

# A Dual-Augmentor Framework for Domain Generalization in 3D Human Pose Estimation

Qucheng Peng, Ce Zheng, Chen Chen

Center for Research in Computer Vision, University of Central Florida

{qucheng.peng, ce.zheng}@ucf.edu, chen.chen@crcv.ucf.edu

## Abstract

3D human pose data collected in controlled laboratory settings present challenges for pose estimators that generalize across diverse scenarios. To address this, domain generalization is employed. Current methodologies in domain generalization for 3D human pose estimation typically utilize adversarial training to generate synthetic poses for training. Nonetheless, these approaches exhibit several limitations. First, the lack of prior information about the target domain complicates the application of suitable augmentation through a single pose augmentor, affecting generalization on target domains. Moreover, adversarial training’s discriminator tends to enforce similarity between source and synthesized poses, impeding the exploration of out-of-source distributions. Furthermore, the pose estimator’s optimization is not exposed to domain shifts, limiting its overall generalization ability.

To address these limitations, we propose a novel framework featuring two pose augmentors: the weak and the strong augmentors. Our framework employs differential strategies for generation and discrimination processes, facilitating the preservation of knowledge related to source poses and the exploration of out-of-source distributions without prior information about target poses. Besides, we leverage meta-optimization to simulate domain shifts in the optimization process of the pose estimator, thereby improving its generalization ability. Our proposed approach significantly outperforms existing methods, as demonstrated through comprehensive experiments on various benchmark datasets. Our code will be released at <https://github.com/davidpengucf/DAF-DG>.

## 1. Introduction

3D human pose estimation (HPE) is the process of predicting the 3D coordinates of human joints from images or videos. It serves as the foundation for various applications including person re-identification [29], action recognition

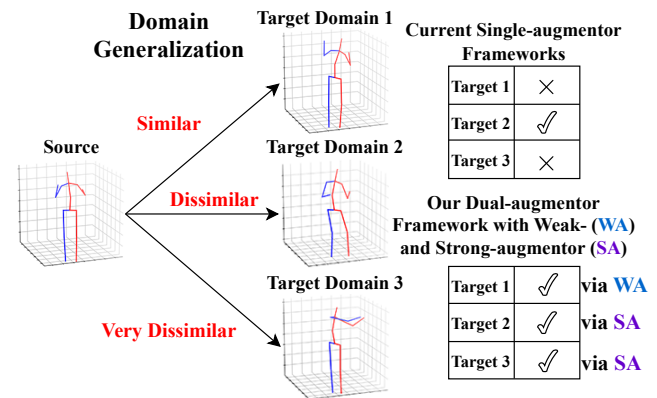


Figure 1. Comparisons between existing single-augmentor frameworks and our proposed dual-augmentor framework on a toy example. Current single-augmentor methods excel at simulating Target Domain 2 but exhibit limitations in simulating Target Domain 1, closely resembling the source, and Target Domain 3, deviating significantly from the source. In our framework, the weak augmentor excels in simulating Target Domain 1, while the strong augmentor effectively imitates both Target Domain 2 and 3.

[10, 20, 34], human mesh recovery [44, 45], virtual reality [9, 35]. However, the annotated 3D data are often collected in controlled laboratory environments for convenience, featuring indoor settings and limited actions performed by few individuals. As a result, pose estimators trained on these labeled datasets face challenges in generalizing to varied in-the-wild scenarios. Hence, the notion of domain generalization (DG) is pivotal in incorporating knowledge from labeled (source) data into a pose estimator that could generalize well on unseen (target) data. Unlike domain adaptation (DA) which involves the training with target data, DG relies solely on the source data as a reference, without any prior information about the target data.

Existing DG approaches for 3D HPE [7, 11, 37] conduct substantial augmentations on the source poses to obtain synthesized poses via adversarial training, and incorporate the synthesized poses as complementary to the source poses for HPE model training. However, these methods have several limitations. *First*, in the context of DG for 3D HPE, there

is a complete lack of information about poses in target domains. If target poses closely resemble the source (as Target Domain 1 in Fig. 1), poses generated by extensive augmentation notably differ from the target, thereby impeding generalization. Conversely, when target poses significantly deviate from the source distributions (as Target Domain 3 in Fig. 1), poses generated via insufficient augmentation may not sufficiently explore out-of-source knowledge to simulate the target. Existing methods only use a single augmentor, making it challenging to simultaneously achieve both objectives. *Second*, the adversarial training between synthesized and source poses constrains the diversity of generation. Existing methods typically employ the generative adversarial network (GAN) [6] structure, which includes one pose generator responsible for pose generation and one discriminator to assist the pose generator by providing feedback. Specifically, the discriminator enforces similarity between synthesized and source poses, aiming to ensure that the generated poses closely resemble the source poses, which harms the exploration of out-of-source knowledge.

To address these limitations, we propose a novel framework featuring two pose augmentors: the weak augmentor and the strong augmentor. The weak augmentor is designed to simulate target poses closely resembling source poses, while the strong augmentor generates target poses that exhibit significant deviations from source distributions. To delineate their characteristics, differential strategies are employed for generation and discrimination processes, as detailed in Sec. 3.3. Notably, our framework alleviates the constraints on strong-augmented poses by traditional adversarial training methods. Instead of enforcing similarity between the source and all the augmented poses, we utilize weak-augmented poses as an intermediary, enabling discrimination between strong- and weak-augmented poses and facilitating discrimination between source and weak-augmented poses. To optimize the utilization of synthesized poses, we introduce meta-optimization among source, weak-augmented, and strong-augmented poses, as elaborated in Sec. 3.4. Our training process exposes the pose estimator to domain shifts during the optimization processes, thereby enhancing its adaptability to handle domain shifts during the inference stage. Our contributions can be summarized in three main aspects:

- We propose a novel framework featuring both the weak and strong pose augmentors, which effectively preserves knowledge related to source poses while simultaneously exploring out-of-source distributions through differential strategies for the generation and discrimination processes of the two augmentors.
- We introduce meta-optimization to enhance the utilization of synthesized poses. By simulating domain shifts among source, weak-augmented, and strong-augmented poses during the optimization processes, the pose estimator's

generalization ability is further improved.

- We conduct comprehensive experiments on several benchmark datasets, and the results demonstrate that our approach significantly outperforms state-of-the-art methods by a considerable margin.

## 2. Related Work

**3D Human Pose Estimation.** The widely adopted two-stage technique in 3D HPE, as demonstrated in [23, 36, 38, 40, 43, 46], initially employs 2D human pose estimators [17, 26, 27, 33] for 2D pose predictions, followed by lifting these predictions to 3D poses. Among these approaches, [38] integrates graph-structured semantic information to enhance the estimation process, while [23] utilizes dilated temporal convolutional layers for temporal information encoding, and [43] presents a purely transformer-based 3D approach. Moreover, [36] effectively models inter-frame correspondences with a mixed sequence-to-sequence encoder, and recent works such as [40] explore the frequency domain to improve inference efficiency, and [46] employs unified representation learning for 3D human poses.

**Domain Generalization.** Current DG methods aim to learn domain-invariant representations and are categorized into three types: domain alignment [21, 41], meta-learning [15, 28], and augmentation strategies [24, 25, 39, 42]. For domain alignment, [41] enhances the conditional invariance of learned features by incorporating an entropy regularization term, leading to improved classifier generalization. [21] iteratively segregates samples into latent domains through clustering. Concerning meta-learning, [15] proposes a model-agnostic training procedure that simulates domain shift during training, whereas [28] applies meta-learning to single-domain generalization. Regarding augmentation strategies, [39] introduces a novel regularization term for adversarial data augmentation derived from the information bottleneck principle, while [42] presents a unique style hallucination module to generate style-diversified samples crucial for generalization.

Differing from the current DG approaches for 3D HPE that focus solely on augmentations, we also incorporate meta-learning-based approaches to enhance generalization.

**Cross-domain Learning for 3D Human Pose Estimation.** Cross-domain learning for 3D HPE is categorized into two types: *domain generalization* [7, 11, 16, 37] and *domain adaptation* [2, 4, 18]. In domain generalization, training processes exclusively utilizes source data, and the resulting model is directly applied to infer target data. [37] adjusts various geometry factors of human poses through differentiable operations. [11] applies DH Forward Kinematics [1] to drive 3D pose augmentation and obtain diverse poses. [7] incorporates Counterfactual Risk Minimization [30] to achieve unbiased learning. [16] addresses with network designs like the interpolation sub-net and body-parts grouping

net. In domain adaptation, labeled source data and unlabeled target data are used simultaneously during the training process. [4] employs generative adversarial network [6] to discriminate between source and target during training. [2] utilizes global position alignment and local pose augmentation to transfer from source to target. [18] employs a multi-hypotheses network along with target-specific source augmentation for the problem.

In this paper, we focus on domain generalization for 3D HPE. In addition to synthesizing novel 3D poses for better generalization, we also simulate domain shifts by using both source and synthesized poses during optimizations.

### 3. Methodology

#### 3.1. Preliminaries

**2D-to-3D lifting Paradigm of 3D HPE.** Current 2D-to-3D lifting paradigm of 3D HPE [23, 36, 38, 43] assumes that  $x_i^s \in \mathbb{R}^{J \times 2}$  represents the 2D coordinates of  $J$  keypoints of a sample in the source domain (2D poses as input), and  $y_i^s \in \mathbb{R}^{J \times 3}$  represents the corresponding 3D positions in the camera coordinate system (3D poses as output), we denote the source domain with  $N$  samples as  $S = \{(x_i^s, y_i^s)\}_{i=1}^N$ , encompassing  $N$  2D-3D pairs. Moreover, we define the pose estimator as  $\mathcal{P} : x_i^s \mapsto \hat{y}_i^s$ , where  $\hat{y}_i^s$  represents the predicted corresponding 3D pose positions. For a fully supervised human pose estimation problem, we aim to achieve an ideal  $\mathcal{P}$  by solving the following optimization objective:

$$\min_{\mathcal{P}} \mathbb{E}_{(x_i^s, y_i^s) \in S} \mathcal{L}_{MSE}(\mathcal{P}(x_i^s), y_i^s), \quad (1)$$

where  $\mathcal{L}_{MSE}$  represents the Mean Squared Error (MSE) loss. However, the objective [23, 43] is designed to achieve optimal performance on source poses, rendering it inadequate for addressing the DG problem, as it does not account for the domain gap between source and target domains.

**DG for 3D HPE.** Within the paradigm of 2D-to-3D lifting HPE, our primary goal is to derive an estimator  $\mathcal{P}$  that demonstrates commendable 3D HPE performance specifically within the target domain  $T$ . Under this scenario, the target domain  $T = \{(x_j^t, y_j^t)\}_{j=1}^M$  with  $M$  samples can only be used for inference and is not involved in the training process. However, when utilizing solely the original source domain, the pose estimator cannot learn out-of-source distributions, which is essential for achieving good performance on the target domain. Existing methods [7, 11, 16, 37] tend to conduct augmentation to the original source poses to enhance data diversity, thereby improving the estimator’s generalization ability. The augmentor is denoted as  $\mathcal{A} : y_i^s \mapsto y_i^a$ , while the projection from 3D to 2D via camera parameters (completely known) is defined as  $\mathcal{R} : y_i^a \mapsto x_i^a$ . Consequently, the min-max optimization objective for domain

generalization can be defined as follows:

$$\min_{\mathcal{P}} \max_{\mathcal{A}} \mathbb{E}_{(x_i^s, y_i^s) \in S} [\mathcal{L}_{MSE}(\mathcal{P}(x_i^s), y_i^s) + \mathcal{L}_{MSE}(\mathcal{P}(\mathcal{R}(\mathcal{A}(y_i^s))), \mathcal{A}(y_i^s))]. \quad (2)$$

The objective is a **min-max game** between the pose augmentor  $\mathcal{A}$  and the pose estimator  $\mathcal{P}$ , encouraging the estimator  $\mathcal{P}$  to learn out-of-source distributions, while conducting augmentations to a significant extent is beneficial to generate more diverse samples, and that is why the loss is minimized with respect to the augmentor  $\mathcal{P}$  and maximized with respect to the augmentor  $\mathcal{A}$  in the optimization.

#### 3.2. Overview of the Proposed Method

Existing methods [7, 11, 37] often apply intricate augmentations to the original poses in the source domain, relying on the discrimination between augmented poses and source poses simultaneously. However, this kind of approach raises two concerns. *First*, as this is a DG problem for 3D HPE, any information about the target domain is entirely unknown. If the target domain bears a striking resemblance to the source domain (like Target Domain 1 in Fig. 1), poses generated by extensive augmentation might hinder the pose estimator’s inference on it. Conversely, in cases where the target domain significantly diverges from the source distributions (like Target Domain 3 in Fig. 1), poses generated by insufficient augmentation may fail to adequately explore out-of-source knowledge for the pose estimator. *Second*, when target domain is distant from the source and needs significant augmentations, the adversarial training between source and synthesized poses limits the diversity of generated poses. Specifically, the discriminator enforces similarity between source and synthesized poses, thereby causing the synthesized poses to remain similar to the source poses.

To tackle these concerns, we propose a novel dual-augmentor framework depicted in Fig. 2. **This framework involves two augmentors that generate weak- and strong-augmented poses**, enabling the handling of diverse unknown target domains. Additionally, the weak-augmentation module serves as a bridge between strong-augmented and source poses. Specifically, **the discrimination between source poses and weak-augmented poses is utilized to update the weak augmentor**, while **the discrimination between weak- and strong-augmented poses is employed to optimize the strong augmentor**. This approach liberates the strong augmentor from heavy reliance on the source domain and enables the exploration of more out-of-source knowledge. Further details regarding the pose augmentation process can be found in Sec. 3.3.

Having elucidated the methodology for synthesizing poses, the subsequent discourse pivots towards the utilization of these synthesized poses. Previous works [7, 11, 37] overlook the interactions between source poses and aug-

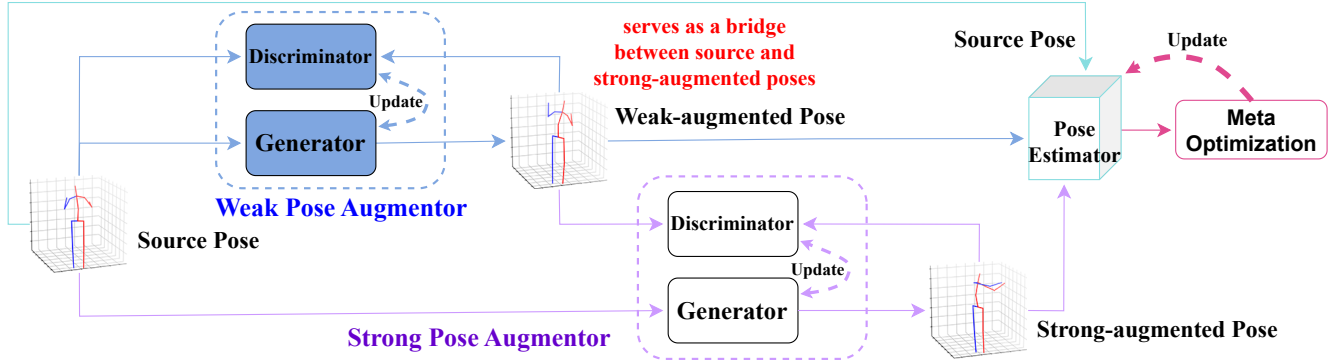


Figure 2. Overall framework of our dual-augmentor method. Initially, the original pose undergoes processing through two pose augmentors, resulting in weak- and strong-augmented poses (See Sec. 3.3). The weak augmentor simulates target domains similar to the source domain, while the strong augmentor emulates target domains that deviate significantly from the source distributions. Subsequently, the original pose and the two augmented poses are input to the pose estimator for further meta-optimization (See Sec. 3.4).

mented poses, dealing with the optimizations of pose estimator on them separately. In contrast, we propose a model-agnostic meta optimization approach that enhances the interactions between source poses and the two types of augmented poses to simulate domain shifts in the optimization and leverage domain-invariant knowledge while maintaining the original 2D-to-3D lifting backbone’s structure unchanged. Further details concerning the meta optimization process can be found in Sec. 3.4.

### 3.3. Pose Augmentation

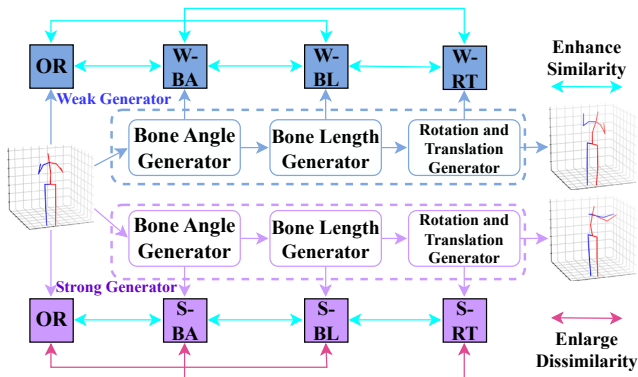


Figure 3. The differentiation of the weak and strong generators. Within each pipeline, denoted as “W-” for weak ones and “S-” for strong ones, there exist four pose states: original (OR), after bone angle operation (BA), after bone length operation (BL) and after rotation and translation operation (RT). For proximate states, similarities are enhanced for both generators. When there is a one-state gap between states, the weak generator continues to enhance similarities, whereas the strong generator enlarges dissimilarities.

The pose augmentation architecture comprises two pose augmentors: the weak augmentor and the strong augmentor. Each augmentor comprises two components: the generator, responsible for producing diverse synthesized poses to facilitate the training of the pose estimator, and the discriminator, which collaborates with the generator to regu-

late the quality of the generated poses. Our objective is to apply *differential strategies* to the generation (named **differential generation**) and discrimination (named **differential discrimination**) of the two augmentors, enabling them to generate weak- and strong-augmented poses.

**Differential Generation:** The generation process, as illustrated in Figure 3, consists of three modules in sequence for each generator pipeline: the Bone Angle Generator, Bone Length Generator, and Rotation and Translation Generator, resulting in four statuses in the pipeline: original (OR), after the bone angle operation (BA), after the bone length operation (BL), and after the rotation and translation operation (RT). Existing approaches such as [16, 37] typically treat the entire generation pipeline in an end-to-end manner and only utilize the  $(OR, RT)$  pair. In contrast, our method deals with the generation in a more fine-grained fashion. We group statuses into pairs based on their relations: proximate pairs as  $PP = \{(OR, BA), (BA, BL), (BL, RT)\}$ , and one-state gap pairs as  $OG = \{(OR, BL), (BA, RT)\}$ .

To begin, we define the measurement of similarity within a pair. Solely relying on conventional absolute position losses, such as the MSE loss, is not adequate in this context for two reasons. First, the three modules within the generator all perform operations on the level of bone vector, not on the joint positions. If one joint undergoes significant position changes after an operation, other joints connected to it will also experience considerable movement, even if the bone vector between them remains stable. In such cases, position-based measurements cannot fully reflect the extent of augmentation based on the bone vector. Second, human poses possess kinematic attributes, and a position-based measurement overlooks the graphical information. Therefore, we introduce the Laplacian weighted similarity measurement. For the human model, it is straightforward to obtain degree matrix  $D$  and adjacency matrix  $A$ , and the normalized Laplacian matrix can be represented as:



$$W_{NL} = I - D^{-\frac{1}{2}} A D^{-\frac{1}{2}}, \quad (3)$$

where  $I$  is the identity matrix, and  $W_{NL}$  is the normalized Laplacian matrix encoding graphical information. Given a pair of statuses  $(st_1, st_2)$  (from either PP or OG), the similarity measurement is defined as:

$$\mathcal{L}_{sim}(st_1, st_2) = \underbrace{\|st_1 - st_2\|}_{\text{MSE Similarity}} + \underbrace{\|W_{NL}st_1 - W_{NL}st_2\|}_{\text{Laplacian Weighted Similarity}}. \quad (4)$$

To differentiate between the two generators, we apply distinctive strategies. For the weak generator, we enhance similarities for its PP and OG sets to maintain a *slight* level of augmentation in the synthesized poses, as indicated by:

$$\begin{aligned} \mathcal{L}_{wg} = & \mathbb{E}_{(st_1, st_2) \in PP} \mathcal{L}_{sim}(st_1, st_2) \\ & + \alpha_1 \mathbb{E}_{(st_1, st_2) \in OG} \mathcal{L}_{sim}(st_1, st_2). \end{aligned} \quad (5)$$

For the strong generator, we enhance similarities within its PP set to ensure the reasonableness of the synthesized output, while enlarging dissimilarities within its OG sets to maintain a *significant* level of augmentation, expressed as:

$$\begin{aligned} \mathcal{L}_{sg} = & \mathbb{E}_{(st_1, st_2) \in PP} \mathcal{L}_{sim}(st_1, st_2) \\ & - \alpha_2 \mathbb{E}_{(st_1, st_2) \in OG} \mathcal{L}_{sim}(st_1, st_2). \end{aligned} \quad (6)$$

$\alpha_1$  and  $\alpha_2$  are trade-off parameters.

**Differential Discrimination:** The discrimination process comprises two min-max games [6, 8]: one between the source pose and the weak-augmented poses, and the other between the weak-augmented pose and the strong-augmented pose. We adopt the WGAN-GP [8] structure here. The discrimination losses regarding the source poses  $y^{or}$ , weak-augmented poses  $y^{wa}$ , and strong-augmented poses  $y^{sa}$  are defined as follows:

$$\begin{aligned} \mathcal{L}_{wd} = & \mathbb{E}[D_{wa}(y^{or})] - \mathbb{E}[D_{wa}(y^{wa})] \\ & + \beta_1 \mathbb{E}(1 - \|\nabla_{\hat{y}^{wa}} D_{wa}(\hat{y}^{wa})\|), \end{aligned} \quad (7)$$

$$\begin{aligned} \mathcal{L}_{sd} = & \mathbb{E}[D_{sa}(y^{wa})] - \mathbb{E}[D_{sa}(y^{sa})] \\ & + \beta_2 \mathbb{E}(1 - \|\nabla_{\hat{y}^{sa}} D_{sa}(\hat{y}^{sa})\|). \end{aligned} \quad (8)$$

Here,  $\mathcal{L}_{wd}$  is the discrimination loss between  $y^{or}$  and  $y^{wa}$ , used to update the weak augmentor, and  $D_{wa}$  is the weak discriminator.  $\mathcal{L}_{sd}$  is the discrimination loss between  $y^{wa}$  and  $y^{sa}$ , used to update the strong augmentor, and  $D_{sa}$  is the strong discriminator.  $\hat{y}^{wa}$  and  $\hat{y}^{sa}$  are built via interpolation, such that  $\hat{y}^{wa} = \epsilon y^{or} + (1 - \epsilon)y^{wa}$  and  $\hat{y}^{sa} = \epsilon y^{wa} + (1 - \epsilon)y^{sa}$ , where  $\epsilon$  is randomly drawn from  $U[0, 1]$ .  $\beta_1$  and  $\beta_2$  are trade-off parameters.

By implementing this discrimination process in two min-max games, the weak augmentor is capable of retaining more source information and alleviating adverse effects

stemming from irrationally synthesized poses. Simultaneously, the strong augmentor can overcome a strong dependency on the source distributions, and explore out-of-source distributions more effectively. With diverse synthesized poses to simulate potential target poses, it is beneficial for further domain generalization in pose estimation.

### 3.4. Meta Optimization

---

#### Algorithm 1 Meta Optimization Pseudo Code

---

**Input:** Original source 2D-3D pose pairs  $(x^{or}, y^{or})$ ; Weak Generator  $G_{wa}$ ; Strong Generator  $G_{sa}$

**Init:** Pose estimator  $\mathcal{P}_t$ , Learning rates  $lr_1$  and  $lr_2$ , inner loop iteration  $k$ , Hyperparameter  $\gamma$

**Output:** Updated pose estimator  $\mathcal{P}_{t+2}$  after two-step meta optimization

- 1: Generate weak-augmented 3D poses  $y^{wa} = G_{wa}(y^{or})$
  - 2: Project  $y^{wa}$  to 2D poses  $x^{wa}$  with camera parameters
  - 3: // Meta-train on Source data:
  - 4: Update  $\mathcal{P}'_t = \mathcal{P}_t - lr_1 \nabla \mathcal{L}_{MSE}(\mathcal{P}_t(x^{or}), y^{or})$
  - 5: **for**  $i \leftarrow 1, \dots, k$  **do**
  - 6: // Meta-test on Weak-augmented data:
  - 7:  $\mathcal{L}_{weak-test} = L_{MSE}(\mathcal{P}'_t(x^{wa}), y^{wa})$
  - 8: **end for**
  - 9: // Meta update on Source and Weak-augmented data:
  - 10:  $\mathcal{L}_{t+1} = \mathcal{L}_{MSE}(\mathcal{P}_t(x^{or}), y^{or}) + \gamma L_{weak-test}$
  - 11:  $\mathcal{P}_{t+1} = \mathcal{P}_t - lr_2 (\partial \mathcal{L}_{t+1} / \partial \mathcal{P}_t)$
  - 12: Generate strong-augmented 3D poses  $y^{sa} = G_{sa}(y^{or})$
  - 13: Project  $y^{sa}$  to 2D poses  $x^{sa}$  with camera parameters
  - 14: // Meta-train on Weak-augmented data:
  - 15: Update  $\mathcal{P}'_{t+1} = \mathcal{P}_{t+1} - lr_1 \nabla \mathcal{L}_{MSE}(\mathcal{P}_{t+1}(x^{wa}), y^{wa})$
  - 16: **for**  $i \leftarrow 1, \dots, k$  **do**
  - 17: // Meta-test on Strong-augmented data:
  - 18:  $\mathcal{L}_{strong-test} = L_{MSE}(\mathcal{P}'_{t+1}(x^{sa}), y^{sa})$
  - 19: **end for**
  - 20: // Meta update on Weak- and Strong-augmented data:
  - 21:  $\mathcal{L}_{t+2} = \mathcal{L}_{MSE}(\mathcal{P}_{t+1}(x^{wa}), y^{wa}) + \gamma L_{strong-test}$
  - 22:  $\mathcal{P}_{t+2} = \mathcal{P}_{t+1} - lr_2 (\partial \mathcal{L}_{t+2} / \partial \mathcal{P}_{t+1})$
- 

For DG problem in 3D HPE, two principal challenges must be addressed. First, there is the issue of synthesizing data, as detailed in Section 3.3. The second challenge revolves around the effective utilization of synthesized data, a facet often overlooked by current methodologies. Existing DG approaches for 3D HPE [11, 37], conduct the optimization of the pose estimator based on source and synthesized data independently. Unfortunately, this approach lacks mechanisms for fostering interactions between these two optimization processes, resulting in a deficiency of simulated domain shifts in the optimization trajectory.

In contrast, our proposed model-agnostic strategy incorporates meta-optimization to bolster interactions among source poses, weak-augmented poses, and strong-augmented poses. This process facilitates the learning of

domain-invariant knowledge during the update of the pose estimator, as outlined in Algorithm 1. The effectiveness of this approach lies in the fact that the objectives in meta-optimization not only aim to minimize losses on source and synthesized poses but also enhance the alignment of optimization directions during training, thus enhancing generalization significantly.

The algorithm can be dissected into two parts: Lines 1-11 manage interactions between source poses and weakly-augmented poses, while Lines 12-22 address interactions between weak- and strong-augmented poses. This step-by-step approach is taken due to the substantial domain gap between source poses and strong-augmented poses. Weak-augmented poses serve as a bridge between source poses and strong-augmented poses, alleviating the challenge of directly aligning source and strong-augmented data. By incorporating all three types of poses in the optimization, the pose estimator can effectively generalize across diverse target domains, avoiding overfitting specific pose data types.

## 4. Experiments

### 4.1. Datasets and Metrics

In this paper, we evaluate our approach using several widely-used 3D human pose benchmarks, including Human3.6M [12], MPI-INF-3DHP [22], and 3DPW [31]. Moreover, following previous works [11, 16, 37], we adopt the 16-keypoint human model with Hip joint as the origin.

**Human3.6M** is an indoor dataset comprising 3.6 million frames and consisting of 7 subjects denoted as S1, S5, S6, S7, S8, S9, and S11. For the cross-dataset evaluation in Tab. 1, Tab. 2 and Tab. 5, we use S1, S5, S6, S7, S8 as the source domain. In the cross-scenario evaluation on the entire Human3.6M dataset in Tab. 3, S1, S5, S6, S7, S8 are the source domain, while S9 and S11 are the target domain. Mean Per Joint Position Error (MPJPE) and Procrustes-Aligned Mean Per Joint Position Error (PA-MPJPE) are employed as evaluation metrics. For the cross-scenario evaluation on partial Human3.6M, we follow previous works [4, 11, 16] to define two tasks as shown in Tab. 4. One task uses S1 as the source and S5, S6, S7, S8 as the target, while the other task uses S1, S5 as the source and S6, S7, S8 as the target. Both tasks utilize MPJPE as the metric.

**MPI-INF-3DHP (3DHP)** is an in-the-wild dataset, and we utilize only its test set for cross-dataset evaluation, as shown in Tab. 1 and Tab. 5, which consists of approximately 3k frames. The results are presented based on three metrics: Percentage of Correct Keypoints (PCK), Area Under the Curve (AUC), and MPJPE.

**3DPW** is another in-the-wild dataset featuring more challenging poses and scenes. We utilize it for cross-dataset evaluation, as shown in Tab. 2. Here PA-MPJPE and MPJPE serve as the evaluation metrics.

### 4.2. Implementation Details

For all the generators and discriminators, we ensure consistency by employing the same fully-connected layers, aligning with the methodology described in [37]. In the *data augmentation process*, the learning rate is set to 1e-4 for generators and 2e-4 for discriminators. We set  $\alpha_1$  and  $\alpha_2$  to 0.50 and 0.35, respectively, while both  $\beta_1$  and  $\beta_2$  are assigned a value of 4. During the *meta optimization process*, we utilize a learning rate of 1e-4 for  $lr_1$  and 5e-4 for  $lr_2$ . The trade-off parameter  $\gamma$  and the inner loop iteration  $k$  are both set to 1.

Moreover, we employ the Adam optimizer [13] for data augmentation and the AdamW optimizer [19] for meta optimization. Our experiments are conducted with a batch size of 1024 over 60 epochs. We initialize the pose estimator with a warm-up phase lasting two epochs for supervised learning on source data. From the third epoch onwards, data augmentation and meta-optimization begin.

### 4.3. Quantitative Results

Table 1. Cross-dataset evaluation on 3DHP dataset.

Method	Venue	DG	PCK $\uparrow$	AUC $\uparrow$	MPJPE $\downarrow$
VPose (1-frame) [23]	CVPR'19	$\times$	80.9	42.5	102.3
EvoSkeleton [17]	CVPR'20	$\checkmark$	81.2	46.1	99.7
RepNet [32]	CVPR'19	$\checkmark$	81.8	54.8	92.5
PoseAug [37]	TPAMI'23	$\checkmark$	88.6	57.3	73.0
DH-AUG [11]	ECCV'22	$\checkmark$	89.5	57.9	71.2
PoseGU [7]	CVIU'23	$\checkmark$	86.3	57.2	75.0
CEE-Net [16]	AAAI'23	$\checkmark$	89.9	58.2	69.7
Ours		$\checkmark$	<b>92.9</b>	<b>60.7</b>	<b>63.1</b>

Table 2. Cross-dataset evaluation on 3DPW dataset.

Method	Venue	DG	PA-MPJPE $\downarrow$	MPJPE $\downarrow$
VPose (1-frame) [23]	CVPR'19	$\times$	94.6	125.7
VIBE [14]	CVPR'20	$\checkmark$	82.3	122.5
PoseAug [37]	TPAMI'23	$\checkmark$	81.6	119.0
DH-AUG [11]	ECCV'22	$\checkmark$	79.3	112.8
PoseGU [7]	CVIU'23	$\checkmark$	92.3	-
CEE-Net [16]	AAAI'23	$\checkmark$	76.8	-
Ours		$\checkmark$	<b>73.2</b>	<b>106.6</b>

Table 3. Cross-scenario evaluation on Entire Human3.6M dataset. S1,S5,S6,S7,S8 are the source and S9,S11 are the target.

Method	Venue	DG	MPJPE $\downarrow$	PA-MPJPE $\downarrow$
VPose (1-frame) [23]	CVPR'19	$\times$	52.7	40.9
EvoSkeleton [17]	CVPR'20	$\checkmark$	50.9	38.0
PoseAug [37]	TPAMI'23	$\checkmark$	50.2	39.1
DH-AUG [11]	ECCV'22	$\checkmark$	49.8	38.3
CEE-Net [16]	AAAI'23	$\checkmark$	47.3	36.8
Ours		$\checkmark$	<b>44.4</b>	<b>34.6</b>

**Cross-dataset evaluation results.** In cross-dataset evaluations, source and target come from different datasets. Following identical paradigms from existing methods [11, 16, 37], we use ground truth 2D keypoints as input, single-frame VPose [23] as the lifting backbone, and Human3.6M as the source dataset. Our method demonstrates notable performance in all metrics, as presented in Tab. 1 and Tab. 2. Notably, our approach outperforms CEE-Net by 3.0% in




Figure 4. Qualitative results on Cross-dataset evaluation. Left is 3DHP dataset, and right is 3DPW dataset.




Figure 5. Results on Cross-scenario evaluation. Left is for task S1,S5,S6,S7,S8 → S9,S11, and right is for task S1,S5 → S6,S7,S8.

Table 4. Cross-scenario evaluation on Partial Human3.6M dataset. For the task "S1", S1 is the source and S5, S6, S7, S8 are the target. For the task "S1+S5", S1 and S5 are the source, and S6, S7, S8 are the target. MPJPE ( $\downarrow$ ) is used for evaluation.

Method	Venue	DG	S1	S1+S5
VPose (1-frame) [23]	CVPR'19	×	65.2	57.9
EvoSkeleton [17]	CVPR'20	✓	61.5	54.6
PoseAug [37]	TPAMI'23	✓	56.7	51.3
DH-AUG [11]	ECCV'22	✓	52.2	47.0
CEE-Net [16]	AAAI'23	✓	51.9	46.7
Ours		✓	<b>50.3</b>	<b>45.4</b>

Table 5. Cross-dataset evaluation with MPJPE ( $\downarrow$ ) on 3DHP with varied 2D predictions and 2D-to-3D backbones (1-frame).

Method	DG	DET [5]	CPN [3]	HR [33]	GT
SemGCN [38]	×	101.9	98.7	95.6	97.4
SemGCN + PoseAug [37]	✓	89.9	89.3	89.1	86.1
SemGCN + CEE-generator [16]	✓	83.6	82.8	82.4	81.3
SemGCN + DH-AUG [11]	✓	79.7	76.7	73.0	71.3
SemGCN + Ours	✓	<b>76.5</b>	<b>74.1</b>	<b>70.7</b>	<b>68.9</b>
VPose [23]	×	92.6	89.8	85.6	86.6
VPose + PoseAug [37]	✓	78.3	78.4	73.2	73.0
VPose + CEE-generator [16]	✓	75.6	75.2	71.2	71.4
VPose + DH-AUG [11]	✓	76.7	74.8	71.1	71.2
VPose + Ours	✓	<b>72.4</b>	<b>70.9</b>	<b>62.4</b>	<b>63.1</b>
PoseFormer [43]	×	91.9	89.2	84.2	85.7
PoseFormer + PoseAug [37]	✓	77.7	77.5	72.1	72.3
PoseFormer + CEE-generator [16]	✓	-	-	-	-
PoseFormer + DH-AUG [11]	✓	75.6	74.8	71.6	72.0
PoseFormer + Ours	✓	<b>72.2</b>	<b>70.5</b>	<b>62.8</b>	<b>63.4</b>
MixSTE [36]	×	90.6	87.4	82.0	84.0
MixSTE + PoseAug [37]	✓	76.1	76.3	71.7	71.6
MixSTE + CEE-generator [16]	✓	-	-	-	-
MixSTE + DH-AUG [11]	✓	74.8	74.4	70.9	70.7
MixSTE + Ours	✓	<b>70.5</b>	<b>68.2</b>	<b>60.4</b>	<b>61.0</b>

PCK and 2.5% in AUC, and reduces MPJPE by 6.6mm in the 3DHP task. In the case of 3DPW, our method shows an improvement of 3.6mm in PA-MPJPE compared to CEE-Net [16]. While CEE-Net [16] and PoseGU [7] do not disclose their codes or report their results on MPJPE, it is evident that our method surpasses DH-AUG [11] by 6.2mm.

**Cross-scenario evaluation results.** In cross-scenario evaluations, source and target come from different subsets of the same dataset. Maintaining consistency with previous works, we utilize ground truth 2D keypoints as input and single-frame VPose [23] as the 2D-to-3D lifting network.

For the situation of using Entire Human3.6M in Tab. 3, our method demonstrates superior performance compared to CEE-Net [16] with a 2.9mm reduction in MPJPE and a 2.2mm improvement in PA-MPJPE. In the case of using partial Human3.6M in Tab. 4, our approach surpasses CEE-Net [16] by 1.6mm in the S1 task and 1.3mm in the S1+S5 task based on the MPJPE metric.

**Results with varied 2D predictions and backbones.** The results presented in Tables 1 to 4 are confined to the usage of ground truth 2D input and single-frame VPose backbone, which may raise concerns about the universality of the proposed method. To address this concern, we assess the performance of our approach with various 2D predictions such as DET [5], CPN [3], HRNet [33], and diverse lifting backbones including SemGCN [38], PoseFormer [43], and MixSTE [36], as displayed in Table 5. In this evaluation, 3DHP serves as the dataset, and MPJPE is the metric used. Notably, all the listed backbones are single-frame versions. As CEE-Net [16] only provides results for its generation part, CEE-generator, and does not offer open-source code, we have included partial results of CEE-generator. From Table 5, it is evident that our method surpasses all the existing methods, demonstrating the robustness of our framework across various settings.

#### 4.4. Qualitative Results

Fig. 4 shows qualitative results on cross-dataset evaluation (3DHP on the left side and 3DPW on the right side), while Fig. 5 displays qualitative results on cross-scenario evaluation (S1,S5,S6,S7,S8 → S9,S11 on the left side and S1,S5 → S6,S7,S8 on the right side). HRNet [33] is applied as the 2D pose estimator and VPose [23] is the 2D-to-3D lifting backbone. We use **Source-only**, **PoseAug** [37], **DH-AUG** [11], **Ours**, and **Ground Truth (GT)** for qualitative comparison. Because CEE-Net does not provide source codes or pretrained models, we cannot generate visual examples from it. It is evident that our method outperforms other baselines significantly.

## 4.5. Ablation Study

**Ablation study on the overall framework.** The ablation study is conducted to illustrate the functions of three proposed modules: differential generation (**DiffGen**) in Sec. 3.3, differential discrimination (**DiffDis**) in Sec. 3.3, and meta optimization (**MetaOpt**) in Sec. 3.4. In Tab. 6, we apply ground truth as 2D predictions and VPose as the backbone. The absence of DiffGen leads to a decrease in PCK and AUC by 2.4% and 2.1% respectively, accompanied by a 5.9mm increase in MPJPE on 3DHP, while it increases PA-MPJPE and MPJPE by 3.1mm and 5.2mm separately on 3DPW. Similarly, the exclusion of DiffDis results in a decrease of 1.7% in both PCK and AUC, with a corresponding 4.1mm increase in MPJPE on 3DHP. As for 3DPW, the removal causes a degradation of 1.9mm in PA-MPJPE and 3.0mm in MPJPE. Besides, the removal of MetaOpt leads to a decline in PCK and AUC by 1.2% and 0.8% respectively, along with a 2.4mm increase in MPJPE on 3DHP, and an increase of 1.3mm and 1.8mm in PA-MPJPE and MPJPE separately. These results show that each module plays a critical role in obtaining better generalization.

Table 6. Overall framework ablation study on 3DHP and 3DPW

Method	3DHP			3DPW	
	PCK $\uparrow$	AUC $\uparrow$	MPJPE $\downarrow$	PA-MPJPE $\downarrow$	MPJPE $\downarrow$
Ours w/o DiffGen	90.5	58.6	69.0	76.3	111.8
Ours w/o DiffDis	91.2	59.0	67.3	75.1	109.6
Ours w/o MetaOpt	91.7	59.9	65.5	74.5	108.4
Ours	<b>92.9</b>	<b>60.7</b>	<b>63.1</b>	<b>73.2</b>	<b>106.6</b>

**Ablation study on the generators.** There exist two generators in our framework, and each with two pair groups. In this part, we discuss the functions of proximate pairs in weak augmentor (**W-PP**), proximate pairs in strong augmentor (**S-PP**), one-state gap pairs in weak augmentor (**W-OG**), and one-state gap pairs in strong augmentor (**S-OG**).

Table 7. Generators' ablation study on 3DHP and 3DPW

Method	3DHP			3DPW	
	PCK $\uparrow$	AUC $\uparrow$	MPJPE $\downarrow$	PA-MPJPE $\downarrow$	MPJPE $\downarrow$
Ours w/o W-PP	88.3	57.5	72.6	81.7	118.8
Ours w/o S-PP	90.8	58.2	71.3	78.1	111.0
Ours w/o W-OG	92.1	59.6	65.8	74.7	108.7
Ours w/o S-OG	91.4	58.9	68.2	75.4	109.5
Ours	<b>92.9</b>	<b>60.7</b>	<b>63.1</b>	<b>73.2</b>	<b>106.6</b>

In Table 7, excluding W-PP or S-PP leads to a significant decline in PCK by 4.6% and 2.1%, and in AUC by 3.2% and 2.5% respectively, accompanied by a notable increase of 9.5mm and 8.2mm in MPJPE separately on 3DHP. These results emphasize the critical role of maintaining similarity in proximate pairs for both weak and strong augmentors, serving as the fundamental basis for generating effective and reasonable synthesized poses. Moreover, the absence of W-OG leads to a decline in PCK and AUC by 0.8% and 1.1% respectively, with a corresponding 1.7mm increase in the MPJPE on 3DHP. The removal of S-OG results in a decrease in PCK and AUC scores by 1.5% and 1.8% respec-

tively, along with a 3.4mm increase in MPJPE on 3DHP. These results highlight the significance of maintaining differentiation between the weak augmentor and the strong augmentor during the generation process, where enlarging dissimilarity in S-OG is more important in discriminating these two generators.

**Ablation study on the number of augmentors.** Comparisons were conducted between our dual-augmentor framework and single-augmentor frameworks. Alongside our proposed framework, two single-augmentor frameworks were considered in the ablation study, utilizing either the weak-augmentor (**WA**) or the strong-augmentor (**SA**). The discrimination and meta-optimization processes exclusively involved source poses and one category of synthesized poses. The results, using ground truth as 2D predictions and VPose as the backbone, are presented in Table 8.

Table 8. Ablation study of number of augmentors on 3DHP and 3DPW

Method	3DHP			3DPW	
	PCK $\uparrow$	AUC $\uparrow$	MPJPE $\downarrow$	PA-MPJPE $\downarrow$	MPJPE $\downarrow$
WA	87.3	56.0	74.5	80.5	117.7
SA	89.8	57.8	71.0	79.1	111.4
Ours	<b>92.9</b>	<b>60.7</b>	<b>63.1</b>	<b>73.2</b>	<b>106.6</b>

From Table 8, it is evident that our proposed framework surpasses both WA and SA significantly, underscoring the superiority of employing two augmentors over a single augmentor in addressing DG for 3D HPE. Furthermore, SA outperforms WA, emphasizing the greater significance of exploring out-of-source distributions compared to retaining source-relevant knowledge in cross-dataset tasks.

## 5. Conclusion

In this paper, we propose a novel dual-augmentor framework designed to enhance domain generalization in 3D human pose estimation. Our framework addresses the critical aspects of data augmentation and the effective utilization of synthesized data. To achieve this, we implement distinctive strategies for the weak and strong generators, ensuring the preservation of source-specific information while simultaneously exploring out-of-source distributions. Moreover, we incorporate meta-optimization techniques to facilitate enhanced interaction among source data, weak-augmented data, and strong-augmented data, thereby simulating domain shifts in the training of pose estimator and fostering the acquisition of domain-invariant knowledge. Extensive experimentation and comprehensive analysis conducted across multiple datasets demonstrate the superior performance of our proposed approach over existing state-of-the-art methods.

**Acknowledgements** This material is based upon work supported by the National Science Foundation under Grant CNS-1910844.



## References

- [1] Maher Baili, Philippe Wenger, and Damien Chablat. A classification of 3r orthogonal manipulators by the topology of their workspace. In *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004*, pages 1933–1938. IEEE, 2004. **2**
- [2] Wenhao Chai, Zhongyu Jiang, Jenq-Neng Hwang, and Gaoang Wang. Global adaptation meets local generalization: Unsupervised domain adaptation for 3d human pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14655–14665, 2023. **2, 3**
- [3] Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, and Jian Sun. Cascaded pyramid network for multi-person pose estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7103–7112, 2018. **7**
- [4] 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*, pages 13075–13085, 2022. **2, 3, 6**
- [5] Ross Girshick, Ilija Radosavovic, Georgia Gkioxari, Piotr Dollár, and Kaiming He. Detectron. <https://github.com/facebookresearch/detectron>, 2018. **7**
- [6] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. **2, 3, 5**
- [7] Shannan Guan, Haiyan Lu, Linchao Zhu, and Gengfa Fang. Posegu: 3d human pose estimation with novel human pose generator and unbiased learning. *Computer Vision and Image Understanding*, 233:103715, 2023. **1, 2, 3, 6, 7**
- [8] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. *Advances in neural information processing systems*, 30, 2017. **5**
- [9] Vladimir Guzov, Aymen Mir, Torsten Sattler, and Gerard Pons-Moll. Human positioning system (hps): 3d human pose estimation and self-localization in large scenes from body-mounted sensors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4318–4329, 2021. **1**
- [10] Yilei Hua, Wenhan Wu, Ce Zheng, Aidong Lu, Mengyuan Liu, Chen Chen, and Shiqian Wu. Part aware contrastive learning for self-supervised action recognition. In *International Joint Conference on Artificial Intelligence*, 2023. **1**
- [11] Linzhi Huang, Jiahao Liang, and Weihong Deng. Dh-aug: Dh forward kinematics model driven augmentation for 3d human pose estimation. In *European Conference on Computer Vision*, pages 436–453. Springer, 2022. **1, 2, 3, 5, 6, 7**
- [12] 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, 2013. **6**
- [13] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. **6**
- [14] Muhammed Kocabas, Nikos Athanasiou, and Michael J Black. Vibe: Video inference for human body pose and shape estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5253–5263, 2020. **6**
- [15] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy Hospedales. Learning to generalize: Meta-learning for domain generalization. In *Proceedings of the AAAI conference on artificial intelligence*, 2018. **2**
- [16] Haolun Li and Chi-Man Pun. Cee-net: Complementary end-to-end network for 3d human pose generation and estimation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(1):1305–1313, 2023. **2, 3, 4, 6, 7**
- [17] Shichao Li, Lei Ke, Kevin Pratama, Yu-Wing Tai, Chi-Keung Tang, and Kwang-Ting Cheng. Cascaded deep monocular 3d human pose estimation with evolutionary training data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6173–6183, 2020. **2, 6, 7**
- [18] Hanbing Liu, Jun-Yan He, Zhi-Qi Cheng, Wangmeng Xiang, Qize Yang, Wenhao Chai, Gaoang Wang, Xu Bao, Bin Luo, Yifeng Geng, et al. Posynda: Multi-hypothesis pose synthesis domain adaptation for robust 3d human pose estimation. In *Proceedings of the ACM International Conference on Multimedia*, 2023. **2, 3**
- [19] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018. **6**
- [20] Zhengzhi Lu, He Wang, Ziyi Chang, Guoan Yang, and Hubert P. H. Shum. Hard no-box adversarial attack on skeleton-based human action recognition with skeleton-motion-informed gradient. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4597–4606, 2023. **1**
- [21] Toshihiko Matsuura and Tatsuya Harada. Domain generalization using a mixture of multiple latent domains. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 11749–11756, 2020. **2**
- [22] 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. **6**
- [23] Dario Pavlo, Christoph Feichtenhofer, David Grangier, and Michael Auli. 3d human pose estimation in video with temporal convolutions and semi-supervised training. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7753–7762, 2019. **2, 3, 6, 7**
- [24] Qucheng Peng. *Multi-source and Source-Private Cross-Domain Learning for Visual Recognition*. PhD thesis, Purdue University, 2022. **2**

- [25] Qucheng Peng, Zhengming Ding, Lingjuan Lyu, Lichao Sun, and Chen Chen. Rain: regularization on input and network for black-box domain adaptation. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pages 4118–4126, 2023. 2
- [26] Qucheng Peng, Ce Zheng, and Chen Chen. Source-free domain adaptive human pose estimation. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4803–4813. IEEE, 2023. 2
- [27] Ekkasit Pinyoanuntapong, Ayman Ali, Kalvik Jakkala, Pu Wang, Minwoo Lee, Qucheng Peng, Chen Chen, and Zhi Sun. Gaitsada: Self-aligned domain adaptation for mmwave gait recognition. In *2023 IEEE 20th International Conference on Mobile Ad Hoc and Smart Systems (MASS)*, pages 218–226. IEEE, 2023. 2
- [28] Fengchun Qiao, Long Zhao, and Xi Peng. Learning to learn single domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12556–12565, 2020. 2
- [29] Chi Su, Jianing Li, Shiliang Zhang, Junliang Xing, Wen Gao, and Qi Tian. Pose-driven deep convolutional model for person re-identification. In *Proceedings of the IEEE international conference on computer vision*, pages 3960–3969, 2017. 1
- [30] Adith Swaminathan and Thorsten Joachims. Counterfactual risk minimization: Learning from logged bandit feedback. In *International Conference on Machine Learning*, pages 814–823. PMLR, 2015. 2
- [31] Timo Von Marcard, Roberto Henschel, Michael J Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3d human pose in the wild using imus and a moving camera. In *Proceedings of the European conference on computer vision (ECCV)*, pages 601–617, 2018. 6
- [32] Bastian Wandt and Bodo Rosenhahn. Repnet: Weakly supervised training of an adversarial reprojection network for 3d human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7782–7791, 2019. 6
- [33] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 43(10):3349–3364, 2020. 2, 7
- [34] Hong Yan, Yang Liu, Yushen Wei, Zhen Li, Guanbin Li, and Liang Lin. Skeletonmae: Graph-based masked autoencoder for skeleton sequence pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5606–5618, 2023. 1
- [35] Hongwei Yi, Chun-Hao P Huang, Shashank Tripathi, Lea Hering, Justus Thies, and Michael J Black. Mime: Human-aware 3d scene generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12965–12976, 2023. 1
- [36] Jinlu Zhang, Zhigang Tu, Jianyu Yang, Yujin Chen, and Junsong Yuan. Mixste: Seq2seq mixed spatio-temporal encoder for 3d human pose estimation in video. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13232–13242, 2022. 2, 3, 7
- [37] Jianfeng Zhang, Kehong Gong, Xinchao Wang, and Jiashi Feng. Learning to augment poses for 3d human pose estimation in images and videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(8):10012–10026, 2023. 1, 2, 3, 4, 5, 6, 7
- [38] Long Zhao, Xi Peng, Yu Tian, Mubbasir Kapadia, and Dimitris N Metaxas. Semantic graph convolutional networks for 3d human pose regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3425–3435, 2019. 2, 3, 7
- [39] Long Zhao, Ting Liu, Xi Peng, and Dimitris Metaxas. Maximum-entropy adversarial data augmentation for improved generalization and robustness. *Advances in Neural Information Processing Systems*, 33:14435–14447, 2020. 2
- [40] Qitao Zhao, Ce Zheng, Mengyuan Liu, Pichao Wang, and Chen Chen. Poseformerv2: Exploring frequency domain for efficient and robust 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8877–8886, 2023. 2
- [41] Shanshan Zhao, Mingming Gong, Tongliang Liu, Huan Fu, and Dacheng Tao. Domain generalization via entropy regularization. *Advances in Neural Information Processing Systems*, 33:16096–16107, 2020. 2
- [42] Yuyang Zhao, Zhun Zhong, Na Zhao, Nicu Sebe, and Gim Hee Lee. Style-hallucinated dual consistency learning for domain generalized semantic segmentation. In *European Conference on Computer Vision*, pages 535–552. Springer, 2022. 2
- [43] Ce Zheng, Sijie Zhu, Matias Mendieta, Taojiannan Yang, Chen Chen, and Zhengming Ding. 3d human pose estimation with spatial and temporal transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11656–11665, 2021. 2, 3, 7
- [44] Ce Zheng, Xianpeng Liu, Guo-Jun Qi, and Chen Chen. Potter: Pooling attention transformer for efficient human mesh recovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 1
- [45] Ce Zheng, Matias Mendieta, Taojiannan Yang, Guo-Jun Qi, and Chen Chen. Feater: An efficient network for human reconstruction via feature map-based transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 1
- [46] Wentao Zhu, Xiaoxuan Ma, Zhaoyang Liu, Libin Liu, Wayne Wu, and Yizhou Wang. Motionbert: A unified perspective on learning human motion representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. 2