

TokenHMR: Advancing Human Mesh Recovery with a Tokenized Pose Representation

Sai Kumar Dwivedi^{1,*} Yu Sun^{2,*} Priyanka Patel¹ Yao Feng^{1,2,3} Michael J. Black¹

¹Max Planck Institute for Intelligent Systems, Tübingen, Germany ²Meshcapade ³ETH Zurich




Figure 1. Existing methods for 3D human pose and shape (HPS) regression (e.g. HMR2.0 [11]) estimate bodies that are either image-aligned or have accurate 3D pose, but not both. We show that this is a fundamental trade-off for existing methods. To address this, TokenHMR, introduces a novel loss, *Threshold-Adaptive Loss Scaling (TALS)*, and a discrete token-based pose representation of 3D pose. With these, TokenHMR achieves state-of-the-art accuracy on multiple in-the-wild 3D benchmarks. The circled regions compare poses.

Abstract

We address the problem of regressing 3D human pose and shape from a single image, with a focus on 3D accuracy. The current best methods leverage large datasets of 3D pseudo-ground-truth (p-GT) and 2D keypoints, leading to robust performance. With such methods, however, we observe a paradoxical decline in 3D pose accuracy with increasing 2D accuracy. This is caused by biases in the p-GT and the use of an approximate camera projection model. We quantify the error induced by current camera models and show that fitting 2D keypoints and p-GT accurately causes incorrect 3D poses. Our analysis defines the invalid distances within which minimizing 2D and p-GT losses is detrimental. We use this to formulate a new loss, “Threshold-Adaptive Loss Scaling” (TALS), that penalizes gross 2D and p-GT errors but not smaller ones. With such a

loss, there are many 3D poses that could equally explain the 2D evidence. To reduce this ambiguity we need a prior over valid human poses but such priors can introduce unwanted bias. To address this, we exploit a tokenized representation of human pose and reformulate the problem as token prediction. This restricts the estimated poses to the space of valid poses, effectively improving robustness to occlusion. Extensive experiments on the EMDB and 3DPW datasets show that our reformulated loss and tokenization allows us to train on in-the-wild data while improving 3D accuracy over the state-of-the-art. Our models and code are available for research at <https://tokenhmr.is.tue.mpg.de>.

1. Introduction

We address the problem of regressing 3D human pose and shape (HPS) from a single image. Recent methods [3, 11, 31, 48] are increasingly accurate on this task.

*Equal contribution

By accurate, we mean two things. A method should correctly regress the 3D pose but it should also align with the image evidence. Unfortunately, current models cannot do both. We observe a seeming paradox, that the more accurate a method is on fitting 2D keypoints, the less accurate it is at predicting 3D pose. We study this problem and identify the common weak-perspective camera assumption as a key culprit. This camera model does not match the true camera used to acquire the images and thus there is a mismatch between the projected 3D joints and the detected 2D ones. Since currently, no reliable method exists to estimate camera parameters from single image, we study and quantify this effect and propose two solutions to address it.

Specifically, we use the synthetic BEDLAM [3] dataset, which has perfect 3D and 2D groundtruth (GT). We project the 3D data into 2D using the camera model from [11] to quantify the 2D error in the best case; it is large. We then also go the other direction and show that low 2D error can result in large 3D error. Even using a full perspective model like [31] does not solve the problem since we lack the precise intrinsic and extrinsic camera parameters.

This analysis highlights the issue of supervising 3D pose regression with a 2D keypoint loss. But such a loss opens up access to large datasets, providing generalization and robustness. Unfortunately, pseudo ground-truth (p-GT) training data suffers the same problem since it is generated by fitting a 3D body to 2D data via optimization using an approximate camera model. How can we leverage the abundant information present in large-scale, in-the-wild, datasets while mitigating the decline in 3D accuracy? Our answer to this is *TokenHMR*, a new HPS regression method that strikes a balance between effectively leveraging 2D keypoints while maintaining 3D pose accuracy, thus leveraging Internet data today without known camera parameters.

TokenHMR has two main components. The first is based on our key insight that supervision from 2D keypoints, while flawed, is valuable for preventing highly incorrect predictions. However, excessive reliance on 2D cues introduces bias. To address this, we define a new loss called *Threshold-Adaptive Loss Scaling (TALS)* that penalizes large 2D and p-GT errors but only minimally penalizes small ones. We use our BEDLAM analysis to define this, so that the network is not encouraged to fit 2D keypoints more accurately than makes sense given the camera model.

This, however, creates a new problem. Predicting 3D pose from 2D keypoints is fundamentally ambiguous. When one relaxes the keypoint matching constraint, even more 3D poses are consistent with the 2D data. To control this, we need to introduce a prior that biases the network to *valid* poses. Unfortunately, existing pose priors based on mixtures of Gaussians [4] or VAEs [41] are biased towards poses that occur frequently in the training data. Instead, we seek an unbiased prior that restricts the network to only out-

put valid poses but does not bias it to any particular pose.

This leads us to the second key component of *TokenHMR*, which gives it its name. Specifically, we convert the problem of continuous pose regression into a problem of token prediction by tokenizing human poses. We use a Vector Quantized-VAE (VQ-VAE) [51] to discretize continuous human poses by pre-training on extensive motion capture datasets, such as AMASS [36] and MOYO [50]. This tokenized representation provides the regressor with a “vocabulary” of valid poses, effectively representing the pose prior as a knowledge bank, *codebook*. Since VQ-VAE’s are designed to represent a uniform prior, we posit that this reduces the biases caused by previous pose priors.

TokenHMR generates discrete tokens through classification, in contrast to regressing continuous pose. When we take a SOTA HPS method and replace the continuous pose with our tokenized pose approach we see consistent improvements in 3D accuracy (all else held the same).

We perform extensive experiments to evaluate different ways of tokenizing pose and their effects on accuracy. Any discretization of pose comes with some loss in accuracy. In our case it results in a loss of 3D accuracy of about 2.5mm, which is 20 times smaller than the accuracy of the state-of-the-art (SOTA) HPS regressors on real data; i.e., the loss in accuracy due to tokenization is negligible.

Finally, we put our two components together and find that they work synergistically. Our new loss does not distort the 3D pose to over-fit the keypoints or p-GT and the tokenization keeps the network from distorting 3D pose for the sake of 2D accuracy. With this combination, we achieve a new state-of-the-art in terms of 3D accuracy. We extensively evaluate *TokenHMR* and other recent methods on EMDB [22] and 3DPW [53], which have accurate 3D ground truth. Using the same data and backbone, *TokenHMR* exhibits a 7.6% reduction in 3D error compared to HMR2.0 [11] on the challenging EMDB dataset. Qualitative results suggest that the *TokenHMR* is robust to ambiguous image evidence and the estimated poses do not suffer from the “bent knees” bias of methods that use p-GT and 2D keypoints (see Fig. 1).

In summary, we make the following key contributions: (1) *Analysis of 3D Accuracy Degradation*: We analyze and quantify the trade-off between 3D and 2D accuracy that current HPS methods face if they use 2D losses. (2) *Threshold-Adaptive Loss Scaling*: To ameliorate the issue, we develop a novel loss function that reduces the influence of 2D and p-GT errors that are less than the expected error due to the incorrect camera model. (3) *Token-Based Pose Representation*: We introduce a token-based representation for human pose and show that it produces more accurate pose estimates. Our models and code are available for research.

2. Related Work

2.1. HPS Regression

Estimating 3D human pose and shape from single images has been studied in great detail from optimization-based approaches to the most recent transformer-based regressors. Optimization approaches fit a parametric model [35, 41, 57] to 2D image cues, including, but not limited to keypoints [4, 41, 57], silhouettes [40], and part segmentations [29]. Some learning-based approaches directly estimate the parametric body model from images [5, 7, 8, 20, 24, 31, 47, 48, 59] and videos [21, 23] and some estimate bodies with a model-free approach either as vertices [27, 33, 45] or as implicit shapes [38, 44, 56]. Recent methods [11, 32] use transformers to estimate 3D bodies, achieving the current best accuracy. To address the challenges of generalization, recent methods like EFT [19], NeuralAnnot [39], HMR2.0 [11] and CLIFF [31] use 2D keypoints and p-GT in the training loss, to produce a good alignment between the projected body and the image. Methods like HuManiFlow [46] and POCO [8] model probabilistic HPS to explicitly address pose ambiguity. The problem of 3D accuracy degradation in pursuit of better 2D alignment has been noted but not extensively quantified before our work. Our statistical analysis highlights this bias in existing HMR methods, offering a new perspective on training strategies for this problem.

Some methods [25, 31, 54] address the 3D-to-2D projection error by estimating the camera from a single image. SPEC [25] uses a network to predict camera parameters but does not generalize well, while CLIFF [31] uses an approximation by providing the network with information about the bounding box coordinates of the person in the image. Estimating the camera from a single image is highly ill-posed so this remains a challenging, unsolved, problem. Our approach reduces the impact of using the wrong camera model and can be applied to any HPS regression method.

2.2. Pose Prior

Human pose priors play a pivotal role in various applications like lifting 2D pose to 3D [4, 41] and estimating human pose from images/videos [19, 28]. Early pose priors focus on learning joint limits [1] to avoid poses that are impossible. Gaussian Mixture Models (GMMs) [4] and Generative Adversarial Networks (GANs) [10, 20] are also used to impose prior knowledge during training. Some recent methods use VAEs [41] and normalizing flows [28] as priors. Many of these methods are biased to commonly occurring poses and this bias is passed on the regressor. Methods like Pose-NDF [49] learn a manifold of plausible poses represented as the zero-level set of a neural implicit function. The mapping of invalid to valid poses involves gradient descent, which is an expensive operation when integrated in HPS training. In contrast to prior work, we learn a dis-

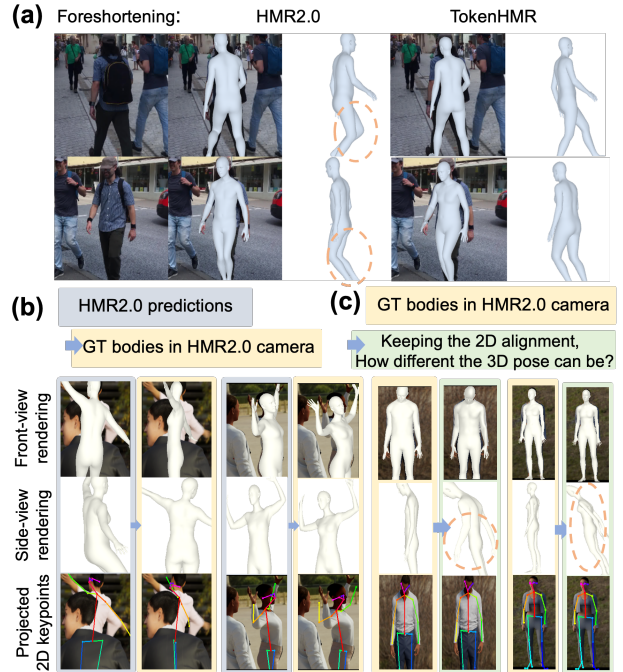


Figure 2. **Visualization of the camera/pose bias issues.** (a) The lack of correct focal length means that foreshortened legs are estimated as bent by methods like HMR2.0. (b) Replacing the predicted body poses with ground truth reveals camera bias; (c) Maintaining 2D alignment, how wrong can the 3D poses be? See Sec. 3 for details.

crete token-based prior over valid SMPL poses, reducing pose bias and improving robustness to occlusion, while being easy to integrate into HPS training.

We use a VQ-VAE [51], which is a variant of VAEs, to learn a discretized prior by quantizing the 3D training poses in a process called “Tokenization” creating a knowledge bank i.e. codebook. Tokenization is widely used in various applications like image synthesis [43, 51], text-to-image generation [42], 2D human pose estimation [9], and learning motion priors [17, 60]. In the context of human pose estimation, tokenization remains relatively unexplored, though it is widely used in human motion generation [13, 60]. Of course, tokenized representations of images and language are widely used for many vision and language problems. Our approach is novel in that it reformulates the regression problem as a pose token classification problem. It thus exploits tokenization to represent valid poses, effectively providing a pose prior.

3. Camera/pose Bias

Methods that estimate 3D HPS typically try to satisfy two goals: accurate 3D pose and accurate alignment with 2D image features. Unfortunately, we observe a trend in all experiments – the better a method does on 2D error, the

worse it does on 3D and vice versa. The key reason for this is that current methods, including those tested here, do not estimate the camera intrinsic parameters (e.g. focal length) or the camera extrinsics (rotation and translation). Instead, current methods estimate the person in camera coordinates using scaled orthographic projection or perspective projection with fixed and incorrect camera parameters. This results in a mismatch between the true 3D joints and their 2D projection. Specifically, since photos are typically taken from roughly eye height, the legs are further away than the upper body. This causes them to be foreshortened. Training models to minimize 2D error forces them to generate incorrect poses in 3D; this is illustrated in Fig. 2 (a). Pseudo ground truth (p-GT) for 2D pose datasets is obtained by minimizing the 2D error with problematic camera parameters. Fully trusting such p-GT and pursuing accurate learning of such annotations will make the problem more prominent. Notice how foreshortening makes the legs appear shorter in the image. The only way to make a human body fit this is to bend the legs at the knees or tilt the body in 3D, making the legs further away. This produces unnatural or unstable poses.

This is a fundamental issue with all current methods and one cannot get low error for both 3D and 2D without knowing the camera. To numerically evaluate the impact of this mismatch, we employ BEDLAM [3], a synthetic dataset where both 3D and 2D data are known exactly along with the camera. This removes any possible noise and allows us to see the effects of using the wrong camera on 2D projection error. Specifically, as shown in Fig. 2 (b), we take ground truth BEDLAM bodies and project them into the image using the camera of HMR2.0 [11].

We evaluate the effect of the incorrect camera in 2D using the standard measure of Percentage of Correct Keypoints (PCK), which we compute for a sequence from the BEDLAM validation set. The 3D bodies computed by HMR2.0b have errors of 0.78 on PCK0.5 and 0.88 on PCK1.0. In contrast, when we use the HMR2.0b camera with the *ground truth* 3D bodies, the PCK scores *decrease* to 0.66 on PCK0.5 and 0.86 on PCK1.0. Ideally, with a correct camera model, both PCK0.5 and PCK1.0 should reach 1.0. The fact that HMR2.0b achieves lower error than the ground truth indicates that its output deviates from the true 3D pose and shape due to camera bias. This demonstrates that methods like HMR2.0b, while obtaining high PCK values, do so at the expense of 3D accuracy. In summary, seeking high PCK values is counterproductive to 3D accuracy unless one has the correct camera model.

We further design experiments to explore how bad the 3D error can be while maintaining good 2D alignment. We modify the loss function of SMPLify [4] to keep the distance between predicted 2D keypoints and GT 2D keypoints J_{2D_g} close, while adding a new loss to *increase* the dis-

tance between predicted 3D keypoints J_{3D} and real 3D keypoints J_{3D_g} , as expressed in the following equation:

$$w_{2D} \|\Pi(J_{3D}, T) - J_{2D_g}\|_2 - w_{3D} \|J_{3D} - J_{3D_g}\|_2 + m \quad (1)$$

where Π represents 3D-to-2D projection using HMR2.0’s camera, $m = 20$ is the margin value, $w_{2D} = 4$ and $w_{3D} = 40.5$ are scalar weights. After 100 iterations of optimization, the Mean Per Joint Position Error (MPJPE) reaches 146mm. As shown in Fig. 2 (c), the projected 3D pose can still maintain a high degree of 2D alignment even with significant errors in the depth direction. When optimized for 200 iterations, the MPJPE exceeds 300mm, and the error continues to increase with further optimization.

Since the field does not currently have a reliable way to estimate the camera parameters from a single image, below we explore the ability of our new methods (TALS and tokenization) to help mitigate the issues caused by approximate camera models. Figure 2 (a) compares results from HMR2.0 and TokenHMR. Note that the effect of foreshortening has less impact on pose with TokenHMR.

4. TokenHMR

4.1. Preliminaries

Our method, TokenHMR, takes an input image, I , and outputs body pose, θ , shape, β and perspective camera, T . We use SMPL [35], a differentiable parametric body model. Its input parameters consist of pose, denoted by $\theta \in \mathbb{R}^{72}$ and shape, denoted as $\beta \in \mathbb{R}^{10}$. As output, it produces a body mesh, \mathcal{M} , and vertices, $V \in \mathbb{R}^{N \times 3}$, where $N = 6890$ is the number of vertices. 3D joints denoted as J_{3D} , are derived through a linear combination of mesh vertices using a pre-trained joint regressor.

4.2. Threshold-Adaptive Loss Scaling: TALS

In Section 3, our analysis reveals a notable impediment to the effective learning using pseudo-ground-truth and 2D keypoints—camera/pose bias. Despite this challenge, the scale provided by such annotations remains integral to achieving optimal generalization and robustness. We assert that, when appropriately utilized, without over-fitting, these annotations significantly enhance the model’s ability to robustly estimate pose. A key insight emerges from our observations: establishing an effective threshold is imperative to discern the error levels that yield no additional benefit as a training signal. When the loss surpasses this threshold, conventional learning mechanisms guide pose estimation. Conversely, when the loss falls below this effective threshold, we minimize its impact to prevent over-fitting to the camera/pose bias.

To determine this effective threshold, we analyze the errors obtained using ground truth (GT) 3D poses and a standard (incorrect) camera model. Again we leverage the 3D

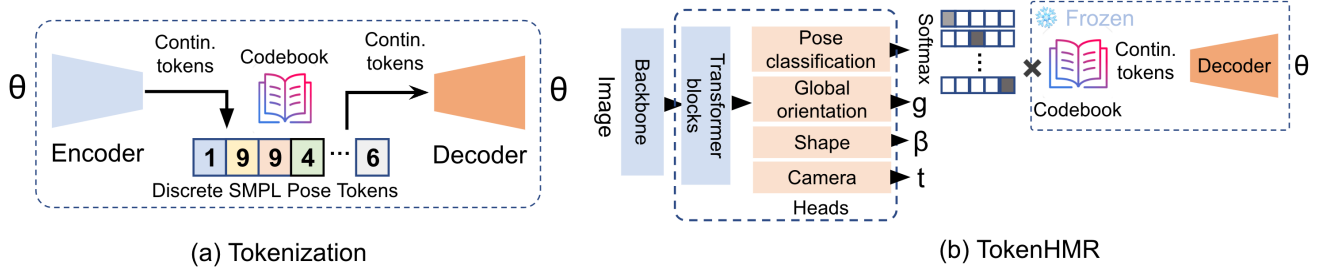


Figure 3. **Framework overview.** Our method has two stages. (a) In the tokenization step, the encoder learns to map continuous poses to discrete pose tokens and the decoder tries to reconstruct the original poses. (b) To train TokenHMR, we replace regression with classification using the pre-trained decoder, which provides a “vocabulary” of valid poses.

GT in BEDLAM [3], this time to establish effective thresholds for both 2D keypoints and SMPL pseudo-ground-truth. For 2D keypoints, we replace the predicted SMPL parameters with ground truth values from BEDLAM to obtain the real 3D human body’s 2D keypoint projections under the HMR2.0 camera. We then calculate the mean L1 norm between these projections and the GT 2D keypoints, and use this as the threshold $\epsilon_{J_{2D}}$ for 2D keypoint supervision. We normalize these 2D keypoints relative to image width and scale the values between -0.5 and 0.5 to mitigate scale-related variances.

Similarly, to establish the effective supervision threshold for SMPL p-GT, we conduct additional experiments. With p-GT, we formulate the pose loss in terms of joint angle error. To set appropriate thresholds on these errors, we evaluate the difference in joint angles between HMR2.0’s predictions on BEDLAM and the ground truth values for each joint in the SMPL model. Specifically, we compute the mean geodesic distance* between the 3D joint rotations on the manifold of rotations in $SO(3)$ predicted by HMR2.0 and the ground-truth rotations in BEDLAM. Please refer to Sec. B.2 in *Sup. Mat.* for specific threshold values.

After establishing the effective thresholds for 2D keypoints and SMPL p-GT, we introduce a new loss called *Threshold-Adaptive Loss Scaling (TALS)*. It scales down the loss only when it goes below the threshold. Specifically, the *TALS* loss terms for p-GT pose and 2D joints are defined as

$$\mathcal{L}_{\theta_{pGT}} = \begin{cases} \|\theta - \theta_g\|^2 & \text{if } \mathcal{L}_{\theta_{pGT}} > \epsilon_\theta \\ \alpha_\theta \cdot \|\theta - \theta_g\|^2 & \text{otherwise} \end{cases} \quad (2)$$

$$\mathcal{L}_{J_{2D}pGT} = \begin{cases} |J_{2D} - J_{2D_g}| & \text{if } \mathcal{L}_{J_{2D}pGT} > \epsilon_{J_{2D}} \\ \alpha_{J_{2D}} \cdot |J_{2D} - J_{2D_g}| & \text{otherwise} \end{cases} \quad (3)$$

where $\alpha_{J_{2D}}$ and α_θ are small scalar multipliers and ϵ_θ is the threshold calculated separately for each pose parameter and $\epsilon_{J_{2D}}$ is the threshold for 2D joints.

*<https://rotations.berkeley.edu/geodesics-of-the-rotation-group-so3/>

4.3. Tokenization

We use a VQ-VAE [51], which learns an encoding of 3D pose in a discrete representation. Specifically, we learn a discrete representation for SMPL body parameters, $\theta = [\theta_1, \theta_2, \dots, \theta_{21}]$, where the θ_i represent each joint’s pose parameters in \mathbb{R}^6 . The process involves encoding and decoding the pose parameters using an autoencoder architecture and a learnable codebook, denoted as $CB = \{c_k\}_{k=1}^K$, with each code $c_k \in \mathbb{R}^{d_c}$, where d_c is the dimension of the codes. The overall architecture of the Pose VQ-VAE is illustrated in Fig. 3 (a). The encoder and decoder of the autoencoder are represented by E and D , respectively. The encoder is responsible for generating discrete pose tokens, while the decoder reconstructs these tokens back to SMPL poses. The latent feature z can be computed as $z = E(\theta)$, resulting in $z = [z_1, z_2, \dots, z_M]$, where $z_i \in \mathbb{R}^{d_c}$ and M is the number of tokens. Each latent feature z_i is quantized using the codebook CB by finding the most similar code element, as expressed in the following equation:

$$\hat{z}_i = \arg \min_{c_k \in CB} \|z_i - c_k\|_2. \quad (4)$$

In training the pose tokenizer, we adopt a strategy similar to previous work [51], which involves three primary loss functions to optimize the tokenizer: the reconstruction loss ($\mathcal{L}_{\mathcal{R}\mathcal{E}}$), the embedding loss ($\mathcal{L}_\mathcal{E}$), and the commitment loss ($\mathcal{L}_\mathcal{C}$). The overall loss (\mathcal{L}_{VQ}) is defined as

$$\begin{aligned} \mathcal{L}_{VQ} &= \lambda_{\mathcal{R}\mathcal{E}} \mathcal{L}_{\mathcal{R}\mathcal{E}} + \lambda_\mathcal{E} \mathcal{L}_\mathcal{E} + \lambda_\mathcal{C} \mathcal{L}_\mathcal{C} \\ &= \lambda_{\mathcal{R}\mathcal{E}} \mathcal{L}_{\mathcal{R}\mathcal{E}} + \lambda_\mathcal{E} \|sg[z] - e\|_2 + \lambda_\mathcal{C} \|z - sg[e]\|_2 \end{aligned} \quad (5)$$

where sg is the stop gradient operator, e is the embedding from the codebook and $\lambda_{\mathcal{R}\mathcal{E}}$, $\lambda_\mathcal{E}$, $\lambda_\mathcal{C}$ are the hyper-parameters of for each term. For reconstruction, we use an \mathcal{L}_1 loss between the ground-truth pose, θ_g , and predicted pose, θ and also on the error between the SMPL ground-truth 3D joints, J_{3D_g} and predicted joints, J_{3D} . So, the $\mathcal{L}_{\mathcal{R}\mathcal{E}}$ loss is defined as

$$\mathcal{L}_{\mathcal{R}\mathcal{E}} = \mathcal{L}_1(\theta_g, \theta) + \mathcal{L}_1(J_{3D_g}, J_{3D}). \quad (6)$$

The original VQ-VAE suffers from codebook collapse, i.e. the codebook is not fully utilized. Following prior work [60], we use the training strategy of exponential moving average (EMA) and codebook reset (Code Reset) for better utilization.

4.4. Architecture

Our architecture exploits the Vision Transformer (ViT) [6], similar to HMR2.0 [11]. The input image, I , is first transformed into input tokens, which are subsequently processed by the transformer to generate output tokens. These output tokens then undergo further processing in the transformer decoder. The transformer decoder has multi-head self-attention that cross-attends a zero input token with an image output token to get features from the transformer block. In contrast to HMR2.0 [11], which employs three linear layers to map the features from transformer block to the SMPL pose, θ , shape, β and camera, T , we propose a novel approach. Our objective is to leverage a tokenizer trained on a significant amount of motion capture (mocap) data, specifically focusing on body pose. To facilitate this, we partition the SMPL pose parameters into body pose and global orientation. We use separate linear layers to predict the global orientation and body pose from the tokenizer.

A straightforward integration of the tokenizer would involve estimating the code index directly from the ViT transformer backbone and selecting embeddings, e , based on the code index from the codebook, CB . However, this poses a challenge as the process of selecting an embedding from the codebook, is non-differentiable. To address this issue, we adopt a logit-based approach. Instead of directly estimating the code index, we output logits, Q for each token. These logits are multiplied with the codebook, resulting in weighted embeddings. Thus, the approximated quantized feature $\bar{z} = [\bar{z}_1, \bar{z}_2, \dots, \bar{z}_M]$ can be calculated as

$$\bar{z} = \sigma(Q_{M \times K}) \times CB_{K \times D} \approx \hat{z} \quad (7)$$

where, Q are the logits estimated by the backbone, σ is the softmax operation, CB is the pretrained codebook, M is the number of tokens, K is the number of entries in the codebook and D is the dimension of each codebook entry. The operation makes it differentiable. The obtained approximated quantized features, \bar{z} are subsequently passed through the tokenizer decoder. This process yields the final pose. In the training of TokenHMR, the learned tokenizer decoder is frozen to take advantage of the prior it has learned from mocap data.

4.5. Losses

Following prior work [11, 24], we define losses on 2D and 3D joints and SMPL pose and shape parameters, i.e. on J_{2D} , J_{3D} , θ , β , respectively. However, following the analysis in Sec. 4.2, we treat data from 2D and 3D datasets dif-

ferently. For 3D ground-truth datasets, we define the standard loss as

$$\mathcal{L}_{GT} = \lambda_{\theta} \mathcal{L}_{\theta}(\theta, \theta_g) + \lambda_{\beta} \mathcal{L}_{\beta}(\beta, \beta_g) + \lambda_{3D} \mathcal{L}_{3D}(J_{3D}, J_{3D_g}) + \lambda_{2D} \mathcal{L}_{2D}(J_{2D}, J_{2D_g}) \quad (8)$$

where \mathcal{L}_{β} is a SMPL shape loss, $\mathcal{L}_{J_{3D}}$ is the 3D joint loss and $\mathcal{L}_{J_{2D}}$ is the joint re-projection loss. λ_{β} , λ_{3D} and λ_{2D} are steering weights for each term. To learn from SMPL pseudo-ground-truth, we use *Threshold-Adaptive Loss Scaling (TALS)* where we scale the loss based on the threshold computed in Sec. 4.2, outlined in Eqs. 2 and 3. Thus, the total loss is defined as

$$\mathcal{L}_{Total} = \mathcal{L}_{GT} + \mathcal{L}_{\theta_{pGT}} + \mathcal{L}_{J_{2D}_{pGT}}. \quad (9)$$

5. Experiments

5.1. Implementation Details

Training of TokenHMR involves two stages: first we train a tokenizer to learn discrete pose representations using AMASS [36] and MOYO [50] mocap data. Then we use the pretrained decoder of the tokenizer as an additional head for regressing body pose. During the training of TokenHMR, the tokenizer is frozen to exploit the prior.

Our tokenizer architecture is inspired from T2M-GPT [60] but instead of learning motion tokens of 3D joints we learn pose tokens of SMPL pose parameters. We use 1 ResNet [14] block and 4 1D convolutions both in the encoder and decoder. The steering weights $\lambda_{\mathcal{R}\mathcal{E}}$, $\lambda_{\mathcal{E}}$, $\lambda_{\mathcal{C}}$ are set at 50.0, 1.0, 1.0, respectively. The model is trained for 150K iterations with batch size of 256 and learning rate of $2e^{-4}$. To train a robust model, we augment random joints with noise starting from $1e^{-3}$, which we progressively increase after every 5K iterations. We choose the best tokenizer model containing 160 tokens and codebook of size 2048×256 for TokenHMR based on reconstruction error on the validation set.

For TokenHMR, we use ViT-H/16 [6] as the backbone and standard transformer decoder [52] following HMR2.0 [11]. We use 4 separate linear layers to map the features of size 1024 from the transformer decoder to the global orientation, hand pose, and body shape of SMPL and one for the camera. However, for body pose, we process the 1024 features through 4 blocks of linear layers, each containing 2 MLPs and an GELU activation function [15]. This gives the final logits, Q , of size 160×2048 for multiplication with a codebook of size 2048×256 , which results in approximate quantized features, \bar{z} ; see Eq. 4. We use ViT-Pose [58] as the pretrained backbone. We train for 100K iterations on 4 Nvidia RTX 6000 GPUs with a batch size of 256 and learning rate of $1e^{-5}$ for about one day. The steering weights, λ_{θ} , λ_{β} , $\lambda_{J_{2D}}$, $\lambda_{J_{3D}}$ are set to $1e^{-3}$, $5e^{-4}$, $1e^{-2}$, $5e^{-2}$, respectively. The loss weights of TALS is set to 1% for both pose α_{θ} and 2D keypoints $\alpha_{J_{2D}}$.

Training Datasets	Method	EMDB [22]			3DPW [53]		
		MVE	MPJPE	PA-MPJPE	MVE	MPJPE	PA-MPJPE
SD	HybrIK [30]	122.2	103.0	65.6	94.5	80.0	48.8
SD	CLIFF [31]	122.9	103.1	68.8	81.2	69.0	43.0
SD	HMR2.0 [11]	120.1	97.8	61.5	84.1	70.0	44.5
BL	BEDLAM-CLIFF [3]	113.2	97.1	61.3	85.0	72.0	46.6
BL	HMR2.0	106.6	90.7	51.3	88.4	72.2	45.1
BL	TokenHMR	104.2	88.1	49.8	86.0	70.5	43.8
SD + ITW	HMR2.0 [11]	140.6	118.5	79.3	94.4	81.3	54.3
SD + ITW	TokenHMR	124.4	102.4	67.5	88.1	76.2	49.3
SD + ITW + BL	HMR2.0	120.7	99.3	62.8	88.4	77.4	47.4
SD + ITW + BL	HMR2.0 + TALS	115.7	96.7	58.5	89.6	73.5	46.8
SD + ITW + BL	HMR2.0 + Token	116.1	95.6	62.2	86.6	75.0	48.0
SD + ITW + BL	HMR2.0 + TALS + VPoser [41]	116.8	97.9	56.4	87.1	73.7	45.7
SD + ITW + BL	TokenHMR	109.4	91.7	55.6	84.6	71.0	44.3

Table 1. 3D human mesh and pose errors on the EMDB and 3DPW datasets. See text.

Method	Crop 30%			Crop 50%		
	MVE	MPJPE	PA-MPJPE	MVE	MPJPE	PA-MPJPE
HMR2.0 [11]	135.24 (+14.98)	113.39 (+14.13)	70.68 (+7.86)	166.71 (+46.45)	137.88 (+38.59)	90.30 (+27.48)
TokenHMR	124.09 (+14.71)	104.72 (+13.01)	62.13 (+6.52)	150.29 (+40.91)	125.99 (+34.28)	78.88 (+23.27)

Table 2. Impact of evenly cropping images at different ratios from the boundaries on the 3D HPS accuracy on the EMDB dataset. The numbers in (parentheses) indicate the changes in performance relative to the non-cropped scenario; smaller is better. All models compared here employ identical backbones and are trained on the same data.

	Method	AMASS [36]		MOYO [50]	
		MVE ↓	MPJPE ↓	MVE ↓	MPJPE ↓
CB	1024 × 256	11.5	4.6	27.1	15.7
	2048 × 128	9.4	3.1	22.5	12.3
	2048 × 256	8.3	2.2	19.9	10.4
Tokens	80	12.5	4.1	24.4	16.7
	160	8.3	2.2	19.9	10.4
	320	8.1	1.9	19.0	10.1
Noise	Yes	8.3	2.2	19.9	10.4
	No	7.9	1.9	21.0	11.5
AMASS + MOYO*		8.7	2.6	16.5	7.6

Table 3. **Tokenizer Ablation.** All methods are trained on the standard training set of AMASS [36] and evaluated on the test set of AMASS and validation set of MOYO [50] except the last row*, which is trained with the MOYO training set. The last model is used as the tokenizer in TokenHMR.

Training Data: For training the tokenizer, we use the standard training split of AMASS [36] and the training data of MOYO [50]. For more details on data preparation of training, please refer to *Sup. Mat.* Following the prior methods [11, 24, 26], we use standard datasets (*SD*) for training which include Human3.6M [16], MPI-INF-3DHP [37], COCO [34], and MPII [2]. Additionally, like HMR2.0b, we also use in-the-wild 2D datasets (*ITW*) like InstaVariety [21], AVA [12], and AI Challenger [55] datasets and their p-GT for training. We also include BEDLAM (*BL*) [3],

a synthetic dataset with accurate ground-truth 3D data. For a fair comparison, we re-train HMR2.0b using a combination of the *SD*, *ITW*, and *BL* datasets. We choose HMR2.0b as a baseline model since the code is open-source and we can reproduce the results.

Evaluation and Metrics: For the tokenizer accuracy, we report the Mean Vertex Error (MVE) and Mean Per Joint Position Error (MPJPE) and evaluate on the standard test split of AMASS and validation set of MOYO. For TokenHMR, we report the Mean Vertex Error (MVE), Mean Per Joint Position Error (MPJPE), and Procrustes-Aligned Mean Per Joint Position Error (PA-MPJPE) between the predictions and the ground-truth. We evaluate on the test set of the 3DPW [53] and EMDB [22] datasets. The former is a standard 3D dataset and the latter is a recently released and more challenging dataset with varying camera motions and varied 3D poses.

5.2. How to Alleviate the 3D Degradation Problem?

Table 1 shows the performance of the HMR2.0 model trained solely with the *SD* dataset and its performance when trained with both *SD* and *ITW* datasets. We observe a significant decrease (over 17%) in 3D accuracy on the EMDB dataset upon the inclusion of the *ITW* data. At the same time, according to the Table 2 in HMR2.0 [11] paper, the *ITW* data improves the method’s 2D performance. A straightforward approach to counter this trend could be to




Figure 4. Qualitative comparisons on challenging poses from the LSP [18] dataset.

integrate more data with precise 3D annotations, such as BEDLAM [3]. Yet, as Table 1 reveals, even with the inclusion of BEDLAM, HMR2.0 (*SD+ITW+BL*) still suffers from noticeable 3D metric degradation. This observation forms our baseline for further investigations.

Employing our novel loss formulation (*TALS*) results in notable performance improvement on both the EMDB and 3DPW datasets, indicating its effectiveness in preventing overfitting to noisy p-GT data. While *TALS* yields improvements, we delve deeper into exploring pose priors to compensate for the diminished supervision. Our evaluation of VPoser [41], a prevalent VAE prior in HPS, yields only marginal improvements, suggesting the need for a more robust alternative. Our VQ-VAE-based pose tokenization approach offers greater improvement. TokenHMR significantly outperforms HMR2.0, with improvements of 9% in MVE, 7.6% in MPJPE, and 11.5% in PA-MPJPE on EMDB. Consistent trends are also observed on 3DPW.

5.3. How does the Token-based Prior Help?

Beyond facilitating more effective learning as we discussed above, we also examine the efficacy of our discrete token-based prior in scenarios with ambiguous image information, such as truncation. We evaluate our method under varying degrees of image cropping on the EMDB dataset. Specifically, we crop 30% and 50% from the image boundaries. As shown in Table 2, compared with HMR2.0 [11], the performance of our approach decreases less in the challenging truncation settings (50% v.s. 30%). Furthermore, the qualitative outcomes, illustrated in Fig. 4, underscore the robustness of the prior embedded within our token-based pose representation. This robustness is crucial for handling real-world scenarios where image truncation is common.

5.4. Ablation Study of Tokenizer

Table 3 presents our ablation study on different tokenizer design choices using AMASS’s standard test set and MOYO’s validation set. To understand the impact of the design choices on out-of-distribution MOYO data, we train solely with AMASS and conduct various ablations. The final tokenizer model (last row in Table 3), however, is

used in TokenHMR and is trained on both the AMASS and MOYO datasets. Our findings indicate that the number of codebook entries has a more significant impact than code dimensions. Although the number of tokens is crucial for an accurate representation, we observe a performance plateau, opting for 160 tokens in our final model. This number strikes a balance between network size and reconstruction accuracy for TokenHMR. Random augmentation of pose parameters with noise builds a more robust tokenizer, slightly reducing performance for in-distribution data but beneficially impacting OOD data.

6. Conclusion

In this paper, we presented a novel approach to 3D human pose estimation from single images. We begin by identifying and quantifying the problem caused by using a 2D key-point loss with an incorrect camera model. This leads to a fundamental tradeoff for current methods – either have high 3D accuracy or 2D accuracy, but not both. Our method, TokenHMR, addresses this problem with two contributions that can easily be used by other methods. TokenHMR adopts a new paradigm for pose estimation based on regressing a discrete tokenized representation of human pose. We combine this with a new loss, *TALS*, which mitigates some of the bias caused by the camera projection error, and biased p-GT, while still allowing the use of in-the-wild training data. Our experiments on the EMDB and 3DPW datasets demonstrate that TokenHMR significantly outperforms existing models like HMR2.0 in terms of 3D accuracy, even with significant occlusion.

7. Acknowledgement

We sincerely thank the department of Perceiving Systems and ML team of Meshcapade GmbH for insightful discussions and feedback. We thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting Sai Kumar Dwivedi. We thank Meshcapade GmbH for supporting Yu Sun and providing GPU resources. This work was partially supported by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039B. **Disclosure:** https://files.is.tue.mpg.de/black/CoI_CVPR_2024.txt

References

- [1] Ijaz Akhter and Michael J. Black. Pose-conditioned joint angle limits for 3D human pose reconstruction. In *Computer Vision and Pattern Recognition (CVPR)*, 2015. 3
- [2] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *Computer Vision and Pattern Recognition (CVPR)*, 2014. 7
- [3] Michael J Black, Priyanka Patel, Joachim Tesch, and Jinlong Yang. Bedlam: A synthetic dataset of bodies exhibiting detailed lifelike animated motion. In *CVPR*, pages 8726–8737, 2023. 1, 2, 4, 5, 7, 8
- [4] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J. Black. Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image. In *European Conference on Computer Vision (ECCV)*, 2016. 2, 3, 4
- [5] Hongsuk Choi, Gyeongsik Moon, JoonKyu Park, and Kyoung Mu Lee. Learning to estimate robust 3d human mesh from in-the-wild crowded scenes. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1475–1484, 2022. 3
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations (ICLR)*, 2021. 6
- [7] Sai Kumar Dwivedi, Nikos Athanasiou, Muhammed Kocabas, and Michael J. Black. Learning to regress bodies from images using differentiable semantic rendering. In *International Conference on Computer Vision (ICCV)*, 2021. 3
- [8] Sai Kumar Dwivedi, Cordelia Schmid, Hongwei Yi, Michael J. Black, and Dimitrios Tzionas. POCO: 3D pose and shape estimation using confidence. In *International Conference on 3D Vision (3DV)*, 2024. 3
- [9] Zigang Geng, Chunyu Wang, Yixuan Wei, Ze Liu, Houqiang Li, and Han Hu. Human pose as compositional tokens. In *Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [10] G. Georgakis, Ren Li, S. Karanam, Terrence Chen, J. Kosecka, and Ziyang Wu. Hierarchical kinematic human mesh recovery. *European Conference on Computer Vision (ECCV)*, 2020. 3
- [11] Shubham Goel, Georgios Pavlakos, Jathushan Rajasegaran, Angjoo Kanazawa, and Jitendra Malik. Humans in 4D: Reconstructing and tracking humans with transformers. In *International Conference on Computer Vision (ICCV)*, 2023. 1, 2, 3, 4, 6, 7, 8
- [12] Chunhui Gu, Chen Sun, David A Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, et al. Ava: A video dataset of spatio-temporally localized atomic visual actions. In *Computer Vision and Pattern Recognition (CVPR)*, pages 6047–6056, 2018. 7
- [13] Chuan Guo, Xinxin Zuo, Sen Wang, and Li Cheng. Tm2t: Stochastic and tokenized modeling for the reciprocal generation of 3d human motions and texts. In *European Conference on Computer Vision (ECCV)*, 2022. 3
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Computer Vision and Pattern Recognition (CVPR)*, 2016. 6
- [15] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv: 1606.08415*, 2016. 6
- [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. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2014. 7
- [17] Biao Jiang, Xin Chen, Wen Liu, Jingyi Yu, Gang Yu, and Tao Chen. Motiongpt: Human motion as a foreign language. *arXiv preprint arXiv: 2306.14795*, 2023. 3
- [18] Sam Johnson and Mark Everingham. Learning effective human pose estimation from inaccurate annotation. In *Computer Vision and Pattern Recognition (CVPR)*, 2011. 8
- [19] 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 (3DV)*, 2020. 3
- [20] Angjoo Kanazawa, Michael J. Black, David W. Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *Computer Vision and Pattern Recognition (CVPR)*, 2018. 3
- [21] Angjoo Kanazawa, Jason Y. Zhang, Panna Felsen, and Jitendra Malik. Learning 3D human dynamics from video. In *Computer Vision and Pattern Recognition (CVPR)*, 2019. 3, 7
- [22] Manuel Kaufmann, Jie Song, Chen Guo, Kaiyue Shen, Tianjian Jiang, Chengcheng Tang, Juan José Zárate, and Otmar Hilliges. EMDB: The Electromagnetic Database of Global 3D Human Pose and Shape in the Wild. In *International Conference on Computer Vision (ICCV)*, 2023. 2, 7
- [23] Muhammed Kocabas, Nikos Athanasiou, and Michael J. Black. VIBE: Video inference for human body pose and shape estimation. In *Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [24] Muhammed Kocabas, Chun-Hao P. Huang, Otmar Hilliges, and Michael J. Black. PARE: Part attention regressor for 3D human body estimation. In *International Conference on Computer Vision (ICCV)*, 2021. 3, 6, 7
- [25] Muhammed Kocabas, Chun-Hao P. Huang, Joachim Tesch, Lea Müller, Otmar Hilliges, and Michael J. Black. SPEC: Seeing people in the wild with an estimated camera. In *International Conference on Computer Vision (ICCV)*, 2021. 3
- [26] Nikos Kolotouros, Georgios Pavlakos, Michael J. Black, and Kostas Daniilidis. Learning to reconstruct 3D human pose and shape via model-fitting in the loop. In *International Conference on Computer Vision (ICCV)*, 2019. 7
- [27] Nikos Kolotouros, Georgios Pavlakos, and Kostas Daniilidis. Convolutional mesh regression for single-image human shape reconstruction. In *Computer Vision and Pattern Recognition (CVPR)*, 2019. 3

- [28] Nikos Kolotouros, Georgios Pavlakos, Dinesh Jayaraman, and Kostas Daniilidis. Probabilistic modeling for human mesh recovery. In *International Conference on Computer Vision (ICCV)*, 2021. 3
- [29] Christoph Lassner, Javier Romero, Martin Kiefel, Federica Bogo, Michael J. Black, and Peter Gehler. Unite the people: Closing the loop between 3D and 2D human representations. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 3
- [30] Jiefeng Li, Chao Xu, Zhicun Chen, Siyuan Bian, Lixin Yang, and Cewu Lu. HybrIK: A hybrid analytical-neural inverse kinematics solution for 3D human pose and shape estimation. In *Computer Vision and Pattern Recognition (CVPR)*, 2021. 7
- [31] 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 (ECCV)*, 2022. 1, 2, 3, 7
- [32] Jing Lin, Ailing Zeng, Haoqian Wang, Lei Zhang, and Yu Li. One-stage 3D whole-body mesh recovery with component aware transformer. In *Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [33] Kevin Lin, Lijuan Wang, and Zicheng Liu. Mesh graphormer. In *International Conference on Computer Vision (ICCV)*, 2021. 3
- [34] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision (ECCV)*, 2014. 7
- [35] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. In *Transactions on Graphics (TOG)*, 2015. 3, 4
- [36] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. AMASS: Archive of motion capture as surface shapes. In *International Conference on Computer Vision (ICCV)*, 2019. 2, 6, 7
- [37] Dushyant Mehta, Helge Rhodin, Dan Casas, Pascal V. Fua, Oleksandr Sotnychenko, Weipeng Xu, and Christian Theobalt. Monocular 3d human pose estimation in the wild using improved cnn supervision. In *International Conference on 3D Vision (3DV)*, 2017. 7
- [38] Marko Mihajlovic, Shunsuke Saito, Aayush Bansal, Michael Zollhoefer, and Siyu Tang. COAP: Compositional articulated occupancy of people. In *Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [39] Gyeongsik Moon, Hongsuk Choi, and Kyoung Mu Lee. Neuralannot: Neural annotator for 3d human mesh training sets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2299–2307, 2022. 3
- [40] Mohamed Omran, Christoph Lassner, Gerard Pons-Moll, Peter Gehler, and Bernt Schiele. Neural body fitting: Unifying deep learning and model based human pose and shape estimation. In *International Conference on 3D Vision (3DV)*, 2018. 3
- [41] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 3, 7, 8
- [42] A. Ramesh, Mikhail Pavlov, Gabriel Goh, S. Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. *International Conference on Machine Learning (ICML)*, 2021. 3
- [43] Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. *Conference on Neural Information Processing Systems (NeurIPS)*, 2019. 3
- [44] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. PIFuHD: Multi-level pixel-aligned implicit function for high-resolution 3D human digitization. In *Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [45] István Sárándi, Timm Linder, Kai O. Arras, and Bastian Leibe. Metric-scale truncation-robust heatmaps for 3D human pose estimation. In *IEEE Int Conf Automatic Face and Gesture Recognition (FG)*, 2020. 3
- [46] Akash Sengupta, Ignas Budvytis, and Roberto Cipolla. HuManiFlow: Ancestor-Conditioned Normalising Flows on SO(3) Manifolds for Human Pose and Shape Distribution Estimation. In *Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [47] Yu Sun, Qian Bao, Wu Liu, Yili Fu, Black Michael J., and Tao Mei. Monocular, One-stage, Regression of Multiple 3D People. In *International Conference on Computer Vision (ICCV)*, 2021. 3
- [48] Yu Sun, Wu Liu, Qian Bao, Yili Fu, Tao Mei, and Michael J Black. Putting People in their Place: Monocular Regression of 3D People in Depth. In *Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 3
- [49] Garvita Tiwari, Dimitrije Antic, Jan Eric Lenssen, Nikolaos Sarafianos, Tony Tung, and Gerard Pons-Moll. Pose-ndf: Modeling human pose manifolds with neural distance fields. In *European Conference on Computer Vision (ECCV)*, 2022. 3
- [50] Shashank Tripathi, Lea Müller, Chun-Hao P. Huang, Taheri Omid, Michael J. Black, and Dimitrios Tzionas. 3D human pose estimation via intuitive physics. In *Computer Vision and Pattern Recognition (CVPR)*, 2023. 2, 6, 7
- [51] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Conference on Neural Information Processing Systems (NeurIPS)*, 2017. 2, 3, 5
- [52] Ashish Vaswani, Noam M. Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Conference on Neural Information Processing Systems (NeurIPS)*, 2017. 6
- [53] 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 (ECCV)*, 2018. 2, 7
- [54] Wenjia Wang, Yongtao Ge, Haiyi Mei, Zhongang Cai, Qingping Sun, Yanjun Wang, Chunhua Shen, Lei Yang, and Taku

- Komura. Zolly: Zoom focal length correctly for perspective-distorted human mesh reconstruction. In *International Conference on Computer Vision (ICCV)*, 2023. 3
- [55] Jiahong Wu, He Zheng, Bo Zhao, Yixin Li, Baoming Yan, Rui Liang, Wenjia Wang, Shiwei Zhou, Guosen Lin, Yanwei Fu, et al. Ai challenger: A large-scale dataset for going deeper in image understanding. *arXiv*, 2017. 7
- [56] Yuliang Xiu, Jinlong Yang, Xu Cao, Dimitrios Tzionas, and Michael J. Black. ECON: Explicit clothed humans optimized via normal integration. In *Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [57] Hongyi Xu, Eduard Gabriel Bazavan, Andrei Zanfir, William T. Freeman, Rahul Sukthankar, and Cristian Sminchisescu. GHUM & GHUML: Generative 3D human shape and articulated pose models. In *Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [58] Yufei Xu, Jing Zhang, Qiming Zhang, and Dacheng Tao. ViTPose: Simple vision transformer baselines for human pose estimation. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2022. 6
- [59] 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 *International Conference on Computer Vision (ICCV)*, pages 11446–11456, 2021. 3
- [60] Jianrong Zhang, Yangsong Zhang, Xiaodong Cun, Shaoli Huang, Yong Zhang, Hongwei Zhao, Hongtao Lu, and Xi Shen. T2m-gpt: Generating human motion from textual descriptions with discrete representations. In *Computer Vision and Pattern Recognition (CVPR)*, 2023. 3, 6