

PGDS: Pose-Guidance Deep Supervision for Mitigating Clothes-Changing in Person Re-Identification

Quoc-Huy Trinh¹, Nhat-Tan Bui³, Dinh-Hieu Hoang^{1,2}, Phuoc-Thao Vo Thi¹
Hai-Dang Nguyen^{1,2}, Debesh Jha⁴, Ulas Bagci⁴, Ngan Le³, Minh-Triet Tran^{1,2}

¹University of Science, Vietnam National University, Ho Chi Minh City, Vietnam

²John von Neumann Institute, Vietnam National University, Ho Chi Minh City, Vietnam

³AICV Lab, University of Arkansas, Fayetteville, Arkansas, USA

⁴Northwestern University, Chicago, Illinois, USA

Abstract

Person Re-Identification (Re-ID) task seeks to enhance the tracking of multiple individuals by surveillance cameras. It supports multimodal tasks, including text-based person retrieval and human matching. One of the most significant challenges faced in Re-ID is clothes-changing, where the same person may appear in different outfits. While previous methods have made notable progress in maintaining clothing data consistency and handling clothing change data, they still rely excessively on clothing information, which can limit performance due to the dynamic nature of human appearances. To mitigate this challenge, we propose the Pose-Guidance Deep Supervision (PGDS), an effective framework for learning pose guidance within the Re-ID task. It consists of three modules: a human encoder, a pose encoder, and a Pose-to-Human Projection module (PHP). Our framework guides the human encoder, i.e., the main re-identification model, with pose information from the pose encoder through multiple layers via the knowledge transfer mechanism from the PHP module, helping the human encoder learn body parts information without increasing computation resources in the inference stage. Through extensive experiments, our method surpasses the performance of current state-of-the-art methods, demonstrating its robustness and effectiveness for real-world applications. Our code is available at <https://github.com/huyquoctrinh/PGDS>.

1. Introduction

Person Re-Identification (Re-ID) involves tracking and distinguishing individuals across various cameras and viewpoints. This technique finds application in surveillance

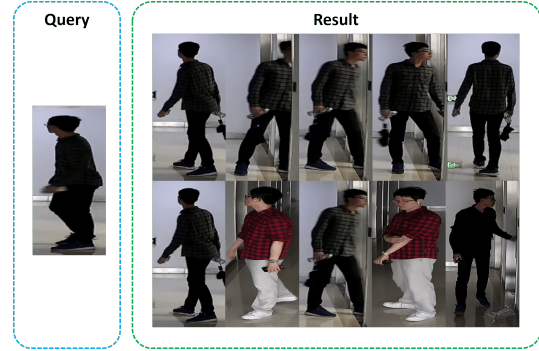


Figure 1: An example query retrieved by our framework under the clothes-changing scenario.

camera systems, enhancing security measures by enabling the search, tracking, and identifying unidentified individuals within specific areas or towns, thereby enhancing security measures.

Despite notable advancements, several challenges persist in the Re-ID task, e.g. occlusion, weather conditions, and other external factors. In such problems, clothing-change presents itself as one of those paramount challenges, first highlighted by Xue et al. [31] and Wu et al. [30]. This challenge becomes particularly crucial in real-world scenarios where individuals always change their clothes.

In recent years, researchers in the Re-ID domain have increasingly focused on mitigating the clothing-change problem. Various methods, such as VC-Clothes [27], LTCC [25], FSAM [11], CAL [7], and GI-ReID [13], have been introduced to address this issue. The primary difficulty in tackling this challenge arises from deep learning models relying heavily on clothing appearance to make predictions rather than considering biometric traits (e.g., face, hair, pose, gait). A growing body of literature, such as PGFL-

KD [33], Pose-transfer [18], Pose-guided re-ID [15], and PSEAMA [14] have demonstrated the effectiveness of pose information in tackling the clothes-changing problem. However, these works often require complex modules to integrate body parts information into the Re-ID model, leading to an increase in computational cost, training time, and inference time.

As a result, we introduce **Pose-Guidance Deep Supervision (PGDS)**, a simple and innovative framework for the clothes-changing Re-ID task. Figure 1 indicates an example query result of our framework under the clothing-change scenario. In essence, our method builds upon three main modules: human encoder, pose encoder, and **Pose-to-Human Projection** module (PHP). The human encoder, i.e., the main Re-ID model, efficiently adapts SOLIDER [5], a self-supervised learning framework for human-centric tasks, to extract the general human representations from the input images. At the same time, the pose encoder is taken from the pose estimation model OpenPose [3] to generate meaningful global features for biometric identification, which includes the pose, body, foot, hand, and facial. Our PHP, which includes several projectors, operates between two encoders to transfer the pose information extracted from the pose encoder to the human encoder across various scales. In practice, we fine-tune the pre-trained human encoder and train our projectors while freezing the pose encoder since we only need the general human pose information from the pose encoder. By incorporating pose information into the multiple layers, the Re-ID model can effectively concentrate on the body parts information, which contain unique information crucial for individual identification. Furthermore, integrating the frozen pre-trained pose estimation model also avoids increasing the computational cost in the inference stage, thus benefiting real-world applications. In summary, this paper’s main contributions are as follows:

- We explore the potential of transferring the human pose structure knowledge from the frozen pose estimation model into the main Re-ID network to overcome the clothes-changing problem. This simple strategy can be the baseline approach to alleviate the clothes-changing effect on the Re-ID task.
- We introduce a Re-ID framework, termed PGDS, incorporating a frozen pre-trained pose estimation model to guide the human representation model across different scales.
- We conduct comprehensive ablation studies on multiple datasets to assess the effectiveness of our proposed PGDS framework.

The content of this paper is organized as follows. In Section 2, we briefly review existing methods for Clothes-

changing Re-Identification problem and similar works that explore the utilization of pose information for Person Re-Identification. Then, we introduce our proposed method in Section 3. Section 4 presents the experiment setup, performance comparison, and ablation study. Finally, we conclude our work and suggest problems for future work in Section 5.

2. Related Work

Clothes-changing Re-Identification: Clothes-changing poses one of the most significant challenges in the Person Re-Identification task. The introduction of the LTCC dataset [25] marked a significant development in addressing the clothes-changing problem. Subsequent research has yielded various methods to tackle this challenge. FSAM [11] aims to obtain coarse ID masks with structure-related details, incorporating ID-relevant information for discriminative structural feature extraction. CAL [7] proposes a novel multi-positive-class classification loss to formulate multi-class adversarial learning. GI-ReID [13] integrates clothes classification and a casual inference model to mitigate bias in clothing information. AIM [32] defines a causality-based auto-intervention model to mitigate clothing bias for robust Clothes-changing Person Re-identification. FIRE² [28] introduces fine-grained feature mining and a fine-grained attribute recombination module to enhance the learning of robust features. Although previous works achieved impressive results on the clothes-changing Re-ID, it is worth noting that these methods lack focusing on the crucial parts of the body, which is the salient information to distinguish a specific person. In contrast, our work focus on prioritizing body parts information by efficiently leverage the pose information from frozen pre-trained pose estimation model to guide the main Re-ID model.

Pose Guide for Clothes-changing Re-Identification: Pose is crucial information utilized to guide Re-Identification models. Several methods have been proposed based on this approach. Among the most notable ones are ABDNet [4], PGR [16], Gated Fusion [1], and Pitr [2]. These methods follow the concept of integrating modules to fuse features extracted from pose heatmaps to the Re-ID models. More recently, PGFL-KD [33], PFD [29], Pose-guided re-ID [15], and PSEAMA [14] have leveraged pose guidance through part matching via global features. These methods demonstrate the effectiveness of pose information in guiding Re-ID models. However, the primary issue noted in these approaches is the substantial computational load imposed by the fusion modules, which are used to incorporate pose information into the main model. Differently from these previous works, PGDS integrates the pre-trained pose estimation model during the training phase, which remains frozen, to guide the main Re-ID model with pose information. This approach prevents an increase in computational

expenses during the inference stage.

3. Proposed PGDS

In this section, we introduce our PGDS framework, which consists of three primary modules: the human encoder, the pose encoder, and the pose-to-human projection module (PHP), as illustrated in Figure 2.

The human encoder takes in the image, learns each person’s feature representations and unique human appearance information, and then returns the image embedding. Meanwhile, the pose encoder (i.e., from a frozen pre-trained pose estimation model) is used to extract the human pose information from the input image and packs this information into a pose heatmap. In the training phase, the PHP encourages the early layers of the human encoder to attain the human representation with additional pose information through a knowledge distillation mechanism. When inferring, only feature representations from the human encoder are used to accomplish the re-identification task. In other words, we leverage a frozen encoder pre-trained for pose estimation as a **teacher** to compel the **student** human encoder to retain pose information into the image embedding. We detail each component in the following subsections to provide more in-depth information.

3.1. Pose Encoder

As mentioned, the proposed PGS aims to transfer the gait, pose, and posture information, which are unique characteristics of each person, to consolidate the model’s capacity to identify the same person wearing different clothes. Thus, we propose a simple strategy that exploits a frozen pre-trained pose estimator to enhance the Re-ID model’s ability to concentrate on body part information instead of cloth appearance.

Our pose encoder is derived from the OpenPose [3], a popular 2D pose estimation model. We hypothesize that the feature representation from OpenPose sufficiently and efficiently covers the pose-related information for our framework since we only want to get the general posture information. As a consequence, we keep the pose encoder unchanged so that we can minimize the number of parameters during training.

The pose encoder takes the input human image to generate the confidence map (i.e. pose feature), which highlights parts of the human, including “human body, hand, facial, and foot keypoints” [3]. The confidence map is then passed through the global average pooling and fully connected layer to produce the meaningful feature maps $f^p \in \mathbb{R}^{1 \times 768}$ that are associated with the overall posture information.

3.2. Human Encoder

Our human encoder is derived from SOLIDER [5], a human-centric self-supervised learning framework. It is important to note that the SOLIDER framework is based on the teacher-student knowledge distillation approach. Thus, it can adapt to a wide range of downstream tasks with different requisitions of semantic information and appearance information. In addition to the teacher and student networks, SOLIDER has a novel semantic controller that can modify the ratio of semantic over appearance information encoded in the feature.

Following the original paper [5], we use SOLIDER’s student network as our human encoder, accompanied by the frozen semantic controller with the ratio $\lambda_{human} = 0.2$. More precisely, the backbone of the human encoder is the Swin Transformer [20], which consists of four main layers. First, the human encoder will encode the image $X \in \mathbb{R}^{W \times H \times 3}$ into five scale features $f_i^h \in \mathbb{R}^{\frac{W}{2^{i+1}} \times \frac{H}{2^{i+1}} \times C_i}$ where $C_i \in \{48, 96, 192, 384, 768\}$, and $i \in \{0, 1, 2, 3, 4\}$. The final scale feature f_4^h undergoes the linear layer to generate the final human embedding, denoted as $f \in \mathbb{R}^{1 \times 768}$. Each human embedding will be utilized to retrieve each individual. We train this human encoder in the contrastive learning manner with the triplet loss so that the model can discriminate between different people. This image embedding will also be aligned with the pose feature f^p by the guide loss function [10], which we will detail in the following subsection 3.4.

3.3. Pose-to-Human Projection Module (PHP)

The primary objective of PHP is to effectively and robustly transfer posture knowledge from the pose encoder to the human encoder. To enhance the consistency of the human encoder throughout all its stages, PHP incorporates three projectors corresponding to the three intermediate layers of the human encoder. At each stage, denoted as i , the projector begins with the input feature from the human encoder, referred to as f_i^h . It then extracts implicit human representations, denoted as f_i^f , through a series of operations, including global average pooling, linear projection, batch normalization, and applying the non-linear ReLU activation function.

It is important to note that the human encoder outputs feature vectors that are more informative and semantically richer compared to the pose encoder, which only contains human pose information. From this understanding, it becomes apparent that directly guiding the features from the human encoder with pose features might result in the loss of other crucial information embedded in the human features. Therefore, we design an indirect transfer knowledge mechanism called a projector. This mechanism enables the human encoder to learn pose information without compro-

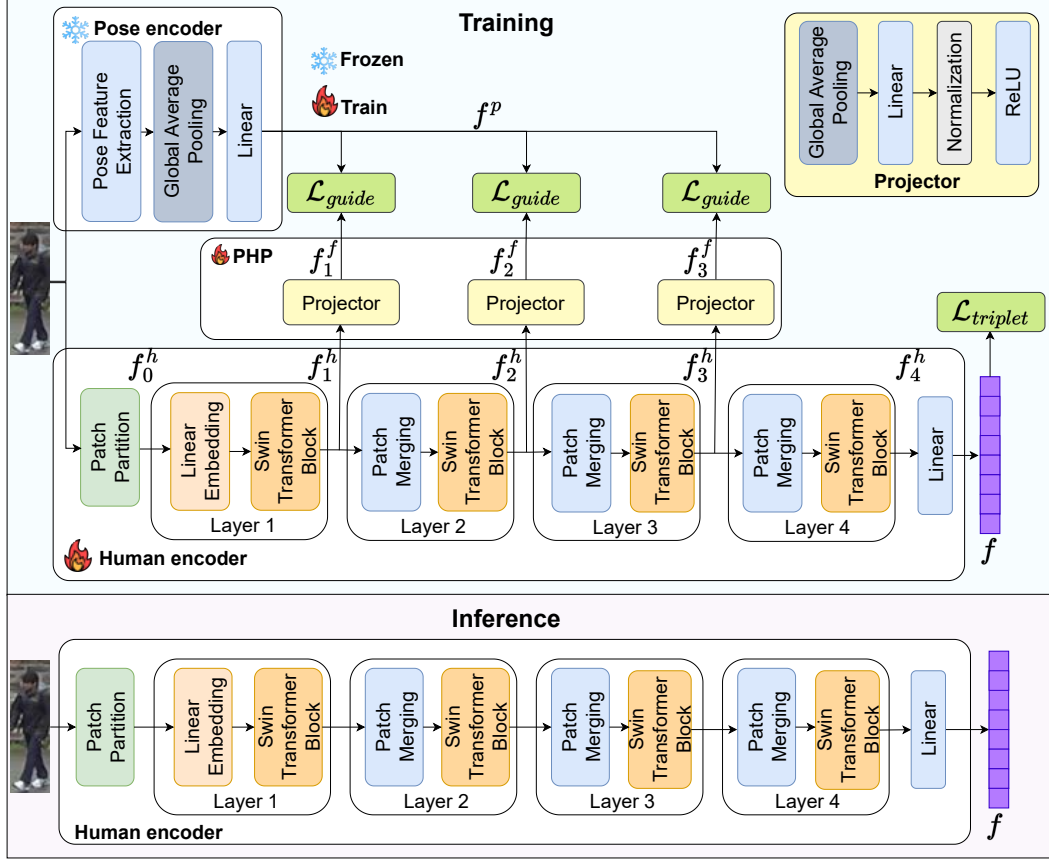


Figure 2: Overall framework of proposed PGDS including three modules: a human encoder, a pose encoder, and a pose-to-human projection module (PHP). The pose encoder module utilizes a frozen pre-trained model while we fine-tune a pre-trained human-centric model for the human encoder module. Our PHP transfers pose knowledge from the pose encoder module to the human encoder module through multiple projectors and guide loss \mathcal{L}_{guide} . H , W , and C denote the height, width, and channel, respectively. $\mathcal{L}_{triplet}$ is the triplet loss [10] which acquires person-centric representations.

missing other essential information. Intuitively, the projector acts as a non-linear, lightweight extractor. It unveils human-specific information within the higher-dimensional image representation of the human encoder and maps it to the lower-dimensional pose heatmap space of the pose encoder. Consequently, the resulting features at each stage, denoted as f_i^f , are incorporated into the guide loss \mathcal{L}_{guide} to align them with the pose feature f^p obtained from the pose encoder.

3.4. Objective Functions

Our framework is trained end-to-end with a combined loss, defined as in Equation 1.

$$\mathcal{L}_{combined} = \mathcal{L}_{triplet} + \lambda \mathcal{L}_{guide} \quad (1)$$

where $\mathcal{L}_{triplet}$ is the triplet loss [10] to learn person-specific representations, and \mathcal{L}_{guide} is the guide loss to transfer the pose knowledge from pose encoder to human encoder. The

hyperparameter λ , which we set to 0.8, controls the contribution and effect of the guide loss on the whole system. See section 4.3 for the ablation study on the λ values.

Triplet Loss: to learn a discriminative embedding in the feature space (i.e., encouraging representations of the same person to be similar to one another), we adapt the popular deep metric learning triplet loss [10], which is defined as follows:

$$\mathcal{L}_{triplet} = \sum_{i=1}^N [||f_i^{an} - f_i^{po}||_2^2 - ||f_i^{an} - f_i^{ne}||_2^2 + \alpha] \quad (2)$$

where N corresponds to the number of data samples in a batch, i denotes the index of the sample, f stands for the human embedding from the human encoder, and an, po, ne symbols individually refer to the anchor, positive, and negative instances. We employ batch hard-triplet mining [9] for the triplet selection. The parameter α signifies the margin, dictating the separation between embeddings. We set α

equal to 0.2 in our work.

Guide Loss: To encourage the model to learn semantically consistent and clustered feature representations, we utilize KL-divergence loss and the cross-entropy defined on the class-probability vectors p^p and p^f of pose embedding f^p and human representation f^f :

$$\mathcal{L}_{KL}(p^p, p^f) = y\mathcal{L}_{sim}(p^p, p^f) + (1 - y)\mathcal{L}_{dis}(p^p, p^f) \quad (3)$$

If p^p and p^f share the same human ID, $y = 1$ and

$$\mathcal{L}_{KL}(p^p, p^f) = \mathcal{L}_{sim}(p^p, p^f) = KL(p^p|p^f) + KL(p^f|p^p) \quad (4)$$

Otherwise, $y = 0$ and

$$\begin{aligned} \mathcal{L}_{KL}(p^p, p^f) = \mathcal{L}_{dis}(p^p, p^f) = \max \left(0, \left(m - KL(p^p|p^f) \right) \right) \\ + \max \left(0, \left(m - KL(p^f|p^p) \right) \right) \end{aligned} \quad (5)$$

where m is user-specified margin and $m = 2$. The KL-divergence is given by $KL(p|q) = \sum_{k=1}^K (p^k \log(\frac{p^k}{q^k}))$

Corresponding to three layers in PHP, the guide loss at each layer i^{th} is computed as $\mathcal{L}_{KL}(p^p, p_i^f)$, where p^p and p_i^f are class-probability vectors of pose embedding f^p and human representation f_i^f .

During the training phase, at each stage, the primary objective of this loss function is to minimize the disparity between the distributions originating from pose embeddings f^p and those obtained from human embeddings f_i^f . This process helps determine the optimal weights for refining human appearance representations by incorporating pose knowledge. Consequently, the model becomes proficient at capturing appearance and pose-related information. Since we aim to introduce a simple framework for transferring biometric knowledge (i.e., pose) to the Re-ID task, we have not conducted an exhaustive search for the best loss functions. This allows other researchers to explore the potential of alternative loss functions for this task.

4. Experiments

4.1. Experimental Setup

Datasets: To assess the effectiveness of the proposed PGDS, we conduct evaluations on five datasets that previous methods have widely used. These datasets cover a wide range of scenarios and closely resemble real-world conditions. Market-1501 [34] contains 1,501 identities observed from 6 camera viewpoints, including 12,936 training images of 751 identities, 19,732 gallery images, and 2,228 queries. Duke-MTMC [26] comprises 34,183 images of 1,404 identities from eight cameras. It contains 16,522 training images, 17,661 gallery images, and 2,228 queries. CuHK03 [17] contains 14,097 pictures of 1,467 identities.

Additionally, the dataset offers 20 sets for training and testing, where 100 identities are reserved for testing in each set while the remaining identities are used for training. LTCC [25] contains 17,119 person images of 152 identities. The training set in this data includes 9,576 images with 77 identities (46 clothes-changing IDs and 31 clothes-consistent IDs). The testing set includes 493 query images and 7,050 test images with 75 identities (45 clothes-changing IDs + 30 clothes-consistent IDs). This dataset includes two types of testing: *Standard* (test cases including clothes-changing and clothes-consistent), and *CC* (test cases include only clothes-changing IDs). VC-Clothes [27] comprises synthesized human images. It encompasses 19,060 images of 512 unique identities in four scenes. In the training phase, 9449 images of 256 identities are employed, leaving the remaining 9611 images of 256 identities for the testing phase.

Implementation Detail: During the training and testing phase, all images are resized to dimensions of 384×128 . Our training and testing for the clothes-consistent follow the setting of Bag of Trick [22], while in the clothes-changing testing, we follow the setup of LTCC [25]. The Pytorch framework is used, and all experiments are done on the Tesla A100 40GB GPU. The batch size for the training and testing is 64; we also use the learning rate scheduler to support our training experiments with a base learning rate of $8e-4$, and the temperature value equals 2. We choose the AdamW [21] as our optimizer, and the best weight has been gotten after 250 epochs.

Baselines: Based on two kinds of Re-ID problems, we select different baselines in two scenarios as follows. For clothes-changing setting, we evaluate our method on LTCC-Standard, LTCC-CC and VC-Clothes datasets. We compare our method with other state-of-the-art methods (i.e., LTCC [25], FSAM [11], CAL [7], GI-ReID [13], AIM [32], and FIRE² [28]). For clothes-consistent setting, we evaluate our method on Market-1501, Duke-MTMC and CuHK03 datasets. We compare our method with several state-of-the-art methods on the Re-ID task, including Bag of Trick [22], ABDNet [4], Auto-ReID+ [6], TransReID [8], PGFL-KD [33], AML [12], PFD [29], and SOLIDER [5]. To further assess the robustness of our model, we conduct the cross-domain setting in two scenarios: the Duke to Market ($D \rightarrow M$) and the Market to Duke ($M \rightarrow D$). We compare our method against several existing methods, e.g., SSL [23], ATNet [19], UDAP [24], and SOLIDER [5].

Evaluation Metrics: For a fair comparison, we employ two metrics, i.e., mean Average Precision (mAP), and Rank-1 accuracy (R1). mAP is calculated via the area under the precision-recall curve while R1 is defined by how many samples are correct in the top-1 prediction.

4.2. Performance Comparisons

The results of the clothes-changing setting, the clothes-consistent setting, and the cross-domain setting are shown in Table 1, Table 2, and the Table 3, respectively.

Table 1: Quantitative results on the clothes-changing datasets. The highest scores are shown in **bold**.

Methods	LTCC-Standard		LTCC-CC		VC-Clothes	
	mAP↑	R1↑	mAP↑	R1↑	mAP↑	R1↑
LTCC [25]	34.3	71.4	11.7	25.2	—	—
FSAM [11]	35.4	73.2	16.2	38.5	78.9	78.6
CAL [7]	40.8	74.2	18.0	40.1	—	—
GI-ReID [13]	29.4	63.2	10.4	23.7	59.0	63.7
AIM [32]	41.1	76.3	19.1	40.6	74.1	73.7
FIRE ² [28]	39.9	75.9	19.1	44.6	—	—
PGDS (Ours)	43.0	77.5	26.7	49.1	84.6	92.5

Clothes-changing: From Table 1, it is evident that on the LTCC-Standard dataset, our model showcases an improvement of over +1.9% in mAP and +1.2% in the R1 metric compared to the second-best method, AIM [32]. In the sole clothes-changing testing scenario (LTCC-CC dataset), our method outperforms the second-best (FIRE²) with a +7.6% improvement in mAP and a substantial margin of +4.5% in the R1 metric. As for the VC-Clothes dataset, our method achieves a notable improvement of +5.7% in mAP and +13.9% in the R1 metric compared to the second-best method, FSAM. These results underscore the efficacy of our method in addressing the clothes-changing problem by prioritizing other body parts over clothing information.

Table 2: Quantitative results on the clothes-consistent datasets. The highest scores are shown in **bold**.

Methods	Market-1501		Duke-MTMC		CuHK03	
	mAP↑	R1↑	mAP↑	R1↑	mAP↑	R1↑
Bag of Trick [22]	94.2	95.4	89.1	90.3	56.6	58.8
ABDNet [4]	88.3	95.6	78.6	89.0	—	—
Auto-ReID+ [6]	88.2	95.8	80.1	90.1	74.2	78.1
PGFL-KD [33]	87.2	95.3	79.5	89.6	—	—
AML [12]	89.5	95.7	81.7	91.1	82.3	85.6
PFD [29]	89.7	95.5	82.2	90.6	—	—
SOLIDER [5]	95.3	96.6	86.1	89.4	71.6	67.4
PGDS (Our)	95.4	96.9	91.4	92.6	89.7	87.9

Clothes-consistent: From Table 2, it is evident that our method leads to an improvement of +0.1% in the mAP metric and +0.3% in the R1 metric on the Market-1501 dataset when compared to the second-best method, SOLIDER [5]. Regarding the Duke-MTMC dataset, our method achieves a better mAP metric by +2.3% and an improvement of +2.3% in the R1 metric compared to the second-best method, Bag of Tricks [22]. In the CuHK03 dataset, our method surpasses the second-best method, AML [12], by a margin of +7.4% in the mAP metric and +2.3% in the R1

metric. These results demonstrate that our method remains competitive with state-of-the-art approaches across all three datasets, emphasizing the efficacy of leveraging pose information for distinguishing individuals. Such outcomes are significant as they highlight the potential robustness of our approach for future studies.

Cross-domain testing: From Table 3, in the $D \rightarrow M$ setting, our method surpasses the second best method by margins of +9.9% and +1.2% in terms of mAP and R1, respectively. While in the $M \rightarrow D$ scenario, we exceed the performance by +3.2% and +3.7% in mAP and R1 compared to the second best method. These results highlight the robustness of our framework in the Re-ID task.

Table 3: Results of the cross-domain testing. D, M represent the Duke-MTMC and Market-1501 datasets. The highest scores are shown in **bold**.

Methods	D \rightarrow M		M \rightarrow D	
	mAP↑	R1↑	mAP↑	R1↑
SSL [23]	37.8	71.7	28.6	52.5
ATNet [19]	25.6	55.7	24.9	45.1
UDAP [24]	53.7	75.8	49.0	68.4
SOLIDER [5]	50.3	59.9	50.3	59.9
PGDS (Our)	63.6	77.0	53.5	72.1

4.3. Ablation Study

To evaluate the effectiveness of the PGDS framework, we conducted four ablation studies. The first experiment assesses the effectiveness of pose guidance by the PHP module across multiple scales in the clothes-consistent problem, compared to the baseline human encoder (i.e., SOLIDER [5]). The second experiment verifies the PHP module’s effectiveness across multiple scales in the clothes-changing problem. The third experiment explores the impact of the parameter λ , which controls the contribution of the guide loss in the total loss, on knowledge transfer.

Effect of PHP across multiple scales in clothes-consistent problem: Table 4 shows that with our proposed PHP module applying to numerous stages, the results are improved over the baseline, even with a relatively small number of added trainable parameters. Since our framework only uses the human encoder in the inference stage, thus the FLOPs value remains the same compared with the baseline.

Table 4: Impact of PHP in clothes-consistent problem compared with baseline in Duke-MTMC dataset. Our method, represented by three distinct versions, i.e., PGDS-1, PGDS-2, and PGDS-3, incorporates one to three PHP projectors for transferring posture knowledge.

Methods	Computational Cost		Metrics	
	Params(M)↓	FLOPs(G)↓	mAP↑	R1↑
SOLIDER [5]	27.51	5.54	86.1	89.4
PGDS-1	28.13(▲0.62)	5.54	88.6(▲2.5)	90.9(▲1.5)
PGDS-2	28.28(▲0.77)	5.54	89.6(▲3.5)	91.4(▲2.0)
PGDS-3	28.57(▲1.06)	5.54	91.4(▲5.3)	92.6(▲3.2)

Effect of PHP across multiple scales in clothes-changing problem: The results are summarized in Table 5 indicate that when integrating the projector of PHP into multiple stages, the Re-ID framework can make more robust prediction since it prioritizes the pose information more than the clothes.

Table 5: Impact of PHP in clothes-changing problem.

Settings	LTCC-CC		VC-Clothes	
	mAP↑	R1↑	mAP↑	R1↑
PGDS-1	23.5	45.1	83.3	91.3
PGDS-2	24.4(▲0.9)	45.6(▲0.5)	83.4(▲0.1)	91.4(▲0.1)
PGDS-3	26.7(▲3.2)	49.1(▲4.0)	84.6(▲1.3)	92.5(▲1.2)

Impact of λ for knowledge transfer: As illustrated in Table 6, the results indicate that the model’s performance first improves and then declines when λ equals 1.0. This phenomenon occurs because a low λ limits the model’s ability to adapt to pose knowledge, while setting λ to 1.0 can make it challenging for the model to optimize both the triplet loss and guide loss simultaneously.

Table 6: Ablation of the λ parameters on LTCC-CC dataset.

Accuracy	$\lambda = 0.2$	$\lambda = 0.4$	$\lambda = 0.6$	$\lambda = 0.8$	$\lambda = 1.0$
mAP↑	25.2	25.6	25.9	26.7	26.2
R1↑	45.8	48.6	48.9	49.1	45.8

4.4. Feature Map Visualization

To gain deeper insights into the behavior of our framework, we conducted heatmap visualization, as depicted in Figure 3. This visualization readily highlights that the resulting heatmaps from the three versions (PGDS-1, PGDS-2, PGDS-3) primarily concentrate on the human’s head and various body parts. In simpler terms, the synergy of the human encoder, pose encoder, and the PHP module equips the model with the ability to accurately prioritize the human body, reducing its dependence on clothing variations.




Figure 3: Heatmap visualization compared with baseline to understand the behavior of our framework. The baseline is the SOLIDER [5].

5. Conclusion

In conclusion, this paper introduces the concept of Pose-Guidance Deep Supervision (PGDS), which transfers the pose knowledge into the main Re-ID model at multiple scales for robust Re-ID in scenarios involving both clothes-consistent and clothes-changing. Our study demonstrates the effectiveness of this approach in enabling the model to learn robust features, as evidenced by its competitive performance compared to current state-of-the-art methods. This simple approach is a promising candidate for real-world camera surveillance applications and provides a solid foundation for future studies to build more robust models. We encourage fellow researchers to explore various methods to improve this framework and address cloth-changing problems through the approach that integrates pose information.

Acknowledgement. Dinh-Hieu Hoang is funded by Vingroup Joint Stock Company and supported by the Domestic Master/ PhD Scholarship Programme of Vingroup Innovation Foundation (VINIF), Vingroup Big Data Institute (VINBIGDATA), code VINIF.2022.ThS.JVN.04. Debesh Jha and Ulas Bagci are supported by the NIH funding: R01-CA246704, R01-CA240639, U01-DK127384-02S1, and U01-CA268808.

References

- [1] A. Bhuiyan, Y. Liu, P. Siva, M. Javan, I. B. Ayed, and E. Granger. Pose guided gated fusion for person re-identification. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2675–2684, 2020.
- [2] A. Bhuiyan, Y. Liu, P. Siva, M. Javan, I. B. Ayed, and E. Granger. Pose guided gated fusion for person re-identification. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2675–2684, 2020.
- [3] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [4] T. Chen, S. Ding, J. Xie, Y. Yuan, W. Chen, Y. Yang, Z. Ren, and Z. Wang. Abd-net: Attentive but diverse person re-identification.

In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019.

- [5] W. Chen, X. Xu, J. Jia, H. Luo, Y. Wang, F. Wang, R. Jin, and X. Sun. Beyond appearance: a semantic controllable self-supervised learning framework for human-centric visual tasks. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [6] H. Gu, G. Fu, J. Li, and J. Zhu. Auto-reid+: Searching for a multi-branch convnet for person re-identification. *Neurocomputing*, 435:53–66, 2021.
- [7] X. Gu, H. Chang, B. Ma, S. Bai, S. Shan, and X. Chen. Clothes-changing person re-identification with rgb modality only. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [8] S. He, H. Luo, P. Wang, F. Wang, H. Li, and W. Jiang. Transreid: Transformer-based object re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021.
- [9] A. Hermans, L. Beyer, and B. Leibe. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*, 2017.
- [10] E. Hoffer and N. Ailon. Deep metric learning using triplet network. In *International Workshop on Similarity-Based Pattern Recognition*, pages 84–92, 2015.
- [11] P. Hong, T. Wu, A. Wu, X. Han, and W.-S. Zheng. Fine-grained shape-appearance mutual learning for cloth-changing person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
- [12] Z. Huang, L. Wang, Y. Li, A. Du, and S. Jiang. Improving performance in person reidentification using adaptive multiple loss baseline. *Information*, 13(10):453, 2022.
- [13] X. Jin, T. He, K. Zheng, Z. Yin, X. Shen, Z. Huang, R. Feng, J. Huang, Z. Chen, and X.-S. Hua. Cloth-changing person re-identification from a single image with gait prediction and regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [14] Y. Junping, Y. Zebin, L. Xiaojun, and G. Yi. Pose-guided self and external attention feature matching and aggregation network for person re-identification. *Displays*, 80:102567, 2023.
- [15] A. Khatun, S. Denman, S. Sridharan, and C. Fookes. Pose-driven attention-guided image generation for person re-identification. *Pattern Recognition*, 137:109246, 2023.
- [16] J. Li, S. Zhang, Q. Tian, M. Wang, and W. Gao. Pose-guided representation learning for person re-identification. *IEEE transactions on pattern analysis and machine intelligence*, 44(2):622–635, 2019.
- [17] W. Li, R. Zhao, T. Xiao, and X. Wang. Deepreid: Deep filter pairing neural network for person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2014.
- [18] J. Liu, B. Ni, Y. Yan, P. Zhou, S. Cheng, and J. Hu. Pose transferable person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.
- [19] J. Liu, Z.-J. Zha, D. Chen, R. Hong, and M. Wang. Adaptive transfer network for cross-domain person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- [20] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021.
- [21] I. Loshchilov and F. Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [22] H. Luo, Y. Gu, X. Liao, S. Lai, and W. Jiang. Bag of tricks and a strong baseline for deep person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2019.
- [23] H. Luo, P. Wang, Y. Xu, F. Ding, Y. Zhou, F. Wang, H. Li, and R. Jin. Self-supervised pre-training for transformer-based person re-identification. *arXiv preprint arXiv:2111.12084*, 2021.
- [24] P. Peng, T. Xiang, Y. Wang, M. Pontil, S. Gong, T. Huang, and Y. Tian. Unsupervised cross-dataset transfer learning for person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016.
- [25] X. Qian, W. Wang, L. Zhang, F. Zhu, Y. Fu, T. Xiang, Y.-G. Jiang, and X. Xue. Long-term cloth-changing person re-identification. In *Proceedings of the Asian Conference on Computer Vision*, 2020.
- [26] E. Ristani, F. Solera, R. Zou, R. Cucchiara, and C. Tomasi. Performance measures and a data set for multi-target, multi-camera tracking. In *European Conference on Computer Vision workshop on Benchmarking Multi-Target Tracking*, 2016.
- [27] F. Wan, Y. Wu, X. Qian, Y. Chen, and Y. Fu. When person re-identification meets changing clothes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020.
- [28] Q. Wang, X. Qian, B. Li, Y. Fu, Y. Fu, and X. Xue. Exploring fine-grained representation and recombination for cloth-changing person re-identification. *arXiv preprint arXiv:2308.10692*, 2023.
- [29] T. Wang, H. Liu, P. Song, T. Guo, and W. Shi. Pose-guided feature disentangling for occluded person re-identification based on transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2022.
- [30] A. Wu, W.-S. Zheng, and J.-H. Lai. Robust depth-based person re-identification. *IEEE Transactions on Image Processing*, 26(6):2588–2603, 2017.
- [31] J. Xue, Z. Meng, K. Katipally, H. Wang, and K. Van Zon. Clothing change aware person identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 2112–2120, 2018.
- [32] Z. Yang, M. Lin, X. Zhong, Y. Wu, and Z. Wang. Good is bad: Causality inspired cloth-debiasing for cloth-changing person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [33] K. Zheng, C. Lan, W. Zeng, J. Liu, Z. Zhang, and Z.-J. Zha. Pose-guided feature learning with knowledge distillation for occluded person re-identification. In *Proceedings of the 29th ACM International Conference on Multimedia*, 2021.
- [34] L. Zheng, L. Shen, L. Tian, S. Wang, J. Wang, and Q. Tian. Scalable person re-identification: A benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2015.