# ADHMR: Aligning Diffusion-based Human Mesh Recovery via Direct Preference Optimization

Wenhao Shen<sup>1</sup> Wanqi Yin<sup>2</sup> Xiaofeng Yang<sup>1</sup> Cheng Chen<sup>1</sup> Chaoyue Song<sup>1</sup> Zhongang Cai<sup>2</sup> Lei Yang<sup>2</sup> Hao Wang<sup>3</sup> Guosheng Lin<sup>1</sup>

#### Abstract

Human mesh recovery (HMR) from a single image is inherently ill-posed due to depth ambiguity and occlusions. Probabilistic methods have tried to solve this by generating numerous plausible 3D human mesh predictions, but they often exhibit misalignment with 2D image observations and weak robustness to in-the-wild images. To address these issues, we propose ADHMR, a framework that Aligns a Diffusion-based HMR model in a preference optimization manner. First, we train a human mesh prediction assessment model, HMR-Scorer, capable of evaluating predictions even for in-the-wild images without 3D annotations. We then use HMR-Scorer to create a preference dataset, where each input image has a pair of winner and loser mesh predictions. This dataset is used to finetune the base model using direct preference optimization. Moreover, HMR-Scorer also helps improve existing HMR models by data cleaning, even with fewer training samples. Extensive experiments show that ADHMR outperforms current state-of-the-art methods. Code is available at: https://github.com/shenwenhao01/ADHMR.

# 1. Introduction

Human mesh recovery (HMR) is a fundamental challenge in computer vision, focused on estimating the 3D human shape and pose from a single RGB image. HMR enables various downstream applications, including clothed human reconstruction (Shuai et al., 2022; Hong & Shen, 2024; Yao et al., 2025), virtual try-on, AR/VR content creation (Xu et al., 2024c; Yang et al., 2024) and etc. Prevailing approaches usually adopt a deterministic style, generating a single prediction for each image (Cai et al., 2024a; Goel et al., 2023; Li et al., 2023; Moon et al., 2022). However, this task faces inherent uncertainty when lifting 2D observations to 3D models, due to depth ambiguity and occlusion. Accordingly, the community is now shifting to probabilistic methods. Probabilistic methods tackle the uncertainty by generating multiple plausible human mesh predictions for each image (Kolotouros et al., 2021; Sengupta et al., 2023). For instance, recent approaches (Foo et al., 2023; Cho & Kim, 2023) frame this task as a denoising diffusion process. However, these probabilistic approaches suffer from limited emphasis on obtaining accurate estimates.

Specifically, the current state-of-the-art probabilistic method ScoreHypo (Xu et al., 2024b) tackles this by designing an additional network for test-time prediction selection after the diffusion-based prediction model. However, we observe that ScoreHypo still exhibits the following shortcomings: (1) misalignment between 3D mesh predictions and 2D image cues, and (2) poor performance on in-the-wild images. This is primarily because end-to-end diffusion models predicting from pure noise typically avoid 3D reprojection loss, as early denoising steps yield unrealistic poses, making such loss ineffective (Huang et al., 2024). Instead, the diffusion loss focuses on generating the target human mesh distribution rather than precisely aligning joints. While this produces plausible poses, it may neglect the alignment between the 3D mesh and the image. Moreover, existing datasets often use optimization-based HMR methods to generate pseudo 3D annotations for in-the-wild images, which inevitably contain some inaccurate or noisy data.

To address the above challenges, we introduce Aligned Diffusion-based Human Mesh Recovery (ADHMR). The key insight is to distill the knowledge of a powerful scorer into the 3D human mesh predictor in a preference optimization manner. Technically, we begin by training the HMR-Scorer, essentially a reward model that assigns a quality score to quantify the human mesh prediction quality. HMR-Scorer gives higher scores to predictions better aligned with the input image ("winners") and lower scores to those poorly aligned ("losers"). In order to increase the sensitivity of

<sup>&</sup>lt;sup>1</sup>Nanyang Technological University <sup>2</sup>SenseTime Research <sup>3</sup>The Hong Kong University of Science and Technology (Guangzhou). Correspondence to: Guosheng Lin <gslin@ntu.edu.sg>, Hao Wang <haowang@hkust-gz.edu.cn>.

Proceedings of the  $42^{nd}$  International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).

HMR-Scorer to nuances in the image cues, we extract multiscale image features as conditions, which provide global and local pixel-aligned features sampled by reprojected human keypoints, enabling HMR-Scorer to identify misalignment between the predicted mesh and 2D image cues.

Then we draw on the concept of direct preference optimization (DPO) for diffusion models (Wallace et al., 2024) to optimize a diffusion-based HMR base model. Traditional joint-wise or pixel-wise losses could overfit noisy labels in real-world data. Besides, a trade-off between 2D reprojection fidelity and 3D accuracy exists due to imprecise camera estimation (Dwivedi et al., 2024). In contrast, DPO focuses on the relative prediction quality, being more robust to imperfect data. However, DPO requires an annotated preference dataset, which is costly to obtain (Rafailov et al., 2024). To this end, we employ HMR-Scorer to evaluate and rank the predictions generated by the base model, resulting in a preference dataset composed of  $\langle winner, loser \rangle$ prediction pairs. Guided by this synthetic dataset, ADHMR refines the HMR base model towards producing human pose predictions that are both more plausible and more closely aligned with 2D image cues. Moreover, ADHMR improves its robustness by finetuning on in-the-wild images without the need for pseudo labels.

Notably, HMR-Scorer can also be leveraged to improve the performance of state-of-the-art HMR models through data cleaning. Many models (Yin et al., 2025; Pang et al., 2024; Sun et al., 2024) incorporate in-the-wild datasets for training to enhance their generalizability. However, as mentioned earlier, the 3D pseudo-labels in these datasets are often unreliable. Prior work relies on expensive manual curation (Lassner et al., 2017) or rigid reprojection-error filtering (Kolotouros et al., 2019) to combat these issues. Instead, we propose to conduct a fully automated data cleaning process to build higher-quality training datasets. We sort the pseudo-labeled images in a dataset based on their scores given by HMR-Scorer and only retain samples with scores above a certain threshold. By filtering out poorly annotated data, we reduce the influence of noisy annotations and boost model performance.

Comprehensive experimental results demonstrate the effectiveness of our approach. The main contributions are summarized below:

- We propose ADHMR, a novel framework for improving existing diffusion-based HMR models by adapting human preference optimization methods to an unlabeled setting, thus outperforming existing state-of-theart probabilistic HMR methods.
- We introduce HMR-Scorer, a robust reward model that effectively quantifies the alignment between human mesh predictions and corresponding input images.

• We show that using HMR-Scorer for data cleaning boosts the performance of state-of-the-art HMR models, even when trained on fewer data.

# 2. Related work

#### 2.1. Human Mesh Recovery from a Single Image

Current HMR approaches can be broadly categorized into two paradigms: deterministic and probabilistic. Deterministic approaches(Goel et al., 2023; Cai et al., 2024a; Moon et al., 2022; Shen et al., 2024; Yin et al., 2024) produce a single estimate for each input. However, due to intrinsic reconstruction ambiguities, probabilistic approaches focus on generating multiple plausible hypotheses or capturing probabilistic distributions. ProHMR (Kolotouros et al., 2021) leverages a conditional normalizing flow to model a conditional probability distribution. Fang et al. (Fang et al., 2023) propose learning probability distributions over human joint rotations by utilizing a learned analytical posterior probability. EgoHMR (Zhang et al., 2023) proposes a 3D scene-conditioned diffusion approach for reconstructing human meshes from egocentric views. ScoreHypo (Xu et al., 2024b) uses a diffusion-based generator to produce a diverse set of plausible estimates, and a separate network is employed to choose from these estimates. Despite their effectiveness, they often require generating numerous candidate poses for selection or averaging. In contrast, we employ direct preference optimization to improve the performance of the prediction model directly.

#### 2.2. Human Preference Optimization

The initial efforts to learn from human preferences originated in training agents (Christiano et al., 2017; Ibarz et al., 2018), later expanding to incorporate human feedback (RLHF) for enhancing tasks like translation (Kreutzer et al., 2018) and summarization (Stiennon et al., 2020; Ziegler et al., 2019). These methods first train a reward model to align with human preferences and then finetune a language model to maximize this reward using reinforcement learning techniques such as PPO (Schulman et al., 2017). Several solutions have been proposed to simplify this complex pipeline: HIVE (Zhang et al., 2024) uses offline reinforcement learning to align instruction editing. Direct Preference Optimization (DPO) (Rafailov et al., 2024) directly optimizes the model using a supervised classification objective on preference data. This approach is now being increasingly adopted across other domains. For instance, ImageReward (Xu et al., 2024a) and Lee et al. (Lee et al., 2023) apply RLHF to text-to-image synthesis models; DreamReward (Ye et al., 2024) and CADCrafter (Chen et al., 2025) use RLHF for text-to-3D generation. Diffusion-DPO (Wallace et al., 2024) adapts the DPO objective for Diffusion Models, improving the performance of models like Stable



*Figure 1.* Overview of ADHMR. We aim to finetune a probabilistic HMR base model that generates multiple human mesh predictions conditioned on the input image. We first train the HMR-Scorer that assesses the reconstruction quality given an image and corresponding human mesh predictions. The reconstruction quality annotations  $Q^*$  are computed using standard HMR metrics, including PVE  $Q^{pve}$ , MPJPE  $Q^{mpjpe}$ , PA-MPJPE  $Q^{pajpe}$ , and PA-PVE  $Q^{papve}$ . Next, we construct a synthetic human preference dataset, where each sample is a (winner, loser) prediction pair rated by the HMR-Scorer. Finally, ADHMR uses this synthetic human preference dataset to finetune the base model to preferentially generate predictions that are more plausible and better aligned with the image cues.

Diffusion for enhanced visual appeal and textual coherence. Our method is inspired by Diffusion-DPO but differs in its implementation. Rather than depending on curated manually labeled human feedback datasets, we devise a method to automatically generate a human preference dataset using HMR-Scorer, offering greater flexibility for our scenario.

#### 3. Preliminary

#### **3.1. HMR Evaluation Metrics**

We use four standard metrics for HMR: Mean Per Vertex Position Error (PVE) and Mean Per Joint Position Error (MPJPE), along with their Procrustes-aligned variants (PA-PVE and PA-MPJPE). All metrics compute the average distance (in mm) between predicted and ground-truth positions, with the pelvis joint aligned as the reference point. We apply the joint regressor of SMPL(-X) to the predicted mesh to obtain 3D joint coordinates.

#### 3.2. Diffusion-based HMR Base Model

Modeling HMR as a reverse diffusion process by noisy samples  $\{\mathbf{x}_t\}_{t=0}^T$ , the base model HypoNet (Xu et al., 2024b)  $\epsilon_{\text{ref}}$  is a denoiser that progressively denoises random pose noise based on the input image *I* to reconstruct the human mesh. *T* is the total number of timesteps. This process is formulated as:

$$p_{\theta}\left(\mathbf{x}_{t-1} \mid \mathbf{x}_{t}\right) = \mathcal{N}\left(\mathbf{x}_{t-1}; \boldsymbol{\epsilon}_{\text{ref}}\left(\mathbf{x}_{t}, t, I\right)\right), \quad (1)$$

where  $p_{\theta}$  is the posterior mean of the forward process.

Specifically, the base model follows (Li et al., 2021) by breaking down the SMPL (Loper et al., 2015) pose param-

eters into two components: swing, derived from 3D joint positions, and twist, representing rotational details for each body part. These two elements are combined into a single data sample, which is then processed through a forward diffusion step to gradually add noise. The noisy samples are mapped to a high-dimensional feature space using a multilayer perceptron. To guide the denoising process, the model incorporates image features extracted through a convolutional neural network backbone. These preprocessed image features are concatenated and fed into a transformerencoder (Vaswani, 2017) based network. The transformer integrates global image context through a cross-attention mechanism, aligning the denoising process with the input image. Finally, the network reconstructs the human pose by removing the added noise. The human shape parameters are estimated by the convolutional backbone.

## 4. Method

#### 4.1. Overview

An overview of ADHMR is presented in Figure 1. Given an input image I, we aim to reconstruct the 3D human mesh in a parameterized way, which is to predict the pose parameters  $\theta \in \mathbb{R}^{24 \times 3}$  and shape parameters  $\beta \in \mathbb{R}^{10}$  of the predefined SMPL model (Pavlakos et al., 2019). We formulate this problem as a generation process conditioned on the input image to tackle the inherent reconstruction ambiguity.

We begin by training a diffusion-based HMR base model (Sec. 3.2). Next, we construct a synthetic human preference dataset (Sec. 4.3), where candidate human mesh predictions are generated by the base model and then paired

based on scores provided by the assessment model HMR-Scorer (Sec. 4.2). To distill knowledge from this synthetic preference dataset, we propose a preference optimization framework ADHMR that finetunes the base model to preferentially generate winner predictions over losers (Sec. 4.4). Furthermore, thanks to the strong capacity of HMR-Scorer to assess mesh predictions, we filter training data to enhance several popular HMR models (Sec. 4.5).

#### 4.2. HMR-Scorer

Given a set of predictions  $\{\mathbf{P}_m = (\theta, \beta, \Pi)_m\}_{m=0}^M$  for an input image *I*, HMR-Scorer aims to assign an estimated quality score  $\{s_m \in \mathbb{R}\}_{i=0}^M$  to each prediction. *M* represents the number of predictions,  $\Pi$  is the predicted camera parameters. Higher scores should be assigned to predictions with higher quality and better aligned with the image.

**Model architecture.** We first introduce the input features into HMR-Scorer. Instead of directly encoding pose parameters, which are prone to ambiguities in representing joint positions and deficient in spatial context, we leverage UVD coordinates as the input to HMR-Scorer. This provides a unified and consistent representation of the 3D human skeleton and preserves the geometric relationships of the skeleton. Specifically, using the camera model, we first project human body keypoints to the input image space to get their UVD coordinates  $\mathbf{J}_{uvd} \in \mathbb{R}^{N\times 3}$ , where N = 29 is the number of keypoints. We use a multilayer perceptron (MLP) to map  $\mathbf{J}_{uvd}$  to a high-dimensional feature vector  $\mathbf{F}^J \in \mathbb{R}^{C^l \times N}$ .

We utilize multi-scale image features as the image condition, denoted as  $\mathbf{c} := {\mathbf{F}^g, \mathbf{F}^l}$ . The input image I is initially divided into fixed-size patches through a patch embedding mechanism, producing a sequence of image tokens. These tokens are subsequently processed using a ViT-Base model (Dosovitskiy et al., 2021) to generate a series of global image feature tokens, denoted as  $\mathbf{F}^{g} \in \mathbb{R}^{C^{g} \times H^{g} \times W^{g}}$ . The global image feature tokens are then passed through a convolutional network to derive the low-channel global features, represented as  $\mathbf{F}^{g} \in \mathbb{R}^{C^{l} \times H^{g} \times W^{g}}$ . A de-convolution head is deployed on the global feature  $\mathbf{F}^{g}$  to obtain the highresolution local feature map  $\mathbf{F}^{l} \in \mathbb{R}^{C^{l} \times H^{l} \times W^{l}}$ .  $C^{*}$  and  $H^* \times W^*$  denote the feature channel and size, respectively. We sample the local feature  $\mathbf{F}^{l}$  according to the re-projected 2D joint positions  $J_{uv}$  and obtain pixel-aligned local image features  $\mathbf{F}^{L} \in \mathbb{R}^{C^{l} \times N}$  for each joint.

The concatenated features of  $\mathbf{F}^J$  and  $\mathbf{F}^L$  are subsequently fed into a transformer-encoder-based network comprising *B* fundamental blocks. Each block integrates a multi-head self-attention (MHSA) mechanism, a cross-attention (CA) unit, and a feed-forward network (FFN). Within the CA unit, the global image feature  $\mathbf{F}^g$  serves as the key and value features, while the query feature is derived from the output of the preceding MHSA unit. Through the crossattention mechanism, the geometric information from the human mesh predictions is effectively aligned with image features, ensuring a coherent integration of structural and visual cues. Finally, a decoder network, implemented as an MLP, is employed to estimate the score s.

**Training.** We construct a training dataset comprising human mesh predictions for corresponding images and their quality labels to train the HMR-Scorer. Specifically, predictions are generated by adding joint-wise Gaussian noise to the ground truth SMPL pose to simulate rotational errors, with magnitudes empirically determined. The reconstruction quality labels are annotated using standard HMR metrics, including PVE  $Q^{pve}$ , MPJPE  $Q^{mpjpe}$ , PA-MPJPE  $Q^{pajpe}$ , and PA-PVE  $Q^{papve}$ . Details of these metrics are provided in Sec. 3.1. To accurately capture subtle quality differences, the training process is designed to learn relative quality preferences among predictions. Inspired by RankNet (Burges et al., 2005), we utilize a probabilistic ranking cost function:

$$\mathcal{L}_{mn}(s_{mn}, y_{mn}) := -y_{mn} \log s_{mn} - (1 - y_{mn}) \log (1 - s_{mn}), \qquad (2)$$

where  $s_{mn} = \text{Sigmoid} (s_m - s_n)$  is the relative quality difference probability between predictions  $\mathbf{P}_m$  and  $\mathbf{P}_n$ , and  $y_{mn}$  is the ground truth quality label representing the quality difference between predictions based on each of the abovementioned four HMR metrics. For instance, for the PVE benchmark, the label is represented as  $y_{mn}^{pve}$  ( $y_{mn}^{pve} = 1$  if  $Q_m^{pve} < Q_n^{pve}$ , and 0 otherwise). The overall training loss for HMR-Scorer is defined as follows:

$$\mathcal{L}^{\text{HMR-Scorer}} = \sum_{m,n=0;n\neq m}^{M} \left( \mathcal{L}_{mn}^{pve} + \mathcal{L}_{mn}^{papve} + \mathcal{L}_{mn}^{pajpe} + \mathcal{L}_{mn}^{mpjpe} \right).$$
(3)

#### 4.3. HMR Preference Dataset Construction

We leverage preference-based optimization rather than traditional supervised training to finetune the base HMR model. However, traditional preference optimization methods require human preference datasets labeled by human annotators, which are currently unavailable for this field.

To this end, we propose to use HMR-Scorer to synthesize an HMR preference dataset  $\mathcal{D} = \{(I, \mathbf{x}_0^w, \mathbf{x}_0^l)\}$ , where each sample contains the input image I and a pair of human mesh predictions generated from the HMR base model  $\epsilon_{\text{ref.}}$ . Specifically, given a set of predictions  $\{\mathbf{P}_m\}_{m=0}^M$ , HMR-Scorer assigns scores  $\{s_m \in \mathbb{R}\}_{m=0}^M$  for these predictions. According to the scores, these predictions undergo ordinal arrangement, which emulates human preference in ranking human mesh reconstructions, from highest to lowest fidelity.

ADHMR: Aligning Diffusion-based Human Mesh Recovery via Direct Preference Optimization

	GTA-Human					DNA-Rendering						
	PVE		MPJPE		PA-MPJPE		PVE		MPJPE		PA-MPJPE	
	PLCC ↑	$\mathbf{SRCC}\uparrow$	PLCC $\uparrow$	SRCC $\uparrow$	PLCC $\uparrow$	SRCC $\uparrow$	PLCC $\uparrow$	SRCC $\uparrow$	PLCC $\uparrow$	$\mathbf{SRCC}\uparrow$	PLCC $\uparrow$	SRCC $\uparrow$
ScoreNet (Xu et al., 2024b)	0.52	0.49	0.52	0.50	0.47	0.43	0.55	0.51	0.55	0.50	0.50	0.46
HMR-Scorer-P	0.30	0.28	0.28	0.26	0.29	0.25	0.34	0.31	0.32	0.29	0.34	0.27
HMR-Scorer-2D	0.59	0.58	0.59	0.58	0.50	0.49	0.62	0.60	0.63	0.61	0.56	0.53
HMR-Scorer (Ours)	0.63	0.62	0.63	0.62	0.57	0.54	0.66	0.64	0.66	0.64	0.62	0.58

*Table 1.* Score prediction results on the GTA-Human (Cai et al., 2024b) and DNA-Rendering (Cheng et al., 2023) dataset. We report the PLCC and SRCC between the predicted scores and the PVE, MPJPE, and PA-PVE ground truth errors, respectively.

This hierarchical organization facilitates the extraction of paired samples  $(\mathbf{P}^w, \mathbf{P}^l)$ , denoting superior (winner) and inferior (loser) predictions respectively, where their associated scores satisfy the preference relation  $(\mathbf{P}^w \succ \mathbf{P}^l \mid I)$ . The pairing process involves stochastic selection of winners from the top K highest-scoring predictions, coupled with losers from the K lowest-scoring predictions, generating K distinct pairs per image. For studio-captured datasets with precise human mesh annotations, the prediction quality ordering is directly determined by computing the reconstruction error against ground truth labels.

#### 4.4. ADHMR

ADHMR aligns the base model  $\epsilon_{ref}$  with the constructed preference dataset  $\mathcal{D} = \{(I, \mathbf{x}_0^w, \mathbf{x}_0^l)\}$  to produce superior predictions. The aligned model  $\epsilon_{\theta}$  is initialized using the parameters of the base model  $\epsilon_{ref}$ , while keeping the base model's parameters frozen throughout the training process. The proposed optimization framework extends the direct preference optimization (DPO) method. The principle of DPO lies in its direct optimization of a conditional distribution  $\epsilon_{\theta}$  ( $\mathbf{x}_0 \mid \mathbf{c}$ ), contrasting with RLHF's approach of optimizing a reward model r ( $\mathbf{c}, \mathbf{x}_0$ ), while simultaneously constraining the KL-divergence from a reference distribution  $\epsilon_{ref}$ :

$$\max_{\boldsymbol{\epsilon}_{\theta}} \mathbb{E}_{\mathbf{c} \sim \mathcal{D}_{c}, \mathbf{x}_{0} \sim \boldsymbol{\epsilon}_{\theta}(\mathbf{x}_{0} | \mathbf{c})}$$

$$[r(c, \mathbf{x}_{0})] - \beta_{\mathrm{KL}} \left[ \boldsymbol{\epsilon}_{\theta} \left( \mathbf{x}_{0} \mid \mathbf{c} \right) \| \boldsymbol{\epsilon}_{\mathrm{ref}} \left( \mathbf{x}_{0} \mid \mathbf{c} \right) \right].$$

$$(4)$$

Following (Wallace et al., 2024), a significant challenge in applying DPO to diffusion models is the intractability of the parameterized distribution  $\epsilon_{\theta}$  ( $\mathbf{x}_0 | \mathbf{c}$ ), which stems from the necessity to marginalize over all possible diffusion trajectories ( $\mathbf{x}_1, \ldots, \mathbf{x}_T$ ) that culminate in  $\mathbf{x}_0$ . Through some mathematical techniques, this challenge is addressed by formulating an objective function that operates on the complete denoising trajectory  $\mathbf{x}_{0:T}$ :

$$\mathcal{L}^{\text{DDPO}}(\theta) = -\mathbb{E}_{(\mathbf{x}_{0}^{w}, \mathbf{x}_{0}^{l}) \sim \mathcal{D}, t \sim \mathcal{U}(0, T), \mathbf{x}_{t}^{w} \sim q(\mathbf{x}_{t}^{w} | \mathbf{x}_{0}^{w}), \mathbf{x}_{t}^{l} \sim q(\mathbf{x}_{t}^{l} | \mathbf{x}_{0}^{l})}}{\log \sigma(-\beta T \omega \left(\lambda_{t}\right) \left(\|\boldsymbol{\epsilon}^{w} - \boldsymbol{\epsilon}_{\theta}\left(\mathbf{x}_{t}^{w}, t\right)\|_{2}^{2} - \|\boldsymbol{\epsilon}^{w} - \boldsymbol{\epsilon}_{\text{ref}}\left(\mathbf{x}_{t}^{w}, t\right)\|_{2}^{2}} - \left(\left\|\boldsymbol{\epsilon}^{l} - \boldsymbol{\epsilon}_{\theta}\left(\mathbf{x}_{t}^{l}, t\right)\right\|_{2}^{2} - \left\|\boldsymbol{\epsilon}^{l} - \boldsymbol{\epsilon}_{\text{ref}}\left(\mathbf{x}_{t}^{l}, t\right)\right\|_{2}^{2}\right)\right)\right)$$
(5)

where  $\mathbf{x}_t^* = \alpha_t \mathbf{x}_0^* + \sigma_t \boldsymbol{\epsilon}^*$  is drawn from  $q(\mathbf{x}_t^* | \mathbf{x}_0^*)$  with  $\boldsymbol{\epsilon}^* \sim \mathcal{N}(0, \mathbf{I})$ . Here,  $\lambda_t = \alpha_t^2 / \sigma_t^2$  denotes the signal-tonoise ratio,  $\omega(\lambda_t)$  is a weighting function, and the constant T is factored into  $\beta$ .

During training, the model improves by comparing points along the diffusion trajectory with examples from the synthetic preference dataset. This helps the model better denoise winner mesh predictions compared to losers, as evaluated by HMR-Scorer. Hence, this methodology guides the model to generate human mesh predictions that not only align closely with the input image but also adhere to a realistic distribution of human poses. By exclusively finetuning the denoiser component within the latent space, the approach achieves more generalized and well-aligned results without overfitting to noisy labels, especially for in-the-wild datasets.

#### 4.5. Data Cleaning

To further evaluate the efficacy of our trained HMR-Scorer, we propose to apply it to the training data cleaning process, aiming to determine whether the proposed scorer can effectively identify noisy data. While many indoor datasets have ground truth labels, in-the-wild datasets often rely on noisy pseudo labels, which hinders model training and generalization. Therefore, we use HMR-Scorer to remove low-quality samples, ensuring a reliable training dataset.

The data cleaning process begins with score computation, where a quality score  $s_i \in [0, 1]$  is assigned to each sample  $(I_i, \hat{\theta}_i, \hat{\beta}_i, \hat{\Pi}_i)$  in the dataset  $\mathcal{X}$  using HMR-Scorer. The score evaluates the alignment of predictions with the input image and the plausibility of the model parameters. Next, a threshold  $\tau$  is applied to filter out low-quality samples, resulting in the cleaned dataset  $\mathcal{X}_{\text{clean}} = \{(I_i, \hat{\theta}_i, \hat{\beta}_i, \hat{\Pi}_i) \mid s_i \geq \tau\}$ . Finally, only high-confidence pseudo-labels are retained for model training.

## 5. Experiments

### 5.1. Setup

**Training.** HMR-Scorer is trained on five datasets, including HI4D (Yin et al., 2023), BEDLAM (Black et al., 2023), DNA-Rendering (Cheng et al., 2023), GTA-Human (Cai

ADHMR: Aligning Diffusion-based Human Mesh Recovery via Direct Preference Optimization

Method			3DPW		Human3.6M			
		$PVE \downarrow$	$MPJPE \downarrow$	$\text{PA-MPJPE} \downarrow$	$PVE \downarrow$	$MPJPE \downarrow$	$\text{PA-MPJPE} \downarrow$	
HMR (Kanazawa et al., 2018)		152.7	130.0	81.3	96.1	88.0	56.8	
HybrIK (Li et al., 2021)	-	86.5	74.1	45.0	65.7	54.4	34.5	
PyMaf (Zhang et al., 2021)	-	110.1	92.8	58.9	-	57.7	40.5	
POTTER (Zheng et al., 2023)	-	87.4	75.0	44.8	-	56.5	35.1	
ImpHMR (Cho et al., 2023)	-	87.1	74.3	45.4	-	-	-	
Zolly (Wang et al., 2023)	-	76.3	65.0	39.8	-	49.4	32.3	
HMR 2.0a (Goel et al., 2023)	-	-	70.0	44.5	-	44.8	33.6	
HMR 2.0b (Goel et al., 2023)	-	-	81.3	54.3	-	50.0	32.4	
ScoreHMR (Stathopoulos et al., 2024)		-	76.8	51.1	-	-		
		-	79.4	56.6	-	59.2	42.2	
Biggs et al. (Biggs et al., 2020a)	25	-	75.8	55.6	-	58.2	42.2	
Sengupta et al. (Sengupta et al., 2021)		-	75.1	47.0	-	-	-	
ProHMR (Kolotouros et al., 2021)		-	-	52.4	-	-	36.8	
HuManiFlow (Sengupta et al., 2023)		-	65.1	39.9	-	-	-	
HMDiff (Foo et al., 2023)		82.4	72.7	44.5	-	49.3	32.4	
	10	79.8	68.5	41.0	52.5	42.4	29.0	
HypoNet (Base Model)		73.4	63.0	37.6	47.5	38.4	26.0	
		71.9	61.8	36.1	46.4	37.4	25.3	
	10	73.8	64.2	38.3	52.1	41.8	28.4	
ADHMR		65.4	57.2	33.5	45.9	36.9	24.8	
		63.5	55.7	32.5	44.6	35.8	24.1	
	10	71.3	61.3	37.1	52.2	41.9	28.3	
ADHMR (ITW)	100	62.6	54.2	32.0	45.9	37.0	24.8	
		60.5	52.6	30.8	44.6	35.9	24.0	

Table 2. Comparison with state-of-the-arts on the 3DPW (Von Marcard et al., 2018) and Human3.6M (Ionescu et al., 2013) dataset. *M* is the number of predictions in probabilistic methods. ADHMR is finetuned on the target benchmark dataset, while ADHMR (ITW) is further finetuned on the preference dataset constructed from an in-the-wild dataset.

et al., 2024b), and SPEC (Kocabas et al., 2021). These datasets contain accurate 3D annotations for human pose, which plays an important role in training an effective scorer. The HMR base model is the current state-of-the-art probabilistic HMR method: HypoNet from (Xu et al., 2024b).

**Evaluation metrics.** We use four standard human mesh reconstruction metrics: PVE, MPJPE, PA-PVE, and PA-MPJPE, as detailed in Sec. 3.1. To evaluate the scorer, we follow score prediction assessment (Zhai & Min, 2020) to employ two standard metrics: the Pearson linear correlation coefficient (PLCC) and Spearman rank correlation coefficient (SRCC). These correlation coefficients quantify the alignment between the predicted scores and ground truth reconstruction error.

#### 5.2. HMR-Scorer Evaluation

**Test benchmark.** Since there are no existing datasets for score prediction of HMR models, we construct a test set from a synthetic dataset GTA-Human (Cai et al., 2024b), which is produced with rendering engines (e.g., Unreal Engine) and contains accurate 3D annotations. We also use DNA-Rendering (Cheng et al., 2023), a large-scale multiview studio-based dataset with an ultra-high resolution, to construct the other test set to show the capacity for common studio-based scenes. We adopt the original test set of the two selected datasets. Then we perturb the ground truth pose labels with random gaussian noise to simulate

Mathad	3DPW							
wiethod	$PVE \downarrow$	$MPJPE \downarrow$	$PA\text{-}MPJPE\downarrow$					
HypoNet (Base Model)	73.4	63.0	37.6					
(a) Finetune	on target b	enchmark dat	aset					
Supervised finetuning	68.0	59.4	35.9					
ADHMR	65.4	57.2	33.5					
(b) Finetu	ne on in-th	e-wild datase	t					
Supervised finetuning	70.2	61.3	36.5					
ADHMR	62.6	54.2	32.0					

Table 3. Ablation of preference finetuning on 3DPW (Von Marcard et al., 2018) dataset. M = 100 for all models.

predictions of HMR models. We save the corresponding reconstruction errors for the noised predictions. We then measure the PLCC and SRCC between the ground truth reconstruction metrics and the predicted scores.

**Results.** Table 1 presents the comparison results for score prediction. We observe that our scorer outperforms all baseline methods in both PLCC and SRCC metrics, showcasing its efficacy. For comparison, we modify the reward model in ScoreHypo as a baseline. We also study two ablations of HMR-Scorer: HMR-Scorer-P accepts SMPL(-X) pose rotation vectors as input, and HMR-Scorer-2D accepts 2D joint positions (without joint depth) as input. Results show that HMR-Scorer outperforms the baselines in aligning the scores with the real reconstruction errors, which underscores its effectiveness indicating human mesh prediction quality.



*Figure 2.* Qualitative comparison of the state-of-the-art probabilistic model ScoreHypo (Xu et al., 2024b) and our ADHMR. Our framework significantly improves image alignment and in-the-wild robustness. (a)  $\sim$  (f) are from the 3DPW (Von Marcard et al., 2018) dataset, and (g)  $\sim$  (h) are challenging in-the-wild images.

Mathad		3DPW	
Method	$PVE \downarrow$	$MPJPE \downarrow$	PA-MPJPE↓
HypoNet (Base Model)	71.9	61.8	36.1
Supervised finetuning	69.9	59.7	35.2
ADHMR (ITW)	60.5	52.6	30.8

Table 4. Ablation of extra training data on 3DPW (Von Marcard et al., 2018) dataset. We use multiple training datasets of the scorer to perform supervised finetuning. M = 200 for all models.

#### 5.3. ADHMR Evaluation

**Comparisons with state-of-the-art methods.** In Table 2, we compare the accuracy of the ADHMR with state-of-the-art methods on two widely used benchmark datasets Human3.6M (Ionescu et al., 2013) and 3DPW (Von Marcard et al., 2018). 3DPW is an in-the-wild dataset. We show results of both deterministic and probabilistic methods. Following the conventions of standard probabilistic approaches (Xu et al., 2024b; Biggs et al., 2020b), we generate multiple estimates and report the min-MPJPE and min-PVE of the M predictions. When finetuned directly on the target benchmark dataset, ADHMR achieves further enhancements, showcasing its strong ability to adapt to and balance domain-specific distribution. To evaluate the effectiveness under the in-the-wild setting, ADHMR (ITW) is

further finetuned on InstaVariety dataset (Kanazawa et al., 2019), which contains various in-the-wild images annotated with noisy pseudo-labels collected from Instagram. Please note that we do not use the original 3D labels but use the HMR-Scorer to construct a preference dataset. Notably, we observe that our finetuned model achieves better performance using M = 10 predictions than the base model using M = 200 predictions on the in-the-wild benchmark 3DPW, showcasing that our finetuning pipeline greatly enhances the generalizability to in-the-wild datasets and its efficiency. Moreover, ADHMR consistently outperforms existing state-of-the-art methods by a substantial margin.

**Qualitative results.** Fig. 2 shows qualitative comparisons between ADHMR and the previous state-of-the-art probabilistic method ScoreHypo. We show randomly selected results of ScoreHypo and ADHMR on 3DPW and internet images with M = 10 candidate predictions. We can see that the finetuned model can produce more accurate results for body pose under challenging cases, such as dense humanenvironment interactions. The base model, however, cannot achieve good image-mesh alignment. For example, in the (a) instance, our method provides more reasonable poses for the occluded right arm. In the (c) instance, ScoreHypo gives erroneous prediction for the person's arms, while ours gives a more accurate body pose prediction, proving the

ADHMR: Aligning Diffusion-based Human Mesh Recovery via Direct Preference Optimization

	31	DPW	Hum	an3.6M	EHF			
Method	$MPJPE \downarrow$	$\text{PA-MPJPE} \downarrow$	$\mathbf{MPJPE}\downarrow$	$\text{PA-MPJPE} \downarrow$	$\text{PA-PVE} \downarrow$	$PVE\downarrow$	PA-MPJPE↓	
Hand4Whole (Moon et al., 2022)	115.2	75.4	78.8	57.7	57.8	89.2	70.2	
+ data cleaning	112.3	73.7	77.9	57.0	56.2	88.8	69.8	
OSX (Base) (Lin et al., 2023)	100.4	66.4	69.5	48.9	54.7	86.6	63.7	
+ data cleaning	99.4	65.2	65.7	46.4	53.1	84.0	62.4	
SMPLer-X-Base (Cai et al., 2024a)	99.5	64.2	59.8	45.8	51.0	82.4	59.7	
+ data cleaning	97.9	62.9	57.5	43.8	51.4	78.6	58.6	

*Table 5.* Quantitative comparisons between several state-of-the-art methods with and without data cleaning on the 3DPW (Von Marcard et al., 2018), Human3.6M (Ionescu et al., 2013) and EHF (Pavlakos et al., 2019) dataset. All methods are trained on four commonly used datasets. After cleaning the training data, these models achieve higher accuracy even when trained on a smaller subset.

efficacy of the proposed finetuning pipeline. In the (g) instance, ScoreHypo produces inaccurate elbow poses, while our prediction fits the input image better. Results show that ADHMR is more robust for challenging internet images than ScoreHypo. Please zoom in to observe our improvement over the base model.

Ablation of preference finetuning. As shown in Table 3, we conduct an ablation study on different finetuning methods to demonstrate the advantages of ADHMR over traditional supervised finetuning. We construct two baselines where the base model is finetuned using the ground truth labels in the datasets. We first finetune the base model on the training sets of the two target benchmarks (3DPW and Human3.6M). We also finetune on the pseudo labels of InstaVariety to simulate training on noisy pseudo-labeled inthe-wild data. Results show that ADHMR consistently outperforms traditional supervised finetuning in both settings. When finetuned directly on the target benchmark dataset, our method achieves further enhancements than supervised finetuning, showcasing its strong ability to adapt to and balance domain-specific distribution. In the meantime, supervised finetuning may overfit one training dataset (3DPW) and the performance on the other test benchmark (Human3.6M) could be corrupted. On the in-the-wild dataset, ADHMR demonstrates robustness under noisy pseudo-label conditions, while supervised finetuning on the noisy labels could overfit the training dataset and harm its performance on 3DPW benchmark.

Ablation of extra training datasets. In Table 4, we aim to confirm that the improvements achieved by ADHMR are contributed to its inherent optimization strategy, rather than including extra information from other datasets used during scorer training. So we additionally finetune the base model on the scorer's training sets. Supervised finetuning on scorer training datasets yields insignificant performance gains, confirming that the improvements are not primarily due to external dataset information. This suggests that simply exposing the base model to more data does not guarantee better performance. In contrast, ADHMR benefits from the ability of HMR-Scorer to implicitly guide the base model towards generating high-quality predictions.

#### 5.4. Data Cleaning Results

Current HMR models are trained using quite different datasets. For a fair and comprehensive comparison, we selected several state-of-the-art HMR methods and retrained them using four commonly used datasets: MSCOCO (Lin et al., 2014), MPII (Andriluka et al., 2014), Human3.6M (Ionescu et al., 2013), and MPI-INF-3DHP (Mehta et al., 2017). We compare training the models on the full training sets of these datasets and training on the filtered datasets obtained after applying data cleaning with the HMR-Scorer. We set the filter threshold  $\tau = 0.6$ .

Quantitative results are in Table 5. The results demonstrate that models trained on the cleaned datasets achieve better performance despite using less training data. This improvement highlights the utility of our method in filtering out low-quality training samples, thereby enabling the models to focus on higher-quality data for learning. This finding provides a new perspective for improving large HMR models by strategically cleaning training data. By leveraging the HMR-Scorer to curate datasets, we can achieve higher-quality model training with even less data, making it a valuable tool for effortless performance gain.

#### 6. Conclusion

In this work, we propose ADHMR, the first framework for aligning diffusion-based HMR models with direct preference optimization. We leverage a trained HMR-Scorer to synthesize a preference dataset automatically without the need for manual annotation. This dataset is then used to align the diffusion-based HMR model through direct preference optimization. Additionally, HMR-Scorer improves the performance of several state-of-the-art HMR models by filtering out low-quality training data. Extensive experiments validate the effectiveness of ADHMR. We believe that our work will pave the way for future advancements in alignment techniques for human mesh recovery.

# Acknowledgements

This study is supported under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s). This research is also supported by A\*STAR under its MTC Programmatic Funds (Grant No. M23L7b0021) and the MoE AcRF Tier 1 grant (RG14/22).

# **Impact Statement**

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

# References

- Andriluka, M., Pishchulin, L., Gehler, P., and Schiele, B. 2d human pose estimation: New benchmark and state of the art analysis. In *Proceedings of the IEEE Conference on computer Vision and Pattern Recognition*, pp. 3686–3693, 2014.
- Biggs, B., Novotny, D., Ehrhardt, S., Joo, H., Graham, B., and Vedaldi, A. 3d multi-bodies: Fitting sets of plausible 3d human models to ambiguous image data. *Advances in neural information processing systems*, 33:20496–20507, 2020a.
- Biggs, B., Novotny, D., Ehrhardt, S., Joo, H., Graham, B., and Vedaldi, A. 3d multi-bodies: Fitting sets of plausible 3d human models to ambiguous image data. *Advances in neural information processing systems*, 33:20496–20507, 2020b.
- Black, M. J., Patel, P., Tesch, J., and Yang, J. Bedlam: A synthetic dataset of bodies exhibiting detailed lifelike animated motion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8726–8737, 2023.
- Burges, C., Shaked, T., Renshaw, E., Lazier, A., Deeds, M., Hamilton, N., and Hullender, G. Learning to rank using gradient descent. In *Proceedings of the 22nd international conference on Machine learning*, pp. 89–96, 2005.
- Cai, Z., Yin, W., Zeng, A., Wei, C., Sun, Q., Yanjun, W., Pang, H. E., Mei, H., Zhang, M., Zhang, L., et al. Smplerx: Scaling up expressive human pose and shape estimation. *Advances in Neural Information Processing Systems*, 36, 2024a.
- Cai, Z., Zhang, M., Ren, J., Wei, C., Ren, D., Lin, Z., Zhao, H., Yang, L., Loy, C. C., and Liu, Z. Playing for 3d human recovery. *IEEE Transactions on Pattern Analysis* and Machine Intelligence, 2024b.

- Chen, C., Wei, J., Chen, T., Zhang, C., Yang, X., Zhang, S., Yang, B., Foo, C.-S., Lin, G., Huang, Q., et al. Cadcrafter: Generating computer-aided design models from unconstrained images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025.
- Cheng, W., Chen, R., Fan, S., Yin, W., Chen, K., Cai, Z., Wang, J., Gao, Y., Yu, Z., Lin, Z., et al. Dna-rendering: A diverse neural actor repository for high-fidelity humancentric rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 19982– 19993, 2023.
- Cho, H. and Kim, J. Generative approach for probabilistic human mesh recovery using diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4183–4188, 2023.
- Cho, H., Cho, Y., Ahn, J., and Kim, J. 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*, pp. 21148–21158, 2023.
- Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., and Amodei, D. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houlsby, N. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021.
- Dwivedi, S. K., Sun, Y., Patel, P., Feng, Y., and Black, M. J. Tokenhmr: Advancing human mesh recovery with a tokenized pose representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1323–1333, 2024.
- Fang, Q., Chen, K., Fan, Y., Shuai, Q., Li, J., and Zhang, W. Learning analytical posterior probability for human mesh recovery. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pp. 8781– 8791, 2023.
- Foo, L. G., Gong, J., Rahmani, H., and Liu, J. Distributionaligned diffusion for human mesh recovery. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 9221–9232, 2023.
- Goel, S., Pavlakos, G., Rajasegaran, J., Kanazawa, A., and Malik, J. Humans in 4d: Reconstructing and tracking humans with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 14783– 14794, 2023.

- Hong, Z. and Shen, W. Free-viewpoint video in the wild using a flying camera. In ECCV 2024 Workshop on Wild 3D: 3D Modeling, Reconstruction, and Generation in the Wild, 2024.
- Huang, B., Li, C., Xu, C., Pan, L., Wang, Y., and Lee, G. H. Closely interactive human reconstruction with proxemics and physics-guided adaption. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024.
- Ibarz, B., Leike, J., Pohlen, T., Irving, G., Legg, S., and Amodei, D. Reward learning from human preferences and demonstrations in atari. *Advances in neural information processing systems*, 31, 2018.
- Ionescu, C., Papava, D., Olaru, V., and Sminchisescu, C. Human3. 6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence*, 36(7):1325–1339, 2013.
- Kanazawa, A., Black, M. J., Jacobs, D. W., and Malik, J. End-to-end recovery of human shape and pose. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 7122–7131, 2018.
- Kanazawa, A., Zhang, J. Y., Felsen, P., and Malik, J. Learning 3d human dynamics from video. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 5614–5623, 2019.
- Kocabas, M., Huang, C.-H. P., Tesch, J., Müller, L., Hilliges, O., and Black, M. J. Spec: Seeing people in the wild with an estimated camera. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 11035– 11045, 2021.
- Kolotouros, N., Pavlakos, G., Black, M. J., and Daniilidis, K. 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*, pp. 2252–2261, 2019.
- Kolotouros, N., Pavlakos, G., Jayaraman, D., and Daniilidis, K. Probabilistic modeling for human mesh recovery. In *Proceedings of the IEEE/CVF international conference* on computer vision, pp. 11605–11614, 2021.
- Kreutzer, J., Uyheng, J., and Riezler, S. Reliability and learnability of human bandit feedback for sequence-tosequence reinforcement learning. In *Proceedings of the* 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 1777– 1788, 2018.

- Lassner, C., Romero, J., Kiefel, M., Bogo, F., Black, M. J., and Gehler, P. V. Unite the people: Closing the loop between 3d and 2d human representations. In *Proceedings* of the IEEE conference on computer vision and pattern recognition, pp. 6050–6059, 2017.
- Lee, K., Liu, H., Ryu, M., Watkins, O., Du, Y., Boutilier, C., Abbeel, P., Ghavamzadeh, M., and Gu, S. S. Aligning textto-image models using human feedback. *arXiv preprint arXiv:2302.12192*, 2023.
- Li, J., Xu, C., Chen, Z., Bian, S., Yang, L., and Lu, C. Hybrik: A hybrid analytical-neural inverse kinematics solution for 3d human pose and shape estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 3383–3393, 2021.
- Li, J., Bian, S., Xu, C., Chen, Z., Yang, L., and Lu, C. Hybrik-x: Hybrid analytical-neural inverse kinematics for whole-body mesh recovery. *arXiv preprint arXiv:2304.05690*, 2023.
- Lin, J., Zeng, A., Wang, H., Zhang, L., and Li, Y. One-stage 3d whole-body mesh recovery with component aware transformer. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pp. 21159– 21168, 2023.
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. Microsoft coco: Common objects in context. In *Computer Vision–ECCV* 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13, pp. 740– 755. Springer, 2014.
- Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., and Black, M. J. Smpl: a skinned multi-person linear model. ACM Transactions on Graphics (TOG), 34(6):1–16, 2015.
- Mehta, D., Rhodin, H., Casas, D., Fua, P., Sotnychenko, O., Xu, W., and Theobalt, C. Monocular 3d human pose estimation in the wild using improved cnn supervision. In 2017 international conference on 3D vision (3DV), pp. 506–516. IEEE, 2017.
- Moon, G., Choi, H., and Lee, K. M. Accurate 3d hand pose estimation for whole-body 3d human mesh estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2308–2317, 2022.
- Pang, H. E., Cai, Z., Yang, L., Tao, Q., Wu, Z., Zhang, T., and Liu, Z. Towards robust and expressive whole-body human pose and shape estimation. *Advances in Neural Information Processing Systems*, 36, 2024.
- Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A. A., Tzionas, D., and Black, M. J. Expressive body capture: 3d hands, face, and body from a single image.

In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 10975–10985, 2019.

- Rafailov, R., Sharma, A., Mitchell, E., Manning, C. D., Ermon, S., and Finn, C. Direct preference optimization: Your language model is secretly a reward model. *Ad*vances in Neural Information Processing Systems, 36, 2024.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Sengupta, A., Budvytis, I., and Cipolla, R. Hierarchical kinematic probability distributions for 3d human shape and pose estimation from images in the wild. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 11219–11229, 2021.
- Sengupta, A., Budvytis, I., and Cipolla, R. Humaniflow: Ancestor-conditioned normalising flows on so (3) manifolds for human pose and shape distribution estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4779–4789, 2023.
- Shen, W., Yin, W., Wang, H., Wei, C., Cai, Z., Yang, L., and Lin, G. Hmr-adapter: A lightweight adapter with dual-path cross augmentation for expressive human mesh recovery. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pp. 6093–6102, 2024.
- Shuai, Q., Geng, C., Fang, Q., Peng, S., Shen, W., Zhou, X., and Bao, H. Novel view synthesis of human interactions from sparse multi-view videos. In *ACM SIGGRAPH 2022 Conference Proceedings*, pp. 1–10, 2022.
- Stathopoulos, A., Han, L., and Metaxas, D. Score-guided diffusion for 3d human recovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 906–915, 2024.
- Stiennon, N., Ouyang, L., Wu, J., Ziegler, D., Lowe, R., Voss, C., Radford, A., Amodei, D., and Christiano, P. F. Learning to summarize with human feedback. *Advances in Neural Information Processing Systems*, 33: 3008–3021, 2020.
- Sun, Q., Wang, Y., Zeng, A., Yin, W., Wei, C., Wang, W., Mei, H., Leung, C.-S., Liu, Z., Yang, L., et al. Aios: Allin-one-stage expressive human pose and shape estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 1834–1843, 2024.
- Vaswani, A. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.

- Von Marcard, T., Henschel, R., Black, M. J., Rosenhahn, B., and Pons-Moll, G. Recovering accurate 3d human pose in the wild using imus and a moving camera. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 601–617, 2018.
- Wallace, B., Dang, M., Rafailov, R., Zhou, L., Lou, A., Purushwalkam, S., Ermon, S., Xiong, C., Joty, S., and Naik, N. Diffusion model alignment using direct preference optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8228–8238, 2024.
- Wang, W., Ge, Y., Mei, H., Cai, Z., Sun, Q., Wang, Y., Shen, C., Yang, L., and Komura, T. Zolly: Zoom focal length correctly for perspective-distorted human mesh reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3925–3935, 2023.
- Xu, J., Liu, X., Wu, Y., Tong, Y., Li, Q., Ding, M., Tang, J., and Dong, Y. Imagereward: Learning and evaluating human preferences for text-to-image generation. *Advances* in Neural Information Processing Systems, 36, 2024a.
- Xu, Y., Ma, X., Su, J., Zhu, W., Qiao, Y., and Wang, Y. Scorehypo: Probabilistic human mesh estimation with hypothesis scoring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 979–989, 2024b.
- Xu, Z., Song, C., Song, G., Zhang, J., Liew, J. H., Xu, H., Xie, Y., Luo, L., Lin, G., Feng, J., et al. High quality human image animation using regional supervision and motion blur condition. *arXiv preprint arXiv:2409.19580*, 2024c.
- Yang, F., Chen, T., He, X., Cai, Z., Yang, L., Wu, S., and Lin, G. Attrihuman-3d: Editable 3d human avatar generation with attribute decomposition and indexing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10596–10605, 2024.
- Yao, N., Zhang, G., Shen, W., Shu, J., and Wang, H. Unify3d: An augmented holistic end-to-end monocular 3d human reconstruction via anatomy shaping and twins negotiating. *arXiv preprint arXiv:2504.18215*, 2025.
- Ye, J., Liu, F., Li, Q., Wang, Z., Wang, Y., Wang, X., Duan, Y., and Zhu, J. Dreamreward: Text-to-3d generation with human preference. *arXiv preprint arXiv:2403.14613*, 2024.
- Yin, W., Cai, Z., Wang, R., Wang, F., Wei, C., Mei, H., Xiao, W., Yang, Z., Sun, Q., Yamashita, A., et al. Whac: Worldgrounded humans and cameras. In *European Conference* on Computer Vision, pp. 20–37. Springer, 2024.

- Yin, W., Cai, Z., Wang, R., Zeng, A., Wei, C., Sun, Q., Mei, H., Wang, Y., Pang, H. E., Zhang, M., et al. Smplest-x: Ultimate scaling for expressive human pose and shape estimation. arXiv preprint arXiv:2501.09782, 2025.
- Yin, Y., Guo, C., Kaufmann, M., Zarate, J. J., Song, J., and Hilliges, O. Hi4d: 4d instance segmentation of close human interaction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 17016–17027, 2023.
- Zhai, G. and Min, X. Perceptual image quality assessment: a survey. *Science China Information Sciences*, 63:1–52, 2020.
- Zhang, H., Tian, Y., Zhou, X., Ouyang, W., Liu, Y., Wang, L., and Sun, Z. Pymaf: 3d human pose and shape regression with pyramidal mesh alignment feedback loop. In *Proceedings of the IEEE/CVF International Conference* on Computer Vision, pp. 11446–11456, 2021.
- Zhang, S., Ma, Q., Zhang, Y., Aliakbarian, S., Cosker, D., and Tang, S. Probabilistic human mesh recovery in 3d scenes from egocentric views. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 7989–8000, 2023.
- Zhang, S., Yang, X., Feng, Y., Qin, C., Chen, C.-C., Yu, N., Chen, Z., Wang, H., Savarese, S., Ermon, S., et al. Hive: Harnessing human feedback for instructional visual editing. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pp. 9026– 9036, 2024.
- Zheng, C., Liu, X., Qi, G.-J., and Chen, C. Potter: Pooling attention transformer for efficient human mesh recovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1611–1620, 2023.
- Ziegler, D. M., Stiennon, N., Wu, J., Brown, T. B., Radford, A., Amodei, D., Christiano, P., and Irving, G. Fine-tuning language models from human preferences. arXiv preprint arXiv:1909.08593, 2019.