

# PoseSyn: Synthesizing Diverse 3D Pose Data from In-the-Wild 2D Data

ChangHee Yang<sup>\*1</sup>    Hyeonseop Song<sup>\*1</sup>    Seokhun Choi<sup>\*1</sup>  
 Seungwoo Lee<sup>1</sup>    Jaechul Kim<sup>1</sup>    Hoseok Do<sup>†1</sup>  
 AI Lab, CTO Division, LG Electronics<sup>1</sup>

{changhee.yang, hyeonseop.song, seokhun.choi, seungwoo5.lee, jaechul1220.kim, hoseok.do}@lge.com

## Abstract

Despite considerable efforts to enhance the generalization of 3D pose estimators without costly 3D annotations, existing data augmentation methods struggle in real-world scenarios with diverse human appearances and complex poses. We propose PoseSyn, a novel data synthesis framework that transforms abundant in-the-wild 2D pose datasets into diverse 3D pose–image pairs. PoseSyn comprises two key components: Error Extraction Module (EEM), which identifies challenging poses from the 2D pose datasets, and Motion Synthesis Module (MSM), which synthesizes motion sequences around the challenging poses. Then, by generating realistic 3D training data via a human animation model–aligned with challenging poses and appearances–PoseSyn boosts the accuracy of various 3D pose estimators by up to 14% across real-world benchmarks including various backgrounds and occlusions, challenging poses, and multi-view scenarios. Extensive experiments further confirm that PoseSyn is a scalable and effective approach for improving generalization without relying on expensive 3D annotations, regardless of the pose estimator’s model size or design.

## 1. Introduction

Estimating 3D human pose [2, 5, 8, 12, 15, 27, 29, 32, 33, 50, 51, 60, 60] from RGB images has been extensively researched due to its applications like action recognition, virtual reality, and sports analytics. Traditional methods rely heavily on large-scale datasets with accurate 3D pose ground truth (GT), but acquiring such 3D annotations often demands multi-camera systems [4, 20, 38] or motion capture setups [28, 54], which are costly and limited to controlled indoor spaces [26]. Consequently, existing 3D pose datasets are mainly biased toward laboratory settings, which hinder them from covering in-the-wild challenging poses

<sup>\*</sup>These authors contributed equally to this work

<sup>†</sup>Corresponding author

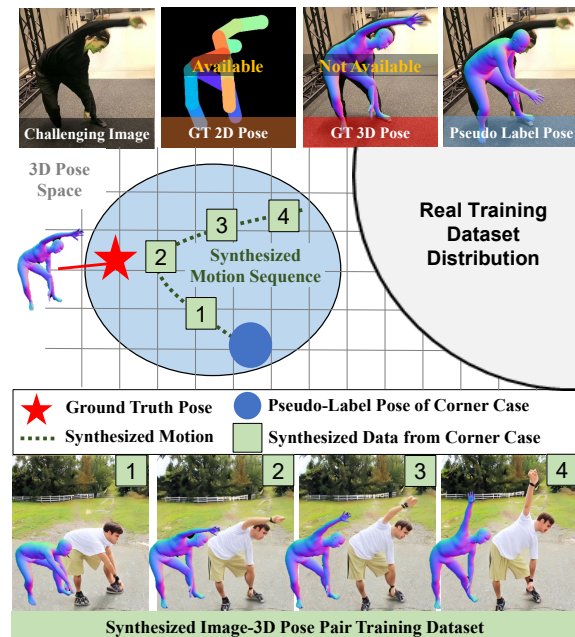


Figure 1. **PoseSyn’s Approach for Corner-Case Pose Generation.** PoseSyn addresses challenging pose cases without direct costly 3D annotations. Starting from a pseudo-labeled pose (blue circle) that is inaccurate compared to the unknown true 3D pose (red star), PoseSyn generates diverse motion sequences (green dotted lines) around the pseudo-labeled pose to produce new image–pose pairs (numbered 1–4). PoseSyn effectively bridges the gap between the inaccurate pseudo-label and the desired challenging pose, expanding the training distribution with hard examples.

and reflecting real-world visual elements such as appearance, background, and lighting. The reliance on indoor data also introduces a domain gap, limiting model robustness in outdoor or uncontrolled environments, thus hindering generalization to real-world scenarios.

To address this gap, prior work has explored data augmentation techniques. Transformation-based methods [13, 16], often powered by Generative Adversarial Networks (GANs) [17], generate new skeletal poses via geometric

transformations. Yet these methods are restricted to modifying keypoints without image-level context. In contrast, synthetic approaches like PoseGen [14] leverage Neural Radiance Fields (NeRF) [39, 53] to render novel views, but they often fail to capture the real-world human appearances and backgrounds—limiting their impact on generalization. Abundant in-the-wild 2D pose datasets offer a cost-effective alternative, but lack corresponding 3D annotations.

We propose PoseSyn, a novel framework that synthesizes diverse 3D pose-image pair data from in-the-wild 2D pose-image pair data to improve the generalization of 3D pose estimators, targeting hard, underrepresented samples. The key insight is that generating plausible variant data that approximates a challenging pose, then retraining the Target Pose Estimator (TPE) with these variants, improves generalization in real-world scenarios. Leveraging abundant 2D pose annotations as seeds, PoseSyn generates diverse 3D training samples tailored to each TPE’s challenging cases. Our Error Extraction Module (EEM) identifies challenging images where TPE underperforms, guided by 2D ground truth. However, GT 3D pose of the challenging image is not available for direct annotation, while the pseudo-labeled 3D pose is often inaccurate. Naively relying on this inaccurate pose or high-level text description extracted from the challenging image, which are approximations of the hidden GT 3D poses, fails to capture full complexity of pose configurations, making it difficult to generate variant data for such challenging cases. To mitigate this, our Motion Synthesis Module (MSM) integrates both text descriptions and the inaccurate pose to synthesize motion sequences near the hidden GT pose. By representing poses as continuous motion rather than isolated frames, MSM reduces ambiguity and produces a range of plausible pose variations that better approximate the desired 3D GT, as shown in Fig. 1. Finally, these generated poses, closely aligned with challenging images, are synthesized into realistic images via a human animation model [64], ensuring well-aligned 3D training data.

Moreover, TPEs can vary in parameter size to meet different hardware constraints and accuracy needs in practical applications. Through experiments across diverse datasets [4, 21, 24–26, 28, 54] and multiple TPE architectures [8, 15, 32], we demonstrate that PoseSyn is a scalable and effective solution for enhancing 3D pose estimation without relying on costly 3D annotations.

In summary, our contributions include:

- We propose PoseSyn, a novel framework that generates diverse 3D pose data from in-the-wild 2D pose datasets to address real-world corner cases.
- We propose Motion Synthesis Module (MSM), which synthesizes motion sequences by leveraging both textual descriptions and initial pose information, extending the coverage of challenging poses.
- We show consistent improvements in generalization

across various datasets and multiple TPE architectures.

## 2. Related Works

### 2.1. Data Augmentation for Pose Estimation

Numerous data augmentation techniques have been developed to enhance the generalization performance of 2D and 3D human pose estimation. Image-level augmentations (*e.g.*, rotation and transformation) diversify training data to some extent [7, 42, 55, 58], yet they do not drastically change the original image. Other approaches utilize human parsing, where body parts are segmented and synthesized into real images [3, 23], but struggle with ambiguities like differentiating front and back views. Pose-guided or appearance-modifying image generation [6, 11, 45, 49] synthesize images from motion capture or transform human appearance; however, the generated data deviate from real-world distribution as they are not derived from real images.

Other approaches rotate or scale 3D poses and project them into 2D [13, 16, 31], though mainly limited to 2D-to-3D lifting tasks. For instance, PoseAug [16] and Adapt-Pose [13] utilize GAN to produce 3D poses via transformations and bone-level augmentations. More recent work, PoseGen [14], generates diverse poses via GAN and renders images [53]. However, challenges still remain, as it lacks background and is constrained to predefined human models. Despite these efforts, the need for augmentation methods that better align with real-world distributions and improve generalization across varied scenarios is evident.

### 2.2. Generative Model-based Data Synthesis

Generative models have recently advanced recognition tasks such as image classification [19, 43, 47], object detection [10, 30, 57], and segmentation [9, 22, 30, 36, 41, 57, 61]. The core idea is to harness pre-trained generative models [46, 52, 63]—which encode rich knowledge from large-scale data [48]—to generate additional training data. For example, DiverGen [9] and FreeMask [61] leverage off-the-shelf text-to-image diffusion models [46, 52, 59] to synthesize diverse training examples with minimal guidance, effectively boosting accuracy on underrepresented classes. Some approaches use auxiliary modules [10, 41, 57] or light fine-tuning [22, 44] for task-specific data generation. A common strategy is to concentrate on corner-cases where standard training data falls short, which is straightforward in object detection and segmentation since these cases can be defined by discrete classes. For example, if “giraffes” yield low accuracy, one can prompt a generative model for more giraffe images, directly enlarging that corner-case category.

In human pose estimation, however, identifying and synthesizing corner cases is considerably more complex. First, the output is a set of continuous keypoint coordinates rather than discrete labels, so corner-case poses cannot be grouped

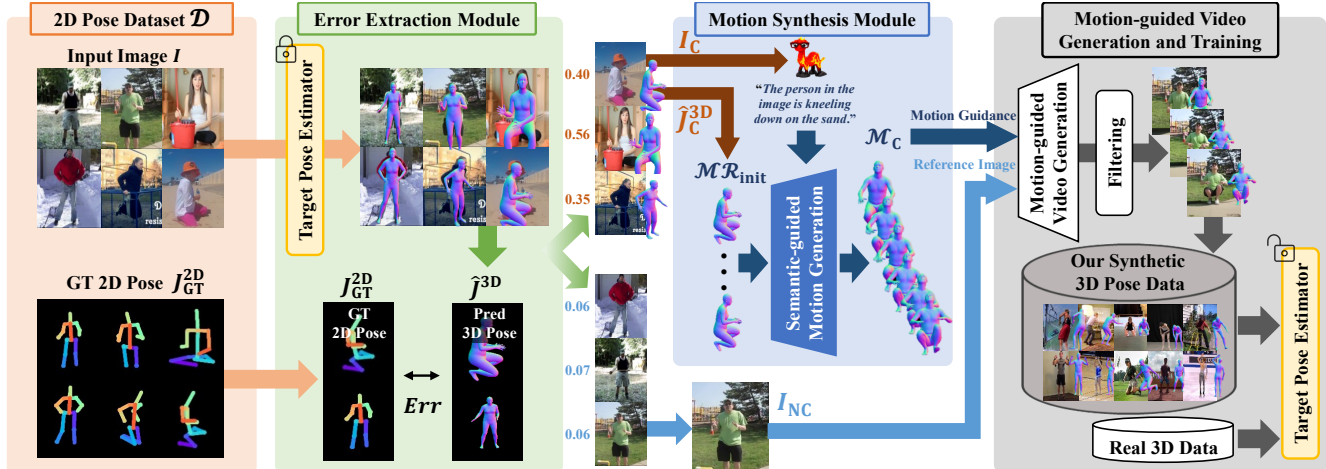


Figure 2. **Proposed Framework for 3D Pose Data Augmentation.** Our framework begins by identifying challenging and non-challenging images in a 2D pose dataset using the Error Extraction Module (EEM). EEM isolates poses with high error from the Target Pose Estimator (TPE) as challenging data, which are then processed in the Motion Synthesis Module (MSM) to generate complex motion sequences. Subsequently, a motion-guided video generation model creates synthetic training data using non-challenging images as references, which then undergoes a filtering process to finally train the TPE.

neatly by class. Instead, each problematic pose must be handled at the sample level. Second, even after identifying the problematic pose, reproducing it via simple text prompts (e.g., “person bending forward with partially occluded right arm”) is nontrivial due to the intricate geometric details in pose. These challenges highlight a gap in current generative data frameworks, underscoring the need for new strategies to effectively pinpoint and synthesize corner-case poses.

### 3. Methods

We propose PoseSyn, a novel framework for augmenting in-the-wild 2D pose dataset into 3D human pose data, enabling 3D pose estimators to generalize effectively across in-the-wild scenes. The framework consists of two main modules: the *Error Extraction Module* (EEM) and the *Motion Synthesis Module* (MSM). First, EEM identifies challenging poses in the 2D dataset  $\mathcal{D}$  by comparing the GT 2D pose  $J_{GT}^{2D}$  with the 2D projection pose  $\hat{J}^{2D}$  of the pseudo-label 3D pose  $\hat{J}^{3D}$  predicted by a target pose estimator (TPE). This process produces a set of challenging images and their mis-predicted pose data  $\mathcal{D}_C = \{I_C, \hat{J}_C^{3D}\}$  (Sec. 3.1). Since the actual GT 3D poses for these challenging samples are not available, MSM generates diverse motion sequences to approximate these problematic cases. To achieve this, MSM first employs a Vision Language Model (VLM) [56] to generate descriptive captions for the challenging images  $I_C$ . Then, using a text-to-motion model under the integrated guidance of both caption and the mis-predicted pose, MSM synthesizes challenging motion sequences  $\mathcal{M}_C$  (Sec. 3.2). Rather than directly estimating the unknown GT pose as a single pose, these motion sequences serve to extend the pose diversity, covering plausible variations of the challeng-

ing scenario. The synthesized motion sequence  $\mathcal{M}_C$  is subsequently utilized in the *Motion-guided Video Generation and Training* stage, where an off-the-shelf human animation model [64] generates image sequences aligned with  $\mathcal{M}_C$ . Finally, these created data (i.e., images and 3D poses) are filtered, retaining only those that improve generalization, and then used to fine-tune TPE (Sec. 3.3). The comprehensive methodology is illustrated in Fig. 2.

#### 3.1. Error Extraction Module

PoseSyn identifies poses within the in-the-wild 2D pose dataset where TPE exhibits high error, termed as challenging poses. Specifically, let’s denote 2D pose dataset as  $\mathcal{D} = \{I, J_{GT}^{2D}\}$ , where  $I$  is an image,  $J_{GT}^{2D} \in \mathbb{R}^{N_{2D} \times 2}$  is the corresponding GT 2D pose, and  $N_{2D}$  is the number of 2D joints. Through TPE, we predict the pseudo-label 3D pose  $\hat{J}^{3D} = \text{TPE}(I)$  and project  $\hat{J}^{3D}$  onto 2D image plane  $\hat{J}^{2D} = \text{Proj}(\hat{J}^{3D}, f, p)$ , where  $f$  is the focal length and  $p$  is the principal point. Next, to identify the challenging poses, we calculate the error for each data sample in 2D dataset  $\mathcal{D}$ :

$$Err = \sum_{n=2}^{N_{2D}} \mathbf{w}_n \left| \left( \hat{j}^{2D,n} - \hat{j}^{2D,1} \right) - \left( J_{GT}^{2D,n} - J_{GT}^{2D,1} \right) \right|, \quad (1)$$

where  $\hat{j}^{2D,n}$  and  $J_{GT}^{2D,n}$  are the predicted and GT coordinates of the  $n$ -th joint, respectively. The pelvis joint is designated as the root joint (i.e.,  $n = 1$ ) for calculating relative joint positions. The parameter  $\mathbf{w}_n$  denotes the weight assigned for the  $n$ -th joint, which allows important joints to be emphasized when determining challenging poses. In our method, we assign greater weight to the arms and legs



Figure 3. **EEM Results.** The proposed EEM identifies (a) challenging and (b) non-challenging data from in-the-wild 2D pose dataset. Challenging data includes intricate and dynamic poses, whereas non-challenging data consists of stationary and static poses, demonstrating the effectiveness of the EEM.

due to their higher variability and complexity in movement, which enhances the accuracy of identifying challenging poses with intricate joint configurations and dynamic orientations. Based on the metric  $Err$ , we obtain the top  $K_C$  poses and their corresponding images as challenging data  $\mathcal{D}_C = \{I_C, \hat{J}_C^{3D}\}$ , using the following operation:

$$\mathcal{D}_C = \text{Top}_{K_C} (Err(d) \mid d \in \mathcal{D}). \quad (2)$$

Similarly, we identify non-challenging data  $\mathcal{D}_{NC} = \{I_{NC}, \hat{J}_{NC}^{3D}\}$ , which consist of the bottom  $K_{NC}$  poses and their corresponding images with respect to the error metric:

$$\mathcal{D}_{NC} = \text{Top}_{K_{NC}} (-Err(d) \mid d \in \mathcal{D}). \quad (3)$$

As shown in Fig. 3, our EEM isolates the challenging data, typically dynamic and complex poses where TPE struggles. The non-challenging data, on the other hand, consists primarily of static poses. Previous methods [14, 35] identify challenging poses from unrealistic synthetic or simulated data, thereby lacking the ability to enhance generalization to real-world intricate poses. However, our EEM identifies the challenging poses in a real-world dataset to create the training data that can improve real-world generalization.

### 3.2. Motion Synthesis Module

Given the challenging data  $\mathcal{D}_C = \{I_C, \hat{J}_C^{3D}\}$  identified by EEM, the next necessary step is to generate training data that approximates the actual challenging poses implied by the images, which are unavailable. One simple approach is to directly utilize the mis-predicted poses  $\hat{J}_C^{3D}$ . However, since  $\hat{J}_C^{3D}$  deviates from the actual problematic poses, using it alone often results in ineffective training data. Another straightforward approach is to extract text descriptions from challenging images  $I_C$  and generate poses from these descriptions. Yet, text alone can be ambiguous, making it difficult to precisely capture complex pose configurations.

To overcome these shortcomings, MSM integrates both information in  $\mathcal{D}_C$  (*i.e.*, textual cues from the challenging

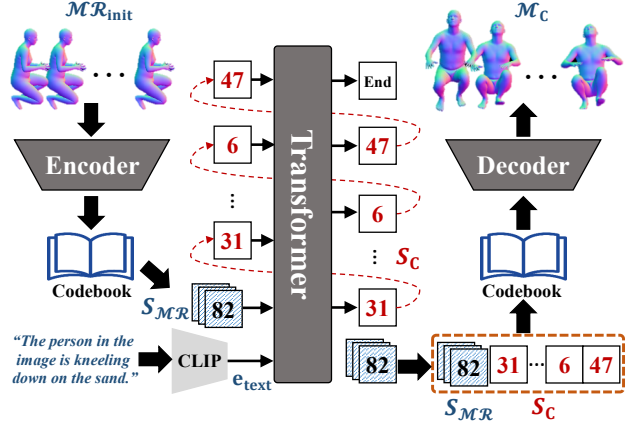


Figure 4. **SMG Architecture.** Proposed Semantic-guided Motion Generation (SMG) augments a mis-predicted pose into motion sequences. An initial motion representation  $\mathcal{MR}_{init}$  is encoded and mapped into the codebook to produce initial motion indices  $\mathcal{S}_{MR}$ . The transformer takes text embeddings  $e_{text}$  and  $\mathcal{S}_{MR}$  both as input and generates motion indices for motion sequence  $\mathcal{M}_C$ .

image  $I_C$  and the mis-predicted pose  $\hat{J}_C^{3D}$ ) and augments the inaccurate pseudo-label  $\hat{J}_C^{3D}$  into motion sequences that encompass plausible pose variations of the challenging scenario. We instantiate this idea by modifying the existing text-to-motion model, T2M-GPT [62], to incorporate both the text description and the mis-predicted pose. Concretely, a VLM [56] provides text descriptions for each challenging image  $I_C$  by answering, “What is the motion of the someone in the image? Please answer similar to {Text-to-Motion prompt examples}” (refer to supplementary material for details). Next, to address the ambiguity of text-based generation, we introduce an initial motion representation,  $\mathcal{MR}_{init}$ , as shown in the SMG architecture in Fig. 4. Our empirical findings, presented in Sec. 4.4, demonstrate that this approach effectively mitigates the ambiguity, surpassing the use of text embedding,  $e_{text}$ , alone.

Specifically, we compute  $\mathcal{MR}_{init} = F(\hat{J}_C^{3D} \otimes T)$  by replicating the pseudo-labeled pose  $\hat{J}_C^{3D}$  over  $T$  time steps and applying a motion representation processing operation  $F$  [18]. This initial motion representation is then encoded via an encoder  $\mathcal{E}$  in Motion VQ-VAE of T2M-GPT:  $\mathcal{Z}_{MR} = \mathcal{E}(\mathcal{MR}_{init})$ , where  $\mathcal{Z}_{MR} = \{z_{MR}^m\}_{m=1}^M$  are the latent features of initial motion sequences. Here,  $M = T/r$  is a sequence length, and  $r$  is a temporal downsampling factor of the encoder. Each latent feature is mapped into a codebook  $\mathcal{C}$  with  $Q$  codes  $\{c_q\}_{q=1}^Q$  in the VQ-VAE to obtain initial motion indices  $\mathcal{S}_{MR} = \{s_{MR}^m\}_{m=1}^M$  as follows:

$$s_{MR}^m = \arg \min_q \|z_{MR}^m - c_q\|_2. \quad (4)$$

The initial motion indices  $\mathcal{S}_{MR}$  provide additional semantic information for guiding the generation process, helping to complement the insufficient information in the text de-

scription alone. Under the guidance of both textual and initial motion inputs, the motion indices are generated in an autoregressive manner by predicting the distribution of possible next motion indices using a transformer as follows:

$$p(\mathcal{S}_C | \mathbf{e}_{\text{text}}, \mathcal{S}_{\mathcal{MR}}) = \prod_{i=0}^{|\mathcal{S}_C|} p(s^i | \mathbf{e}_{\text{text}}, \mathcal{S}_{\mathcal{MR}}, s^{<i}). \quad (5)$$

Finally, the generated motion indices  $\mathcal{S}_C$ , along with  $\mathcal{S}_{\mathcal{MR}}$ , are input into the codebook and decoded through Motion VQ-VAE, producing the augmented challenging motion sequence  $\mathcal{M}_C = \{J_{C,l}^{3D}\}_{l=1}^L$  consisting of  $L$  3D joint poses.

### 3.3. Motion-guided Video Generation and Training

To create images corresponding to the generated motion  $\mathcal{M}_C$ , we utilize Champ [64], an off-the-shelf human animation model that applies motion guidance to reference RGB images. Using RGB images  $I_{NC}$  from non-challenging data  $\mathcal{D}_{NC}$ , we generate a human animated video  $\mathcal{V}_C = \{I_{C,l}\}_{l=1}^L$  based on the generated motion  $\mathcal{M}_C$ . To align the global orientation of  $\mathcal{M}_C$  with the human orientation in  $I_{NC}$ , we preprocess the parameters to ensure natural human animation. Specifically, among the camera parameters, SMPL pose, and SMPL shape parameters from  $I_{NC}$ , we replace and adjust the pose parameters with those from generated motion  $\mathcal{M}_C$  (refer to supplementary materials for details). Unlike previous method [14], which utilizes NeRF [53] to render novel-view human images from generated 3D poses—often resulting in artifacts like background absence and limited human appearance diversity—our approach leverages diverse reference images  $I_{NC}$ . This approach produces images of varied backgrounds and human appearances, making them more suitable as training data for image-to-3D pose estimators.

While these animated videos exhibit visually appealing qualities that align with the generated motions, artifacts such as the blending between human figures and their backgrounds can still occur. To filter out noisy samples, each generated image  $I_{C,l} \in \mathcal{V}_C$  is processed by TPE to predict the 3D joint pose  $\hat{J}_l^{3D} = \text{TPE}(I_{C,l})$ . We then compute the error  $\text{Err}_{3D,l}$  for each image  $I_{C,l}$  between the predicted 3D joint pose  $\hat{J}_l^{3D,n}$  and the corresponding generated 3D joint pose  $J_{C,l}^{3D}$  (as derived in Sec. 3.2), as follows:

$$\text{Err}_{3D,l} = \sum_{n=2}^{N_{3D}} \left| \left( \hat{J}_l^{3D,n} - \hat{J}_l^{3D,1} \right) - \left( J_{C,l}^{3D,n} - J_{C,l}^{3D,1} \right) \right|, \quad (6)$$

where  $J_{C,l}^{3D,n}$  and  $J_{C,l}^{3D,1}$  denote the generated 3D joint pose for the  $n$ -th joint and root joint, respectively, and  $N_{3D}$  denotes the number of 3D joints. Any images with  $\text{Err}_{3D,l}$  above the threshold  $\tau$  are discarded. Finally, the remaining high-quality 3D pose data are combined with the original pre-training real data, to re-train TPE, enhancing its in-the-wild generalization capability.

## 4. Experiments

### 4.1. 3D Pose Estimators and Baselines

We evaluated three image-to-3D pose estimation models as our target pose estimators (TPEs): Hybrik [32], 3DCrowdNet [8], and 4DHumans [15]. Hybrik, which bridges the gap between 3D keypoint estimation and body mesh estimation via a hybrid analytical-neural inverse kinematics, is widely used for its strong performance in 3D human pose estimation. 3DCrowdNet, designed to address the domain gap by leveraging a crowded scene-robust image feature, is commonly employed for its robust 3D pose estimation performance in in-the-wild crowded scenes. 4DHumans, which achieves leading 3D human mesh recovery through a transformer architecture, is distinguished by its exceptional performance on challenging poses. To ensure the broad applicability of our approach, we conduct experiments across these diverse TPEs, which represent different architectures and target scenarios, reflecting the varying model sizes and constraints encountered in real-world applications.

We compared our method against PoseGen [14] and *Ours-N*, two baselines focused on augmenting 3D pose data. Unlike *Ours*, which measures TPE performance on each real-world image to identify challenging poses where TPE underperforms, PoseGen generates synthetic poses with a GAN and renders them via NeRF, identifying challenging poses by measuring TPE performance on these artificially created images. *Ours-N* employs the same challenging-pose identification strategy (EEM) and motion synthesis (MSM) as *Ours*, but replaces the human animation model [64] with NeRF rendering—mirroring PoseGen’s image-generation method. Thus, comparing *Ours* and *Ours-N* reveals the importance of realistic image rendering, while comparing *Ours-N* and PoseGen highlights the advantage of our key modules (*i.e.*, EEM and MSM).

### 4.2. Datasets

Before applying our approach, we pre-trained each TPE (*i.e.*, Hybrik, 3DCrowdNet, and 4DHumans) on five datasets: Human3.6M [20], MuCo [38], MPI-INF-3DHP [37], MPII [1], and MSCOCO [34]. The first three datasets (*i.e.*, Human3.6M, MuCo, and MPI-INF-3DHP) provide accurate 3D pose annotations captured with multi-view camera systems, but they either lack background diversity or have synthetic backgrounds. On the other hand, MPII and MSCOCO offer more diverse and complex backgrounds, but only contain pseudo-labeled [40] 3D pose annotations. After pre-training TPE on these datasets, our framework synthesized an additional 3D pose dataset from the MPII dataset, an in-the-wild 2D pose dataset. Specifically, we identified the top 500 challenging samples ( $K_C = 500$ ) and the bottom 200 non-challenging samples ( $K_{NC} = 200$ ) from single-person in-the-wild 2D pose




Figure 5. **Qualitative Comparison with Baselines.** 3DCrowdNet trained with only real data exhibits inaccurate 3D pose predictions with limited generalization capability. Baselines (*i.e.*, PoseGen and Ours-N) show limited performance gains, but our approach achieves more accurate pose predictions across diverse real-world datasets with various backgrounds and occlusions, challenging poses, and multi-views. Red boxes highlight areas of incorrect predictions in models trained with real-only data and baseline methods.

TPE	Method	3DPW		EMDB		CMU_171204		CMU_171026		HuMMan		Mean		LSPET	JHMDB	Mean
		MPJPE $\downarrow$	PA-MPJPE $\downarrow$	MPJPE $\downarrow$	PA-MPJPE $\downarrow$	MPJPE $\downarrow$	PA-MPJPE $\downarrow$	MPJPE $\downarrow$	PA-MPJPE $\downarrow$	MPJPE $\downarrow$	PA-MPJPE $\downarrow$	MPJPE $\downarrow$	PA-MPJPE $\downarrow$	PCK $^{0.6\uparrow}$	PCK $^{0.9\uparrow}$	PCK $^{0.9\uparrow}$
3DCrowdNet	Real-only	81.7	51.1	115.8	71.2	108.8	72.5	110.7	70.4	98.9	65.8	103.2	66.2	69.5	80.5	75.0
	PoseGen	80.0	50.0	113.1	70.5	104.0	68.6	106.8	68.3	94.5	64.0	99.7	64.3	70.0	79.6	74.8
	Ours-N	78.6	49.7	111.8	69.7	103.1	68.8	105.8	68.2	<b>93.1</b>	63.9	98.5	64.1	70.5	81.0	75.8
	<b>Ours</b>	<b>77.4</b>	<b>48.9</b>	<b>111.0</b>	<b>68.3</b>	<b>101.0</b>	<b>67.3</b>	<b>105.0</b>	<b>67.9</b>	<b>93.1</b>	<b>62.3</b>	<b>97.5</b>	<b>62.9</b>	<b>70.7</b>	<b>81.2</b>	<b>76.0</b>
Hybrik	Real-only	88.0	48.6	155.4	104.1	117.5	79.8	125.6	82.2	119.7	75.1	121.2	78.0	70.7	78.7	74.7
	PoseGen	84.8	47.6	146.9	101.0	120.8	83.2	124.8	90.9	111.2	73.3	117.9	79.6	71.1	78.7	74.9
	Ours-N	81.1	46.5	141.1	99.2	115.7	80.1	127.0	86.4	100.4	<b>70.4</b>	113.1	76.5	71.3	78.7	75.0
	<b>Ours</b>	<b>78.4</b>	<b>46.2</b>	<b>129.9</b>	<b>90.2</b>	<b>100.3</b>	<b>68.9</b>	<b>119.1</b>	<b>77.1</b>	<b>95.3</b>	71.1	<b>104.6</b>	<b>70.7</b>	<b>71.7</b>	<b>78.8</b>	<b>75.3</b>
4DHumans	Real-only	81.3	54.3	116.3	79.1	115.1	82.0	115.3	74.2	106.1	73.8	106.8	72.7	86.3	88.3	87.3
	PoseGen	81.1	54.1	114.5	78.3	111.3	79.2	114.5	71.3	105.3	73.8	105.3	71.3	86.1	88.4	87.3
	Ours-N	80.6	53.0	112.6	77.1	106.6	74.9	110.1	68.2	101.0	70.4	102.1	68.7	86.4	88.5	87.5
	<b>Ours</b>	<b>77.0</b>	<b>52.1</b>	<b>108.6</b>	<b>74.8</b>	<b>104.1</b>	<b>72.9</b>	<b>107.8</b>	<b>65.9</b>	<b>98.0</b>	<b>68.5</b>	<b>99.1</b>	<b>66.8</b>	<b>86.6</b>	<b>88.8</b>	<b>87.7</b>

Table 1. **Quantitative comparison with baselines.** Pose estimation performance is evaluated with MPJPE and PA-MPJPE metrics for 3D pose datasets (*i.e.*, 3DPW, EMDB, CMU, and HuMMan), and with PCKh metric for 2D pose datasets (*i.e.*, LSPET and JHMDB). For three TPEs (*i.e.*, 3DCrowdNet, Hybrik, and 4DHumans), our approach outperforms other baselines (*i.e.*, PoseGen and Ours-N) across each dataset, demonstrating the highest generalization improvement in real-world scenarios.

data in the MPII dataset. Then, using our proposed methodology including the filtering based on Eq. 6, we synthesized 27,000 3D pose data samples from the source 2D data. For evaluation, we utilized six datasets: 3DPW [54], EMDB [28], CMU [26], HuMMan [4], LSPET [24, 25], and JHMDB [21]. These datasets encompass various human appearances and poses in in-the-wild scenarios [21, 24, 25, 28, 54] or multi-view camera setups [4, 26]. For 3D pose evaluation on the 3DPW, EMDB, CMU, and HuMMan datasets, we used two metrics: Mean Per Joint Position Error (MPJPE [20]) and Procrustes-aligned MPJPE (PA-MPJPE [20]). For 2D pose evaluation on the LSPET and JHMDB datasets, we used PCKh [1] metric.

### 4.3. Comparison with Baselines

We evaluated PoseSyn through both qualitative and quantitative comparisons against two baselines across three TPEs. As shown in Fig. 5, our methodology, along with PoseGen and Ours-N, improved the performance of TPE (*i.e.*, 3DCrowdNet). However, PoseGen underperformed both Ours-N and Ours on the sample from the JHMDB dataset in the first row of Fig. 5. This discrepancy occurs

because PoseGen obtains challenging poses from NeRF-rendered images, as shown in Fig. 6 (a), rather than directly from in-the-wild images. Furthermore, even when challenging poses were obtained from in-the-wild images, Ours surpassed Ours-N on the sample from the EMDB dataset in the second row. This improvement stems from the fact that the 3D pose data generated by Ours-N primarily consists of NeRF-rendered images, limiting diversity in both backgrounds and human appearances—an important element for training image-to-3D pose estimators—as illustrated in Fig. 6 (b). In contrast, as depicted in Fig. 6 (c), PoseSyn not only acquires challenging poses from real-world data, but also leverages a human animation model to generate images that showcase various human appearances and intricate poses with realistic backgrounds. Consequently, PoseSyn demonstrated significantly improved performance, as illustrated in Fig. 5, aligning with quantitative metrics in Tab. 1. Our method yields 6–14% improvements in MPJPE and 5–9% in PA-MPJPE on 3D datasets, along with a 1–2% accuracy gain on 2D datasets across various TPEs. These results confirm that PoseSyn provides robust data augmentation benefits, consistently enhancing performance across




Figure 6. **Qualitative Comparison of Augmented Data.** (a) **PoseGen** generates 3D pose data where poses differ from real-world problematic poses due to the way of acquiring challenging poses from synthetic images rendered by NeRF. (b) **Ours-N** extracts challenging poses from in-the-wild dataset but lacks diversity in rendered images. (c) **Ours** not only acquires problematic poses from real-world but also generates images with various human appearances and realistic backgrounds. Note that each column in (b) and (c) displays images generated from the same poses.

various real-world scenarios.

#### 4.4. Effectiveness of MSM

*Does MSM effectively augment identified challenging poses into motion sequences?* To evaluate whether MSM effectively augments challenging poses identified by EEM into motion sequences, we conducted a direct evaluation of how closely MSM-generated motions align with the GT 3D poses—without training the TPE on these synthesized data. We first selected challenging samples from EMDB dataset, which contains GT 3D pose annotations, by applying TPE (3DCrowdNet) and choosing the top 100 samples with the largest error. On these challenging samples, the average PA-MPJPE between the TPE-predicted and GT poses was 181.7 mm (Tab. 2 (a)). We then augmented each of 100 TPE-predicted poses into motion sequences using two methods: (b) MSM without initial motion representation (*i.e.*, w/o  $\mathcal{MR}_{init}$ ) and (c) our full MSM with both initial motion and text description. Within each sequence, we computed the PA-MPJPE between every synthesized pose and the GT pose, then calculated the mean, standard deviation (std), and minimum (min) of these PA-MPJPE values. Finally, we averaged these three statistics across all 100 sequences. Since (a) is just a single TPE-predicted pose for each hard case, its mean and min are the same and its std is zero. As shown in Tab. 2, though motion synthesis approaches (*i.e.*, (b) and (c)) have higher mean error than mis-prediction (a), these approaches achieved lower min error. This indicates that, by synthesizing plausible pose variants, they produce at least one pose that is closer to the true problematic pose than the single naive mis-predicted pose. Notably, (b) omitting the initial pose representation raises both mean (222.3 mm) and min (151.1 mm) error, emphasizing the vital role of  $\mathcal{MR}_{init}$  to accurately target challenging poses.

*Is the effectiveness of MSM essential for TPE’s per-*

Method	Mean ( $\downarrow$ ) $\pm$ Std ( $\uparrow$ )	Min ( $\downarrow$ )
(a) $\hat{j}^{3D}$	<b>181.7</b> $\pm$ 0. mm	181.7 mm
(b) w/o $\mathcal{MR}_{init}$	222.3 $\pm$ 36.4 mm	151.1 (-16.8%) mm
(c) Ours	209.3 $\pm$ <b>36.5</b> mm	<b>140.8</b> (-22.5%) mm

Table 2. **Analysis on Effectiveness of MSM.** We assessed how well the generated poses reflect the identified challenging poses. This evaluation was conducted using PA-MPJPE between the GT 3D pose in EMDB dataset and the generated poses. Results show (a) mis-prediction, (b) without  $\mathcal{MR}_{init}$ , (c) our full MSM.


Method	Metrics	3DPW	EMDB	CMU 171204	CMU 171206	HuMMan	Mean
Real-only	MPJPE $\downarrow$	81.7	115.8	108.8	110.7	98.9	103.2
	PA-MPJPE $\downarrow$	51.1	71.2	72.5	70.4	65.8	66.2
$\hat{j}^{3D}$	MPJPE $\downarrow$	78.7	115.8	102.7	<b>104.8</b>	96.2	99.6
	PA-MPJPE $\downarrow$	51.1	72.1	68.4	69.6	65.6	65.4
w/o $\mathcal{MR}_{init}$	MPJPE $\downarrow$	78.2	112.7	102.6	106.1	<b>93.1</b>	98.5
	PA-MPJPE $\downarrow$	50.1	70.2	68.6	68.9	64.0	64.4
Ours	MPJPE $\downarrow$	<b>77.4</b>	<b>111.0</b>	<b>101.0</b>	105.0	<b>93.1</b>	<b>97.5</b>
	PA-MPJPE $\downarrow$	<b>48.9</b>	<b>68.3</b>	<b>67.3</b>	<b>67.9</b>	<b>62.3</b>	<b>62.9</b>

Table 3. **Effectiveness of MSM in TPE Training.** We assessed the impact of MSM on TPE performance by training 3DCrowdNet using data augmented with three different MSM configurations (mis-prediction, without  $\mathcal{MR}_{init}$ , and our full MSM).

*performance improvement?* MSM aims to augment the mis-predicted pose for challenging image into motion sequences better approximate the actual challenging pose, followed by human animation model to generate images under condition of the augmented motion sequences. Without MSM, images are generated using only the mis-predicted pose as a condition, as shown in Fig. 7 (a). Since this mis-predicted pose differs from the actual challenging pose, 3D pose data synthesized in this manner does not properly reflect such challenging scenarios. Furthermore, this approach has no ability to produce any diverse pose variations for these difficult cases, leading to minimal performance gains when fine-tuning with such monotonous pose data, as indicated in Tab. 3 (*i.e.*,  $\hat{j}^{3D}$ ). Conversely, when images are generated using augmented motion that aligns more closely with the challenging images—thanks to MSM—the performance improvement was maximized (See Ours in Tab. 3). This is because the human animation model generates images conditioned on motions close to pose where TPE shows low performance, effectively targeting those challenging scenarios. Lastly, without initial motion representation, generated motions failed to reflect the actual problematic pose’s configuration as shown in Fig. 7 (b), leading to limited performance gains (*i.e.*, w/o  $\mathcal{MR}_{init}$  in Tab. 3).

#### 4.5. Ablation Study

To assess the effectiveness of EEM, we first ablated EEM in our framework. As shown in Tab. 4, without EEM, the performance gain was minimal, highlighting the importance of identifying and synthesizing challenging samples where TPE underperforms. Additionally, we evaluated how using a *different* TPE inside EEM would affect final performance on the actual TPE (*i.e.*, 3DCrowdNet). Specifically, we replaced 3DCrowdNet with Hybrik or 4DHu-



SMG text input: “The person in the image is kneeling down and chopping with an axe.”

Figure 7. **Effectiveness of MSM in Data Synthesis.** The first row shows the challenging image and pose in  $\mathcal{D}_C$  and reference image in  $\mathcal{D}_{NC}$ . (a) Without MSM, images are generated using the repeated mis-predicted pose as a condition. (b) Without  $\mathcal{M}\mathcal{R}_{init}$ , MSM-generated motion only reflects “chopping with an axe” missing “kneeling down”, the actual problematic pose. (c) Ours effectively augments the actual problematic pose into 3D pose data.

Model (EEM)	Training method	Metrics	3DPW	EMDB	CMU 171204	CMU 171206	HuMMan	Mean
w/o EEM	FT	MPJPE ↓	80.2	114.2	107.1	109.0	94.9	101.1
		PA-MPJPE ↓	50.5	71.1	72.0	71.1	65.1	66.0
Hybrik	FT	MPJPE ↓	79.3	113.5	103.5	106.2	94.9	99.5
		PA-MPJPE ↓	50.0	70.4	69.4	68.8	64.6	64.6
4DHumans	FT	MPJPE ↓	78.6	112.7	103.0	106.6	93.2	98.8
		PA-MPJPE ↓	50.0	69.9	68.4	68.4	64.2	64.2
3DCrowdNet	FS	MPJPE ↓	78.8	112.4	103.0	106.2	93.5	98.8
		PA-MPJPE ↓	49.5	69.9	68.2	67.9	63.7	63.8
3DCrowdNet	FT	MPJPE ↓	<b>77.4</b>	<b>111.0</b>	<b>101.0</b>	<b>105.0</b>	<b>93.1</b>	<b>97.5</b>
		PA-MPJPE ↓	<b>48.9</b>	<b>68.3</b>	<b>67.3</b>	<b>67.9</b>	<b>62.3</b>	<b>62.9</b>

Table 4. **Ablation Study on EEM.** We conducted an ablation study to examine the effect of different pose estimators within EEM on TPE performance. We compared ours with three settings: (1) excluding EEM entirely, (2) using a pose estimator different from the TPE (*i.e.*, Hybrik and 4DHumans) in EEM, and (3) using the same TPE model, 3DCrowdNet, in EEM but training estimator from scratch. “FT” denotes a fine-tuning approach and “FS” represents an approach of training estimator from scratch.

mans for identifying challenging poses. Although both variants showed improved performance upon the without EEM baseline, neither matched the performance achieved by using 3DCrowdNet in EEM, which is our original method. This demonstrates that each TPE has its own problematic poses with lower performance, so focusing data synthesis on these poses yields better results. We also tested an approach called, 3DCrowdNet (FS), which uses a pre-trained 3DCrowdNet to identify hard samples, but trains a randomly initialized 3DCrowdNet using synthesized data from those hard samples. While 3DCrowdNet (FS) surpassed different model architectures, it still underperformed compared to using the exact same TPE weights. Overall, these results confirm that using the same pre-trained TPE within EEM is the most effective way to generate tailored training data and maximize generalization gains. Lastly, we conducted an ablation study for various threshold values in the

Filtering $\tau$	Pass	3DPW		EMDB		Mean	
		MPJPE ↓	PA-MPJPE ↓	MPJPE ↓	PA-MPJPE ↓	MPJPE ↓	PA-MPJPE ↓
200	72.6%	79.9	50.2	114.5	70.8	97.2	61.0
160	64.0%	78.7	49.7	113.0	70.3	95.9	60.0
120	45.7%	<b>77.4</b>	<b>48.9</b>	<b>111.0</b>	<b>68.3</b>	<b>94.2</b>	<b>58.6</b>
80	28.2%	78.6	49.6	111.9	69.8	95.3	59.7
40	6.6%	79.5	49.8	112.4	70.1	96.0	60.0

Table 5. **Ablation Study on Filtering.** We experimented with different threshold values ( $\tau$ ) in filtering process to examine their impact on TPE (*i.e.* 3DCrowdNet) performance when training with our synthesized 3D pose data. “Pass” denotes the percentage of data retained after filtering step.

filtering step, allowing us to assess its impact on training the TPE with our created 3D pose data, as shown in Tab 5. Though our method utilizes an off-the-shelf human animation model for creating images aligned with synthesized motion, editing by motion can sometimes lead to blending artifacts between the human and background. Using the degraded data blindly (high  $\tau$ ) can lower model performance, while excessive filtering (low  $\tau$ ) could leave only simpler samples, restricting performance improvement. We set  $\tau = 120$  to ensure high-quality images for training the image-to-3D pose estimator, while still retaining problematic pose samples where TPE struggles.

## 5. Conclusion

We present PoseSyn, a novel framework to improve TPE generalization by effectively augmenting in-the-wild 2D pose data into enriched 3D pose training data. Our EEM effectively identifies challenging poses and images, while our MSM generates motion sequences aligned with the identified problematic poses under the guidance of both text description and initial motion representation. These motion sequences are then used to generate images of various human appearances and poses, enriching 3D pose data and enhancing the generalization capability of TPEs. Experimental results confirm that PoseSyn achieves notable accuracy improvements across various benchmarks and TPE architectures. We believe this framework provides a promising solution for enhancing generalization of pose estimator regardless of its model characteristic without the need for costly acquisition of 3D pose data.

**Limitation** Our framework utilizes an off-the-shelf human animation model [64] for motion-guided video generation, editing a human in a reference image based on generated motions. However, since the animation model struggles with multi-person scenarios (*e.g.*, human interaction), our framework is also limited to single-person scenario. Extending our framework would require advancements in multi-person animation model with the integration of interaction-aware motion generation models into our MSM. Future work could focus on these developments to support multi-person 3D pose data synthesis.

## References

- [1] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 5, 6
- [2] Fabien Baradel\*, Matthieu Armando, Salma Galaoui, Romain Brégier, Philippe Weinzaepfel, Grégory Rogez, and Thomas Lucas\*. Multi-hmr: Multi-person whole-body human mesh recovery in a single shot. In *ECCV*, 2024. 1
- [3] Yanrui Bin, Xuan Cao, Xinya Chen, Yanhao Ge, Ying Tai, Chengjie Wang, Jilin Li, Feiyue Huang, Changxin Gao, and Nong Sang. Adversarial semantic data augmentation for human pose estimation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIX 16*, pages 606–622. Springer, 2020. 2
- [4] Zhongang Cai, Daxuan Ren, Ailing Zeng, Zhengyu Lin, Tao Yu, Wenjia Wang, Xiangyu Fan, Yang Gao, Yifan Yu, Liang Pan, Fangzhou Hong, Mingyuan Zhang, Chen Change Loy, Lei Yang, and Ziwei Liu. HuMMan: Multi-modal 4d human dataset for versatile sensing and modeling. In *17th European Conference on Computer Vision, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VII*, pages 557–577. Springer, 2022. 1, 2, 6
- [5] Junuk Cha, Muhammad Saqlain, GeonU Kim, Mingyu Shin, and Seungryul Baek. Multi-person 3d pose and shape estimation via inverse kinematics and refinement. In *ECCV*, 2022. 1
- [6] Wenzheng Chen, Huan Wang, Yangyan Li, Hao Su, Zhenhua Wang, Changhe Tu, Dani Lischinski, Daniel Cohen-Or, and Baoquan Chen. Synthesizing training images for boosting human 3d pose estimation. In *2016 Fourth International Conference on 3D Vision (3DV)*, pages 479–488. IEEE, 2016. 2
- [7] Yu Chen, Chunhua Shen, Xiu-Shen Wei, Lingqiao Liu, and Jian Yang. Adversarial posenet: A structure-aware convolutional network for human pose estimation. In *Proceedings of the IEEE international conference on computer vision*, pages 1212–1221, 2017. 2
- [8] Hongsuk Choi, Gyeongsik Moon, JoonKyu Park, and Kyoung Mu Lee. Learning to estimate robust 3d human mesh from in-the-wild crowded scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1475–1484, 2022. 1, 2, 5
- [9] Chengxiang Fan, Muzhi Zhu, Hao Chen, Yang Liu, Weijia Wu, Huaqi Zhang, and Chunhua Shen. Divergen: Improving instance segmentation by learning wider data distribution with more diverse generative data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3986–3995, 2024. 2
- [10] Haoyang Fang, Boran Han, Shuai Zhang, Su Zhou, Cuixiong Hu, and Wen-Ming Ye. Data augmentation for object detection via controllable diffusion models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1257–1266, 2024. 2
- [11] Yongtao Ge, Wenjia Wang, Yongfan Chen, Hao Chen, and Chunhua Shen. 3d human reconstruction in the wild with synthetic data using generative models. *arXiv preprint arXiv:2403.11111*, 2024. 2
- [12] Yongtao Ge, Wenjia Wang, Yongfan Chen, Hao Chen, and Chunhua Shen. 3d human reconstruction in the wild with synthetic data using generative models. *arXiv preprint arXiv:2403.11111*, 2024. 1
- [13] Mohsen Gholami, Bastian Wandt, Helge Rhodin, Rabab Ward, and Z. Jane Wang. Adaptpose: Cross-dataset adaptation for 3d human pose estimation by learnable motion generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13075–13085, 2022. 1, 2
- [14] Mohsen Gholami, Rabab Ward, and Z Jane Wang. Posegen: Learning to generate 3d human pose dataset with nerf. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1905–1913, 2024. 2, 4, 5
- [15] Shubham Goel, Georgios Pavlakos, Jathushan Rajasegaran, Angjoo Kanazawa, and Jitendra Malik. Humans in 4D: Reconstructing and tracking humans with transformers. In *ICCV*, 2023. 1, 2, 5
- [16] Kehong Gong, Jianfeng Zhang, and Jiashi Feng. Poseaug: A differentiable pose augmentation framework for 3d human pose estimation. In *CVPR*, 2021. 1, 2
- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 1
- [18] Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5152–5161, 2022. 4
- [19] Ruifei He, Shuyang Sun, Xin Yu, Chuhui Xue, Wenqing Zhang, Philip Torr, Song Bai, and XIAOJUAN QI. Is synthetic data from generative models ready for image recognition? In *The Eleventh International Conference on Learning Representations*. 2
- [20] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, 2014. 1, 5, 6
- [21] H. Jhuang, J. Gall, S. Zuffi, C. Schmid, and M. J. Black. Towards understanding action recognition. In *International Conf. on Computer Vision (ICCV)*, pages 3192–3199, 2013. 2, 6
- [22] Yuru Jia, Lukas Hoyer, Shengyu Huang, Tianfu Wang, Luc Van Gool, Konrad Schindler, and Anton Obukhov. Dginstyle: Domain-generalizable semantic segmentation with image diffusion models and stylized semantic control. In *European Conference on Computer Vision*, pages 91–109. Springer, 2024. 2
- [23] Wentao Jiang, Sheng Jin, Wentao Liu, Chen Qian, Ping Luo, and Si Liu. Posetrans: A simple yet effective pose transformation augmentation for human pose estimation. In *European Conference on Computer Vision*, pages 643–659. Springer, 2022. 2

- [24] Sam Johnson and Mark Everingham. Clustered pose and nonlinear appearance models for human pose estimation. In *bmvc*, page 5. Aberystwyth, UK, 2010. 2, 6
- [25] Sam Johnson and Mark Everingham. Learning effective human pose estimation from inaccurate annotation. In *CVPR 2011*, pages 1465–1472. IEEE, 2011. 6
- [26] Hanbyul Joo, Tomas Simon, Xulong Li, Hao Liu, Lei Tan, Lin Gui, Sean Banerjee, Timothy Scott Godisart, Bart Nabbe, Iain Matthews, Takeo Kanade, Shohei Nobuhara, and Yaser Sheikh. Panoptic studio: A massively multiview system for social interaction capture. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017. 1, 2, 6
- [27] 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. 1
- [28] 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. 1, 2, 6
- [29] 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 *ICCV*, 2019. 1
- [30] Orest Kupyn and Christian Rupprecht. Dataset enhancement with instance-level augmentations. In *European Conference on Computer Vision*, pages 384–402. Springer, 2024. 2
- [31] Haolun Li and Chi-Man Pun. Cee-net: complementary end-to-end network for 3d human pose generation and estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1305–1313, 2023. 2
- [32] Jiefeng Li, Chao Xu, Zhicun Chen, Siyuan Bian, Lixin Yang, and Cewu Lu. Hybrik: A hybrid analytical-neural inverse kinematics solution for 3d human pose and shape estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3383–3393, 2021. 1, 2, 5
- [33] 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 *ECCV*, 2022. 1
- [34] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. 5
- [35] Qihao Liu, Adam Kortylewski, and Alan L Yuille. Pose-examiner: Automated testing of out-of-distribution robustness in human pose and shape estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 672–681, 2023. 4
- [36] Thibaut Loiseau, Tuan-Hung Vu, Mickael Chen, Patrick Pérez, and Matthieu Cord. Reliability in semantic segmentation: Can we use synthetic data? In *European Conference on Computer Vision*, pages 442–459. Springer, 2024. 2
- [37] Dushyant Mehta, Helge Rhodin, Dan Casas, Pascal Fua, Oleksandr Sotnychenko, Weipeng Xu, and Christian Theobalt. Monocular 3d human pose estimation in the wild using improved cnn supervision. In *2017 international conference on 3D vision (3DV)*, pages 506–516. IEEE, 2017. 5
- [38] Dushyant Mehta, Oleksandr Sotnychenko, Franziska Mueller, Weipeng Xu, Srinath Sridhar, Gerard Pons-Moll, and Christian Theobalt. Single-shot multi-person 3d pose estimation from monocular rgb. In *2018 International Conference on 3D Vision (3DV)*, pages 120–130. IEEE, 2018. 1, 5
- [39] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 2
- [40] Gyeongsik Moon, Hongsuk Choi, Sanghyuk Chun, Jiyoung Lee, and Sangdoon Yun. Three recipes for better 3d pseudo-gts of 3d human mesh estimation in the wild. In *Computer Vision and Pattern Recognition Workshop (CVPRW)*, 2023. 5
- [41] Quang Nguyen, Truong Vu, Anh Tran, and Khoi Nguyen. Dataset diffusion: Diffusion-based synthetic data generation for pixel-level semantic segmentation. *Advances in Neural Information Processing Systems*, 36:76872–76892, 2023. 2
- [42] Xi Peng, Zhiqiang Tang, Fei Yang, Rogerio S Feris, and Dimitris Metaxas. Jointly optimize data augmentation and network training: Adversarial data augmentation in human pose estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2226–2234, 2018. 2
- [43] Maan Qraitem, Kate Saenko, and Bryan A Plummer. From fake to real: Pretraining on balanced synthetic images to prevent spurious correlations in image recognition. In *European Conference on Computer Vision*, pages 230–246. Springer, 2024. 2
- [44] Fazle Rahat, M Shifat Hossain, Md Rubel Ahmed, Sumit Kumar Jha, and Rickard Ewetz. Data augmentation for image classification using generative ai. *arXiv preprint arXiv:2409.00547*, 2024. 2
- [45] Grégory Rogez and Cordelia Schmid. Mocap-guided data augmentation for 3d pose estimation in the wild. *Advances in neural information processing systems*, 29, 2016. 2
- [46] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 2
- [47] Mert Bülent Sarıyıldız, Karteek Alahari, Diane Larlus, and Yannis Kalantidis. Fake it till you make it: Learning transferable representations from synthetic imagenet clones. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8011–8021, 2023. 2
- [48] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021. 2
- [49] Yi-Ting Shen, Hyungtae Lee, Heesung Kwon, and Shuvra S Bhattacharyya. Diversifying human pose in syn-

- thetic data for aerial-view human detection. *arXiv preprint arXiv:2405.15939*, 2024. 2
- [50] Karthik Shetty, Annette Birkhold, Srikrishna Jaganathan, Norbert Strobel, Markus Kowarschik, Andreas Maier, and Bernhard Egger. Pliks: A pseudo-linear inverse kinematic solver for 3d human body estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 574–584, 2023. 1
- [51] Soyong Shin, Juyong Kim, Eni Halilaj, and Michael J. Black. Wham: Reconstructing world-grounded humans with accurate 3d motion. In *Computer Vision and Pattern Recognition (CVPR)*, 2024. 1
- [52] Alex Shonenkov, Misha Konstantinov, Daria Bakshandaeva, Christoph Schuhmann, Ksenia Ivanova, and Nadiia Klokova. Deepfloyd-iff. <https://github.com/deep-floyd/IF>, 2023. 2
- [53] Shih-Yang Su, Frank Yu, Michael Zollhöfer, and Helge Rhodin. A-nerf: Articulated neural radiance fields for learning human shape, appearance, and pose. *Advances in neural information processing systems*, 34:12278–12291, 2021. 2, 5
- [54] Timo von Marcard, Roberto Henschel, Michael Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3d human pose in the wild using imus and a moving camera. In *European Conference on Computer Vision (ECCV)*, 2018. 1, 2, 6
- [55] Jiahang Wang, Sheng Jin, Wentao Liu, Weizhong Liu, Chen Qian, and Ping Luo. When human pose estimation meets robustness: Adversarial algorithms and benchmarks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11855–11864, 2021. 2
- [56] Weizhi Wang. Llava-llama-3-8b: A reproduction towards llava-3 based on llama-3-8b llm backbone, 2024. 3, 4
- [57] Weijia Wu, Yuzhong Zhao, Hao Chen, Yuchao Gu, Rui Zhao, Yefei He, Hong Zhou, Mike Zheng Shou, and Chunhua Shen. Datasetdm: Synthesizing data with perception annotations using diffusion models. *Advances in Neural Information Processing Systems*, 36:54683–54695, 2023. 2
- [58] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *Proceedings of the European conference on computer vision (ECCV)*, pages 466–481, 2018. 2
- [59] Han Xue, Zhiwu Huang, Qianru Sun, Li Song, and Wenjun Zhang. Freestyle layout-to-image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14256–14266, 2023. 2
- [60] ChangHee Yang, Kyeongbo Kong, SungJun Min, Dongyoon Wee, Ho-Deok Jang, Geonho Cha, and SukJu Kang. Sefd: learning to distill complex pose and occlusion. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 14941–14952, 2023. 1
- [61] Lihe Yang, Xiaogang Xu, Bingyi Kang, Yinghuan Shi, and Hengshuang Zhao. Freemask: Synthetic images with dense annotations make stronger segmentation models. *Advances in Neural Information Processing Systems*, 36:18659–18675, 2023. 2
- [62] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 4
- [63] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023. 2
- [64] Shenhao Zhu, Junming Leo Chen, Zuozhuo Dai, Yinghui Xu, Xun Cao, Yao Yao, Hao Zhu, and Siyu Zhu. Champ: Controllable and consistent human image animation with 3d parametric guidance. In *European Conference on Computer Vision (ECCV)*, 2024. 2, 3, 5, 8