, $diag(KCS)$ is a function of position and body scale. On the contrary $KCS - diag(KCS)$ is a function of the camera viewpoint and scale of the person. Thus, the KCS matrix disentangles different parameters that the motion generator requires to learn. For the 3D discriminator, in order not to condition the 3D discriminator on the source domain, we first apply a random perturbation of β degrees to the input bone vectors $\beta < 10^\circ$ and then feed the perturbed 3D to a part-wise KCS branch [11] and the original 3D to a KCS branch. Further details about the 3D discriminator are provided in the supplementary material.

4.4. Selection

In order to stabilize the training of the lifting network we introduce a selection step by evaluating samples via the lifting network N . In this step, the lifting network receives $(X_{2D}^{src}, X_{3D}^{src})$ and $(X_{2D}^{fake}, X_{3D}^{fake})$ which are source and generated samples, respectively. We exclude samples that are either too simple or too hard using the following rule

$$\text{selection} = \begin{cases} \text{yes} & \text{if } (\frac{\mathcal{L}(N(X_{2D}^{fake}))}{\mathcal{L}(N(X_{2D}^{src}))} - a)^2 < b^2 \\ \text{no} & \text{otherwise} \end{cases}, \quad (14)$$

where \mathcal{L} is an L_2 loss.

5. Training

In each epoch we generate 1.5 million synthetic samples followed by fine-tuning the lifting network.

Motion Generator. Our adversarial framework is trained using three losses for the motion generator and for the discriminators which are defined as

$$\mathcal{L}_{\mathcal{D}_{3D}} = \frac{1}{2} \mathbb{E}[(\mathcal{D}(X_{3D}^{src}) - 1)^2] + \frac{1}{2} \mathbb{E}[\mathcal{D}(X_{3D}^{fake})^2], \quad (15)$$

$$\mathcal{L}_{\mathcal{D}_D} = \frac{1}{2} \mathbb{E}[(\mathcal{D}(X_{2D}^{tar}) - 1)^2] + \frac{1}{2} \mathbb{E}[\mathcal{D}(X_{2D}^{fake})^2], \quad (16)$$

$$\mathcal{L}_{G_{adv}} = \frac{1}{2} \mathbb{E}[(\mathcal{D}(X_{2D}^{fake}) - 1)^2], \quad (17)$$

where $(X_{3D}^{src}, X_{3D}^{fake})$ are 3D samples from the source dataset and synthetic generated samples, respectively.

$(X_{2D}^{tar}, X_{2D}^{fake})$ are 2D keypoints from the target dataset and the generated synthetic data, respectively. The generator also receives a feedback loss from the lifting network. The feedback loss has two components: 1) reprojection loss of the estimated 3D keypoints of the target domain 2) fixed hard ratio feedback loss adapted from [11]. The lifting network N accepts X_{2D}^{tar} from the target dataset and predicts X_{3D}^{tar} . We define the reprojection loss as

$$\mathcal{L}_{proj} = \left\| \frac{X_{proj}^{fake}}{\|X_{proj}^{fake}\|} - \frac{X_{2D}^{fake}}{\|X_{2D}^{fake}\|} \right\|_1, \quad (18)$$

where $\|\cdot\|_1$ is the L_1 norm and

$$X_{proj}^{fake} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} N(X_{2D}^{tar}). \quad (19)$$

The fixed hard ratio loss provides feedback depending on the difficulty of generated sample relative to the source samples as follows:

$$f = \left(\frac{\mathcal{L}(N(X_{2D}^{fake}))}{\mathcal{L}(N(X_{2D}^{src}))} - c \right)^2, \quad (20)$$

$$\mathcal{L}_{hr} = \begin{cases} f & \text{if } f < d^2 \\ 0 & \text{otherwise} \end{cases}, \quad (21)$$

where \mathcal{L} is L_2 loss. The summation of the above mentioned losses is our generator loss

$$\mathcal{L}_G = \mathcal{L}_{G_{adv}} + \mathcal{L}_{proj} + \mathcal{L}_{hr}. \quad (22)$$

Lifting Network. The lifting network N is trained using $(X_{2D}^{src}, X_{3D}^{src})$ and $(X_{2D}^{fake}, X_{3D}^{fake})$ which gives the lifting loss

$$\mathcal{L}_N = \|X_{3D}^{src} - N(X_{2D}^{src})\|_2 + \|X_{3D}^{fake} - N(X_{2D}^{fake})\|_2. \quad (23)$$

6. Experiments

We perform extensive experiments to evaluate the performance of AdaptPose for cross-dataset generalization. We further conduct ablation studies on the different elements of our network. In the following, we discuss different datasets and subsequently baselines and metrics.

- **Human3.6M (H3.6M)** contains 3D and 2D data from seven subjects captured in 50 fps. We use the training set of H3.6M (S1, S5, S6, S7, S8) as our source dataset for cross-dataset evaluations. While performing experiments on the H3.6M dataset itself we will use S1 as the source dataset and S5, S6, S7, and S8 as the target.
- **MPI-INF-3DHP (3DHP)** contains 3D and 2D data from 8 subjects and covers 8 different activities. We will use the 2D data from the training set of 3DHP [28]

as our target dataset when evaluating 3DHP. The test set of 3DHP includes more than 24K frames. However, some of the previous work use a subset of test data which includes 2,929 frames for evaluation [11, 20]. The 2,929 version has temporal inconsistency which is fine for the single-frame networks. We use the official test set of 3DHP and compare our results against the previous work’s results on the official test set of 3DHP for a fair comparison.

- **3DPW** contains 3D and 2D data captured in an outdoor environment. The camera is moving in some of the trials. 3DPW [37] is captured in 25fps and has more variability than 3DHP and H3.6M in terms of camera poses. We use the training set of 3DPW as our target dataset when experimenting on this dataset.
- **Ski-Pose PTZ-Camera (Ski)** includes 3D and 2D labels from 5 professional ski athletes in a ski resort. The dataset is captured in 30 fps and frames are cropped in 256×256 . The cameras are moving and there is a major domain gap between Ski and previous datasets in terms of the camera pose/position.

Evaluation Metrics. We use mean per joint position error (MPJPE) and Procrustes aligned MPJPE (P-MPJPE) as our main evaluation metrics. P-MPJPE measures MPJPE after performing Procrustes alignment of the predicted pose and the target pose. We also report the percentage of correct keypoint (PCK) with a threshold of 150 mm and area under the curve (AUC) for evaluation on 3DHP following previous arts.

Baseline (Lifting Network). We use VideoPose3D [32] (VPose3D) as the baseline pose estimator model. VPose3D is a lifting network that regresses 3D keypoints from input 2D keypoints. We use 27 frames as the input in our experiments. As preprocessing for H3.6M, 3DHP, and 3DPW datasets we normalize image coordinate such that $[0, w]$ is mapped to $[-1, 1]$. Note that the 3DPW dataset has some portrait frames with a height greater than width. In these cases, we pad the width so that height is equal to width to avoid the 2D keypoint coordinates being larger than the image frame after normalization. Our experiments show that this preprocessing has lower cross-dataset error compared with root centering and Frobenius normalization of 2D keypoints. While performing experiments on the Ski dataset we use root centering and Frobenius normalization of 2D keypoints since the image frames are already cropped to 256×256 with the person in the center of the image. Since there is an fps difference and also motion speed difference between our source dataset and target datasets, we also perform random downsampling in our data loader for training the baseline network. Specifically, our data loader samples $\{x_{r(t-n)}, \dots, x_{r(t+n)}\}$ from the source dataset and

r is a random number sampled from a uniform distribution of $[2, 5]$. Table 5 shows that the baseline model has a cross-dataset MPJPE of 96.4 mm using 3DHP as the target dataset.

6.1. Quantitative Evaluation

H3.6M. We compare our results with previous semi-supervised learning methods that only use 3D labels from S1 and 2D annotations from the remaining subjects for training [32] as well as data augmentation methods. Our results improve upon the previous state-of-the-art by 16%. We use ground truth 2D keypoints and therefore compare with previous work with the same setting. Since the camera pose does not change much between subjects, we hypothesize that the comparison in the current setting compares our bone generation method against previous work.

3DHP. Table 2 gives MPJPE, AUC, and PCK on test set of 3DHP. We report the results of PoseAug’s released pre-trained model on the complete test set of 3DHP. Our results have a 14% margin in terms of MPJPE compared with previous methods that report cross-dataset evaluation results [11, 13, 23, 42, 45]. This includes the comparison to [46] that uses information from the target test data to perform test-time optimization.

3DPW. Table 3 provides MPJPE and PA-MPJPE on the test set of 3DPW. Our method outperforms previous methods by 12 mm in PA-MPJPE. This includes previous methods that particularly were designed for cross-dataset generalization [8, 11, 13] and those that use temporal information [13, 19]. In comparison with test-time optimization methods [13, 46], ours also has an advantage of fast inference.

SKI. Table 4 gives the cross-dataset results on the Ski dataset. Skiing is fast and sequences of the Ski dataset are as short as 5s. This provides little training data for temporal models and, therefore, we use a single-frame input model. We report the performance of VPose3D with single-frame input in a cross-dataset scenario to compare as a baseline model. Moreover, our results compared with Rhodin *et al.* [33] and CanonPose [40] that use multi-view data from the training set of Ski show 28mm improvement in MPJPE and 2mm in PA-MPJPE.

6.2. Qualitative Evaluation

Figure 4 shows qualitative evaluation on Ski, 3DHP, and 3DPW datasets. The predictions of the baseline and AdaptPose are depicted vs. the ground truth. We observe that AdaptPose successfully enhances the baseline predictions. Figure 5 provides some examples of the generated motion and the input 3D keypoints. Generated motions are smooth and realistic. We provide further qualitative examples in the supplementary material.

Table 1. Cross-scenario learning on H3.6M. Source: S1. Target: S5, S6, S7, S8

Method	3D	PA-MPJPE	MPJPE
Martinez <i>et al.</i> [27]	Full	-	45.5
Pavlo [32]	Full	27.2	37.2
Lui <i>et al.</i> [25]	Full	-	34.7
Wang [41]	Full	-	25.6
PoseAug [11]	S1	-	56.7
Pavlo [32]	S1	-	51.7
Li <i>et al.</i> [23]	S1	-	50.5
Ours	S1	34.0	42.5

Table 2. Cross-dataset (CD) evaluation on 3DHP dataset. Source: H3.6M-target:3DHP

Method	CD	PCK	AUC	MPJPE
Mehta <i>et al.</i> [28]		76.5	40.8	117.6
VNet [30]		76.6	40.4	124.7
MultiPerson [29]		75.2	37.8	122.2
OriNet [26]		81.8	45.2	89.4
BOA [13]	✓	90.3	-	117.6
Wang <i>et al.</i> [42]	✓	76.1	-	109.5
SRNET [45]	✓	77.6	43.8	-
Li <i>et al.</i> [23]	✓	81.2	46.1	99.7
PoseAug [11]	✓	82.9	46.5	92.6
Zhang <i>et al.</i> [46]	✓	83.6	48.2	92.2
Ours	✓	88.4	54.2	77.2

Table 3. Cross-dataset (CD) evaluation on 3DPW dataset. Source: H3.6M-target:3DPW

Method	CD	PA-MPJPE	MPJPE
EFT [17]		55.7	-
Vibe [18]		51.9	82.9
Lin <i>et al.</i> [24]		45.6	74.7
Sim2real [8]	✓	74.7	-
Zhang <i>et al.</i> [46]	✓	70.8	-
Wang <i>et al.</i> [42]	✓	68.3	109.5
SPIN [20]	✓	59.2	96.9
PoseAug [11]	✓	58.5	94.1
VIBE [18]	✓	56.5	93.5
BOA [13]	✓	49.5	77.2
Ours	✓	46.5	81.2

Table 4. Cross-dataset (CD) evaluation on Ski dataset. Source: H3.6M-target:Ski

Method	CD	PA-MPJPE	MPJPE
Rhodin <i>et al.</i> [33]		85	-
CanonPose [40]		89.6	128.1
Pavlo <i>et al.</i> [32]	✓	88.1	106.0
PoseAug [11]	✓	83.5	105.4
Ours	✓	83.0	99.4

6.3. Ablation Studies

Ablation on Components of AdaptPose. We ablate components of our framework including bone generation, camera generation, domain discriminator, and selection.

Table 5 provides the performance improvements by adding any of the components starting from the baseline. All of the components have a major contribution to the results. Comparing bone generation and camera generation, the latter has larger effects on the performance. However, in contrast to PoseAug [11], our bone generation method is significantly contributing to the results (10 mm vs 1 mm). A3 shows that a combination of bone and camera generation is as good as camera generation alone. Therefore, A4 excludes bone generation from the pipeline that causes a 9 mm performance drop in MPJPE. A3 and A5 give the role of domain adaptation that is 10 mm improvements.

Table 5. Ablation study on supervision elements of the proposed model. Source: H3.6M-target:3DHP

Index	BG	Cam	DD	Select	PMPJPE	MPJPE
Baseline					66.5	96.4
A1	✓				61.7	90.1
A2		✓			62.0	88.2
A3	✓	✓			61.8	88.1
A4		✓	✓	✓	59.3	86.5
A5	✓	✓	✓		54.0	78.6
AdaptPose	✓	✓	✓	✓	53.6	77.2

Ablation on bone generation methods. In this section we compare the performance of three different bone generation methods that were explained in Section 4.1. Table 6 gives performance of BG1, BG2, and BG3 while performing cross-dataset evaluation on 3DHP. We observe that using an axis-angle representation for rotating bone vectors is superior to generating bone directions. We hypothesize that learning $\Delta\vec{B}$ is a harder task since there are infinitely many $\Delta\vec{B}$ that can generate $[\vec{B}'_{t=0}]^N$ from \vec{B}_t . On the contrary, there are only two axis-angles that map \vec{B}_t to $[\vec{B}'_{t=0}]^N$.

Table 6. Ablation study on bone generation strategies

Method	PMPJPE	MPJPE
BG1	59.3	85.1
BG2	56.2	80.0
BG3	53.6	77.2

Ablation on camera generation methods. In this section we perform analysis on three different camera generation methods that were introduced in Section 4.2. In terms of rotation representation, axis-angle outperforms quaternions and Euler-angles. Euler-angles are sensitive to the order of rotations and can lead to degenerate solutions. Comparing probabilistic and deterministic methods, the former obtains 5 mm more accurate results.

Ablation on temporal information. Table 8 shows the performance of the network while excluding temporal information from the input and generating single 2D-3D pairs. Our cross-dataset MPJPE is 86.4 mm which still improves over previous methods (86.4 mm vs. 92.2). Therefore, although using temporal information is highly contributing

Table 7. Ablation study on camera generation strategies

Method	Representation	PMPJPE	MPJPE
Deterministic	Axis-Angle	58.0	82.8
Probabilistic	Axis-Angle	53.6	77.2
Probabilistic	Quaternion	58.7	83.5
Probabilistic	Euler-Angle	60.9	85.3

Table 8. Ablation study on temporal information

Input	PCK	AUC	MPJPE
1 frame	84.6	50.3	86.4
27 frames	88.4	54.2	77.2

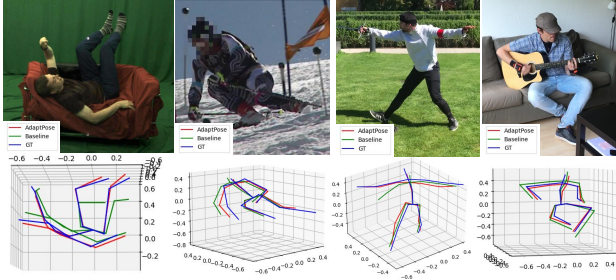


Figure 4. 3D human pose predictions (red) vs. ground truth (blue) for samples of Ski and 3DPW.

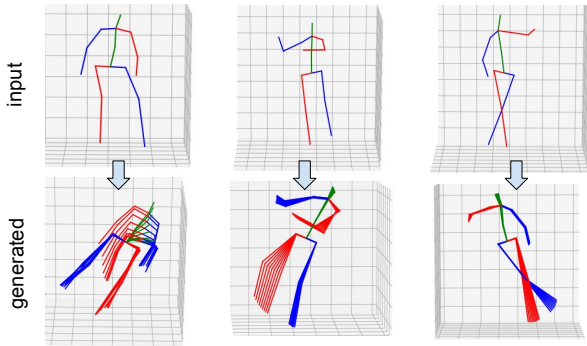


Figure 5. Samples of generated motions and the corresponding input 3D keypoints. Motions are smooth and realistic.

to our framework, our network still excels in non-temporal settings.

6.4. Are we really adapting to new datasets?

To evaluate our claim that we are adapting poses and camera views to the target dataset we visualize some samples of generated motions for 3DHP and 3DPW datasets in Figure 6. The ceiling viewpoint in the first row is from 3DHP that is out of the distribution of our source dataset. While the 2D input is from a chest view camera the generated sample is from a ceiling view, similar to the target samples. We observe that our approach generates qualitatively similar camera poses. The second and third rows also provide examples of new poses that are out of the distribution of source poses and similar to samples in the target dataset.

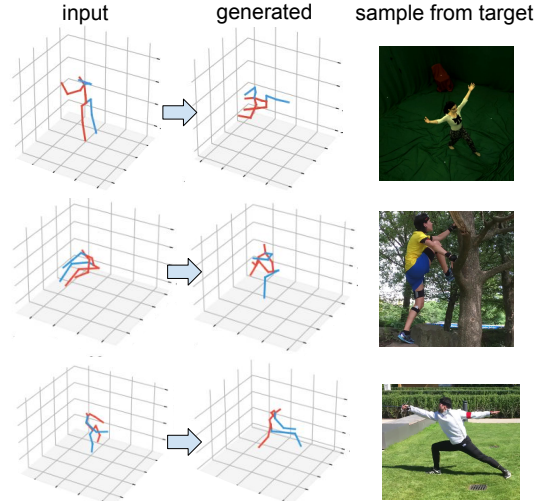


Figure 6. Sample of input images from the source dataset and the generated 3D keypoints. For visualization purposes, we only plot the middle frame from the sequence of generated frames. We manually select the images on the right from the target that matched the generated.

We provide further qualitative examples in the supplementary material. Table 5 also provides numbers regarding the importance of domain discriminators in our framework (A5 vs A3). It is important to mention that we substitute the domain discriminator with a 2D discriminator from the source dataset when excluding the domain discriminator in Table 5. Thus, the performance drop while excluding the domain discriminator is essentially attributed to the lack of adaptation to the target space and not because of excluding the 2D discriminator. The supplementary material provides further experiments on the domain adaption.

7. Conclusion

We proposed an end-to-end framework that adapts a pre-trained 3D human pose estimation model to any target dataset by generating synthetic motions by only looking at 2D target poses. AdaptPose outperforms previous work on four public datasets by a large margin ($> 10\%$). Our proposed solution can be applied to applications where limited motion data is available. Moreover, our method is able to generate synthetic human motion for other tasks such as human action recognition. The major limitation of our work is that it underperforms when there is a large body scale difference between source and train set. Although we have defined a parameter that learns to adjust the body bone lengths we observe a 10mm difference between normalized MPJPE and actual MPJPE when there is a large scale difference between source and target body scales (cross-dataset on 3DPW). Future work should address the scale ambiguity between source and target domains.

References

- [1] Anurag Arnab, Carl Doersch, and Andrew Zisserman. Exploiting temporal context for 3d human pose estimation in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. **1**
- [2] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. **3**
- [3] Yujun Cai, Lihao Ge, Jun Liu, Jianfei Cai, Tat-Jen Cham, Junsong Yuan, and Nadia Magnenat Thalmann. Exploiting spatial-temporal relationships for 3d pose estimation via graph convolutional networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. **1**
- [4] Ching-Hang Chen, Amrith Tyagi, Amit Agrawal, Dylan Drover, Rohith MV, 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 (CVPR)*, June 2019. **2**
- [5] Wenzheng Chen, Huan Wang, Yangyan Li, Hao Su, Zhenhua Wang, Changhe Tu, Dani Lischinski, Daniel Cohen-Or, and Baoquan Chen. Synthesizing training images for boosting human 3d pose estimation. In *3D Vision (3DV)*, 2015. **2**
- [6] Yu Cheng, Bo Yang, Bo Wang, and Robby T. Tan. 3d human pose estimation using spatio-temporal networks with explicit occlusion training. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(07):10631–10638, Apr. 2020. **1**
- [7] Yu Cheng, Bo Yang, Bo Wang, Wending Yan, and Robby T. Tan. Occlusion-aware networks for 3d human pose estimation in video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. **2**
- [8] Carl Doersch and Andrew Zisserman. Sim2real transfer learning for 3d human pose estimation: motion to the rescue. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. **6, 7**
- [9] Dylan Drover, Rohith MV, Ching-Hang Chen, Amit Agrawal, Amrith Tyagi, and Cong Phuoc Huynh. Can 3d pose be learned from 2d projections alone? In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, September 2018. **2**
- [10] Mohsen Gholami, Ahmad Rezaei, Helge Rhodin, Rabab Ward, and Z. Jane Wang. Self-supervised 3d human pose estimation from video. *Neurocomputing*, 488:97–106, 2022. **2**
- [11] 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 (CVPR)*, pages 8575–8584, June 2021. **2, 5, 6, 7**
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Commun. ACM*, 63(11):139–144, Oct. 2020. **3**
- [13] Shanyan Guan, Jingwei Xu, Yunbo Wang, Bingbing Ni, and Xiaokang Yang. Bilevel online adaptation for out-of-domain human mesh reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10472–10481, June 2021. **1, 2, 6, 7**
- [14] Mir Rayat Imtiaz Hossain and James J. Little. Exploiting temporal information for 3d human pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. **1**
- [15] 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, jul 2014. **1**
- [16] Umar Iqbal, Pavlo Molchanov, and Jan Kautz. Weakly-supervised 3d human pose learning via multi-view images in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. **2**
- [17] Hanbyul Joo, Natalia Neverova, and Andrea Vedaldi. Exemplar fine-tuning for 3d human model fitting towards in-the-wild 3d human pose estimation, 2021. **7**
- [18] Muhammed Kocabas, Nikos Athanasiou, and Michael J. Black. Vibe: Video inference for human body pose and shape estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. **7**
- [19] Muhammed Kocabas, Salih Karagoz, and Emre Akbas. Self-supervised learning of 3d human pose using multi-view geometry. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. **2, 6**
- [20] 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*, 2019. **6, 7**
- [21] Jogendra Nath Kundu, Siddharth Seth, Varun Jampani, Mugalodi Rakesh, R. Venkatesh Babu, and Anirban Chakraborty. Self-supervised 3d human pose estimation via part guided novel image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. **2**
- [22] Jogendra Nath Kundu, Rahul M V, Jay Patravali, and Venkatesh Babu RADHAKRISHNAN. Unsupervised cross-dataset adaptation via probabilistic amodal 3d human pose completion. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, March 2020. **2**
- [23] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. **1, 2, 6, 7**

- [24] Kevin Lin, Lijuan Wang, and Zicheng Liu. Mesh graphormer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12939–12948, October 2021. 7
- [25] Ruixu Liu, Ju Shen, He Wang, Chen Chen, Sen-ching Cheng, and Vijayan Asari. Attention mechanism exploits temporal contexts: Real-time 3d human pose reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 7
- [26] Chenxu Luo, Xiao Chu, and Alan Loddon Yuille. Orinet: A fully convolutional network for 3d human pose estimation. In *BMVC*, 2018. 7
- [27] 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)*, Oct 2017. 1, 7
- [28] 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 *3D Vision (3DV), 2017 Fifth International Conference on*. IEEE, 2017. 5, 7
- [29] Dushyant Mehta, Oleksandr Sotnychenko, Franziska Mueller, Weipeng Xu, Srinath Sridhar, Gerard Pons-Moll, and Christian Theobalt. Single-shot multi-person 3d pose estimation from monocular rgb. In *2018 International Conference on 3D Vision (3DV)*, pages 120–130, 2018. 7
- [30] Dushyant Mehta, Srinath Sridhar, Oleksandr Sotnychenko, Helge Rhodin, Mohammad Shafiei, Hans-Peter Seidel, Weipeng Xu, Dan Casas, and Christian Theobalt. Vnect: Real-time 3d human pose estimation with a single rgb camera. volume 36, 2017. 7
- [31] 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
- [32] Dario Pavllo, Christoph Feichtenhofer, David Grangier, and Michael Auli. 3d human pose estimation in video with temporal convolutions and semi-supervised training. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 6, 7
- [33] Helge Rhodin, Jörg Spörri, Isinsu Katircioglu, Victor Constantin, Frédéric Meyer, Erich Müller, Mathieu Salzmann, and Pascal Fua. Learning monocular 3d human pose estimation from multi-view images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 1, 2, 6, 7
- [34] Grégory Rogez and Cordelia Schmid. Mocap-guided data augmentation for 3d pose estimation in the wild. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS’16*, page 3116–3124, 2016. 2
- [35] Jörg Spörri. Research dedicated to sports injury prevention – the ‘sequence of prevention’ on the example of alpine ski racing. *Habilitation with Venia Docendi in “Biomechanics”*. Department of Sport Science and Kinesiology, University of Salzburg., 2016. 1
- [36] Gul Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J. Black, Ivan Laptev, and Cordelia Schmid. Learning from synthetic humans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 2
- [37] 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. 6
- [38] Bastian Wandt, Hanno Ackermann, and Bodo Rosenhahn. A kinematic chain space for monocular motion capture. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, September 2018. 5
- [39] Bastian Wandt and Bodo Rosenhahn. Repnet: Weakly supervised training of an adversarial reprojection network for 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2, 5
- [40] Bastian Wandt, Marco Rudolph, Petrisa Zell, Helge Rhodin, and Bodo Rosenhahn. Canonpose: Self-supervised monocular 3d human pose estimation in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13294–13304, June 2021. 2, 6, 7
- [41] Jingbo Wang, Sijie Yan, Yuanjun Xiong, and Dahua Lin. Motion guided 3d pose estimation from videos. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 764–780, Cham, 2020. Springer International Publishing. 7
- [42] Zhe Wang, Daeyun Shin, and Charless C. Fowlkes. Predicting camera viewpoint improves cross-dataset generalization for 3d human pose estimation. In Adrien Bartoli and Andrea Fusiello, editors, *Computer Vision – ECCV 2020 Workshops*, pages 523–540, Cham, 2020. Springer International Publishing. 1, 2, 6, 7
- [43] 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 (CVPR)*, June 2018. 2
- [44] Frank Yu, Mathieu Salzmann, Pascal Fua, and Helge Rhodin. Pcls: Geometry-aware neural reconstruction of 3d pose with perspective crop layers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9064–9073, June 2021. 5
- [45] 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 Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 507–523, Cham, 2020. Springer International Publishing. 1, 2, 6, 7
- [46] Jianfeng Zhang, Xuecheng Nie, and Jiashi Feng. Inference stage optimization for cross-scenario 3d human pose estimation. In *NeurIPS*, 2020. 1, 2, 6, 7
- [47] Jianfeng Zhang, Dongdong Yu, Jun Hao Liew, Xuecheng Nie, and Jiashi Feng. Body meshes as points. In *Proceed-*

*ings of the IEEE/CVF Conference on Computer Vision and
Pattern Recognition (CVPR), pages 546–556, June 2021. 2*