

Watch Where You Head: A View-biased Domain Gap in Gait Recognition and Unsupervised Adaptation

Gavriel Habib* Noa Barzilay* Or Shimshi Rami Ben-Ari Nir Darshan
 OriginAI, Israel
 {gavrielh, noab, ors, ramib, nir}@originai.co

Abstract

Gait Recognition is a computer vision task aiming to identify people by their walking patterns. Although existing methods often show high performance on specific datasets, they lack the ability to generalize to unseen scenarios. Unsupervised Domain Adaptation (UDA) tries to adapt a model, pre-trained in a supervised manner on a source domain, to an unlabelled target domain. There are only a few works on UDA for gait recognition proposing solutions to limited scenarios. In this paper, we reveal a fundamental phenomenon in adaptation of gait recognition models, caused by the bias in the target domain to viewing angle or walking direction. We then suggest a remedy to reduce this bias with a novel triplet selection strategy combined with curriculum learning. To this end, we present Gait Orientation-based method for Unsupervised Domain Adaptation (GOUDA). We provide extensive experiments on four widely-used gait datasets, CASIA-B, OU-MVLP, GREW, and Gait3D, and on three backbones, GaitSet, GaitPart, and GaitGL, justifying the view bias and showing the superiority of our proposed method over prior UDA works.

1. Introduction

Gait recognition is a biometric method that analyzes and identifies individuals based on their walking patterns. It involves the analysis of various characteristics such as stride, body shape, motion postures, and limb movements. Gait recognition has been widely used in different application areas such as sport [45], health monitoring [34], forensic analysis [31] as well as security and surveillance [33]. However, compared to different biometrics, like face and fingerprint recognition, gait is a challenging task due to complex conditions such as clothes changing, object carrying, and diverse camera views. Existing methods achieve high performance on specific datasets [4, 9, 29]. GaitSet [4] learns

from a set of gait silhouettes, GaitPart [9] proposes body part-based model, and GaitGL [29] utilizes global and local visual information. However, the ability of these methods to generalize to new datasets is limited due to domain gaps. Domain gaps can be caused by *e.g.* different camera settings or clothing conditions (changed between seasons), making the gait sequences appear differently across use cases.

Unsupervised Domain Adaptation (UDA) aims to transfer a model learned from a *source* domain with labeled data to an unseen *target* domain without labels. UDA is common in different computer vision tasks, such as object classification [42, 50], semantic segmentation [5, 16], and person re-identification [3, 21, 40].

UDA for Gait Recognition addresses the case in which the source and target domains include different identities. One main concern for gait recognition task is the ability to recognize an individual over different camera views and walking directions. The interplay between camera positioning and an individual’s walking direction within the 3D realm defines what is referred to as the *viewing angle* in literature [8, 35, 44]. Gait recognition models are trained to generalize over different viewing angles, particularly tested by robustness on such scenarios [35, 44].

In this paper, we consider a case where a trained model on a labeled source domain should be adapted to a target domain without any access to the target labels or the source data (that could be unavailable due to privacy issues). To this end, we first use the source model to generate an initial representation of the target domain followed by an adaptation scheme to adapt the pre-trained backbone to the new domain. Such adaptation is often carried out by applying a contrastive loss such as triplet loss [14]. Without access to the labels in the target domain, prior works used different cues to distinguish (approximately) between the samples [39, 47]. In this paper, we propose GOUDA - Gait Orientation-based method for Unsupervised Domain Adaptation. Our method is driven mainly by a strong observation of the gait view distribution in the target domain (we consider yaw angle here). Fig. 1 shows an example of view distribution on OU-MVLP dataset [35]. The depicted UMAPs

* The authors contributed equally

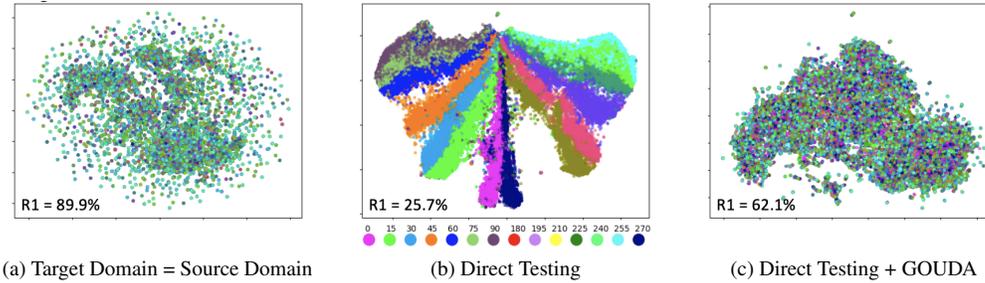


Figure 1. Gait embedding visualization (UMAP [32]) for domain transfer between GREW as source and OU-MVLP as the target domain. Points are color coded by viewing angle. Fig. 1a showcases the single domain scenario, in which source and target domains are the same (OU-MVLP). Fig. 1b presents the direct testing scenario, in which the model was trained on GREW and tested on OU-MVLP. The Rank-1 performance is dropped and a strong viewing angle bias is observed. Fig. 1c presents the direct testing scenario after applying GOUDA using the viewing angles provided in the dataset (just for sake of illustration). Our method is capable of disentangling the viewing angle bias to improve the recognition performance, showing a relatively uniform view distribution similar to the source domain.

show the distribution of gait representations color coded by corresponding viewing angles. While in the source domain (Fig. 1a) the distribution is nearly uniform indicating that relations in the embedding space are *invariant* to the view, the clustered colors in the target domain (Fig. 1b) imply a strong bias toward the viewing angle. In the target domain, nearby points are likely to be from similar view and not necessarily the same identity (see Rank-1 accuracies in the figure). This view bias causes a significant drop in recognition. Fig. 1c shows that GOUDA is capable of disentangling the view bias, modifying the embedding space in the target domain to improve the recognition accuracy. We show that this fundamental observation and correction method is a key contributor to gait UDA methods as evaluated over various gait backbones and cross domain datasets. Essentially, we suggest a novel triplet selection strategy, based on this key observation. Our method also uses a self-supervised method and curriculum learning, to further boost the results. To summarize, these are our key contributions:

- We reveal a fundamental bias in cross-domain gait recognition models toward viewing angles.
- We propose GOUDA, a UDA method with a novel view-based Triplet Selection strategy, mitigating the observed viewing angle bias and leading to a significant improvement in gait UDA task.
- We conduct extensive experiments with three backbones and with 14 source and target combinations, outperforming prior state-of-the-art by a large gap.

2. Related Work

Existing gait recognition methods can be grouped into two primary categories: model-based approaches and appearance-based approaches.

Model-based approaches [24,27,36,43] focus on capturing the underlying structure and kinematic features of the human body to learn the walking patterns for recognition. These methods utilize structural models, such as 2D or 3D skeleton keypoints [24,27,36,49] as their input representation, aiming to achieve robustness to clothing and object carrying conditions. However, they often face challenges in estimating body parameters in low-resolution videos. Appearance-based methods extract gait representations directly from RGB images or binary silhouette sequences without explicitly modeling the human body. In this context, Silhouette-based recognition is widely used in the literature e.g., template-based methods [12,23,38,38] that treat entire gait sequences as templates, sequence-based methods that consider gait sequences as videos [7,17,18,26,28,29] or as an unordered image set [4,15]. In our study, we use appearance-based gait recognition models, specifically employing silhouette sequences or unordered sets as input data.

2.1. Unsupervised Domain Adaptation (UDA)

In general UDA task [30,41], domain-invariant feature learning aims to align the source and target distributions by minimizing divergence measures [20] or by utilizing adversarial training [11,37]. Other UDA approaches employ ensemble methods [10] or adapt batch normalization statistics from source to target domain [25]. A different approach applies clustering using pseudo-labels [3,21,40], in the target domain. Our approach draws inspiration from aligning the source and target distributions.

UDA has been employed in the re-identification task coping with variations in view, camera settings, and environmental conditions, over different domains. In this respect, UNRN [48] is a clustering-based method with uncertainty-guided optimization driven by a teacher-student framework. SECRET [13] utilizes global and local features for a self-consistent pseudo-label refinement. A triplet strat-



Figure 2. An example of positive sample selection. We use the setting with GaitSet and OU-MVLP (source) to CASIA-B (target). TraND [47] usually selects positive samples with the exact same view as the anchors due to its strategy to bring adjacent samples closer. Contrary to that, GOUDA selects positives with different views by default. Here, the anchor view is 54° , and the positive sample selected by GOUDA is of view 108° . In both cases, the positives are correct identity-wise. By bringing closer samples of different views, our method improves the cross-view benchmarks and outperforms TraND, as shown in Table 2.

egy was used by Bertocco et al. [3] based on the diversity of cameras within a cluster. However, these studies address a different task of re-identification. In contrast to the person re-identification task, gait recognition focuses on videos of binary silhouettes or human skeletons rather than RGB images, to be invariant to the RGB appearance information.

Despite the necessity of UDA in Gait Recognition, only a few studies have addressed this specific task [39,47]. Both of these methods attempt for a fine-tuning adaptation on the target domain. From Indoor to Outdoor (FI2O) [39] conducts an adaptation step by using a pseudo-labeling strategy in the target domain, to adapt a model trained on an indoor environment to an outdoor using noisy labels. Their method relies on a clustering approach that assumes the existence of multiple sequences per person in the target dataset, which may not always be available. TraND [47] attempts to close the domain gap by a neighborhood selection strategy (for training) that relies on an entropy-based confidence measure. Effectively, its main effort is to keep bringing closer adjacent samples, with high confidence of belonging to the same identity. As a result, TraND tends to favor positive samples with similar views. Due to this bias, TraND struggles to deal with the common case of different views. In contrast to TraND, GOUDA adopts a different strategy by selecting positive samples from diverse views (see Fig. 2). It enriches the training set and contributes to the development of a model that is robust against variations in views. We also show that adding self-supervised learning on top of our triplet selection method, boosts the results, reducing the gap toward the upper bound of fine-tuning with labels.

3. Method

Figure 3 presents an overview of our method with three main stages. First, a gait model, pre-trained on a source domain, is used as a feature extractor for an unseen target domain. In the embedding space, each sequence is described by its evaluated viewing angle, while its label (iden-

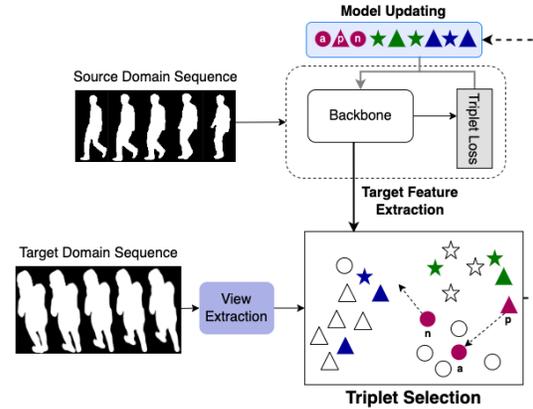


Figure 3. An overview of GOUDA. First, a pre-trained backbone, trained on the source domain, is used to extract gait features of sequences from an unseen target domain. Each sequence in the figure is characterized by a distinct shape, representing its corresponding view extracted using View Extraction module. Second, the Triplet Selection algorithm is applied, aiming to push away sequences of similar views and bring closer different ones in the embedding space. The colored shapes depict the selected triplets. The triplets are then used to update the backbone using triplet loss.

tity) is unknown. Second, our Triplet Selection approach is applied, aiming to push away sequences of similar views while bringing closer different views. In order to increase the probability that the triplets are identity-wise correct, we further adopt a Curriculum Learning framework [2] based on a certainty measure. Lastly, the selected triplets are used to fine-tune the network using standard triplet loss.

3.1. View Extraction

In gait recognition, the view is defined as the walking direction of the person relative to the camera. To estimate the view of a gait sequence in the target domain, we use the provided 2D keypoints (representing the posture of human body parts). First, we apply MHFormer [22] to regress the 3D keypoints per frame. Then, we estimate the sequence view (yaw angle) using the median 3D keypoints of relevant body parts - hips and shoulders. For real application uses, the 2D keypoints can be estimated from RGB images.

3.2. Triplet Selection

This section provides a description of the Triplet Selection algorithm, which aims to improve performance on the target domain by disentangling the view bias. Ideally, we aim to find correct triplets such that, given an anchor, the positive sample has the same identity but a different view, and the negative sample belongs to a different identity but of similar view.

Negative Sample Selection. Based on the observation that gait recognition models strongly emphasize view-based

features in the target domain (Fig. 1), we assume that within a neighborhood of samples with a similar view, there are likely to exist samples of different identities.

Positive Sample Selection. Gait models demonstrate a certain level of learned general gait features. Hence, we assume that among all samples with different views, the ones with the same identities are reasonable to be the closest¹. In Fig. 1b, it is implied by the cross-view Rank-1 accuracy (25.7%), which is much higher than random ($\pm 0.02\%$).

In the embedding space, we prioritize pushing away similar-view samples and bringing closer together cross-view samples. We use various techniques, such as curriculum learning, to increase the identity-wise certainty.

The input to the algorithm (pseudo-code presented in Alg. 1) is an unlabeled target domain set $S := \{(x_r, v_r)\}_{r=1}^R$, where R is the number of gait sequences, $x_r \in \mathbb{R}^{W \times H \times T}$ is the r -th sequence with T frames of size $W \times H$, and $v_r \in \mathbb{R}$ is its estimated view. Each gait sequence x_r is represented by its embedding $e_r = f(x_r)$, where $f(\cdot)$ is a gait model pre-trained in a supervised manner on the source domain. In addition, we define the following hyper-parameters: margin m , and view thresholds T_s, T_c such that $T_s \leq T_c$. We use v_j as an indicator (in the range of $[0, 360)$), for separation of the data for each anchor a into two mutually exclusive sets. Considering an anchor a , the sample j belongs to the set $\chi_{similar}^a$, if $|v_a - v_j| < T_s$, and belongs to χ_{cross}^a , if $|v_a - v_j| > T_c$. The triplets $\{a, p, n\}$ are selected such that $n \in \chi_{similar}^a$ and $p \in \chi_{cross}^a$.

Given a set of embeddings and views $\{(e_r, v_r)\}_{r=1}^R$, we now calculate the distances matrix $D \in \mathbb{R}^{R \times R}$, such that $D_{i,j} := d(e_i, e_j)$, where d is the cosine distance metric. In order to increase the certainty of triplets to be identity-wise correct, we select them as follows. The positive p is defined as the closest sample to a in χ_{cross}^a (line 4 in Alg. 1). We choose $n \in \chi_{similar}^a$ to be the farthest sample from a within a fixed margin m from p (line 5 in Alg. 1). Eventually, the triplet $\{a, p, n\}$ is “valid” if exists at least one sample in $\chi_{similar}^a$ that is closer to a than n (line 6 in Alg. 1). Both restrictions on n are used to increase the probability that a and n do not share the same identity, assuming that the closest similar-view sample is more likely to share the same identity. The margin m is incorporated into the negative sample selection (line 5 in Alg. 1) in order to align with the objective of triplet loss. An illustration of the Triplet Selection algorithm is shown in Fig. 4, and described in Algorithm 1. The algorithm is applied iteratively during the training process (Section 3.3).

3.3. Curriculum Learning

Inspired by [2, 47], we apply the Triplet Selection algorithm in four different stages during training, in a from-

¹In the supplementary materials, we present further intuition and justification, including a shared visualization of views and identities.

Algorithm 1 Triplet Selection

```

0: function TRIPLET SELECTION( $D, \{v_r\}_{r=1}^R$ )
1: for  $a \leftarrow 1$  to  $R$  do
2:    $\chi_{similar}^a \leftarrow \{j \mid |v_a - v_j| < T_s, a \neq j\}$ 
3:    $\chi_{cross}^a \leftarrow \{j \mid |v_a - v_j| > T_c\}$ 
4:    $p \leftarrow \chi_{cross}^a[\arg \min_{j \in \chi_{cross}^a} \{D_{a,j}\}]$ 
5:    $n \leftarrow \chi_{similar}^a[\arg \max_{j \in \chi_{similar}^a} \{D_{a,j} \mid D_{a,j} < D_{a,p} + m\}]$ 
6:   if  $\exists j \in \chi_{similar}^a$  s.t:  $D_{a,j} < D_{a,n}$  then
7:      $\mathcal{V}$ : Valid Triplets  $\leftarrow \{a, p, n\}$ 
8:      $\mathcal{C}$ : Confidence Scores  $\leftarrow 1 - D_{a,p}$ 
9:   Selected Triplets  $\leftarrow \mathcal{V}[\arg \max_{\text{top } q\%} \mathcal{C}]$ 
10: return Selected Triplets

```

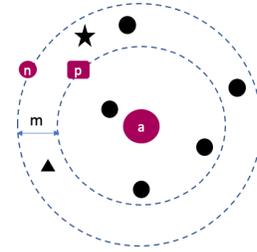


Figure 4. Triplet selection illustration (shape describes view). Given an anchor a , the positive p is the closest sample with a different view. The negative n is the farthest sample (up to margin m from p) with a similar view as a . A triplet $\{a, p, n\}$ is “valid” if exists at least one similar-view sample closer to a than n .

easy-to-hard manner. We define a *triplet confidence score* as the cosine similarity between the anchor and the positive samples ($1 - D_{a,p}$). We select the top $q\%$ valid triplets \mathcal{V} according to their confidence scores \mathcal{C} , while q gets higher in each stage. We keep improving target domain performance by starting from updating the model with a small amount but high-quality data and carefully extending the dataset size.

3.4. Stopping Criteria

In contrast to previous works [39, 47], we introduce an indicative stopping criteria (SC) using a validation set. Although training with a fixed number of iterations is simple, it lacks control over the distribution obtained in the embedding space, due to an over-fitting risk. Our SC is driven by the desire to achieve greater diversity in view within the neighborhoods. Given a validation set of size N_{val} , for each sample a we compute its K nearest neighbors, $KNN(a)$, using the cosine distance metric. Then, we count how many of the K samples are of similar view as a . The final SC value is the average of all counts (Eq. 1). Eventually, we

choose the checkpoint in which SC is minimal.

$$SC = \frac{1}{N_{val}} \sum_{a=1}^{N_{val}} \sum_{i \in KNN(a)} \mathbb{1}[i \in \mathcal{X}_{similar}^a] \rightarrow \min \quad (1)$$

3.5. Model Updating

In the adaptation stage, we optimize the backbone f with two triplet loss objectives: GOUDA and self-supervision loss. Suppose $\{a_i, p_i, n_i\}$ is a selected triplet, GOUDA loss is given by:

$$\mathcal{L}_{GOUDA} = \sum_{i=1}^N [D_{a_i, p_i} - D_{a_i, n_i} + m]_+ \quad (2)$$

where N represents the number of triplets, m is the triplet loss margin and $[x]_+$ is $\max(x, 0)$. Note that a and n share a similar-view, different from p . Our self-supervision loss (SSL) is defined by:

$$\mathcal{L}_{SSL} = \sum_{i=1}^N [D_{a_i, \tilde{a}_i} - D_{a_i, n_i} + m]_+ \quad (3)$$

In the self-supervision loss we replace p with \tilde{a} , an augmented version of a . The augmentations are two random sub-sequences of the same gait sequence. Eventually, the losses are linearly combined to our domain adaptation loss:

$$\mathcal{L} = \mathcal{L}_{GOUDA} + \mathcal{L}_{SSL} \quad (4)$$

The selected triplets are randomly extracted for training.

4. Experiments

4.1. Datasets

In this paper we evaluate the proposed GOUDA on four well-known gait datasets, CASIA-B [44], OU-MVLP [35], Gait3D [46], and GREW [51].

CASIA-B [44] is a popular indoor gait dataset. It contains 13,640 gait sequences, including 124 subjects with 11 views. Each subjects' videos include 3 different settings, normal walking (NM), walking with a bag (BG), and walking with a coat (CL). The first 74 subjects are used for the training set and the rest 50 subjects are used for the test set.

OU-MVLP [1, 35] is a large indoor gait dataset. It contains 10,307 subjects with 14 views. According to the provided protocol [4], 5,153 subjects are used for the training set, and the remaining 5,154 subjects are used for testing.

Gait3D [46] is a 3D representation-based gait dataset captured in the supermarket. The dataset contains 4,000 subjects and over 25,000 sequences extracted from 39 cameras. The training set includes 3,000 subjects, and the test set includes the remaining 1,000 subjects.

GREW [51] is a large-scale in-the-wild gait recognition dataset. The dataset contains 26,345 subjects and 128,671 sequences extracted from 882 cameras. We adopt 2 protocols to evaluate the dataset [51], [39]. According to the first protocol [51], the training set contains 20,000 subjects and the test set includes 6,000 subjects. The second protocol [39] splits the original training set, the first 2,000 subjects are utilized for training and the last 2,000 subjects are utilized for test. We refer to the latter partition as GREW*.

4.2. Implementation details

We use three different well-known gait backbones - GaitSet [4], GaitPart [9], and GaitGL [29] to evaluate the proposed GOUDA. The input silhouettes in all datasets are resized to 64×44 . Each backbone was first trained on the source domain data, using the implementations provided in OpenGait code [8]. The training set in each target domain was partitioned into separate training and validation sets. Specifically, we created the validation set by selecting the last 10% of subjects from the training set.

In the model updating stage, all backbones were trained using the loss function as described in Eq. 4, with $m = 0.2$. View thresholds T_s , and T_c are set to 10° and 20° respectively. We adopt Adam optimizer [19] with a learning rate of $1e-5$ and weight decay of $5e-4$. In every training iteration, a batch size of 32×4 was employed, where the number 4 corresponds to the inclusion of two different augmentations applied to the anchor sample (a, \tilde{a}) , a positive sample (p) , and a negative sample (n) . The final weights of the model were selected according to the optimal results of the stopping criteria (cf. Sec. 3.4) on the validation set. K value of the SC was set to 5. The curriculum learning strategy includes four stages, where each stage involves selecting a higher percentage of top $q\%$ valid triplets: 10, 25, 50, and 100^2 . The amount of training iterations after each stage is proportional to the number of selected triplets. Specifically, we move to the next stage after each selected triplet appears in training 10 times on average.

5. Evaluation

5.1. Results

Comparison to Direct Testing. Our method does not assume special architectural conditions and therefore can be integrated with different kinds of deep backbones³. We evaluate the performance of GOUDA when combined with various backbones, including GaitSet [4], GaitPart [9], and GaitGL [29], as shown in Table 1. These results are then compared to Direct Testing performances, *i.e.* models that are trained solely on the source domain in a supervised man-

²Ablation study on q values is shown in the Supplementary Materials.

³In the Supplementary Materials, we show that the domain gap presented in Fig. 1 is consistent across different gait backbones.

Source Dataset	Backbone	Target Dataset											
		CASIA-B						OU-MVLP		GREW		Gait3D	
		NM		BG		CL		GO	DT	GO	DT	GO	DT
		GO	DT	GO	DT	GO	DT	GO	DT	GO	DT	GO	DT
CASIA-B	GaitSet							27.9	9.6	9.4	9.4	9.3	7.7
	GaitPart							27.5	10.8	15.3	9.3	11.5	6.6
	GaitGL							34.0	16.2	14.6	11.6	10.9	7.6
OU-MVLP	GaitSet	87.0	74.0	68.0	55.5	27.2	16.4			11.7	12.3	13.3	13.0
	GaitPart	91.5	73.9	78.4	56.9	38.7	20.7	-		21.0	14.8	18.8	11.2
	GaitGL	92.8	81.7	80.9	71.5	44.6	28.8			25.7	21.4	21.6	16.0
GREW	GaitSet	76.9	65.6	56.8	44.9	26.0	20.8	38.8	21.8			16.3	19.0
	GaitPart	81.4	69.2	68.3	52.1	33.9	25.4	43.6	23.9	-		19.5	19.3
	GaitGL	82.7	69.8	73.2	61.1	44.6	31.9	44.2	25.7			18.4	15.6
Gait3D	GaitSet	73.4	62.8	51.9	45.8	20.4	11.9	41.9	20.8	14.9	19.2		
	GaitPart	76.7	61.8	61.1	47.0	24.5	13.7	32.2	19.1	14.9	14.2		
	GaitGL	79.0	63.5	65.4	51.2	24.9	16.3	37.0	23.7	10.7	14.3		

Table 1. Rank-1 accuracy of GOUDA (referred to as GO) compared to the results of Direct Testing (DT) for various backbones [4], [9], [29]. This table describes 12 settings of source and target datasets.

ner. It is evident that the direct testing results exhibit a certain level of competence in learning general gait features, which contributes to achieving reasonably decent results in the target domains. Nevertheless, these results are still unsatisfactory. Referring to Table 1, it becomes apparent that our method yields substantial improvements in Rank-1 accuracy across almost all settings. Notably, GOUDA demonstrates improvements across multiple backbones, underscoring the algorithm’s robustness independent of any specific architecture. This finding reinforces the observation presented in Section 1. GOUDA demonstrates superior performance compared to Direct Testing results in 49 out of 54 cases. In some cases, our method improves the Rank-1 accuracy by a significant margin, for example from 23.9% to 43.6%, when using GaitPart trained on GREW (source) to OU-MVLP (target). In the supplementary materials, we present even higher improvements when using ground truth views provided as meta-data in CASIA-B and OU-MVLP datasets, avoiding the noise produced by the View Extraction module. It may suggest that the potential of GOUDA is higher, if using an improved 3D keypoints estimator.

It is apparent that the Rank-1 improvements achieved by GOUDA are considerably higher when employing CASIA-B or OU-MVLP datasets as the target domain, compared to GREW or Gait3D. This disparity can be attributed to the nature of the datasets themselves. CASIA-B and OU-MVLP are indoor datasets captured in controlled environments, whereas GREW and Gait3D were collected in-the-wild. Consequently, the latter datasets present more challenging conditions for gait recognition, potentially leading to lower gains in performance compared to the former datasets. Furthermore, our method relies on estimating a

single view through the entire sequence. However, such an approach does not accurately reflect the conditions encountered in datasets captured in the wild. In practice, one can decompose the gait sequence into snippets with piece-wise constant views and analyze them separately.

Comparison with State-of-The-Art. We compare GOUDA to state-of-the-art UDA gait recognition methods [39, 47] and present them in Table 2. To ensure a fair comparison, we employ the same backbone in each comparison, GaitGL when comparing with [39] and GaitSet when comparing with [47]. In order to compare to TraND [47], we trained their model using the implementations provided by the authors. As a reference, we further report the results of UDA methods for person re-identification, UNRN [48] and SECRET [13], applied by [39]. In these cases, the entire gait sequence was compressed to a single template (GEI [12]) as these models get a single image as an input.

As shown in Table 2, GOUDA significantly outperforms prior works in 9 out of 10 settings. The difference in results between GOUDA and TraND [47] can be attributed to the fundamental principle on which the TraND approach relies. This approach is centered around an entropy-first neighbor selection strategy focusing solely on the nearest neighbor sample that is likely to share the same view. We believe that due to this reason, TraND is less robust to different views.

5.2. Ablation Study

To demonstrate the efficacy of the proposed method and assess the influence of various key components, we carried out an ablation study using GaitGL with OU-MVLP (source) to CASIA-B (target) setting. The experimental results are reported in Table 3.

Method	CASIA-B → GREW*			OU-MVLP → GREW*		
	R1	R5	R10	R1	R5	R10
UNRN	20.2	32.1	39.0	12.7	22.6	27.9
SECRET	12.3	19.6	24.8	20.1	30.9	37.0
GaitGL	30.0	40.5	46.6	42.1	56.2	63.0
FI2O	37.4	49.2	53.8	47.0	60.2	65.7
GOUDA	35.6	49.9	56.0	50.4	65.0	71.0

Method	CASIA-B → OU-MVLP	OU-MVLP → CASIA-B		
		NM	BG	CL
GaitSet	9.6	74.0	55.5	16.4
TraND	15.8	74.7	56.1	16.2
GOUDA	27.9	87.0	68.0	27.2

Table 2. Comparison with the state-of-the-art UDA gait recognition methods, FI2O [39] and TraND [47], and with UDA methods for person re-identification, UNRN [48] and SECRET [13]. Results are reported by Rank-1 (R1), Rank-5 (R5) and Rank-10 (R10). GOUDA is compared to each method in two different settings. GREW* uses the protocol presented in [39]. Direct testing and domain adaptation results of the UDA gait methods are shown with relevant backbones (as used in the corresponding papers).

Analysis of Various Key Components. We trained various models that emphasize the impact of key components in GOUDA (rows 2-8 in Table 3). Row 2 highlights the impact of incorporating m in the Triplet Selection process. In this experiment, we used $m = 0$ when selecting the negative samples (line 5 in Alg. 1). Interestingly, the results are even lower than those achieved with Direct Testing. This outcome can be attributed to the fact that excluding the margin in the Triplet Selection algorithm leads to a significantly reduced number of valid triplets available for training. Thus, it may potentially result in an overfitted model, as it lacks the necessary diversity and complexity in the training data. In row 3, we demonstrate the effectiveness of incorporating a curriculum learning strategy. In this training process, the model undergoes four stages of triplet selection as well. However, unlike GOUDA, all valid triplets are selected during training ($q = 100\%$ in all four stages).

Loss Components. The models in rows 4-5 were trained without \mathcal{L}_{SSL} and \mathcal{L}_{GOUDA} , respectively, *i.e.* all triplets related to each one of these loss terms were excluded from \mathcal{L} . We observe the significant impact of these loss terms on the results, especially of \mathcal{L}_{GOUDA} , which is the key component of the loss.

Hyperparameters. In row 6, we explore a different Stopping Criteria (SC) for GOUDA. Instead of using 5 neighbors ($K = 5$), we examine the SC with a single nearest neighbor ($K = 1$). In this case, the results are comparable to GOUDA. However, selecting an environment of 5 neighbors might be less prone to noise compared to solely

#	Method	OU-MVLP → CASIA-B		
		NM	BG	CL
1	GaitGL Direct Testing	81.7	71.5	28.8
2	GOUDA w/o m	68.2	43.6	20.3
3	GOUDA w/o curriculum learning	88.2	72.0	36.8
4	GOUDA w/o \mathcal{L}_{SSL}	90.8	76.7	39.0
5	GOUDA w/o \mathcal{L}_{GOUDA}	87.5	72.9	41.3
6	GOUDA w/ SC $K = 1$	92.1	79.6	45.9
7	GOUDA w/ $T_c = 10^\circ$	87.8	70.2	39.8
8	GOUDA w/ $T_s = 20^\circ$	92.1	80.2	46.8
9	Triplet Selection Oracle	96.3	90.3	55.6
10	Supervised fine-tuning	98.3	94.7	86.7
11	GOUDA	92.8	80.9	44.6

Table 3. Ablation experiments with domain adaptation from OU-MVLP (source) to CASIA-B (target). Average Rank-1 accuracies across all 11 views are shown, excluding identical-view cases.

considering the nearest neighbor. Rows 7-8 present experiments with different values of T_c and T_s (defined in lines 2, 3 of Alg. 1), where in each study only one threshold was changed. These results are inferior or comparable to those achieved with the original values.

GOUDA vs. Triplet Selection Oracle. To evaluate the effectiveness of our novel Triplet Selection algorithm, we conducted an experiment comparing GOUDA with a Triplet Selection Oracle (row 9). In this particular experiment, we trained the model using solely the correct triplets from the set of valid triplets mentioned in Algorithm 1. Specifically, the oracle utilizes the target identity labels l and exclusively involves a valid triplet $\{a, p, n\}$ only if $l_a = l_p \neq l_n$. By comparing the results of the oracle with those of GOUDA, it is evident that despite the ability of our method to generate high-quality triplets, there is still a gap to overcome.

Supervised Fine-tuning. In row 10, we present the results of supervised fine-tuning, as an upper bound, where target labels are available. In this experiment, a standard identity triplet loss [14] and cross entropy loss [6] were employed to fine-tune the GaitGL (without adaptation) on the target domain for 50K iterations. It is evident that there is still a large potential remained to achieve fully supervised results, especially in BG and CL scenarios.

6. Analysis

In this section, we analyze our method, its consequences, and limitations.

Curriculum Learning. To retrospectively analyze the effect of curriculum learning (Section 3.3), we evaluate the rate of correct triplets by determining triplets as *correct* if $l_a = l_p \neq l_n$ (based on true labels). Fig. 5 shows that in

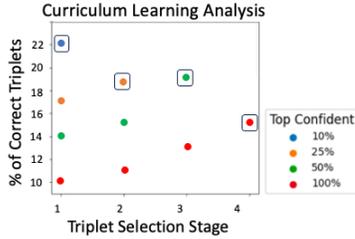


Figure 5. Curriculum Learning Analysis. Triplet Selection is applied 4 times during training (x-axis). Y-axis shows how many of the selected triplets are correct considering person identities. The colors represent the top confident triplets, and in each stage the chosen percentage is marked by a rectangle (see Section 3.3).

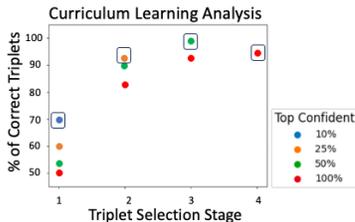


Figure 6. Curriculum Learning Analysis (as described in Fig. 5), using the ground-truth views instead of estimated views. The trend remains, but the percentage of correct triplets is higher.

each stage of Triplet Selection, selecting a smaller number of triplets with the highest confidence scores is better in terms of triplets’ correctness. Moreover, the percentage of correct triplets increases with stages, as the model performance on the target domain improves. Both insights emphasize the importance of learning in an easy-to-hard manner. Although the percentage of correct triplets is approximately 15%-20% (Fig. 5), GOUDA significantly improves direct testing results, due to its ability to reduce the target domain view bias. Higher percentage (up to 90%) is achieved when applying GOUDA using ground-truth views (see Fig. 6), highlighting the importance of quality view estimations.

View Analysis. Here we show the recognition accuracy as a function of view. Let us consider a triplet $\{a, p, n\}$ from our triplet selection algorithm. As described in Section 3.2, p is the closest sample to a out of all *cross view* samples. This results in a limited diversity among the selected triplets, as each view is closely associated with a specific set of views based on its gait features (see Fig. 7). Despite that, we achieve significant improvements on the target domain test set. Fig. 8 presents the accuracy distribution on the test set according to view, before and after applying GOUDA. It is evident that the results improved for most of the cross view combinations, due to the indirect effect of changing the target domain view distribution (brighter colors in the right figure, implying higher accuracies). Furthermore, our method maintains the initial direct testing results

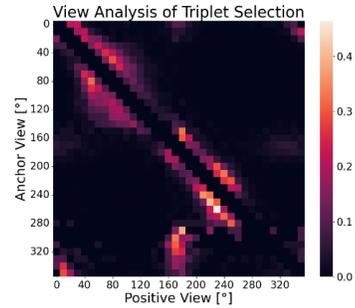


Figure 7. View analysis of Triplet Selection algorithm, presented on GaitGL from GREW (source) to OU-MVLP (target). Rows represent the anchor view and columns the selected positive view. Each row sums up to 1, and the measure presented is the percentage of each view of the selected positive samples. We suffer from a lack of diversity, due to the similarity of views. For example, for anchors with view 260, we mostly choose positives within the view range [210, 240].

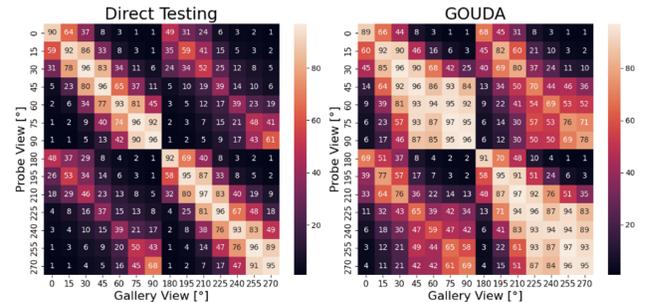


Figure 8. The internal Rank-1 accuracies on the target domain (OU-MVLP) test set. “Direct Testing” is GaitGL pretrained in a supervised manner on the source domain (GREW). Each element in the matrix is the Rank-1 accuracy for the setting in which the probe view is α and the gallery set includes sequences of view β . Brighter color represents higher Rank-1. The final Rank-1 is the average of all Rank-1 accuracies such that $\alpha \neq \beta$. We present improvements in most of these cases after applying GOUDA.

on the exact same view combinations (diagonal), avoiding any degradation effects.

7. Summary

In this paper, we introduce a core problem of Gait Recognition models, often exhibiting a strong bias toward viewing angles when transferred to a new target domain. To reduce the presented gap causing a significant performance drop, we propose GOUDA - an unsupervised domain adaptation method, based on a new triplet selection strategy. We evaluate our method on multiple gait backbones and datasets, showing significant improvements and demonstrating superiority over prior works. We view our work as a first step toward exploring approaches to adapt gait models to target domains by looking at their view distributions.

References

- [1] Weizhi An, Shiqi Yu, Yasushi Makihara, Xinhui Wu, Chi Xu, Yang Yu, Rijun Liao, and Yasushi Yagi. Performance evaluation of model-based gait on multi-view very large population database with pose sequences. *IEEE transactions on biometrics, behavior, and identity science*, 2(4):421–430, 2020. 5
- [2] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48, 2009. 3, 4
- [3] Gabriel C. Bertocco, Fernanda Andaló, and Anderson Rocha. Unsupervised and self-adaptive techniques for cross-domain person re-identification. *IEEE Transactions on Information Forensics and Security*, 16:4419–4434, 2021. 1, 2, 3
- [4] Hanqing Chao, Yiwei He, Junping Zhang, and Jianfeng Feng. Gaitset: Regarding gait as a set for cross-view gait recognition. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 8126–8133, 2019. 1, 2, 5, 6
- [5] Mu Chen, Zhedong Zheng, Yi Yang, and Tat-Seng Chua. Pipa: Pixel-and patch-wise self-supervised learning for domain adaptive semantic segmentation. *arXiv preprint arXiv:2211.07609*, 2022. 1
- [6] Pieter-Tjerk De Boer, Dirk P Kroese, Shie Mannor, and Reuven Y Rubinfeld. A tutorial on the cross-entropy method. *Annals of operations research*, 134:19–67, 2005. 7
- [7] Chao Fan, Saihui Hou, Jilong Wang, Yongzhen Huang, and Shiqi Yu. Learning gait representation from massive unlabelled walking videos: A benchmark. *arXiv preprint arXiv:2206.13964*, 2022. 2
- [8] Chao Fan, Junhao Liang, Chuanfu Shen, Saihui Hou, Yongzhen Huang, and Shiqi Yu. Opengait: Revisiting gait recognition towards better practicality. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9707–9716, 2023. 1, 5
- [9] Chao Fan, Yunjie Peng, Chunshui Cao, Xu Liu, Saihui Hou, Jiannan Chi, Yongzhen Huang, Qing Li, and Zhiqiang He. Gaitpart: Temporal part-based model for gait recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14225–14233, 2020. 1, 5, 6
- [10] Geoffrey French, Michal Mackiewicz, and Mark Fisher. Self-ensembling for visual domain adaptation. *arXiv preprint arXiv:1706.05208*, 2017. 2
- [11] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016. 2
- [12] J. Han and Bir Bhanu. Individual recognition using gait energy image. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(2):316–322, 2006. 2, 6
- [13] Tao He, Leqi Shen, Yuchen Guo, Guiguang Ding, and Zhenhua Guo. Secret: Self-consistent pseudo label refinement for unsupervised domain adaptive person re-identification. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pages 879–887, 2022. 2, 6, 7
- [14] Alexander Hermans, Lucas Beyer, and Bastian Leibe. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*, 2017. 1, 7
- [15] Saihui Hou, Chunshui Cao, Xu Liu, and Yongzhen Huang. Gait lateral network: Learning discriminative and compact representations for gait recognition. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IX*, pages 382–398. Springer, 2020. 2
- [16] Lukas Hoyer, Dengxin Dai, and Luc Van Gool. Hrda: Context-aware high-resolution domain-adaptive semantic segmentation. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXX*, pages 372–391. Springer, 2022. 1
- [17] Xiaohu Huang, Duowang Zhu, Hao Wang, Xinggong Wang, Bo Yang, Botao He, Wenyu Liu, and Bin Feng. Context-sensitive temporal feature learning for gait recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12909–12918, 2021. 2
- [18] Zhen Huang, Dixiu Xue, Xu Shen, Xinmei Tian, Houqiang Li, Jianqiang Huang, and Xian-Sheng Hua. 3d local convolutional neural networks for gait recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14920–14929, October 2021. 2
- [19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proc. Int. Conf. Learn. Represent.*, page 1–41, 2015. 5
- [20] Limin Li and Menglan Cai. Cross-species data classification by domain adaptation via discriminative heterogeneous maximum mean discrepancy. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 18(1):312–324, 2019. 2
- [21] Shihua Li, Mingkuan Yuan, Jie Chen, and Zhilan Hu. Adadc: Adaptive deep clustering for unsupervised domain adaptation in person re-identification. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(6):3825–3838, 2022. 1, 2
- [22] Wenhao Li, Hong Liu, Hao Tang, Pichao Wang, and Luc Van Gool. Mhformer: Multi-hypothesis transformer for 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13147–13156, 2022. 3
- [23] Xiang Li, Yasushi Makihara, Chi Xu, Yasushi Yagi, and Mingwu Ren. Gait recognition via semi-supervised disentangled representation learning to identity and covariate features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13309–13319, 2020. 2
- [24] Xiang Li, Yasushi Makihara, Chi Xu, Yasushi Yagi, Shiqi Yu, and Mingwu Ren. End-to-end model-based gait recognition. In *Proceedings of the Asian conference on computer vision*, 2020. 2
- [25] Yanghao Li, Naiyan Wang, Jianping Shi, Xiaodi Hou, and Jiaying Liu. Adaptive batch normalization for practical domain adaptation. *Pattern Recognition*, 80:109–117, 2018. 2

- [26] Junhao Liang, Chao Fan, Saihui Hou, Chuanfu Shen, Yongzhen Huang, and Shiqi Yu. Gatedge: Beyond plain end-to-end gait recognition for better practicality. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part V*, pages 375–390. Springer, 2022. 2
- [27] Rijun Liao, Shiqi Yu, Weizhi An, and Yongzhen Huang. A model-based gait recognition method with body pose and human prior knowledge. *Pattern Recognition*, 98:107069, 2020. 2
- [28] Beibei Lin, Shunli Zhang, and Feng Bao. Gait recognition with multiple-temporal-scale 3d convolutional neural network. In *Proceedings of the 28th ACM international conference on multimedia*, pages 3054–3062, 2020. 2
- [29] Beibei Lin, Shunli Zhang, and Xin Yu. Gait recognition via effective global-local feature representation and local temporal aggregation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14648–14656, 2021. 1, 2, 5, 6
- [30] Xiaofeng Liu, Chaehwa Yoo, Fangxu Xing, Hyejin Oh, Georges El Fakhri, Je-Won Kang, Jonghye Woo, et al. Deep unsupervised domain adaptation: A review of recent advances and perspectives. *APSIPA Transactions on Signal and Information Processing*, 11(1), 2022. 2
- [31] Ioana Macoveciuc, Carolyn J Rando, and Hervé Borrión. Forensic gait analysis and recognition: standards of evidence admissibility. *Journal of forensic sciences*, 64(5):1294–1303, 2019. 1
- [32] Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018. 2
- [33] Yang Ran, Qinfen Zheng, Rama Chellappa, and Thomas M Strat. Applications of a simple characterization of human gait in surveillance. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 40(4):1009–1020, 2010. 1
- [34] Fangmin Sun, Weilin Zang, Raffaele Gravina, Giancarlo Fortino, and Ye Li. Gait-based identification for elderly users in wearable healthcare systems. *Information fusion*, 53:134–144, 2020. 1
- [35] Noriko Takemura, Yasushi Makihara, Daigo Muramatsu, Tomio Echigo, and Yasushi Yagi. Multi-view large population gait dataset and its performance evaluation for cross-view gait recognition. *IPSJ transactions on Computer Vision and Applications*, 10:1–14, 2018. 1, 5
- [36] Torben Teepe, Ali Khan, Johannes Gilg, Fabian Herzog, Stefan Hörmann, and Gerhard Rigoll. Gaitgraph: Graph convolutional network for skeleton-based gait recognition. In *2021 IEEE International Conference on Image Processing (ICIP)*, pages 2314–2318. IEEE, 2021. 2
- [37] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 2
- [38] Chen Wang, Junping Zhang, Liang Wang, Jian Pu, and Xiaoru Yuan. Human identification using temporal information preserving gait template. *IEEE transactions on pattern analysis and machine intelligence*, 34(11):2164–2176, 2011. 2
- [39] Likai Wang, Ruize Han, Wei Feng, and Song Wang. From indoor to outdoor: Unsupervised domain adaptive gait recognition. *arXiv preprint arXiv:2211.11155*, 2022. 1, 3, 4, 5, 6, 7
- [40] Pengfei Wang, Changxing Ding, Wentao Tan, Mingming Gong, Kui Jia, and Dacheng Tao. Uncertainty-aware clustering for unsupervised domain adaptive object re-identification. *IEEE Transactions on Multimedia*, pages 1–1, 2022. 1, 2
- [41] Garrett Wilson and Diane J Cook. A survey of unsupervised deep domain adaptation. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 11(5):1–46, 2020. 2
- [42] Tongkun Xu, Weihua Chen, Pichao Wang, Fan Wang, Hao Li, and Rong Jin. Cdtrans: Cross-domain transformer for unsupervised domain adaptation. *arXiv preprint arXiv:2109.06165*, 2021. 1
- [43] Wenqian Xue, Ming Yang, Rujie Liu, Yamamoto Takuma, Yoshioka Takahiro, and Konno Takeshi. Spatial-temporal graph convolutional network for skeleton-based gait recognition. In *2022 3rd International Conference on Pattern Recognition and Machine Learning (PRML)*, pages 77–82, 2022. 2
- [44] Shiqi Yu, Daoliang Tan, and Tieniu Tan. A framework for evaluating the effect of view angle, clothing and carrying condition on gait recognition. In *18th International Conference on Pattern Recognition (ICPR'06)*, volume 4, pages 441–444, 2006. 1, 5
- [45] Zhaoxiang Zhang. A gait recognition method for a moving target image in sports based on a decision tree. *International Journal of Biometrics*, 13(2-3):165–179, 2021. 1
- [46] Jinkai Zheng, Xinchun Liu, Wu Liu, Lingxiao He, Chenggang Yan, and Tao Mei. Gait recognition in the wild with dense 3d representations and a benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20228–20237, June 2022. 5
- [47] Jinkai Zheng, Xinchun Liu, Chenggang Yan, Jiyong Zhang, Wu Liu, Xiaoping Zhang, and Tao Mei. Trand: Transferable neighborhood discovery for unsupervised cross-domain gait recognition. In *2021 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–5. IEEE, 2021. 1, 3, 4, 6, 7
- [48] Kecheng Zheng, Cuiling Lan, Wenjun Zeng, Zhizheng Zhang, and Zheng-Jun Zha. Exploiting sample uncertainty for domain adaptive person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 3538–3546, 2021. 2, 6, 7
- [49] Haidong Zhu, Zhaocheng Zheng, and Ram Nevatia. Gait recognition using 3-d human body shape inference. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 909–918, 2023. 2
- [50] Jinjing Zhu, Haotian Bai, and Lin Wang. Patch-mix transformer for unsupervised domain adaptation: A game perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3561–3571, 2023. 1
- [51] Zheng Zhu, Xianda Guo, Tian Yang, Junjie Huang, Jiankang Deng, Guan Huang, Dalong Du, Jiwen Lu, and Jie Zhou.

Gait recognition in the wild: A benchmark. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14789–14799, October 2021. 5