

MoCap-to-Visual Domain Adaptation for Efficient Human Mesh Estimation from 2D Keypoints

Bedirhan Uguz¹, Ozhan Suat¹, Batuhan Karagoz¹, Emre Akbas^{1,2}

¹Department of Computer Engineering and ²METU ROMER Robotics Center
Middle East Technical University, Ankara, Turkey

{bedirhan.uguz, ozhan.suat, batuhan.karagoz, eakbas}@metu.edu.tr

Abstract

This paper presents Key2Mesh, a model that takes a set of 2D human pose keypoints as input and estimates the corresponding body mesh. Since this process does not involve any visual (i.e. RGB image) data, the model can be trained on large-scale motion capture (MoCap) datasets, thereby overcoming the scarcity of image datasets with 3D labels. To enable the model’s application on RGB images, we first run an off-the-shelf 2D pose estimator to obtain the 2D keypoints, and then feed these 2D keypoints to Key2Mesh. To improve the performance of our model on RGB images, we apply an adversarial domain adaptation (DA) method to bridge the gap between the MoCap and visual domains. Crucially, our DA method does not require 3D labels for visual data, which enables adaptation to target sets without the need for costly labels. We evaluate Key2Mesh for the task of estimating 3D human meshes from 2D keypoints, in the absence of RGB and mesh label pairs. Our results on widely used H3.6M and 3DPW datasets show that Key2Mesh sets the new state-of-the-art by outperforming other models in PA-MPJPE for both datasets, and in MPJPE and PVE for the 3DPW dataset. Thanks to our model’s simple architecture, it operates at least 12× faster than the prior state-of-the-art model, LGD [44]. Additional qualitative samples and code are available on the project website: <https://key2mesh.github.io/>.

1. Introduction

Accurate estimation of 3D human pose and shape (HPS) from a single-view image is a challenging problem in computer vision with many applications in human motion understanding [56] and generation [37], AR/VR [50], human-computer interaction [15], and medical field [3, 10, 39]. Parametric human body models like SMPL [30] have played a pivotal role in driving recent advancements in this domain. A typical HPS estimation method [5, 18, 22, 26,

28, 31, 46] takes an image (or video) as input and estimates SMPL parameters. Mostly, such methods follow a fully supervised approach and rely on the availability of image datasets with 3D annotations. However, obtaining such “paired” annotations is both costly and challenging, particularly in the wild. One usually needs to set up an expensive motion capture system that is extensively tailored for the environment. Consequently, existing datasets for direct 3D supervision are small and do not cover the complete range of variations in human body (shape, size, pose, appearance), and scene (lighting, environment). In addition to technical challenges, the practical application of traditional data collection methods faces significant challenges in acquiring comprehensive visual data, particularly in sensitive fields like the medical domain, where privacy concerns restrict full-image collection.

Several prior work have addressed the shortage of 3D labels by generating pseudo-labels through multiview geometry [21], optimization [24], and fine-tuning a model [17, 27]. Another set of approaches explores the usage of 2D labels as auxiliary information sources. The availability of large image and video datasets made it possible to reliably estimate and train image-based HPS models with 2D cues such as 2D keypoints [4, 7, 18, 22, 24, 27, 28, 45, 46], silhouettes [36, 52], optical flow [47], and body part segmentations [53, 54]. Unlike 2D annotations, large-scale MoCap datasets like AMASS [32] provide extensive 3D details, spanning body pose, shape distributions, and mesh vertex configurations. However, incorporating such a dataset in HPS estimation is not trivial since it does not contain any RGB data (hence, “unpaired”).

Some image-based methods leverage unpaired MoCap data to supervise their models with adversarial losses [18, 22] or establish human body priors [25, 38]. Moreover, as a further step, recent work employs various proxy representations such as 2D keypoints [8, 44], silhouettes [40, 42], edge images [41] and dense correspondences [9] to train models with unpaired MoCap data. Common to these methods is the generation of “proxy representation (e.g.

2D keypoints)”—“human mesh” training pairs from MoCap samples. During inference, obtaining proxy representations is straightforward through readily available estimators trained on the aforementioned large 2D datasets. Also, using proxy representations offers a strategic workaround for data collection challenges in sensitive fields by avoiding the need to store identifiable image data, thus addressing privacy concerns. However, as this process only involves MoCap data, it falls short of accurately capturing visual data attributes, disregarding crucial factors such as estimation errors, occlusions, and noise introduced by the estimators.

In this paper, we introduce Key2Mesh, a model estimating a human body mesh from 2D keypoints, applied to RGB images after extracting keypoints via an off-the-shelf 2D human pose estimator (Fig. 1 inference). We pre-train Key2Mesh on the body mesh instances obtained from a large-scale unpaired MoCap dataset. Specifically, we extract 2D keypoints from a body mesh in the dataset using a random virtual camera, which serves as input for our model. Then, Key2Mesh is pre-trained on such 2D keypoint - body mesh pairs (Fig. 1 training).

To bridge the gap between the source (MoCap) and target (Visual) domains, we introduce a domain adaptation process that utilizes detected 2D keypoints from the target domain and further train our Key2Mesh model. We employ a domain critic and train the Key2Mesh model in an adversarial manner similar to ADDA [48], along with a re-projection loss and a feature regularization to support HPS estimation task on the target domain (Fig. 1 training). Notably, our method does not require 3D labels from the target domain, making it easy to adapt the pre-trained Key2Mesh model from the source domain to different target domains.

We evaluate our approach on two widely-used indoor and outdoor 3D human pose and shape estimation datasets, H3.6M [16] and 3DPW [49]. In the task of reconstructing human meshes using 2D detections without depending on paired 3D labels, Key2Mesh outperforms prior models [8, 9, 44] in terms of PA-MPJPE performance on both datasets, as well as in MPJPE and PVE on the 3DPW dataset. In addition, thanks to Key2Mesh’s simple architecture, which implements a single forward pass, our model is up to 33× faster compared to LGD [44] – the only model that (i) solves exactly the same task as Key2Mesh, (ii) is trained on an unpaired dataset, and (iii) has a public implementation.

2. Related Work

2.1. Human Pose and Shape Estimation

Considering the existing work on HPS estimation, our model Key2Mesh has four critical properties: (i) For inference, 2D keypoints are used as input, and no RGB data is required. (ii) During training, it does not need a “paired”

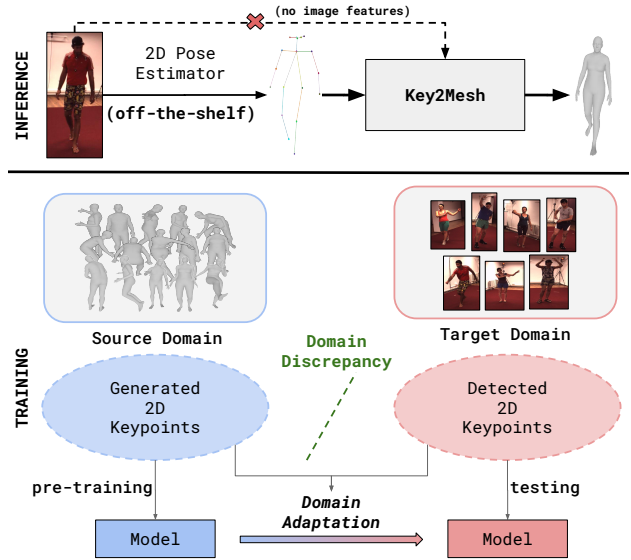


Figure 1. **An overview of Key2Mesh’s inference and training process.** (Top) **Inference:** For a given input image, we run an off-the-shelf 2D pose estimator to obtain 2D keypoints, which are input to our Key2Mesh model to estimate a body mesh. (Bottom) **Training:** Source domain consists of 3D body meshes obtained from MoCap data. We generate 2D keypoints from these body meshes using virtual cameras and a range of augmentations. We then pre-train our model using these generated 2D keypoints and corresponding meshes, without incorporating any RGB images. Finally, we adapt our model to target RGB images by using an off-the-shelf 2D pose estimator and bridging the gap between the “2D keypoints obtained from MoCap” domain and the “2D keypoints obtained from the target RGB images” domain.

dataset, i.e., {image, 3D label} pairs. It only requires a dataset with 3D body meshes without any corresponding visual data. MoCap datasets provide such data. (iii) During inference, the model runs only a single forward pass, without needing any recurrent or iterative operations. (iv) It uses an adversarial domain adaptation method to adapt the model to visual data. Still, no 3D labels (i.e. paired data) are required.

Existing methods that have the first two properties above are all iterative: SMPLify [1], LGD [44] and NeuralFitter [8], while Key2Mesh runs in a single-forward pass. Due to this difference, Key2Mesh is up to 33x faster than LGD, which is reported to be faster than SMPLify. We cannot do any inference time comparison with NeuralFitter as they do not provide code. Another key difference is domain adaptation, which shows its effect in the end result: Key2Mesh outperforms SMPLify, LGD and NeuralFitter in terms of PA-MPJPE on both H3.6M and 3DPW, as well as in MPJPE and PVE on 3DPW.

Several studies utilize unpaired 3D data for training or none, aligning with property (ii). Among these, some re-

quire supplementary 2D observations, such as 2D keypoint sequences (PoseNet3D [45]), dense correspondences (CRA [9]), silhouettes (STRAPS [40], Skeleton2Mesh [52]), and edge images (STRAPS V3 [41]). In contrast, Key2Mesh only requires 2D keypoints as input, making it easy to use and more flexible by eliminating the need for additional 2D estimation. Alternatively, some studies rely on RGB images as inputs such as HMR [18] and SPIN [24] and report results both using paired and unpaired data. They exhibit a significant performance degradation in unpaired settings, whereas Key2Mesh outperforms them under such conditions.

A related line of research employs paired data for training, utilizing an RGB image-based pipeline to estimate human pose and shape [18, 23, 25, 27, 31, 33, 55]. However, obtaining these paired datasets poses challenges. Pose2Mesh [6], and MPT [29] adopts a hybrid approach, aiming to estimate pose and shape from 2D keypoints similar to (i), combining paired datasets with 2D ground truth and unpaired datasets. Also, MPT utilizes a pre-training strategy with MoCap data similar to ours, but implements an end-to-end fine-tuning pipeline using image-3D ground truth pairs through a 2D pose estimation backbone. In contrast, Key2Mesh leverages only unpaired datasets, implementing domain adaptation to mitigate the discrepancies between 2D ground truths and detections.

2.2. Unsupervised Domain Adaptation in HPS

In the realm of HPS estimation, the integration of domain adaptation techniques has recently commenced to address domain gaps. Mugaludi *et al.* [34] proposed a silhouette based technique to alleviate domain shifts. BOA [11], along with its extension DynaBOA [12], introduced an online domain adaptation method for deploying HPS estimation models in out-of-domain scenarios. They employ a bi-level video-based update rule to adapt the initial model trained on paired 3D datasets. In a different approach, CycleAdapt [34] presented a test-time adaptation framework that does not rely on 2D observations, in contrast to the BOA family. Unlike all these methods, Key2Mesh utilizes 2D keypoints as a proxy representation, relies solely on MoCap data for training the initial model, and implements an efficient single-forward pass estimation pipeline without the need for updating model parameters during inference.

3. Method

This section describes how we train our Key2Mesh model. The training is divided into two phases: “Pre-Training” and “Domain Adaptation”, as illustrated in Fig. 2. In the first phase, we train Key2Mesh using a large MoCap dataset. Then, we introduce a domain adaptation pipeline to adapt our trained model to the testing domain (target).

3.1. Problem Setting

Following previous work [8, 44], we define our objective as the estimation of the full 3D mesh of human bodies exclusively from 2D keypoint observations, without reliance on any paired datasets. The 3D mesh of a human body is encoded by Skinned Multi-Person Linear (SMPL) [30] model, which is a parametric body model that generates a mesh denoted as $M \in \mathbb{R}^{6890 \times 3}$ based on the provided pose parameters $\theta \in \mathbb{R}^{24 \times 3}$ that include global rotation parameters, and shape parameters $\beta \in \mathbb{R}^{10}$. The 3D keypoints $X \in \mathbb{R}^{k \times 3}$, where k is the number keypoints, are defined as a linear combination of the mesh vertices and can be computed as $X = WM$ where W is a learned weight matrix $W \in \mathbb{R}^{k \times 6890}$. Our goal can be summarized as estimating full 3D human body meshes encoded with SMPL parameters from the 2D keypoint detections obtained by an off-the-shelf 2D pose estimator, *e.g.* OpenPose [2].

3.2. Training Data Generation from Unpaired 3D Human Body Data

Here we describe the data generation process employed in our first, pre-training, phase. We only utilize publicly available ground truth 3D human body data, obtained by commercial MoCap systems and specialized algorithms like Mosh++ [32], without requiring any corresponding images. We generate pairs of 2D keypoints $x \in \mathbb{R}^{k \times 2}$ and SMPL parameters (θ, β) on the fly by projecting 3D SMPL keypoints to 2D using a fixed camera c which has a fixed focal length f , identity global rotation matrix $R = I^{3 \times 3}$ looking towards to origin, and fixed translation.

Since the 2D keypoints are generated from the ground truth 3D human body, they are free from any defects. However, in the RGB domain, 2D keypoint estimators suffer from problems such as occlusions, noise, image quality, small size of body in pixel space, etc. This leads to significant differences between 2D keypoints obtained from MoCap (training) and those obtained from RGB (test) images. To mitigate this disparity, we use a range of augmentation techniques during the training pair generation process. These techniques simulate real-world scenarios observed in RGB data, incorporating random camera rotations, keypoint occlusions, and the addition of random noise to emulate jitter effects caused by the 2D pose estimators.

3.3. Pre-Training

In the pre-training phase, we train Key2Mesh to regress human pose and shape from 2D keypoints generated from the unpaired dataset, which we refer to as the source dataset. The model consists of a feature extractor and an SMPL head.

Given 2D keypoints x , the feature extractor learns a function $F_{pt} : \mathbb{R}^{k \times 2} \rightarrow \mathbb{R}^d$ that maps the keypoints to a d -dimensional representation ϕ . Then, SMPL head takes ϕ

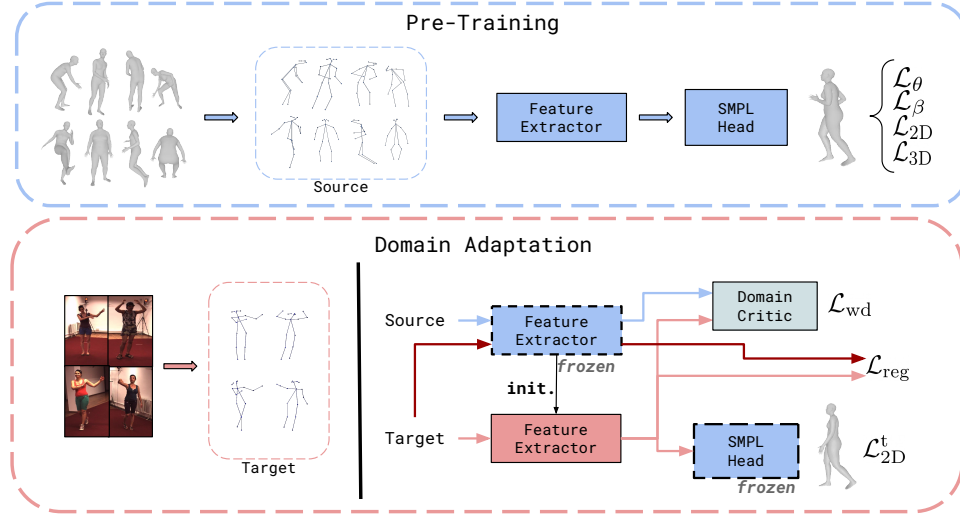


Figure 2. **Overview of our pre-training and domain adaptation phases.** First, we begin with “Pre-Training” phase where we train a feature extractor and an SMPL head using only an unpaired MoCap dataset. Second, we introduce a “Domain Adaptation” phase to close the gap between the source domain (MoCap) and the target domain (Visual) by using 2D detections obtained by an off-the-shelf 2D pose estimator in the target domain.

as input and learns a function $H_{pt} : \mathbb{R}^d \rightarrow \mathbb{R}^{|\Theta|}$ where Θ is $[\theta, \beta]$. Following [24], we use a 6D continuous rotation representation [57] in the estimation of SMPL pose parameters θ , instead of axis-angle representation. Finally, for a given x , we end up with estimated SMPL parameters $\hat{\theta}$ that includes global rotation parameters, and $\hat{\beta}$.

We compute the estimated 3D keypoints \hat{X} by using the estimated SMPL vertices with learned weight matrix W based on $\hat{\theta}$ and $\hat{\beta}$. The estimated 3D keypoints are then projected onto 2D using the fixed camera c from the training data generation process, yielding the 2D keypoints \hat{x} . We further compute losses for SMPL parameters (θ, β), 2D re-projection (x), and 3D (X) keypoints and train our overall pre-trained (F_{pt}, H_{pt}) model. Specifically, the overall loss function for this phase can be written as follows:

$$\mathcal{L}_{PT} = w_1 \mathcal{L}_\theta + w_2 \mathcal{L}_\beta + w_3 \mathcal{L}_{2D} + w_4 \mathcal{L}_{3D} \quad (1)$$

where each term is calculated as:

$$\begin{aligned} \mathcal{L}_\theta &= \|\hat{\theta} - \theta\|_1, \mathcal{L}_\beta = \|\hat{\beta} - \beta\|_1, \\ \mathcal{L}_{2D} &= \|\hat{x} - x\|_1, \mathcal{L}_{3D} = \|\hat{X} - X\|_1. \end{aligned}$$

3.4. Domain Adaptation

The pre-trained model, trained solely on ground-truth 2D keypoints-SMPL pairs, struggles with domain changes when tested on real indoor and outdoor RGB datasets. Even though we employed a range of augmentation techniques in the pair generation process to mimic the RGB 2D pose estimator, we observe that the pose variances, cameras, occlusion scenarios, and jitter introduced by the 2D pose estimator are not fully covered by our augmented source domain.

Taking inspiration from previous domain adaptation work, ADDA [48] and WDGR [43], we propose an adversarial domain adaptation technique that builds upon our pre-trained model. This technique aims to mitigate the discrepancies between data domains without relying on any labeled data from the target domain. We apply an off-the-shelf pose estimator to the target domain images that are available for use in the training to obtain 2D keypoints which are referred to as the target dataset and use this dataset to adapt our pre-trained model to the target domain.

In this phase, we freeze the parameters of the feature extractor F_{pt} and SMPL head H_{pt} trained in the previous phase. Next, we introduce a new feature extractor F which has the same configuration with F_{pt} and is initialized with the weights of F_{pt} . Additionally, we introduce a domain critic that learns a mapping $D : \mathbb{R}^d \rightarrow \mathbb{R}$ that takes feature vectors from the source and target domain as input and outputs a real number. The domain critic assists us in training F to generate target domain features ϕ^t that are indistinguishable from the features ϕ^s generated by the frozen F_{pt} in the source domain, as proposed in WDGR [43]. The final Key2Mesh model closely resembles the pre-trained model, with one key distinction. While the SMPL head remains unchanged, we replace F_{pt} with F that has undergone an adaptation process to the target domain. This adaptation process involves utilizing domain adaptation losses, described in detail in Sec. 3.6. By incorporating the adapted feature extractor F , we enhance the pipeline’s ability to generate accurate SMPL predictions for the given 2D keypoints in the target domain, as the SMPL Head en-

counters features in the target domain that resemble those from its training set, i.e., the source domain.

3.5. Domain Critic

Instead of ADDA’s domain discriminator, we follow WD-GRL [43] and employ a domain critic and calculate Wasserstein distance between the source features ϕ^s generated by F_{pt} taken from the pre-trained model and target features ϕ^t generated by F introduced in the adaptation process. The special case of Wasserstein distance between two representation distributions \mathbb{P}_{ϕ^s} and \mathbb{P}_{ϕ^t} proposed in [43] can be formulated here as:

$$W(\mathbb{P}_{\phi^s}, \mathbb{P}_{\phi^t}) = \sup_{\|D\|_L \leq 1} \mathbb{E}_{\mathbb{P}_{\phi^s}}[D(\phi)] - \mathbb{E}_{\mathbb{P}_{\phi^t}}[D(\phi)], \quad (2)$$

where $\|\cdot\|_L$ denotes the Lipschitz constant. When the learned critic function D is enforced to be 1-Lipschitz via the gradient penalty [14], we can approximate Eq. (2) with features generated from source samples x^s and target samples x^t as:

$$\mathcal{L}_{\text{wd}}(\Phi^s, \Phi^t) = \frac{1}{|x^s|} \sum_{\phi^s \in \Phi^s} D(\phi^s) - \frac{1}{|x^t|} \sum_{\phi^t \in \Phi^t} D(\phi^t), \quad (3)$$

where $\Phi^s = \{F_{\text{pt}}(x) \mid x \in x^s\}$ and $\Phi^t = \{F(x) \mid x \in x^t\}$. The gradient penalty we use is defined as

$$\mathcal{L}_{\text{grad}}(\hat{\phi}) = (\|\nabla_{\hat{\phi}} D(\hat{\phi})\|_2 - 1)^2, \quad (4)$$

where $\hat{\phi}$ is sampled uniformly along straight lines between source and target feature pairs, similar to [14]. While training the domain critic, we solve the problem

$$\max\{\mathcal{L}_{\text{wd}} - \gamma \mathcal{L}_{\text{grad}}\} \quad (5)$$

where γ represents the balancing coefficient. Following WDGRL [43], we first train D to optimality with k steps for each domain adaptation training step. Then, by fixing the parameters of optimal D , we calculate the domain adaptation losses and update the F .

3.6. Domain Adaptation Losses

We apply 2D re-projection loss $\mathcal{L}_{2\text{D}}^t$, domain loss \mathcal{L}_{wd} and regularization on domain-adapted features \mathcal{L}_{reg} in the adaptation process. Below, we describe these loss and regularization terms.

2D re-projection loss on the target domain. F takes 2D keypoint detections x^t from the target dataset as inputs, which allows us to calculate a 2D re-projection loss between the final SMPL predictions and x^t . This 2D re-projection loss serves as a supervision signal during the adaptation without requiring human-labeled or 3D ground truth data.

Similar to $\mathcal{L}_{2\text{D}}$, we use the fixed camera c to project predicted 3D SMPL keypoints and obtain 2D keypoints (\hat{x}^t). The calculation of $\mathcal{L}_{2\text{D}}^t$ is as follows:

$$\mathcal{L}_{2\text{D}}^t = \|\hat{x}^t - x^t\|_1. \quad (6)$$

Domain loss. By utilizing D , we enforce ϕ^t to be inseparable from ϕ^s . To do this, we solve the problem $\min\{\mathcal{L}_{\text{wd}}\}$. Practically, we re-calculate the \mathcal{L}_{wd} after fixing the optimal parameters of the domain critic.

Regularization on domain-adapted features. We apply a regularization term on ϕ^t . We compute the corresponding feature vector $\bar{\phi}^t = F_{\text{pt}}(x^t)$ and apply the following loss:

$$\mathcal{L}_{\text{reg}} = \|\phi^t - \bar{\phi}^t\|. \quad (7)$$

We combine the loss functions mentioned in this section using specific weights to obtain an overall loss function to train F :

$$\mathcal{L}_{\text{DA}} = w_5 \mathcal{L}_{2\text{D}}^t + w_6 \mathcal{L}_{\text{wd}} + w_7 \mathcal{L}_{\text{reg}}. \quad (8)$$

4. Experiments

4.1. Implementation Details

We employ standard MLP architectures: 5 blocks for feature extractors, 2 for the SMPL head. Each block includes a linear layer (1024 hidden dimensions), batch normalization, parametric ReLU, and 0.2 dropout. The Domain Critic uses 2 blocks without normalization layers. For pre-training, we set hyperparameters ($w_1 = w_2 = 100$, $w_3 = w_4 = 50$) for \mathcal{L}_{PT} , optimizing it via Adam optimizer [20] (lr: 1e-3) over ten epochs, batch size 256. During domain adaptation, both the domain critic and feature extractor are trained for five epochs, using lr 1e-4, batch size 256 per domain (512 for critic). The domain critic trains $k = 3$ times per adaptation step. For \mathcal{L}_{DA} , we set hyperparameters ($w_5 = 10$, $w_6 = 0.1$, $w_7 = 20$). Experiments are conducted using PyTorch [35] on one RTX3080ti GPU for all experiments.

4.2. Datasets and Evaluation Metrics

4.2.1 Training Data

During the pre-training of Key2Mesh (Fig. 2), we solely utilize the AMASS dataset [32]. This extensive 3D MoCap dataset consists of a wide range of human poses and shapes (and motions), with SMPL parameters. However, it does not include RGB images corresponding to the 3D data.

In the domain adaptation phase, we leverage the InstaVariety dataset [19], which offers a large set of in-the-wild 2D keypoint detections. Moreover, our domain adaptation process incorporates the Human3.6M dataset [16] (H3.6M) and the 3DPW dataset [49]. **We do not use any ground-truth labels available in these datasets.** Instead, we employ OpenPose [2] to obtain detected 2D keypoints for H3.6M [16] and the detections available in 3DPW [49]. We use these 2D detections in the adaptation phase.

Method	Input	PA-MPJPE↓	MPJPE↓	Time↓
SMPLify [1]	K	82.3	-	-
HMR (u) [18]	I	66.5	-	-
SPIN (u) [24]	I	62.0	-	-
STRAPS [40]	K, S	55.4	-	-
LGD [44]	K	55.0	102.4	42.4
CRA [9]	K, D	53.9	81.0	(D : 110.0)+
Key2Mesh [†]	K	51.4	108.1	3.4
Key2Mesh [‡]	K	51.0	107.1	3.4

Table 1. Comparison of models trained on unpaired data similar to ours, assessed on H3.6M [16] for PA-MPJPE and MPJPE (both in mm) under P2, and processing time (ms) with a single batch size. Inputs include (D)ense Correspondences, (I)mage, 2D (K)eypoints, or (S)ilhouette. As CRA does not have a publicly available code, our time evaluation focused on measuring the time needed to acquire dense correspondences (D) (DensePose [13]), an input used alongside keypoints (K) unlike Key2Mesh and LGD. (u) denotes unpaired. Our models marked with (\dagger) and (\ddagger) were obtained by adapting our PT model to training subjects and evaluation subjects of the H3.6M dataset, respectively, without employing any ground truth labels. Our Key2Mesh obtains the best PA-MPJPE result and is $12\times$ faster than LGD and at least $32\times$ faster than CRA.

4.2.2 Evaluation Data

We quantitatively and qualitatively evaluate our approach on widely used indoor and in-the-wild datasets, H3.6M [16] (subjects S9, S11) and 3DPW [49] datasets.

We use OpenPose [2] on the H3.6M test subjects to obtain 2D keypoint detections, and following the evaluation protocol P2, we only used a frontal camera (camera 3). Following the approach of prior works [8, 44], we discard frames in which fewer than 6 keypoints were detected. For the 3DPW evaluation, we use the detections available in the dataset. We report PA-MPJPE and MPJPE for the H3.6M dataset and PA-MPJPE, MPJPE and PVE for the 3DPW dataset, all measured in mm.

4.3. Quantitative Results

We first pre-train Key2Mesh by following the data generation process detailed in Sec. 3.2. Then, we apply our domain adaptation method by using OpenPose detections from H3.6M subjects, 3DPW, and InstaVariety. For H3.6M, we report the performance of the domain-adapted model, in which we used the detections from all H3.6M training subjects to adapt our pre-trained model. For the 3DPW dataset, we report our InstaVariety adapted model since adaptation to InstaVariety leads to better results compared to the 3DPW train split. Our adaptation process doesn’t require target labels, allowing direct adaptation of the pre-trained model to the test sets using the respective OpenPose detections. Consequently, we present evaluations for models adapted to both H3.6M evaluation subjects and the 3DPW test split.

Method	Input	PA-MPJPE↓	MPJPE↓	PVE↓
SMPLify [1]	K	106.1	-	-
PoseNet3D* [45]	K	73.6	-	-
STRAPS V3 [41]	K, E	59.2	-	-
CRA [9]	K, D	55.9	89.1	115.3
LGD [44]	K	54.2	91.3	104.8
NeuralFitter [8]	K	52.2	-	-
Key2Mesh [†]	K	50.1	89.0	101.7
Key2Mesh [‡]	K	49.8	86.7	99.5

Table 2. Comparison of our models with others that utilize 2D proxy representations and trained on unpaired data similar to our model like ours, assessed on the 3DPW [49] dataset for PA-MPJPE, MPJPE and PVE (all in mm). The input for the presented methods includes (D)ense Correspondences, (E)dge Images, and 2D (K)eypoints. (*) denotes methods that take temporal data as input. (\dagger) and (\ddagger) were achieved through domain adaptation of the PT model to the InstaVariety dataset and the test split of the 3DPW dataset, respectively, without utilizing ground-truth labels.

We present detailed evaluations on the target dataset in Sec. 4.4.2.

Evaluation on H3.6M. We compare our models with others on H3.6M benchmark under Protocol 2 (P2). Since we only use 2D keypoint observations without relying on images and utilize unpaired 3D training data, it is only possible to compare Key2Mesh with other models following a similar approach. As presented in Tab. 1, our model outperforms existing state-of-the-art methods in terms of PA-MPJPE without relying on additional 2D observations like dense correspondences or silhouettes used in CRA [9] and STRAPS [40]. However, our method falls short in MPJPE compared to CRA, largely because CRA leverages dense correspondences for estimation. This advantage is more pronounced in controlled environments like H3.6M, where obtaining high-quality dense correspondences is less challenging. For a fair comparison, we evaluate an iterative baseline LGD [44] on our 2D detections, as the original study did not utilize OpenPose [2] keypoint detections (Tab. 1 row 5). Note that we achieve improved PA-MPJPE for LGD using our 2D detections compared to what its authors report. While our model improves upon iterative methods such as LGD and SMPLify [1] in PA-MPJPE, it demonstrates relatively lower performance in MPJPE when compared to LGD. Our model is $12\times$ faster than LGD and at least $32\times$ faster than CRA. In addition, we compared our model with some image-based baselines, HMR [18] and SPIN [24], under the unpaired training setting. our model outperforms these baselines.

Evaluation on 3DPW. In Tab. 2, we report PA-MPJPE, MPJPE and PVE on 3DPW test set. We compare our model with others that utilize 2D proxy representations such as dense correspondences, 2D keypoints, and edge images to

	Losses	PA-MPJPE↓	MPJPE↓
	Pre-Trained	53.9	114.7
	\mathcal{L}_{2D}^t	53.8	109.0
DA	$\mathcal{L}_{2D}^t + \mathcal{L}_{wd}$	52.2	109.0
	$\mathcal{L}_{2D}^t + \mathcal{L}_{wd} + \mathcal{L}_{reg}$	51.4	108.1

Table 3. **Ablation on domain adaptation losses.** We conduct experiments to assess the impact of loss terms during the domain adaptation phase using the H3.6M [16] dataset. We start with a pre-trained model and add (D)omain (A)daptation losses.

recover SMPL meshes. Also, similar to ours, these models only rely on unpaired data during training. Noticeably, prior works such as CRA [9], STRAPSV3 [41] and Posenet3D [45], as depicted in the Tab. 2, require additional inputs or temporal keypoint detections to regress SMPL meshes. Our model achieves better performance without such dependencies for in-the-wild scenarios. We also compare our model with iterative ones such as SMPLify [1], LGD [44], and NeuralFitter [8], by evaluating LGD in MPJPE and PVE (metrics initially unreported), we test LGD on the 3DPW dataset. This evaluation yields a slight improvement in LGD’s PA-MPJPE compared to its original reported values, and we obtain MPJPE and PVE metrics for comparison. Across all metrics, our model outperforms these iterative methods.

Processing time comparison. We also compared our model’s processing time with LGD and CRA in Tab. 1. As CRA lacks available code, our time evaluation focused on measuring the time needed to acquire dense correspondences (DensePose [13] with ResNet-101 backbone), an extra input used alongside keypoints, unlike Key2Mesh and LGD. Our inference pipeline is highly efficient, operating at 3.4 ms with a batch size of 1, compared to LGD’s 42.4 ms where DensePose requires 110.0 ms for Dense Correspondences on the same hardware. With a larger batch size of 512, our method achieves 3.6 ms, a speed advantage of $33\times$ faster than LGD, which slows down to 120 ms due to its iterative nature. These comparisons highlight our method’s crucial rapid processing for speed-focused applications.

4.4. Ablation Studies

4.4.1 Ablation on Domain Adaptation Loss Terms

We conduct experiments to measure the effects of the loss terms during the domain adaptation phase. As shown in Tab. 3, starting from a pre-trained model on the AMASS dataset, we employ our domain adaptation method using varying combinations of loss terms such as \mathcal{L}_{2D}^t , \mathcal{L}_{wd} , \mathcal{L}_{reg} . Applying domain adaptation with only 2D re-projection loss on the target (\mathcal{L}_{2D}^t) helps the model to better align its predictions on the target domain. This approach leads to a slight enhancement in the PA-MPJPE metric and a noticeable improvement in MPJPE. We obtain further im-

Method	Target Dataset	Size	PA-MPJPE↓	MPJPE↓
Ours - PT	-	-	53.9	114.7
Ours - DA	InstaVariety	2.19M	52.0	111.3
Ours - DA	H3.6M - S1	23.5K	52.8	108.9
Ours - DA	H3.6M - Train*	160K	51.4	108.1
Ours - DA	H3.6M - Val.†	12.4K	51.0	107.1

Table 4. **Ablation on target dataset and size using H3.6M [16].** PT and DA denote the pre-trained model and the domain-adapted model, respectively. (*) represents training subjects (S1, S5-S8) and (†) indicates validation subjects (S9, S11).

Method	Target Dataset	Size	PA↓	MPJPE↓	PVE↓
Ours - PT	-	-	57.5	93.6	109.8
Ours - DA	3DPW-Train	22.8K	57.4	91.7	107.3
Ours - DA	InstaVariety	2.19M	57.1	91.7	107.1
Ours - DA	3DPW-Val.	33.5K	56.4	88.6	104.4

Table 5. **Ablation on target dataset and size using 3DPW [49] validation set.** PT and DA represent the pre-trained model and the domain-adapted model, respectively, with PA referring to PA-MPJPE.

provement in PA-MPJPE when we use a combination of \mathcal{L}_{2D}^t and domain critic loss (\mathcal{L}_{wd}). This observation supports our intuition that adapting the features computed on the target dataset helps in reducing pose estimation errors. Also, by introducing a regularization term on the domain-adapted features (\mathcal{L}_{reg}), we achieved the best PA-MPJPE and MPJPE. This indicates that applying regularization to the domain-adapted features leads to a more effective feature extraction for the target dataset.

4.4.2 Ablation on the Target Dataset and Size

We conduct measurements to assess the impact of the target dataset and size during our domain adaptation phase, using H3.6M and 3DPW validation sets. By adapting our pre-trained model to different target datasets, including validation sets, we calculate the evaluation metrics for each adapted model.

Results on H3.6M. For H3.6M, we adapt our pre-trained model to four different datasets separately: InstaVariety, a small dataset that includes only a single subject (S1) from the H3.6M training subjects, the entire H3.6M training subjects, and validation subjects. Tab. 4 row 2 shows that despite InstaVariety being an in-the-wild 2D dataset, different from H3.6M, our domain adaptation process reduces pose alignment errors. In row 3, even using a single subject shows a reduction in error metrics thanks to our domain adaptation process. Row 4 shows the model adapted using the entire H3.6M training subjects, demonstrating better performance than the model adapted using only a single subject. Lastly, by employing our label-independent domain adaptation framework, we adapt our pre-trained model




Figure 3. **Qualitative comparison of our pre-trained and domain-adapted model on H3.6M dataset [16].** The first column shows the input image and keypoint detections. The second and third columns display outputs from the pre-trained model, while the last two columns present results from the domain-adapted model, which outperforms our pre-trained model, especially on complex poses.

to the validation subjects, resulting in the best PA-MPJPE and MPJPE compared to all other adaptation strategies. Interestingly, the model adapted using the H3.6M training subjects performs comparably to the one adapted specifically to validation subjects. This underlines the potential of using pre-collected data for controlled datasets like H3.6M. **Results on 3DPW.** We leverage InstaVariety, 3DPW-Train, and 3DPW-Validation datasets in the adaptation process for the 3DPW dataset. In Tab. 5, row 2 demonstrates that incorporating the 3DPW-train dataset, which is relatively small, leads to improvements in the evaluation metrics. In addition, Tab. 5 row 3 shows that using the InstaVariety dataset, which is also an in-the-wild dataset, further reduces PA-MPJPE. Finally, we employed the 3DPW-validation dataset to enhance our model, leveraging our label-independent adaptation framework, resulting in the best metrics (Tab. 5 row 4). This implies that in a challenging in-the-wild dataset like 3DPW, direct adaptation to the test data results in more accurate pose alignment.

4.5. Qualitative Results

Our method does not involve estimating camera translation. To produce qualitative results, we compute the best camera translation using a least squares optimization to overlay meshes onto images. This optimization minimizes the difference between the 3D keypoints of the SMPL meshes, projected onto the image plane, and their corresponding 2D keypoint inputs.



Figure 4. **Qualitative results on the 3DPW dataset [49].** Each sample displays the input image and its corresponding 2D keypoint detections in the first column, the SMPL mesh overlaid on the image in the second column, and the side view of the SMPL mesh in the third column, all generated using a domain-adapted model.

In Fig. 3, we qualitatively compare our pre-trained and adapted models. The results show that domain adaptation helps our model to produce more aligned results. Further qualitative results showcasing our domain-adapted models for the 3DPW dataset are presented in Fig. 4.

5. Conclusion

In this study, we present Key2Mesh for recovering 3D human mesh parameters from 2D keypoints. We address the lack of image-to-3D, “paired” training data and utilize an “unpaired” training setting. We introduce an adversarial domain adaptation method to adapt models trained on unpaired data to visual data. We comprehensively evaluate our Key2Mesh model on widely used indoor and in-the-wild benchmarks, H3.6M and 3DPW datasets. Our model outperforms similar methods, showcasing superior PA-MPJPE performance on both datasets, along with superior MPJPE and PVE performance on the 3DPW dataset. Comparatively, our model is significantly faster (up to 33×) than the previous state-of-the-art model LGD [44], emphasizing its efficiency for applications where processing time is critical.

Acknowledgements Dr. Akbas is supported by the “Young Scientist Awards Program (BAGEP)” of Science Academy, Turkey.

References

- [1] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *European Conference on Computer Vision*, pages 561–578. Springer, 2016. [2](#), [6](#), [7](#)
- [2] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7291–7299, 2017. [3](#), [5](#), [6](#), [2](#)
- [3] Kenny Chen, Paolo Gabriel, Abdulwahab Alasfour, Chenghao Gong, Werner K Doyle, Orrin Devinsky, Daniel Friedman, Patricia Dugan, Lucia Melloni, Thomas Thesen, et al. Patient-specific pose estimation in clinical environments. *IEEE journal of translational engineering in health and medicine*, 6:1–11, 2018. [1](#)
- [4] Hanbyel Cho, Yooshin Cho, Jaesung Ahn, and Junmo Kim. Implicit 3d human mesh recovery using consistency with pose and shape from unseen-view. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21148–21158, 2023. [1](#)
- [5] Junhyeong Cho, Kim Youwang, and Tae-Hyun Oh. Cross-attention of disentangled modalities for 3d human mesh recovery with transformers. In *European Conference on Computer Vision (ECCV)*, 2022. [1](#)
- [6] Hongsuk Choi, Gyeongsik Moon, and Kyoung Mu Lee. Pose2mesh: Graph convolutional network for 3d human pose and mesh recovery from a 2d human pose. In *European Conference on Computer Vision*, pages 769–787. Springer, 2020. [3](#)
- [7] Hongsuk Choi, Gyeongsik Moon, Ju Yong Chang, and Kyoung Mu Lee. Beyond static features for temporally consistent 3d human pose and shape from a video. In *Conference on Computer Vision and Pattern Recognition*, 2021. [1](#)
- [8] Vasileios Choutas, Federica Bogo, Jingjing Shen, and Julien Valentin. Learning to fit morphable models. In *European Conference on Computer Vision*, pages 160–179. Springer, 2022. [1](#), [2](#), [3](#), [6](#), [7](#)
- [9] Xuan Gong, Meng Zheng, Benjamin Planche, Srikrishna Karanam, Terrence Chen, David Doermann, and Ziyang Wu. Self-supervised human mesh recovery with cross-representation alignment. In *European Conference on Computer Vision*, pages 212–230, 2022. [1](#), [2](#), [3](#), [6](#), [7](#)
- [10] Robert Grimm, Sebastian Bauer, Johann Sukkau, Joachim Hornegger, and Günther Greiner. Markerless estimation of patient orientation, posture and pose using range and pressure imaging: For automatic patient setup and scanner initialization in tomographic imaging. *International journal of computer assisted radiology and surgery*, 7:921–929, 2012. [1](#)
- [11] 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 Conference on Computer Vision and Pattern Recognition*, pages 10472–10481, 2021. [3](#)
- [12] Shanyan Guan, Jingwei Xu, Michelle Zhang He, Yunbo Wang, Bingbing Ni, and Xiaokang Yang. Out-of-domain human mesh reconstruction via dynamic bilevel online adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):5070–5086, 2022. [3](#)
- [13] Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7297–7306, 2018. [6](#), [7](#)
- [14] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. *Advances in Neural Information Processing Systems*, 30, 2017. [5](#)
- [15] Vladimir Guzov, Aymen Mir, Torsten Sattler, and Gerard Pons-Moll. Human positioning system (hps): 3d human pose estimation and self-localization in large scenes from body-mounted sensors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4318–4329, 2021. [1](#)
- [16] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, 2014. [2](#), [5](#), [6](#), [7](#), [8](#), [1](#)
- [17] Hanbyul Joo, Natalia Neverova, and Andrea Vedaldi. Exemplar fine-tuning for 3d human pose fitting towards in-the-wild 3d human pose estimation. In *International Conference on 3D Vision*, 2020. [1](#)
- [18] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7122–7131, 2018. [1](#), [3](#), [6](#)
- [19] Angjoo Kanazawa, Jason Y. Zhang, Panna Felsen, and Jitendra Malik. Learning 3d human dynamics from video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. [5](#)
- [20] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [5](#)
- [21] Muhammed Kocabas, Salih Karagoz, and Emre Akbas. Self-supervised learning of 3d human pose using multi-view geometry. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. [1](#)
- [22] Muhammed Kocabas, Nikos Athanasiou, and Michael J. Black. Vibe: Video inference for human body pose and shape estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020. [1](#)
- [23] Muhammed Kocabas, Chun-Hao P Huang, Otmar Hilliges, and Michael J Black. Pare: Part attention regressor for 3d human body estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 11127–11137, 2021. [3](#)
- [24] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019. [1](#), [3](#), [4](#), [6](#)

- [25] Nikos Kolotouros, Georgios Pavlakos, Dinesh Jayaraman, and Kostas Daniilidis. Probabilistic modeling for human mesh recovery. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 11605–11614, 2021. 1, 3
- [26] Jiefeng Li, Chao Xu, Zhicun Chen, Siyuan Bian, Lixin Yang, and Cewu Lu. Hybrik: A hybrid analytical-neural inverse kinematics solution for 3d human pose and shape estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3383–3393, 2021. 1
- [27] Zhihao Li, Jianzhuang Liu, Zhensong Zhang, Songcen Xu, and Youliang Yan. Cliff: Carrying location information in full frames into human pose and shape estimation. In *European Conference on Computer Vision*, pages 590–606. Springer, 2022. 1, 3
- [28] Kevin Lin, Lijuan Wang, and Zicheng Liu. End-to-end human pose and mesh reconstruction with transformers. In *CVPR*, 2021. 1
- [29] Kevin Lin, Chung-Ching Lin, Lin Liang, Zicheng Liu, and Lijuan Wang. Mpt: Mesh pre-training with transformers for human pose and mesh reconstruction. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3415–3425, 2024. 3
- [30] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, 2015. 1, 3
- [31] Xiaoxuan Ma, Jiajun Su, Chunyu Wang, Wentao Zhu, and Yizhou Wang. 3d human mesh estimation from virtual markers. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 534–543, 2023. 1, 3
- [32] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. AMASS: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5442–5451, 2019. 1, 3, 5
- [33] Gyeongsik Moon and Kyoung Mu Lee. I2l-meshnet: Image-to-lixel prediction network for accurate 3d human pose and mesh estimation from a single rgb image. In *European Conference on Computer Vision*, pages 752–768. Springer, 2020. 3
- [34] Ramesha Rakesh Mugaludi, Jogendra Nath Kundu, Varun Jampani, et al. Aligning silhouette topology for self-adaptive 3d human pose recovery. *Advances in Neural Information Processing Systems*, 34:4582–4593, 2021. 3
- [35] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, 32, 2019. 5
- [36] 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
- [37] Mathis Petrovich, Michael J Black, and Gül Varol. Action-conditioned 3d human motion synthesis with transformer vae. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10985–10995, 2021. 1
- [38] Davis Remppe, Tolga Birdal, Aaron Hertzmann, Jimei Yang, Srinath Sridhar, and Leonidas J Guibas. Humor: 3d human motion model for robust pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 11488–11499, 2021. 1
- [39] Vinicius F Rodrigues, Rodolfo S Antunes, Lucas A Seewald, Rodrigo Bazo, Eduardo S dos Reis, Uelison JL dos Santos, Rodrigo da R Righi, Luiz G da S Junior, Cristiano A da Costa, Felipe L Bertollo, et al. A multi-sensor architecture combining human pose estimation and real-time location systems for workflow monitoring on hybrid operating suites. *Future Generation Computer Systems*, 135:283–298, 2022. 1
- [40] Akash Sengupta, Ignas Budvytis, and Roberto Cipolla. Synthetic training for accurate 3d human pose and shape estimation in the wild. *arXiv preprint arXiv:2009.10013*, 2020. 1, 3, 6
- [41] Akash Sengupta, Ignas Budvytis, and Roberto Cipolla. Hierarchical kinematic probability distributions for 3d human shape and pose estimation from images in the wild. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 11219–11229, 2021. 1, 3, 6, 7
- [42] Akash Sengupta, Ignas Budvytis, and Roberto Cipolla. Probabilistic 3d human shape and pose estimation from multiple unconstrained images in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 16094–16104, 2021. 1
- [43] Jian Shen, Yanru Qu, Weinan Zhang, and Yong Yu. Wasserstein distance guided representation learning for domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018. 4, 5
- [44] Jie Song, Xu Chen, and Otmar Hilliges. Human body model fitting by learned gradient descent. In *European Conference on Computer Vision*, pages 744–760. Springer, 2020. 1, 2, 3, 6, 7, 8
- [45] Shashank Tripathi, Siddhant Ranade, Amrith Tyagi, and Amit Agrawal. Posenet3d: Learning temporally consistent 3d human pose via knowledge distillation. In *International Conference on 3D Vision*, pages 311–321. IEEE, 2020. 1, 3, 6, 7
- [46] Shashank Tripathi, Lea Müller, Chun-Hao P Huang, Omid Taheri, Michael J Black, and Dimitrios Tzionas. 3d human pose estimation via intuitive physics. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4713–4725, 2023. 1
- [47] Hsiao-Yu Tung, Hsiao-Wei Tung, Ersin Yumer, and Katerina Fragkiadaki. Self-supervised learning of motion capture. *Advances in Neural Information Processing Systems*, 30, 2017. 1
- [48] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7167–7176, 2017. 2, 4
- [49] Timo Von Marcard, Roberto Henschel, Michael J 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*, pages 601–617, 2018. [2](#), [5](#), [6](#), [7](#), [8](#)
- [50] Chung-Yi Weng, Brian Curless, and Ira Kemelmacher-Shlizerman. Photo wake-up: 3d character animation from a single photo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5908–5917, 2019. [1](#)
- [51] Yufei Xu, Jing Zhang, Qiming Zhang, and Dacheng Tao. ViTPose: Simple vision transformer baselines for human pose estimation. In *Advances in Neural Information Processing Systems*, 2022. [2](#)
- [52] Zhenbo Yu, Junjie Wang, Jingwei Xu, Bingbing Ni, Chenglong Zhao, Minsi Wang, and Wenjun Zhang. Skeleton2mesh: Kinematics prior injected unsupervised human mesh recovery. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8619–8629, 2021. [1](#), [3](#)
- [53] Andrei Zanfir, Eduard Gabriel Bazavan, Hongyi Xu, William T Freeman, Rahul Sukthankar, and Cristian Sminchisescu. Weakly supervised 3d human pose and shape reconstruction with normalizing flows. In *European Conference on Computer Vision*, pages 465–481. Springer, 2020. [1](#)
- [54] Andrei Zanfir, Eduard Gabriel Bazavan, Mihai Zanfir, William T Freeman, Rahul Sukthankar, and Cristian Sminchisescu. Neural descent for visual 3d human pose and shape. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 14484–14493, 2021. [1](#)
- [55] Hongwen Zhang, Yating Tian, Xinchu Zhou, Wanli Ouyang, Yebin Liu, Limin Wang, and Zhenan Sun. Pymaf: 3d human pose and shape regression with pyramidal mesh alignment feedback loop. In *Proceedings of the IEEE International Conference on Computer Vision*, 2021. [3](#)
- [56] Siwei Zhang, Yan Zhang, Federica Bogo, Marc Pollefeys, and Siyu Tang. Learning motion priors for 4d human body capture in 3d scenes. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 11343–11353, 2021. [1](#)
- [57] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5745–5753, 2019. [4](#)