

# Learning Gait Representation from Massive Unlabelled Walking Videos: A Benchmark

Chao Fan, Saihui Hou, Jilong Wang, Yongzhen Huang, and Shiqi Yu

**Abstract**—Gait depicts individuals' unique and distinguishing walking patterns and has become one of the most promising biometric features for human identification. As a fine-grained recognition task, gait recognition is easily affected by many factors and usually requires a large amount of completely annotated data that is costly and insatiable. This paper proposes a large-scale self-supervised benchmark for gait recognition with contrastive learning, aiming to learn the general gait representation from massive unlabelled walking videos for practical applications via offering informative walking priors and diverse real-world variations. Specifically, we collect a large-scale unlabelled gait dataset GaitLU-1M consisting of 1.02M walking sequences and propose a conceptually simple yet empirically powerful baseline model GaitSSB. Experimentally, we evaluate the pre-trained model on four widely-used gait benchmarks, CASIA-B, OU-MVLP, GREW and Gait3D with or without transfer learning. The unsupervised results are comparable to or even better than the early model-based and GEI-based methods. After transfer learning, GaitSSB outperforms existing methods by a large margin in most cases, and also showcases the superior generalization capacity. Further experiments indicate that the pre-training can save about 50% and 80% annotation costs of GREW and Gait3D. Theoretically, we discuss the critical issues for gait-specific contrastive framework and present some insights for further study. As far as we know, GaitLU-1M is the first large-scale unlabelled gait dataset, and GaitSSB is the first method that achieves remarkable unsupervised results on the aforementioned benchmarks. The source code of GaitSSB and anonymous data of GaitLU-1M is available at <https://github.com/ShiqiYu/OpenGait>.

**Index Terms**—Gait Recognition, Self-Supervised, Contrastive Learning, GaitSSB, GaitLU-1M.

## 1 INTRODUCTION

GAIT recognition aims to capture human walking patterns by computer vision technologies for individual identification. Compared with other vision-based biometrics, *e.g.*, face, fingerprint and iris, gait is hard to disguise and can be easily depicted at a long distance in non-intrusive ways without cooperative subjects. These advantages make gait recognition particularly suitable for various security applications, such as suspect tracking, crime prevention, identity verification, and so on [1].


Recently, with the rising of deep learning [2], the gait recognition community has continuously attracted increasing interest from both academic and industrial fields since the promising performance achieved on multiple popular gait benchmarks, such as CAISA-B [3] and OU-MVLP [4]. However, evidence from [5] and [6] indicate that existing methods [7], [8], [9] are still far away from practical applications. One of the primary reasons is that almost all the gait datasets are collected by fixed cameras under a fully-controlled environment, *e.g.*, indoor laboratory or

specific areas in the wild, and therefore struggle to simulate the fully-unconstrained challenges, such as diverse camera view and height, complex occlusion and background, various carrying and dressing, unpredictable illumination and weather, *et al.* One straightforward solution is collecting a large amount of real gait data to fill the massive gap between theory and practice. However, large-scale gait data annotation in the wild is particularly expensive. Specifically, unlike the face and pedestrian that can be easily identified by just one image, the gait characteristics are generally represented as the combination of static appearance and dynamic movement patterns inherent in the walking videos, thus requiring much time and effort for annotations. More importantly, as shown in Fig. 1(a), if there are no other apparent soft-biometric features or auxiliary information, such as gender, age, dress, location, background, and so on, it would be almost impossible for an annotator to identify individuals from massive unlabelled videos.

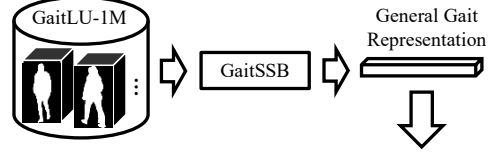
Nowadays, the fact becomes that, despite gait data with manual labels being costly, there are massive unlabelled walking videos on the web and apps, such as the prevailing YouTube and TikTok. Furthermore, pre-training on unlabelled data by self-supervised approaches has been believed to effectively explore the data-specific prior knowledge so as to alleviate the 'out-of-distribution' problem. These advantages make self-supervised methods highly successful in both natural language processing [10], [11] and computer vision fields in recent years. Specifically, the impressive BERT [11] can be fine-tuned with just one additional output layer to achieve state-of-the-art for eleven language tasks. For visual tasks, reducing the distance between representations of two different augmented views of the same image (positive pairs), and meanwhile increasing the dis-

- Chao Fan and Shiqi Yu are with Department of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen 518055, China. E-mail: 12131100@mail.sustech.edu.cn and yusq@sustech.edu.cn.
- Saihui Hou and Yongzhen Huang are with School of Artificial Intelligence, Beijing Normal University and also with Watrix Technology Limited Co. Ltd., China. E-mail: {huangyongzhen, housaihui}@bnu.edu.cn.
- Jilong Wang is with the University of Science and Technology of China and also with Center for Research on Intelligent Perception and Computing (CRIPAC), National Laboratory of Pattern Recognition (NLPR), Institute of Automation, Chinese Academy of Sciences (CASIA). Email: jilongw@mail.ustc.edu.cn.
- Corresponding author: Shiqi Yu.

Manuscript received June 08, 2022.



(a) Gait data annotation in the wild.



Unsupervised Performance Under Cross-View Conditions				
Accuracy	CASIA-B(NM)	OU-MVLP	GREW	Gait3D
Rank@1	83.8	37.2	16.6	24.7

(b) The proposed self-supervised benchmark for gait recognition.

Fig. 1. (a) It is generally hard for an annotator to identify individuals according to their walking patterns without the help of extra attributes in practical scenarios. (b) Our prospective and comprehensive self-supervised benchmark with contrastive learning for gait recognition, including a large-scale unlabelled gait dataset GaitLU-1M and a conceptually simple yet empirically powerful baseline model GaitSSB.

tance between representations of different images (negative pairs) [12], *i.e.*, the contrastive learning framework has been steady progress in self-supervised representation learning, with encouraging results on multiple visual tasks [12], [13], [14], [15]. However, existing gait recognition methods are still immersed in relatively small-scale supervised datasets. More importantly, we thus lack a well pre-trained model providing knowledgeable walking priors obtained by dealing with substantial real-world factors.

To address these issues, we develop a pioneering and comprehensive self-supervised benchmark to learn the general gait representation from massive unlabelled walking videos. Firstly, a new dataset is required. Hence we collect a **Large-scale Unlabelled** gait dataset and name it as **GaitLU-1M**. Inspired by [16], [17], [18], GaitLU-1M is extracted from 1.6K publicly available videos and consists of 1.02 million gait sequences, whose training set is about 10 times larger than that of the largest gait datasets, *e.g.*, OU-MVLP [4] and GREW [5]. Moreover, the collected videos are filmed by different cameras worldwide, covering a wide range of capturing scenes, individual attributes, photography conditions, *etc.*

In addition, we introduce contrastive learning, one of the most popular visual self-supervised methods, to gait recognition and point out two gait-specific issues. The first is about *input modality*. Human gait patterns captured by walking videos involve dynamic motion information along the temporal dimension, which differs a lot from a single RGB image, indicating that some crucial and well-explored data augmentation strategies for common contrastive framework [12] are no longer applicable to gait recognition. In this paper, we propose a silhouette augmentation operator  $\pi(\cdot)$  and find that the gait-specific augmentation strategy is vital for contrastive gait recognition. Another issue is about the *task granularity*. The phenomenon that two different images from semantically near categories own relatively lower visual distance and vice versa has been considered helpful for contrastive framework [19]. However, gait recognition is a fine-grained identification task aiming to identify individuals, whose one of the most prominent features is that intra-class distances are often larger than inter-class, especially for cross-view cases. Through experimental and theoretical analysis, we find that the negative samples play an essential role in learning view-invariant gait characteristics. Taking the above insights together, we propose a conceptually simple and empirically powerful model, **GaitSSB**, to build a

comprehensive **Self-Supervised Benchmark** for gait. Overall, this paper makes the following contributions.

- We successfully leverage the massive unlabelled walking videos to become beneficial for gait recognition. Specifically, we propose a pioneering and comprehensive self-supervised benchmark with contrastive learning, including a large-scale unlabelled gait dataset GaitLU-1M and a conceptually simple yet empirically powerful model GaitSSB.
- As far as we know, GaitSSB is the first method achieving noteworthy unsupervised performance on four of the most popular supervised gait benchmarks, *i.e.*, 83.8%, 35.0%, 16.6% and 24.7% rank-1 average accuracy under cross-view conditions on CASIA-B (normal walking), OU-MVLP, GREW and Gait3D, respectively.
- With transfer learning, GaitSSB outperforms existing methods by a large margin on two practical gait datasets (GREW and Gait3D). On two indoor gait datasets (OU-MVLP and CASIA-B), GaitSSB can also achieve state-of-the-art or fully comparable performances. Moreover, under the challenging cross-domain settings, GaitSSB further exhibits its superior capacity of generalization. Experiments also show that the pre-training can save about 50% and 80% annotation costs of GREW and Gait3D.
- We discuss the crucial issues for introducing the contrastive framework into gait recognition from both experimental and theoretical perspectives, convincingly boosting the further explorations of learning the general gait representation for practical gait recognition applications.

In the remaining part of this paper, after introducing more related works on gait recognition and contrastive learning in Sec. 2, we will present the process of building GaitLU-1M and show the details of GaitSSB in Sec. 3. Experimental results, theoretical discussions and brief conclusions are respectively given in Sec. 4, Sec. 5 and Sec. 6. In Sec. 7, we state the concerns about privacy protection and biometrics abuse.

## 2 RELATED WORKS

### 2.1 Gait Recognition

Almost all the existing gait recognition methods are based on supervised learning with labeled datasets, and can be

generally grouped into two distinct categories, involving the model-based and appearance-based methods.

*Model-based* methods [20], [21], [22], [23], [24] aim to construct the underlying structure of the human body first and then extract gait features for individual identification. Thanks to its compactness, the skeleton data is a kind of human body model usually used for gait representation learning [24], [25], [26]. In the latest literature, the human mesh that preserves rich 3D characteristics is becoming increasingly popular [27], [28], [29], [30]. Due to the nature of directly representing the dynamic changes of body parts and being robust to real-world factors like dressing and carrying, these body models that exclude the visual clues by merely maintaining the structural information are considered a clean representation of gait in theory. In practice, lacking the description of body shape usually makes the performance of skeleton-based methods relatively unsatisfactory [24], [25], [26]. On the other hand, the mesh-based methods representing human body surfaces have achieved more competitive recognition accuracy [27], [28], [29], [30].

*Appearance-based* methods usually extract the gait features from silhouettes directly and thus benefit from the informative shape features. According to the input class, these methods can be roughly divided into two categories: template-based [31], [32], [33], [34], [35], [36] and sequence-based [7], [8], [9], [37], [38], [39], [40], [41], [42], where the former usually compresses the gait sequence into one template along the temporal dimension while the latter tends to regard the walking video as an ordered sequence [8], [9], [42] or an unordered set [7]. With the rising of deep learning, this kind of method mainly focuses on the design of backbone structures (*e.g.*, set block proposed by [7], focal convolution proposed by [8], lateral connection proposed by [37], and 3D global-local convolution proposed by [9], [39].), the neck of network (*e.g.*, horizontal pyramid pooling introduced by [7], micro-motion module proposed by [8], generalized-mean pooling introduced by [9]), and the head of network (*e.g.*, triplet loss introduced by [7] and quintuplet loss proposed by [36]).

Additionally, there raise a series of end-to-end based gait recognition methods directly taking RGB images as input [29], [42], [43], [44]. They usually eliminate the gait-unrelated features by carefully designing the intermediate modality, such as float mask [43], SMPL model [29] and synthetic silhouettes [44]. Despite achieving impressive performances, they have not attracted enough attention due to the lack of public RGB-based gait datasets.

Besides, like other vision tasks, the development of gait datasets has greatly promoted gait recognition research. Specifically, CASIA-B [3] and OU-MVLP [4] are two of the most widely-used gait datasets. They were proposed in 2006 and 2018, respectively, and captured by requiring the subjects to walk around the laboratory, differing significantly from the real-world scenarios. With facing more practical applications, GREW [5] and Gait3D [45] are collected in the wild in 2021 and 2022 respectively. Though being collected from outdoor areas, they are still inferior in both scale and diversity compared to other biometrics recognition datasets, *e.g.*, MegaFace [46] for face recognition.

## 2.2 Contrastive Methods for Self-Supervised Learning

The success of self-supervised learning is mainly owing to its ability to explore valuable prior knowledge from massive unlabelled data. The contrastive learning [47], whose core idea is to attract two different augmented views<sup>1</sup> of the same image and simultaneously enlarge the distance between representations of two different images, has been popularized for self-supervised representation learning due to its state-of-the-art performance and excellent capacity of generalization.

In practice, current contrastive learning methods mainly involve several aspects of designs, *i.e.*, data augmentation to form positive pairs, various encoders to embed input images, and pretext task as a training goal for learning good data representations. Specifically, Chen *et al.* demonstrate that the composition of data augmentations plays a critical role in contrastive learning, and most of the follow-on works agree with this conclusion [12], [15], [48], [49]. For the encoder type, with mostly taking ResNet [50] or ViT [51] as backbone, some works [14], [15] employ the widely-used siamese networks [52] to compare entities, and others tend to maintain an extra momentum memory bank to remain massive negative samples [12], [13], [48] or clustering centers [49] to build the large and informative training batch.

Pretext tasks, where *pretext* implies that the solved task is not of genuine interest but is solved only for the true purpose of learning a good representation [15], can distinguish the categories of self-supervised methods clearly, *e.g.*, pixel-level image reconstruction for denoising autoencoder [53], patch-level image reconstruction for masked autoencoder [54], and language probability model for word vector embedding [55]. There are many pretext tasks usually employed in contrastive framework, *e.g.*, context auto-encoding in [56] and contrastive multiview coding in [57].

## 2.3 Gait Recognition with Contrastive Learning

Before the boost of deep learning, there are many classical works extracting the gait representation via unsupervised approaches, such as Principal Components Analysis (PCA) [58] and Canonical Analysis [59]. The main focus of them is to capture the valuable features used as the input of the downstream gait pattern classifier, rather than directly pre-train the model on massive unlabelled walking videos. In the latest literature, some works [60], [61] have mentioned the idea of contrastive learning but have not performed significantly better than other progressive methods. Besides, instead of the large-scale practical unlabelled datasets, these works [60], [61] are developed on the relatively small-scale and fully-constrained CAISA-B [3] or OU-MVLP [4]. Hence, they are believed to deviate from the primary intention of pre-training, that is to explore the prior knowledge from massive unlabelled walking videos for the excellent capacity of representation and generalization.

In this paper, we constructively point out two gait-specific issues distinct from canonical contrastive frameworks, involving input modality (sequence of silhouettes *v.s.* RGB image) and task granularity (fine-grained recognition

1. A view generally means a glance at the augmented image in contrastive learning, but indicates the camera angle in gait recognition.




Fig. 2. (a) shows the example of collected public videos worldwide by different devices. (b) and (c) present the statistics of GaitLU-1M.

TABLE 1

The comparison between widely-used gait datasets with GaitLU-1M. #Id, #Seq and #Cam refer to the count of identities, sequences and cameras, respectively. NM, BG and CL indicate normal, carrying and dressing factors respectively.  $\approx$  indicates the approximate number.

Dataset	Source	Labelled	#Id	#Seq	#Cam	Environment	Viewpoint	Variations
CASIA-B	ICPR2006	Yes	124	13,640	11	laboratory	fixed	NM, BG, CL
OU-MVLP	CVA2018	Yes	10,307	288,596	14	laboratory	fixed	NM
GREW	ICCV2021	Yes	26,345	128,671	882	particular area	free	real world
Gait3D	CVPR2022	Yes	4,000	25,309	39	supermarket	free	real world
GaitLU-1M	ours	No	1,035,309 $\approx$	1,035,309	1,379	open world	free	real world

*v.s.* common object classification). These new issues make it uncertain whether contrastive learning would still work well for gait recognition. Through careful designs, we build a pioneer and comprehensive benchmark for contrastive gait recognition, including a large-scale unlabelled gait dataset GaitLU-1M and a conceptually simple yet empirically powerful model GaitSSB.

## 2.4 Contrastive Learning on Related Tasks

Here we discuss two related tasks, *i.e.*, person re-identification and spatiotemporal representation learning.

*Person Re-ID.* Similar to gait recognition, person Re-ID aims to find the target person presented in CCTV cameras. Recently, SpCL [62] fine-tunes the network using pseudo labels generated from reliable clustering results trained with contrastive loss. Fu *et al.* [63] present a large-scale unlabelled person re-identification dataset extracted from public videos and discover some unique factors specific to Re-ID task, including that color distortion is harmful, Random Erasing [64] is needed, and a proper temperature parameter for contrastive loss is important.

*Spatiotemporal representation learning.* Feichtenhofer *et al.* [65] present a large-scale study on contrastive spatiotemporal representation learning and find that temporally-persistent features in the same video work surprisingly well across different contrastive frameworks, pre-training datasets, downstream datasets and encoder architectures.

Gait recognition differs from the above-related tasks in both the input modality, *i.e.*, sequence of silhouettes *v.s.* a single or sequence of RGB images, and task granularity, *i.e.*, fine-grained biometric recognition *v.s.* coarse-grained human appearance or action recognition. But this paper gets some consistent conclusions from them, *e.g.*, gait-specific data augmentation strategies are necessary, and consistency features along temporal dimensions are crucial.

## 3 PROPOSED BENCHMARK

### 3.1 GaitLU-1M

Collecting a large-scale unlabelled gait dataset depicting various practical factors is a prerequisite for constructing the self-supervised gait benchmark. However, existing gait datasets either capture the walking videos in an indoor laboratory, *e.g.*, CASIA-B [3] and OU-MVLP [4], or in particular areas in the wild, *e.g.*, GREW [5] and Gait3D [45], struggling from simulating the fully open real-world challenges. Moreover, the scale of gait datasets is also far less than that of related fields, *e.g.*, MegaFace [46] that consists of millions of images used for face recognition.

In this work, we build a Large-scale Unlabelled gait dataset, **GaitLU-1M**. It consists of over one million walking sequences averaging 92 frames, namely about 92M silhouettes, 76 times larger than the number of images in the famous ImageNet-1k (1.2M images) [66]. Inspired by [16], [17] and [18], GaitLU-1M is extracted from public videos shot around the world, making it cover a wide range of variations, involving the capturing scenes (streets, subway and airport, *etc.*), individual attributes (gender, age and race, *etc.*), photography conditions (camera view, resolution, illumination and weather, *etc.*), and so on. This great diversity and scale offer an excellent chance to learn general gait representation in a self-supervised manner.

#### 3.1.1 Data Collection and Processing

By using queries like "city name + walking", we choose 16 big cities around the world. There are 3 American cities (New York, Sao Paulo and Mexico City), 3 European cities (London, Paris and Moscow), 9 Asian cities (Beijing, Hong Kong, Shanghai, Tokyo, Bangkok, Delhi, Mumbai, Kuala Lumpur and Dubai) and 1 African city (Cairo). The top 100 searched videos will be collected for each city, that is, we crawled about 1.6K videos in total. Besides, the videos appearing with no person in the first 2 minutes are skipped for efficiency, and 221 videos are ignored in this way.



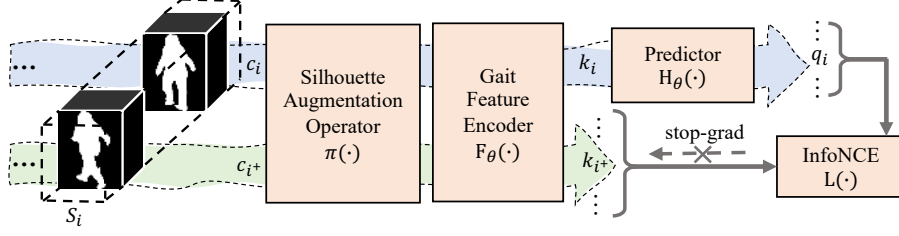


Fig. 3. The overall of GaitSSB.  $c_i^+$  and  $c_i$  denotes two non-overlap clips of the gait sequence  $S_i$  where  $i \in \{1, \dots, n\}$ . And  $\pi(\cdot)$ ,  $F_\theta(\cdot)$ ,  $H_\theta(\cdot)$  and  $L(\cdot)$  indicates the silhouette augmentation operator, gait feature encoder, predictor and InfoNCE loss, respectively.

Then the pedestrian tracking algorithm, BYTETracker [67] is performed to provide initial sequences of human bounding boxes. Since most videos are filmed in crowded public places, such as streets, squares, airports, and subway stations, we obtain over 3M raw walking sequences even if the average duration of collected videos is just a half hour. However, we find a significant portion of the gained boxes suffers from severe occlusion because of being in crowds and not looking like the human appearance. Therefore, we discard those small bounding boxes in each image if they overlap others over an IOU-measured threshold for data cleaning. After performing the pedestrian segmentation [68], [69] frame by frame, we get 1,035,309 walking sequences containing 92,627,263 silhouettes. Some example videos are given in Fig. 2 (a), showing that GaitLU-1M has pretty diverse variations, such as camera angle, background and individual race. *To protect personal privacy, the obtained gait videos are anonymous with random index and exhibited as the sequences of binary silhouette and body skeleton. The pre-processing is based on automatic algorithms, meaning that nobody has manually processed personal visual data. The raw data, including website links, RGB videos and other intermediate results, have been deleted.*

### 3.1.2 Statistics and Comparison

Statistics about GaitLU-1M are illustrated in Fig. 2 (b) and (c). Among them, Fig. 2 (b) shows the amounts of sequences collected from different cities, and more sequences can be collected from the cities with relatively large populations. Fig. 2 (c) tells over 80% of gait sequences contain over 25 frames, namely including about one gait cycle at least.

Compared with existing gait datasets, such as CASIA-B [3], OU-MVLP [4], GREW [5] and Gait3D [45], GaitLU-1M is superior in the following aspects.

a. *Scale*. GaitLU-1M makes one of the first attempts toward large-scale self-supervised learning for gait recognition, and its size is completely comparable to other visual tasks using contrastive methods, *e.g.*, ImageNet [66] for image classification.

b. *Diversity*. Since the chosen 1,379 videos are filmed around the world, varying from the city, place to time, it is almost impossible for the same person appears in two different videos in GaitLU-1M. Hence we can regard the 1M walking sequences as 1M identities, implying that GaitLU-1M possesses millions of individual walking patterns. Similarly, the 1,379 videos can be considered captured from 1,379 different cameras, varying from device to photographer. These diversities offer excellent potential for our baseline model GaitSSB to learn the general gait representation.

c. *Massive Variations*. Though the latest Gait3D [45] and GREW [5] are collected in complex environments (supermarkets and other particular areas), we consider that their diversities of variations are still less than GaitLU-1M since the latter is captured in worldwide practical scenarios.

Besides, two limitations of GaitLU-1M are worth mentioning: a. Since all the gait sequences are anonymous with the random index, GaitLU-1M has no weak annotations like filming cities. Also, the representations of silhouette and skeleton lack visual clues, making it hard for further attribute annotations. b. Most videos are captured by handheld phones/cameras, which may cause the view distribution gap between GaitLU-1M and typical CCTVs, especially for tilt and pitch angles. In Sec. 3.2.2, the perspective transformation strategy can alleviate this issue to some extent.

## 3.2 GaitSSB

Two gait-specific issues are making GaitSSB distinct from other contrastive frameworks, *i.e.*, the input modality and task granularity. The first one requires a new data augmentation strategy concentrating on gait silhouette sequences, and the latter reminds us to learn the subtle yet discriminative gait's unique features. In this subsection, we first describe the high-level design of GaitSSB. Then, we present more details involving the silhouette augmentation operator and gait feature encoder. Moreover, we also introduce the prevailing predictor and stop-gradient trick. In the end, we compare GaitSSB with other contrastive frameworks from the perspectives of fine-grained and view-invariant gait representation learning.

### 3.2.1 Description of GaitSSB

Our goal is to learn the general gait representation without labeling samples. Inspired by successful contrastive learning, the core idea is to reduce the distance between two different augmented views of the same gait sequence (positive pairs) and simultaneously repulse the negative pairs, where the assumption is that data augmentation can simulate the practical variations to some extent without damaging identity information.

As shown in Fig. 3, we sample two of the non-overlap clips ( $c_i$  and  $c_{i+}$ ) in the sequence  $S_i$  as two augmented views of identical walking patterns, where  $i \in \{1, \dots, n\}$ . Hence, there are  $2n$  gait clips in a training batch with  $n$  positive pairs. After being randomly augmented by silhouette augmentation operator (SAO), these clips will be fed into the gait feature encoder  $F_\theta(\cdot)$  to extract the feature vectors, that is  $k_i = F_\theta(c_i)$  and  $k_{i+} = F_\theta(c_{i+})$ . Notably, the encoder  $F_\theta(\cdot)$

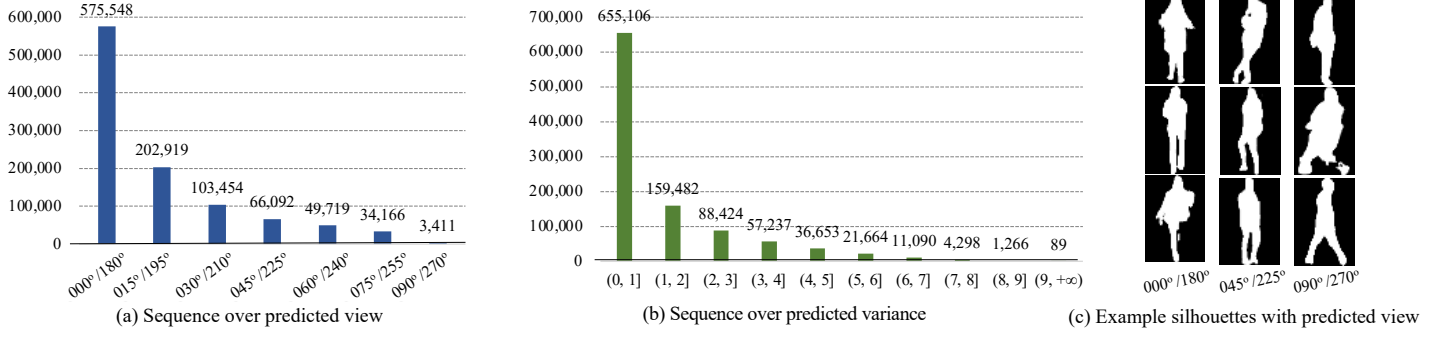


Fig. 4. (a) shows the statistics of sequences number over the predicted views. (b) shows the statistics of sequences number over the predicted variances. (c) shows example silhouettes with the predicted view.




Fig. 5. Example silhouettes with spatial augmentation. From left to right: raw image, horizontal flipping, affine transformation including rotation, perspective transformation and dilation.

is a part-based model, meaning that the output  $k_i$  is a group of parallel feature vectors responding to different parts of the human body. Since these part vectors will be processed independently, the following formulation loosely treats  $k_i$  as a single feature vector for brevity.

Then, the predictor head, denoted as  $H_\theta(\cdot)$ , transforms the output of one view and matches it to the another view, denoted as  $q_i = H_\theta(k_i) \rightarrow k_{i+}$ . Consider  $(q_i, k_{i+})$  as the positive pair and  $(q_i, k_{j+})$  with  $j \neq i$  as negative pairs, we formulate the InfoNCE loss [70]:

$$L(q_i, k_{i+}) = -\log \frac{\exp(\text{sim}(q_i, k_{i+})/\tau)}{\sum_{j=1}^n \exp(\text{sim}(q_i, k_{j+})/\tau)} \quad (1)$$

where  $\text{sim}(q_i, k_{j+}) = q_i^\top k_{j+} / (\|q_i\|_2 \|k_{j+}\|_2)$  is the cosine similarity, and  $\tau > 0$  denotes the temperature hyper-parameter set to 16.

Following [15], we use a symmetrized loss as:

$$\mathcal{L} = \frac{1}{2} L(q_i, k_{i+}) + \frac{1}{2} L(q_{i+}, k_i) \quad (2)$$

where  $q_{i+} = H_\theta(k_{i+})$ . This loss is defined for each gait clip, and the total loss is averaged over all clips in the training batch. Moreover, the stop-gradient operation alternatively treats  $k_i$  and  $k_{i+}$  as constant to prevent collapsing solutions [15]. Fig. 3 only shows the latter case for brevity.

### 3.2.2 Silhouette Augmentation Operator

Data augmentation plays a vital role in contrastive framework because it can partially simulate the practical variations without damaging identity information. However,

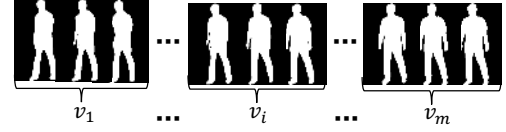


Fig. 6. A gait sequence may involve several clips captured from visually different camera angles.

since our input modality is the sequence of silhouettes, the well-explored data augmentation methods mostly centralized in RGB images no longer apply to our gait contrastive framework. In this paper, we propose a novel silhouette augmentation operator (**SAO**,  $\pi(\cdot)$ ) with making following three-fold effects: spatial, intra-sequence and sampling augmentations.

*Spatial Augmentation.* Performed on each frame, spatial augmentation simulates clothing/carrying changing factors. Here we list several common operations centralizing in RGB image augmentation and discuss whether they still work for silhouette data.

- *Inapplicable: Color distort.* The silhouettes are binary masks with no RGB information.
- *Unnecessary: Crop and resize, and Cutout.* Raw videos are detected and segmented frame by frame, suggesting GaitLU-1M dominates enormous natural occlusion and detection box drift.
- *Available: Flip, Rotate, Affine and Perspective.* These transformations with appropriate settings are still indispensable for silhouette augmentation.

Based on the above analysis, we adopt several augmentation methods, including horizontal flipping, slight rotation, affine and perspective transformations. In addition, we introduce an extra dilation operator trying to simulate the most challenging clothing-changing factors. Its core idea is to dilate a random consecutive region of the human body by randomly choosing structuring elements and kernel size, making the appearance look like wearing an additional coat. Fig. 5 shows the instances illustrating the spatial augmentation used for GaitSSB.

*Intra-sequence Augmentation.* As mentioned in Sec. 3.2.1, we consider two non-overlap silhouette clips inside the same walking sequence as two augmented views of identical walking patterns. This strategy enables the model to learn gait feature consistency hidden in the temporal dimension [65], [71], [72], [73]. Ablation studies in Sec. 4 and

discussions in Sec. 5 will declare the significance of temporal consistencies for the self-supervised gait pre-training.

*Sampling Augmentation.* We observe that frontal-view gait sequences possess almost no viewpoint transitions while others own relatively remarkable viewpoint changes inside the sequence. Hence, to enhance the learning of view-invariant gait features, we propose a sampling strategy to choose these gait sequences possessing view-changing factors with a higher probability.

As shown in Fig. 6, consider all the clips extracted from a walking sequence, we let the average of their camera angle,  $\bar{v} = \frac{\sum_{i=1}^m v_i}{m}$ , present the view of the entire sequence, and use the variance to quantitatively measure the degree of view transitions inside this sequence, denoted as  $\sigma^2 = \frac{\sum_{i=1}^m (v_i - \bar{v})^2}{m-1}$ . The larger  $\sigma^2$  indicates the more salient view changes and vice versa. Specifically, a small multiple layers perceptron composed of a stack of four fully-connected layers (equipped with the batch normalization and ReLU activation layers), and an extra classification layer is trained on OU-MVLP to serve as the view classifier. We use the softmax loss with label smoothing to drive the training process. It is worth noting that we tag some view pairs close to being horizontally symmetric with the same label to reduce the difficulty of convergence caused by the highly similar appearances. Specifically, we tag both the original  $000^\circ$  and  $180^\circ$  as #0,  $015^\circ$  and  $195^\circ$  as #1, and so on<sup>2</sup>.

Consequently, we can get more statistical analysis about GaitLU-1M. As shown in Fig. 4 (a), we observe that about one-half sequences in GaitLU-1M are captured in frontal-view cases, *i.e.*, those sequences under  $000^\circ$ ,  $015^\circ$ ,  $180^\circ$  and  $195^\circ$  camera angles. On the other hand, there are over one-half of sequences in GaitLU-1M that own almost no view transitions, namely, those sequences with their  $\sigma^2 \in (0, 1]$  as shown in Fig. 4 (b). We name them by dumb gait sequences for convenience since they tell us no view-changing information. Further statistics show that 85.6% dumb sequences come from frontal-view cases. Based on the above empirical analysis, we take a simple mining strategy at the pre-training phase: sample non-dumb sequences with a large probability to enhance view-invariant gait feature learning.

### 3.2.3 Gait Feature Encoder

Gait feature encoder  $F_\theta(\cdot)$  aims to transform the input sequence into gait feature vectors. In practice, it can be any existing methods, such as GaitSet [7], GaitPart [8] and GaitGL [9]. However, we fail to introduce them and meet the under-fitting cases. One possibility is that their networks are too shallow to fit millions of gait sequences. Therefore, we build a deep network for gait feature extraction, although that is not the main target of this paper.

Based on the review of recent typical works [7], [8], [9], we summarize that, to extract a good representation of gait sequence, four essential components are necessary for the encoder, *i.e.*, convolution backbone, temporal pooling (TP) module, horizontal pooling (HP) module and projection head. We construct a deep ResNet-like [50] network as the backbone. Besides that, to make it concise enough, we decide not to utilize any other well-designed modules

TABLE 2

Architecture of gait feature encoder  $F_\theta(\cdot)$ . The strings following each layer present the size and channel of the convolution kernels or the dimension of the separate fully-connected layers.

Component	Layer	Architecture
Backbone	ConvLayer	$3 \times 3, 64, \text{stride}=1$
	ResBlock 1	$\begin{bmatrix} 3 \times 3, 64 \\ 3 \times 3, 64 \end{bmatrix} \times 1, \text{stride} = 2$
	ResBlock 2	$\begin{bmatrix} 3 \times 3, 128 \\ 3 \times 3, 128 \end{bmatrix} \times 1, \text{stride} = 2$
	ResBlock 3	$\begin{bmatrix} 3 \times 3, 256 \\ 3 \times 3, 256 \end{bmatrix} \times 1, \text{stride} = 1$
	ResBlock 4	$\begin{bmatrix} 3 \times 3, 512 \\ 3 \times 3, 512 \end{bmatrix} \times 1, \text{stride} = 1$
TP	Temporal Pooling	
HP	Horizontal Pooling	
Head	Separate FC 0	$512 \times 512$
	Separate FC 1	$512 \times 512$

proposed in existing approaches [7], [8], [9], such as Multilayer Global Pipeline in GaitSet [7], Focal Convolution layer in GaitPart [8] and Global/Local Convolutional layer in GaitGL [9].

More specifically, as illustrated in Table 2, the backbone is composed of one initial convolution layer followed by a stack of four basic residual blocks, devoting to transforming each silhouette of the input gait clip into a frame-level feature map. And then, the temporal pooling (TP) module compresses the clip of feature maps along the temporal dimension by maximum function, thus outputting a feature map dominating clip-level understandings. Next, the horizontal pooling (HP) module slices the obtained clip-level feature map horizontally into pre-defined parts, *i.e.*, 16 partial feature map uniform in size. Each part-level feature map will be down-sampled into a part-informed vector by global average and max pooling operations. Finally, those extracted vectors will be mapped into the metric space by a projection head, who consists of two separate fully-connected layers and treats each part vector independently. Notably, all the layers in Table 2 are equipped with the batch normalization and ReLU activation layers except the last output layer, and Table 2 skips them for conciseness.

Totally, these designs enable our model to hold the fitting of million gait sequences, making it structurally concise yet experimentally strong to serve as the baseline model for learning the general gait representation from massive unlabelled walking videos.

### 3.2.4 Predictor and Stop-gradient

The predictor head  $H_\theta(\cdot)$  transforms the output of one view and matches it to another view, that is  $q_i = H_\theta(k_i) \rightarrow k_{i+}$ . Its structure consists of two separate fully-connected layers, with just the first one equipped with the batch normalization and ReLU activation layers, which is similar to the mentioned projection head in the gait feature encoder. Besides, a stop-gradient operation in Fig. 3 is introduced in [12], [15], and aims at preventing the model's outputs from collapsing to constant vectors.

Empirically, some works [12], [15] formulate the learning process of predictor  $H_\theta(\cdot)$  as an optimization problem that is being solved underlying. By definition, the optimal solution

2. The experimental results present over 95% accuracy in predicting the camera angle of sequences with an error of  $\pm 15^\circ$  on OU-MVLP.

TABLE 3  
The amount of the identities (#ID) and sequences (#Seq) covered by the employed datasets.

Dataset	Train Set		Test Set		Condition
	#Id	#Seq	#Id	#Seq	
CASIA-B	74	8,140	50	5,500	NM, BG, CL
OU-MVLP	5,153	144,284	5,154	144,312	NM
GREW	20,000	102,887	6,000	24,000	Diverse
Gait3D	3,000	18,940	1,000	6,369	Diverse

to  $H_\theta(\cdot)$  should satisfy:  $q_i = H_\theta(k_i) = \mathbb{E}_\pi(k_{i+})$ , indicating that the predictor head  $H_\theta(\cdot)$  is an approximation of the expectation  $\mathbb{E}_\pi(\cdot)$  over the data augmentation involved by  $\pi(\cdot)$ . This hypothesis and other similar interpretations [12], [15], [48], [49] can declare the necessity of the predictor  $H_\theta(\cdot)$ , that is to overcome the variations caused by  $\pi(\cdot)$ .

### 3.2.5 Comparison with Other Contrastive Frameworks

By ignoring details in each particular component, here we compare our GaitSSB with other contrastive frameworks in design concepts.

Interestingly, if we remove the negative samples, in other words, replace InfoNCE [56] with cosine similarity loss [12], [15], GaitSSB would be similar to the SimSiam [15] equipped with some modifications for fitting the gait-specific issues. However, we argue that the task granularity makes it different a lot. More specifically, gait recognition is a fine-grained task whose main challenge is that intra-class distances are often larger than inter-class, especially for cross-view cases. Once we discard the explicit signal of increasing the distance between representations of different gait sequences, those sequences owning similar appearance because of the similar camera angle would be automatically clustered together since the overall visual similarity often dominates in contrastive frameworks [19], [74]. This may cause a massive obstacle in identifying individuals because the extracted features take the view characteristics as centers. Hence, we think that the InfoNCE loss, which not only reduce the distance between representations of positive pairs but also enlarge that of negative pairs, is crucial for the gait contrastive framework. The ablation study in Sec. 4.5.2 and the hypothesis in Sec. 5 stand for this point experimentally and theoretically.

In addition, compared with those contrastive learning frameworks using InfoNCE loss, *e.g.*, MoCo [13] and SwAV [49], we take a more concise siamese network [75] as the encoder, instead of the relatively complex momentum or on-line clustering encoder. The primary motivation comes from the philosophy that the simple one is enough. In a nutshell, after absorbing the advantages of existing works, GaitSSB is conceptually simple and empirically powerful enough to be a self-supervised benchmark for gait recognition.

## 4 EXPERIMENTS

This section first introduces the datasets in use, including our unlabelled GaitLU-1M for pre-training and four public supervised datasets for transfer learning. Then, the implementation details about self-supervised pre-training, supervised transfer learning, and training from scratch will be reported. In the end, we conduct systematical comparisons

and comprehensive ablation studies, adequately exhibiting the superiority of GaitSSB and its components.

### 4.1 Datasets

GaitLU-1M collected by this paper is an unlabelled dataset used for pre-training. As for transfer learning, we employ four public supervised gait datasets, involving the most widely-used CASIA-B [3], the largest indoor OU-MVLP [4], the largest dataset in the wild GREW [5], and the most up-to-date popular Gait3D [45]. Here we describe implementation details that strictly follow the official protocols.

*CASIA-B* [3] is composed of 124 subjects, and each subject contains 3 walking conditions, including the normal (NM#1-6), bag-carrying (BG#1-2) and clothing-change (CL#1-2) cases. Each walking video is taken from 11 views uniformly distributed in  $[0^\circ, 180^\circ]$ . Hence, there are  $11 \times (6 + 2 + 2) = 110$  sequences per subject. In our experiments, the first 74 subjects are grouped into the training set, while the remaining 50 are grouped for evaluation. At the test stage, for each subject, the first 4 sequences under the NM walking (NM#1-4) are grouped as the gallery set, and the remaining 6 are divided into 3 probe subsets according to walking conditions, *i.e.*, NM, BG and CL.

*OU-MVLP* [4] is one of the largest indoor gait datasets. This dataset comprises 10,307 subjects, and each subject only contains one walking condition, namely normal walking (NM#0-1). Each walking video is taken from 14 views uniformly distributed inside  $[0^\circ, 90^\circ]$  and  $[180^\circ, 270^\circ]$ . Therefore, there are  $2 \times 14 = 28$  sequences per subject ideally. Following the official partition strategy, we take 5,153 subjects as the training set and 5,154 subjects as the test set. At the test stage, each subject's first sequence is regarded as the probe, and another is used for the gallery.

*GREW* [5] is the largest gait dataset in the wild up to date to our best knowledge. Its raw videos are collected from 882 cameras in a large public area, containing nearly 3,500 hours of  $1,080 \times 1,920$  streams. Except for the identities, some other human attributes have been annotated as well, *e.g.*, the gender, 14 age groups, 5 carrying conditions and 6 dressing styles. Therefore, GREW is believed to include adequate and diverse practical variations. In addition, this dataset is divided into 4 parts: a train set with 20,000 identities and 102,887 sequences, a validation set with 345 identities and 1,784 sequences, a test set with 6,000 identities and 24,000 sequences, a distractor set with 233,857 anonymous sequences. At the testing stage, each subject in the test set has 4 sequences, 2 for the probe and another 2 for the gallery. We use no valuation and distractor set for convenience.

*Gait3D* [45] was collected in a large supermarket and contains 1,090 hours of videos with  $1,920 \times 1,080$  resolution and 25 FPS. It consists of 4,000 subjects with over 25,000 walking sequences, in which 3,000/1,000 IDs are grouped into train/test subsets. For the test set, only one sequence from each ID will be selected to build the probe set with 1,000 sequences, while the rest becomes the gallery set with 5,369 sequences. For clarity, we list several key indexes for employing the above datasets in Table 3.

### 4.2 Implementation Details

Unless otherwise stated, a) The input silhouettes are aligned by the size-normalization method used in [4] and resized to



TABLE 4

Rank-1 accuracy on four authoritative gait datasets: unsupervised performance of GaitSSB v.s. supervised performance of some early methods.

Mode	Method	Srouce	Dataset						
			CASIA-B			OU-MVLP	OU-MVLP <sup>†</sup>	GREW	Gait3D
			NM	BG	CL				
Supervised	PoseGait [76]	PR2020	63.8	42.5	32.0	-	-	0.2	0.2
	GaitGraph [25]	ICIP2021	87.7	74.8	<b>66.3</b>	-	-	1.3	6.2
	CNN-LB [77]	PAMI2015	<b>94.1</b>	72.4	54.0	-	-	13.6	-
	GEINet [78]	ICB2016	-	-	-	-	<b>35.8</b>	6.8	7.0
Self-Supervised	GaitSSB	ours	83.3	<b>75.6</b>	28.7	<b>37.2</b>	35.0	<b>16.6</b>	<b>24.7</b>

TABLE 5

Common training settings on various datasets at different stages. The batch size of  $(q, k)$  indicates  $q$  subjects with  $k$  sequences per subject.

Stage	DataSet	Batch Size	Milestones	Steps
Pre-training	GaitLU-1M	512	(80K, 120K)	150K
Transfer Learning	CASIA-B	(8, 16)	(10K)	12K
	OU-MVLP	(32, 16)	(50K, 60K, 70K)	80K
	GREW	(128, 4)	(50K, 60K, 70K)	80K
	Gait3D	(64, 4)	(6K, 8K, 10k)	12K
Training from Scratch	CASIA-B	(8, 16)	(10K, 20K, 30K)	40K
	OU-MVLP	(32, 16)	(60K, 80K, 100K)	120K
	GREW	(128, 4)	(60K, 80K, 100K)	120K
	Gait3D	(64, 4)	(20K, 40K, 50K)	60K

64 × 44. b) SGD optimizer with an initial learning rate of 0.1 and momentum of 0.9 is employed for training. c) The multiple-step scheduler for dropping the learning rate with milestones shown in Table 5 is adopted. d) The triplet loss [79] with a margin of 0.3 is employed at supervised learning stages. e) The batch size and total iterations are listed in Table 5. f) The length of the input sequence is set to 16 frames in the pre-training phase and 30 frames for the other phases. g) Several state-of-the-art methods are reproduced, involving the GaitSet [7], GaitPart [8], GaitGL [9] and CSTL [80] (with 128 × 88 input), because their performance on GREW [5] and Gait3D [81] are missing. h) Our code is based on the popular OpenGait [82] codebase available at <https://github.com/ShiqiYu/OpenGait>.

#### 4.2.1 Pre-training

The particular settings for pre-training mainly center on the hyper-parameters of silhouette augmentation operator  $\pi(\cdot)$ .

For spatial augmentation, we randomly perform the horizontal flipping, affine, perspective and dilation transformations with a probability of 0.5. The affine transformation consists of a rotation with a random angle from  $[-10^\circ, +10^\circ]$ , and a shear with its level ranging on  $[-5 \times 10^{-3}, 5 \times 10^{-3}]$ . The transform axes of perspective are generated by randomly skewing the source axes within 10 pixels. Besides, dilation operations are conducted on a random contiguous region in the silhouette with a random kernel shape (rectangle, cross, or ellipse) and size (3 or 5).

For sampling augmentation, we sample the dumb sequences with a small probability of 0.1, while the rest are for others to enhance the learning of view-invariant features.

#### 4.2.2 Transfer Learning

For transfer learning, we import an additional separate fully-connected layer, whose parameters are  $l_2$ -normalized over channels dimension, to map the feature vectors ( $l_2$ -normalized) output by gait feature encoder  $F_\theta(\cdot)$  into the metric space. Several common fine-tuning tricks are employed to avoid catastrophic forgetting, *i.e.*, using the small initial learning rate, fixing the batch normalization layers, and removing the weight decay. Specifically, the initial learning rate of backbone, projection and additional head is set to  $1.0 \times 10^{-3}$ ,  $1.0 \times 10^{-2}$  and  $1.0 \times 10^{-1}$ , respectively.

#### 4.2.3 Training from Scratch

To intuitively show the effectiveness of pre-training, we take GaitSSB directly trained with random initialization into comparison scope. As for the training settings, there are only two slight differences between this case with the transfer learning, *i.e.*, starting with the learning rate of 0.1 and training with a new scheduler shown in Table 5.

We take the widely-used average rank-1 accuracy as the metric. Moreover, a) For CASIA-B, there are three metrics responding to the performance on its three probe subsets, *i.e.*, NM, BG and CL. b) For CASIA-B and OU-MVLP, the identical-view cases are excluded. c) For OU-MVLP, those probe sequences owning no gallery are skipped by default during the evaluation. Also, we make OU-MVLP<sup>†</sup> present the results without this skipping process.

### 4.3 Unsupervised Performance Comparison

We compare the unsupervised performance of GaitSSB with the supervised performance of four early representative methods, including the model-based PoseGait [24] and GaitGraph [25], and the GEI-based CNN-LB [1] and GEINet [83], on four authoritative gait benchmarks, namely CASIA-B [3], OU-MVLP [4], GREW [5] and Gait3D [45].

#### 4.3.1 Compared with Early Supervised Methods

As shown in Table 4, GaitSSB achieves nearly comparable or even better performance than the model-based (PoseGait and GaitGraph) and early GEI-based methods (CNN-LB and GEINet) in most cases, *i.e.*, on the OU-MVLP, GREW and Gait3D, and in the NM, BG walking situations of CASIA-B. However, for the most challenging CL cases on CASIA-B, GaitSSB encounters an evident degradation, implying that the dress-changing factors stimulated by the silhouette augmentation operator  $\pi(\cdot)$  are far from enough compared with those in the supervised datasets.

TABLE 6

Rank-1 accuracy on four authoritative gait datasets: unsupervised performance of GaitSSB v.s. cross-domain performance of recently state-of-the-art methods. The numbers in [ ] indicate the supervised results, and others present the cross-domain or unsupervised results.

Mode	Source Dataset	Method	Target Dataset					
			NM	BG	CL	OU-MVLP	GREW	Gait3D
On-diagonal: Supervised	CASIA-B	GaitSet	95.8	90.0	75.4	10.2	9.4	7.7
		GaitPart	96.1	90.7	78.7	11.5	9.3	6.6
		GaitGL	97.4	94.5	83.8	17.3	11.6	7.6
	OU-MVLP	GaitSet	74.0	55.5	16.4	94.1	12.3	13.0
		GaitPart	73.9	56.9	20.7	95.1	14.8	11.2
		GaitGL	81.7	71.5	28.8	96.2	21.4	16.0
Off-diagonal: Cross Domain	GREW	GaitSet	65.6	44.9	20.8	23.2	48.4	19.0
		GaitPart	69.2	52.1	25.4	25.4	47.6	19.3
		GaitGL	69.8	61.1	31.9	27.5	41.5	15.6
	Gait3D	GaitSet	62.8	45.8	11.9	22.1	19.2	42.6
		GaitPart	61.8	47.0	13.7	20.3	14.2	29.9
		GaitGL	63.5	51.2	16.3	25.2	14.3	23.5
Self-Supervised	GaitLU-1M	GaitSSB	83.3	75.6	28.7	37.2	16.6	24.7

#### 4.3.2 Compared with Cross-Domain Testings

To our knowledge, there are no reports about the unsupervised gait recognition performance evaluated on the aforementioned datasets. For a comprehensive comparison, we take the cross-domain testings into contrast scope due to its similarities with our evaluation protocol, *i.e.*, train the model in supervised *v.s.* self-supervised mode on the source dataset first and then test it on another target dataset.

As shown off-diagonal in Table 6, our unsupervised performances remarkably outperform the cross-domain performance of existing supervised methods over most of the source-target dataset pairs, convincingly suggesting the robustness and generalization of the gait representation learned by our contrastive framework. We think that the underlying reasons are two-fold: a) *Diverse data*. The large-scale unlabelled gait dataset, GaitLU-1M, is far beyond the traditional gait datasets in terms of scale and variety, strongly supporting this impressive performance in the case of cross-domain. b) *Strong baseline*. GaitSSB can explore the general and discriminative representations for gait recognition without requiring labeling information.

#### 4.3.3 Cross-View Performance Analysis

Furthermore, we are highly concerned about the unsupervised performance facing one of the most crucial and essential issues for practical gait recognition, namely the cross-view cases. Hence, we draw the rank-1 accuracy heat map by taking the probe-gallery view pairs as coordinates based on the unsupervised results on CASIA-B and OU-MVLP. Here we only consider the normal cases in CASIA-B to exclude the influence of other non-view factors.

As shown in Fig. 7, the results on the diagonal present the unsupervised performance in identical-view cases. Encouragingly, GaitSSB makes a great achievement in these cases, *i.e.* 99.5% and 98.5% rank-1 accuracy average over camera angles on CASIA-B and OU-MVLP, respectively, declaring that our unsupervised GaitSSB can nearly solve the identical-view gait recognition problem.

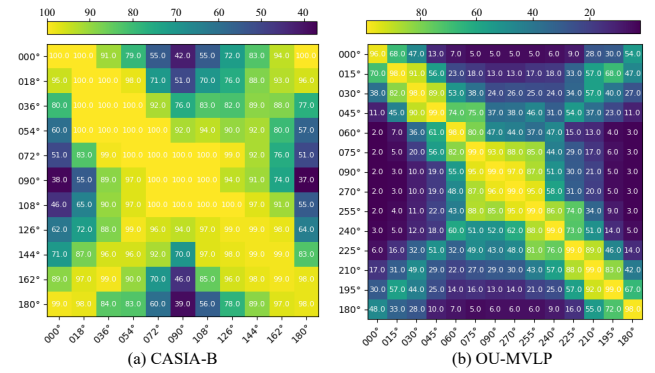


Fig. 7. GaitSSB's unsupervised rank-1 accuracy on (a) CASIA-B (NM walking) and (b) OU-MVLP over probe-gallery view pairs.

For the more challenging cross-view cases, whose results are displayed off-diagonal in Fig. 7, there exist only two subcases for the unsupervised GaitSSB to reach an impressive over 50% rank-1 accuracy on OU-MVLP. a) The view difference between probe and gallery is relatively small, *e.g.*, 15° and 30°, *i.e.*, the results shown around the diagonal of Fig. 7 (b). b) The appearance of the probe and gallery is close to being horizontally symmetric, *i.e.* the results around the anti-diagonal of Fig. 7 (b). But even so, GaitSSB exhibits competitive and promising results in some hard cases as well, *e.g.*, 42.0% rank-1 accuracy in the 000°-090° case for CASIA-B, and 24.0% rank-1 accuracy in the 030°-090° case for OU-MVLP. In Sec. 5, we will provide an instructive hypothesis about why and how our framework conquers such difficult view-changing factors.

Compared with the supervised results in the diagonal of Table 6, there remains a significant gap in front GaitSSB's unsupervised performance compared with the supervised performance of existing state-of-the-art methods, especially for the dramatic dress/view-changing conditions. However, instead of exceeding the supervised methods on some particular datasets, the primary goal of GaitSSB is to learn the informative prior knowledge from the large-scale unlabelled

TABLE 7  
Rank-1 accuracy on five authoritative gait datasets: GaitSSB with transfer learning v.s. some recently state-of-the-art methods.

Method	Source	Dataset									
		CASIA-B			CASIA-B*			OU-MVLP	OU-MVLP†	GREW	Gait3D
		NM	BG	CL	NM	BG	CL				
GaitSet [7]	AAAI 2019	95.8	90.0	75.4	92.3	86.1	73.4	94.1	87.2	48.4	42.6
GaitPart [8]	CVPR2020	96.1	90.7	78.7	93.1	86.0	75.0	95.1	88.6	47.6	29.9
KoopmanGait [41]	CVPR2021		-			-		-	74.7	-	-
GaitGL [84]	ICCV2021	97.4	94.5	83.8	94.1	90.0	81.4	96.2	89.9	47.3	23.5
CSTL [80]	ICCV2021	98.0	95.4	87.0	90.1	85.6	79.2	-	90.2	50.6	-
3DLocal [39]	ICCV2021	<b>98.3</b>	<b>95.5</b>	84.5		-		96.5	90.9	-	-
GaitMPL [85]	TIP2022	97.5	94.5	<b>88.0</b>		-		-	90.6	-	-
SMPLGait [45]	CVPR2022		-			-		-	-	53.2	-
LagrangeGait [86]	CVPR2022	97.5	94.6	85.1		-		-	90.0	-	-
GaitStrip [87]	ACCV2022	97.6	95.2	86.2		-		97.0	-	-	-
GaitSSB	w/o pretraining	94.0	87.6	68.0	96.6	93.1	79.9	97.3	91.0	55.9	54.2
	w/ pretraining	97.8	93.9	77.4	<b>96.8</b>	<b>94.5</b>	<b>81.8</b>	<b>97.8</b>	<b>91.8</b>	<b>61.7</b>	<b>63.6</b>

walking videos for the more generalizable performance achieved in downstream tasks.

#### 4.4 Transfer Learning Performance

After pre-training GaitSSB on GaitLU-1M, we fine-tune it on five target datasets, namely CASIA-B, CASIA-B\*, OU-MVLP, GREW and Gait3D, to check its effectiveness and generalization for transfer learning. Among these datasets, CASIA-B\* is a re-segmented version [44]. The extra import of CASIA-B\* owes to the fact that the background subtraction algorithm that CASIA-B uses for generating silhouette data tends to produce much noise and is outdated for real-world applications nowadays. In addition, to directly reveal the effectiveness of pre-training, we also evaluate the GaitSSB without pre-training, in other words, with being trained from scratch to exclude the backbone influence.

As shown in Table 7, except on the CASIA-B, our GaitSSB reaches the new state-of-the-art performances on other four datasets, namely CASIA-B\* (NM: +2.7%, BG: +4.5%, CL: +0.4%), OU-MVLP (+0.5%), GREW (+9.7%) and Gait3D (+10.4%). It is worth noting that GaitSSB exceeds other methods by a larger margin on the datasets captured from real-world scenarios (GREW and Gait3D) instead of the controlled laboratory (CASIA-B\* and OU-MVLP), illustrating two issues: a) *Generalization*. It is hard for other methods to transfer their superiority to practical cases. b) *Practicability*. GaitSSB is suitable for both indoor and outdoor cases.

For the confusing results on CASIA-B, we think that the salient cross-domain problem may cause it since the pedestrian segmentation algorithm that CASIA-B uses differs a lot from that of GaitLU-1M's, namely background subtraction *v.s.* deep-learning based segmentation [69]. For further analysis, we provide another two experiments to verify the effectiveness of our framework on CASIA-B: a) *w/ v.s. w/o pre-training*. The experiment that initializes GaitSSB from the pre-trained model performs better than that from scratch. b) *CASIA-B\* v.s. CASIA-B*. Once we re-segment the videos, *i.e.*, experiments on CASIA-B\*, the GaitSSB exceeds other methods simultaneously.

Notably, we also observe that excluding the CASIA-B and CASIA-B\*, the performances of training GaitSSB from scratch, *i.e.*, without pre-training, are distinctly beyond most of the current state-of-the-art methods on various

TABLE 8  
Rank-1 accuracy under cross-domain settings.

Method	Source	Cross-domain: Source→Target	
		GREW→Gait3D	Gait3D→GREW
GaitSet	AAAI2019	19.0	19.2
GaitPart	CVPR2020	19.3	14.2
GaitGL	ICCV2021	15.6	14.3
GaitSSB	w/o pre-training	20.1	20.6
	w/ pre-training	<b>27.2</b>	<b>27.6</b>

benchmarks, demonstrating that exploring basic network structures for gait recognition is far from enough. The reason may be that the gait recognition community always relies too much on a gait dataset proposed before the boom of deep learning, namely, the CASIA-B, whose scale is pretty small and the pretreatment is out-of-date, from now view.

GaitSSB is a strong baseline model for gait recognition thanks to its simple structure and exceptional performance, even if we train it from scratch. More importantly, the large-scale self-supervised pre-training makes our GaitSSB forward a step to bring the gait recognition an informed initialization, namely, the prior knowledge capacity to describe the gait patterns.

##### 4.4.1 Generalization Capacity Validation

To directly validate the generalization capacity brought by the pre-training, we introduce Table 8 showing that, under the challenging cross-domain settings, the pre-training provides a noticeable performance gain for GaitSSB to outperform prior representative state-of-the-art methods by a margin. As we can see, the superior generalization capacity of GaitSSB mainly comes from the pre-training process instead of its architectural designs.

##### 4.4.2 Annotation Cost Saving

Fig. 8 shows the performance of GaitSSB *w/ v.s. w/o* pre-training when the fine-tuning (training) scale varies. In practice, the first 10%, 20%, 50% and 100% subjects of the training set will be employed to fine-tune the pre-trained GaitSSB or directly train GaitSSB from scratch. Due to scale

TABLE 9  
Ablation studies: w/ and w/o the spatial, intra-sequence and sampling augmentation. w/ and w/o the negative samples.

Index	Control Variables				CASIA-B*			OU-MVLP	GREW	Gait3D
	Spatial Aug	Intra-Sequence Aug	Sampling Aug	Negative Samples	NM	BG	CL			
(a)	✓	✓	✓	✓	<b>75.0</b>	<b>65.3</b>	<b>28.7</b>	<b>37.2</b>	16.6	<b>24.7</b>
(b)	×	✓	✓	✓	67.4	56.7	21.6	26.7	10.9	18.5
(c)	✓	×	×	✓	6.9	5.6	3.4	0.1	0.1	0.1
(d)	✓	✓	×	✓	67.5	58.6	25.5	33.5	<b>17.0</b>	24.2
(e)	✓	✓	✓	×	43.4	32.1	13.2	14.6	8.8	15.5

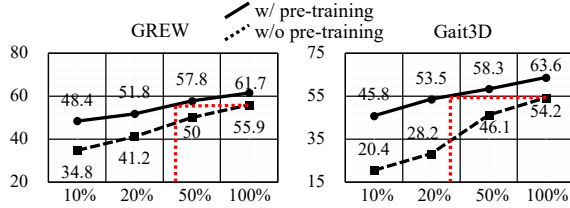


Fig. 8. Rank-1 accuracy of GaitSSB w/ v.s. w/o pre-training when the fine-tuning (training) scale varies.

differences, we proportionally reduce the iterations used for dropping the learning rate and stopping training with keeping other settings shown in Table 5 unchanged. Consequently, Fig. 8 illustrates that: a) The smaller fine-tuning (training) scale, the more significant improvements brought by pre-training. b) According to the red dotted lines, compared with the full performance of GaitSSB without pre-training, GaitSSB with pre-training reaches the competitive performance using only 50% and 20% training data of GREW and Gait3D, namely, the pre-training on the GaitLU dataset can save about 50% and 80% annotation costs of GREW and Gait3D.

#### 4.5 Ablation Study

Unlike previous methods that only conduct the ablative studies on the fairly small-scale CASIA-B [7], [8], [9], [39], [40], here we perform the experiments on four datasets, including the CASIA-B\*, OU-MVLP, GREW and Gait3D, to make our conclusions more solid. Moreover, various controlling settings have been taken into consideration, involving w/ and w/o spatial, intra-sequence and sampling augmentations, w/ and w/o negative samples in the training batch, and different scales of pre-training datasets. As shown in Table 9, we number the experiments for description convenience, where experiment (a) denotes the complete GaitSSB.

##### 4.5.1 Impact of Silhouette Augmentation

In this paper, the proposed silhouette augmentation operator  $\pi(\cdot)$  consists of three parts, involving the spatial, intra-sequence and sampling augmentations. Next, we will discuss their roles for GaitSSB, respectively.

If we remove the spatial augmentation from GaitSSB, *i.e.*, experiment (a) *v.s.* (b) shown in Table 9, the performance would drop dramatically over all the datasets. This phenomenon shows that the proposed spatial augmentation strategy, which is composed of the horizontal flipping, affine, perspective and dilation transformations, can

simulate the spatial variations usually seen in silhouettes adequately and thus boost the pre-training performance remarkably.

The comparison between experiment (a) and (c) in Table 9 indicates that the performance would collapse once we discard the intra-sequence and sampling augmentation. In other words, GaitSSB cannot depict discriminative gait features by only taking the spatial augmentation. By taking two non-overlap clips inside the same walking sequence as two different augmented views, the intra-sequence augmentation gives GaitSSB the ability against frame-level pretreatment errors, pose and appearance changes, and cross-view factors along the sequence dimension:

- Since the natural walking videos are processed frame by frame, learning the intra-sequence consistencies enables the model to overcome the massive noises caused by pedestrian tracking and segmentation.
- Two non-overlap clips inside the same walking sequence present identical gait patterns with two different views of appearance and pose. Therefore, the intra-sequence consistencies imply the uniqueness of an individual's gait patterns to some extent, which is fundamental and crucial for gait recognition.
- Walking is a dynamic course. The angle between the individual's walking route and the camera's orientation is not constant. Hence, reducing the distance between representations of two non-overlap clips can benefit the learning of view-invariant features.

As for the sampling augmentation, its goal is to enhance the learning of view-invariant gait features via sampling the gait sequences possessing view-changing factors with a higher probability and vice versa. By comparing experiment (a) with (d) in Table 9, we find this strategy benefits GaitSSB on various benchmarks observably, *i.e.*, CASIA-B\* (NM: +7.5%, BG: +6.7%, CL: +3.2%), OU-MVLP (+3.7%) and Gait3D (+0.5%). However, we also investigate that the sampling augmentation leads to a slight degradation in GREW. Due to its largest scale and in-the-wild settings, identifying individuals on GREW is more difficult than other employed gait datasets. This situation may make the improvement brought by sampling augmentation during the pre-training hard to be transferred to the downstream unsupervised evaluation on GREW. We consider that this strategy can boost pre-training performance in most cases, and meanwhile, a better sampling strategy is needed.

##### 4.5.2 Impact of Negative Sample

Some contrastive frameworks argue that the negative samples are crucial [12], [13], [48] while others say they are unnecessary [15]. Hence here, we discuss the role of negative

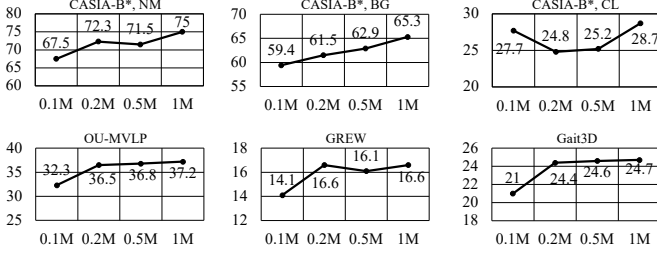


Fig. 9. Ablation studies: impact of the pre-training scale (w/o fine-tune).

samples in GaitSSB. Specifically, we remove the negative samples by simply replacing the InfoNCE with cosine similarity loss, where the latter aims to attract the positive pairs and do nothing for the negative pairs.

The comparison between experiment (a) and (e) in Table 9 declares that, if we remove the negative samples at the pre-training phase, GaitSSB would meet a prominent degeneration over all the benchmarks, *i.e.*, CASIA-B\* (NM: -31.6%, BG: -33.2%, CL: -15.5%), OU-MVLP (-22.6%), GREW(-7.5%) and Gait3D (-9.2%).

As far as we know, most of the contrastive frameworks are developed on coarse-grained tasks, *e.g.*, the image recognition and person re-identification, where a noticeable feature is that two different images from semantically near categories naturally own relatively more minor visual distance. However, it is the similar camera angle instead of the same identity, being more likely to make two different walking sequences alike in visual appearance for gait data. Hence, increasing the distance between the representation of negative pairs, especially those with a similar view, is crucial for GaitSSB to build a discriminative metric space. The hypothesis proposed in Sec. 5 will offer a further methodological explanation for this issue.

#### 4.5.3 Impact of Pre-training Scale

To study the impact of the pre-training scale, we build the GaitLU-0.5M, 0.2M and 0.1M by randomly selecting walking sequences from GaitLU-1M, where the GaitLU-0.5 is a subset of GaitLU-1M, GaitLU-0.2M is a subset of GaitLU-0.5M, and so on.

As shown in Fig. 9, the unsupervised performance of our GaitSSB achieved on OU-MVLP, GREW and Gait3D mainly benefits from the larger pre-training scale. However, for CASIA-B\*, only the testings under NM and BG conditions are generally consistent with this observation, while the accuracy oscillates with the pre-training scale in the most challenging CL case. The underlying reason may be that the unlabelled GaitLU-1M lacks the natural dress-changing factors, giving rise to the invalidity of enlarging the pre-training scale. Therefore, forming positive pairs by effectively simulating the cloth-changing cases should be critical for further study.

Moreover, Fig. 10 presents the counterpart of Fig. 9 after fine-tuning the pre-trained GaitSSB on GREW and Gait3D. As we can see, the GaitSSB pre-trained on varying scales consistently outperforms the GaitSSB w/o pre-training, adequately showing the GaitSSB's robustness to the pre-training scale. Moreover, we note that the performance on GREW gradually increases with the larger pre-

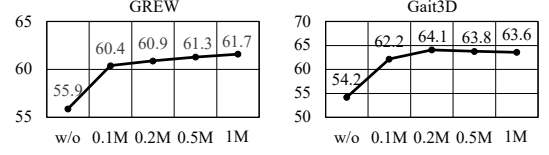


Fig. 10. Ablation studies: impact of the pre-training scale (w/ fine-tune).

training scale, and the growth curve on Gait3D quickly reaches saturation. Taking the scale issue into consideration, *i.e.*, GaitLU-1M is about 55 times larger than the training set of Gait3D and only 10 times larger than that of GREW, we think that the self-supervised data mining capacity of GaitSSB still needs to be improved to explore continual knowledge from the increasing pre-training scale. Therefore, further study majoring in gait data mining efficiency is highly desired.

## 5 HYPOTHESIS AND DISCUSSION

Previous contrastive frameworks have provided some pioneer intuitions on the behavior of the pre-training process. For example, Grill *et al.* consider that the two-stream network in BYOL [12] is similar to the discriminator and generator in GANs [88] respectively. Chen *et al.* hypothesize that SimSiam [15] is an implementation of an Expectation-Maximization-like algorithm that aims at solving two underlying sub-problems. In this section, we pay more attention to gait-specific issues, *i.e.*, proposing the hypothesis on why the intra-sequence augmentation and negative samples are crucial for GaitSSB. Also, we propose some methodological suggestions for further research.

Let a map  $\Pi(\cdot)$  present the corruptions caused by various noisy factors. For a sample's feature vector  $x$  generated by our GaitSSB,  $\Pi(x)$  denotes a set of the samples owning the identical class as  $x$ , covering a wide range of practical factors, *e.g.*, carrying, dressing and camera views, *etc.* We define the intra-class and inter-class distance of sample  $x$  against the noisy factors in  $\Pi(\cdot)$  as:

$$\begin{aligned} d_{\Pi}^{+}(x) &= \max_{x^{+} \in \Pi(x)} \|x - x^{+}\|_2 \\ d_{\Pi}^{-}(x) &= \min_{x^{-} \notin \Pi(x)} \|x - x^{-}\|_2 \end{aligned} \quad (3)$$

Since all the vectors are  $l_2$ -normalized in the obtained metric space, here we use the Euclidean distance  $\|\cdot\|_2$  that is equivalent to the cosine similarity as a measure of the relative distance. Obviously, we say that the sample  $x$  can be identified correctly against  $\Pi(\cdot)$  if  $d_{\Pi}^{+}(x) < d_{\Pi}^{-}(x)$ .

For our self-supervised benchmark, since GaitLU-1M is unlabeled, we introduce the silhouette augmentation operator  $\pi(\cdot)$  as a substitute to simulate noisy factors within  $\Pi(\cdot)$ . Differently,  $\pi(\cdot)$  only contains three categories of hand-designed noisy factors, *i.e.*, the spatial augmentations for gait appearance changes, the intra-sequence augmentations for pre-treatment errors and camera angle changes, and the sampling augmentations for enhancing the view-invariant representation learning. It is worth noting that the view changes inside a walking video are primarily slight, indicating that  $\pi(\cdot)$  usually contains subtle view-changing variations. Similarly, we define the intra and inter-class



distance of sample  $x$  against  $\pi(\cdot)$  as  $d_{\pi}^{+}(x)$  and  $d_{\pi}^{-}(x)$ . Next, we concentrate on one of the most critical scenarios for gait recognition, namely the cross-view cases, and ignore other noisy factors for simplicity.

In identical-view cases, there are almost no view differences between the walking videos of the probe and gallery, indicating  $\Pi(x) \rightarrow \pi(x)$  at this time. Therefore, the proposed GaitSSB can maintain much pre-training performance when tested on other supervised benchmarks. The results in the diagonal of Fig. 7 have verified this point.

For more challenging cross-view cases,  $\Pi(\cdot)$  covers more dramatic view-changing factors than that within  $\pi(\cdot)$ , i.e.,  $\pi(x) \subset \Pi(x)$ , intuitively meaning that it is hard for GaitSSB to maintain its pre-training performance on this occasion. But encouragingly, Fig. 7 shows that GaitSSB can perform still promising performance in the most challenging cross-view case, such as 42.0% rank-1 (68.0% rank-5) average accuracy for the most challenging 000°-090° (probe-gallery) view pair on the widely-used CASIA-B [3]. This observation encourages us to explore the further understanding of the view-invariant feature learning in our GaitSSB. Next, we start by analyzing the inter/intra-class distance.

Since  $\pi(x) \subset \Pi(x)$ , we can get  $\bar{\Pi}(x) \subset \bar{\pi}(x)$  where  $\bar{\Pi}(x) = \{x^- | x^- \notin \Pi(x)\}$  and  $\bar{\pi}(x) = \{x^- | x^- \notin \pi(x)\}$ . Consequently,

$$d_{\pi}^{-}(x) \leq d_{\Pi}^{-}(x) \quad (4)$$

indicates GaitSSB can bring much of the ability to increase inter-class distance into the evaluation stage regardless of the challenging cross-view factors. However, we notice that some contrastive frameworks [12], [15] have achieved amazing performance without the design of enlarging the distance between representations of negative pairs. We argue that the task granularity makes it different a lot:

- *Data Granularity.* For some coarse-grained tasks, their pre-training dataset, e.g., the ImageNet [66] for image recognition, usually has millions of images distributed over only thousands of classes. This relatively dense distribution implies that a training batch may contain two different images from the same class, but the contrastive model will still enlarge their distance and sequentially make a mistake. As for our benchmark, it is almost impossible for an individual to appear twice in GaitLU-1M since the videos are collected around the world, varying in locations, dates and photographers. Therefore, we consider that GaitLU-1M contains only one walking sequence for each subject, meaning that GaitSSB would not mistakenly repulse two different sequences owning the same identity away. Considering Eq. 4, we assume that the ability to enlarge the distance within the negative pair learned at the pre-training stage is not misleading and, thus, pretty useful for the downstream gait task.
- *Overall Similarity.* The overall visual similarity often dominates in the final representation for contrastive frameworks [19]. However, for gait recognition, it is a similar camera view rather than the same identity that usually makes two silhouette sequences alike in overall appearance. Hence, using negative samples is a good choice since it tells the network needs to

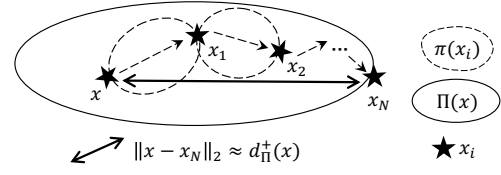


Fig. 11. Approximate  $\Pi(x)$  by repetitively performing  $x_i \in \pi(x_{i-1})$ .

increase the distance within the negative pair, even if two sequences are similar in the camera angle.

Combining with the comparison between experiment (a) with (e) in Table 9, we argue that learning to enlarge the distance within negative pair (using the negative samples) is crucial for contrastive gait pre-training.

As to the intra-class distance, inspired by calculus thinking, we regard  $\pi(\cdot)$  as the differential element of  $\Pi(\cdot)$ . As shown in Fig. 11, consider the sample  $x$  can be augmented  $N > 1$  times by  $\pi(\cdot)$ , i.e.,  $x_1 \in \pi(x)$ ,  $x_2 \in \pi(x_1)$ , ...,  $x_N \in \pi(x_{N-1})$  with

$$x_i = \arg \max_{x^+ \in \pi(x_{i-1})} \|x_{i-1} - x^+\|_2 \quad (5)$$

where  $i \in \{1, \dots, N\}$ , and  $x_0$  here denotes  $x$  (the same below). This process goes like transferring the camera angle by a maximal offset again and again. Then, we can approximate  $d_{\Pi}^{+}(x)$  by finding a large enough  $N$  to make:

$$\begin{aligned} d_{\Pi}^{+}(x) &\approx \|x - x_N\|_2 \\ &\leq \sum_{i=1}^N \|x_{i-1} - x_i\|_2 \\ &= \sum_{i=1}^N d_{\pi}^{+}(x_i) \end{aligned} \quad (6)$$

Above inequality holds for Euclidean-metric space. Obviously, the larger  $N$  implies the larger camera angle difference between the probe ( $x$ ) and gallery ( $x_N$ )<sup>3</sup>. Eq. 6 displays the distance transitivity law in Euclidean-metric space, and implies that we can pull  $x$  towards  $x_N$  (reduce intra-distance  $d_{\Pi}^{+}(x)$ ) by minimizing the distance between  $x_i$  and  $x_{i-1}$  where  $x_i \in \pi(x_{i-1})$  and  $i \in \{1, \dots, N\}$ .

Next, we make a general assumption that the well-pre-trained GaitSSB can ensure  $d_{\pi}^{+}(x_i) \leq a \leq b \leq d_{\pi}^{-}(x_i)$  for sample  $x_i$  against  $\pi(\cdot)$  where  $i \in \{1, \dots, N\}$ . Taking Eq. 4 and 6 into consideration, we know that:

$$\begin{aligned} b &\leq d_{\Pi}^{-}(x) \\ d_{\Pi}^{+}(x) &\leq Na \end{aligned} \quad (7)$$

showing that the lower bound of the inter-class distance of sample  $x$  against the practical noisy factors in  $\Pi(\cdot)$  is  $b$  while the upper bound of the intra-class distance is  $Na$ .

For the convenience of analysis, here we fix  $N$  first. If GaitSSB can increase the distance between representations of negative pairs by a large margin at the pre-training stage, i.e., make  $a \ll b$  until  $N < b/a$ , we can get  $d_{\Pi}^{+}(x) < d_{\Pi}^{-}(x)$  at this time. Our self-supervised framework can rightly recognize the sample  $x$  against the challenging cross-view

3. We consider that the view differences between the similar or symmetric camera angle pairs are relatively small, e.g., 000° with 015° or 180°. The same to below.

factors in  $\Pi(\cdot)$  at the supervised testing stage. Therefore, increasing  $b$  and meanwhile decreasing  $a$  via enhancing the recognition capacity of the backbone is vital for GaitSSB.

On the other hand, the larger cross-view angle implies the larger  $N$ , showing that it is harder to ensure  $N < b/a$ , especially for those cross-large-view cases. The salient performance degradation around the diagonal/anti-diagonal of Fig. 7 can verify this point experimentally. Moreover, the degree of view-changing transition involved by  $\pi(\cdot)$ , which is regarded as the differential element of  $\Pi(\cdot)$  as shown in Fig. 11, make it different as well. Intuitively, the more dramatic view-changing factors that  $\pi(\cdot)$  contain, the smaller  $N$  we need to approximate  $\Pi(\cdot)$ . If we remove the intra-sequence and sampling augmentation strategies from  $\pi(\cdot)$ , it will include almost no view-changing factors as discussed in Sec. 4.5.1. At this time, it is impossible to ensure  $N < b/a$  since  $N \rightarrow \infty$ . The performance collapse in Table 9 (experiment (a) *v.s.* (c)) can stand for it point experimentally. Therefore, exploring the more advanced augmentation strategy dominating view-changing variations dramatically should be vital for further research.

Last but not least, the most challenging bag-carrying and dress-changing cases, which have been neglected for clarity analysis, are worth being discussed seriously as well. Intuitively, the spatial augmentation strategies should matter in learning the clothing-invariant representations. However, though its effectiveness has been verified by experiment (b) in Table 9, our silhouette augmentation operator is far from effectively simulating the enormous variety of practical appearance changes. This issue still needs further systematic explorations.

## 6 CONCLUSION

Since manual labeling is costly and insatiable, the real-world applications of gait recognition have been subject to small-scale supervised datasets for a long time. Inspired by the development of self-supervised learning, this paper makes an attempt to verify the possibility of learning the generalizable and discriminative gait representations from massive unlabelled walking videos from the aspects of both experimental and theoretical analysis. Technically, we propose a comprehensive and systematical benchmark, involving a large-scale unlabelled gait dataset GaitLU-1M, and a structurally concise yet experimentally powerful baseline model GaitSSB. We hope this work can serve as a stepping stone to learning the general gait representation by self-supervised approaches, and inspire future works to narrow the gap between theory and practice for gait recognition.

## 7 ETHICAL STATEMENTS

The improper use or abuse of gait recognition will significantly threaten personal privacy. As of now, gathering web data on a large scale has become a common practice that significantly promotes AI research and applications. However, the collection inevitably involves personal data since most Internet data is created by and related to individuals. To protect personal privacy, most laws from different countries impose restrictions on the use of personal data by private businesses, while providing derogations for security and

research purposes. For instance, according to General Data Protection Regulation passed by the European Parliament, in the version of the OJ L 119, 04.05.2016 [89], personal data can be processed for scientific research purposes without demanding for the data subject's consent if adequately minimized and anonymized. Though this practice is generally perceived as legally acceptable, it introduces an unsolved ethical conundrum at the intersection of privacy protection and advancements in AI. On this background, the distributor of the GaitLU-1M dataset would like to try the best to protect personal data by taking the following measures.

- a) Data Minimization. The data of GaitLU-1M is exhibited as the sequence of binary silhouettes and body skeletons. No visual RGB information will be available.
- b) Expiration Date. GaitLU-1M can only be used for 20 years since the corresponding paper is published. After this date, all data will be deleted and not allowed to be used.
- c) Opt-out. For each gait sequence in GaitLU-1M, its data will be deleted once identified by some algorithms or people with providing confirmed evidence.
- d) Unattended Pre-treatment. The pre-processing of GaitLU-1M is based on automatic algorithms, meaning that nobody has manually processed personal visual data.
- e) Reasonable Anonymization. The raw data of GaitLU-1M, including video links and other intermediate results, have been deleted. In addition, weak annotations such as the date and city information have been erased by randomly numbering the sequences.
- f) Purpose Limitation. GaitLU-1M can be distributed only for non-commercial research purposes with a case-by-case access application.
- g) Ethical Requirements. Applicants are required to sign an additional agreement about data security and privacy protection. Conducts like identity inferring, manual labeling, private distribution, inversion attack, and biased analysis are prohibited explicitly.

We hope the above measures can inspire further study to balance AI development and privacy protection.

## ACKNOWLEDGMENT

We would like to thank the helpful discussion with Mrs. Yunjie Peng, Dr. Chunshui Cao, Dr. Xu Liu and Dr. Hanyang Peng. This work was supported in part by the National Natural Science Foundation of China under Grant 61976144, Grant 62276025 and Grant 62206022, in part by the National Key Research and Development Program of China under Grant 2020AAA0140002, and in part by the Shenzhen Technology Plan Program (Grant No. GJHZ20220913142611021 and KQTD20170331093217368).


## REFERENCES

- [1] Z. Wu, Y. Huang, L. Wang, X. Wang, and T. Tan, "A comprehensive study on cross-view gait based human identification with deep cnns," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 2, pp. 209–226, 2016.
- [2] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, pp. 436–444, 2015.

- [3] 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)*, vol. 4, 2006, pp. 441–444.
- [4] N. Takemura, Y. Makihara, D. Muramatsu, T. Echigo, and Y. Yagi, "Multi-view large population gait dataset and its performance evaluation for cross-view gait recognition," *IPSN Transactions on Computer Vision and Applications*, vol. 10, 12 2018.
- [5] Z. Zhu, X. Guo, T. Yang, J. Huang, J. Deng, G. Huang, D. Du, J. Lu, and J. Zhou, "Gait recognition in the wild: A benchmark," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 14 789–14 799.
- [6] S. Zhang, Y. Wang, T. Chai, A. Li, and A. K. Jain, "Real-gait: Gait recognition for person re-identification," *arXiv preprint arXiv:2201.04806*, 2022.
- [7] H. Chao, K. Wang, Y. He, J. Zhang, and J. Feng, "Gaitset: Cross-view gait recognition through utilizing gait as a deep set," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 7, pp. 3467–3478, 2022.
- [8] C. Fan, Y. Peng, C. Cao, X. Liu, S. Hou, J. Chi, Y. Huang, Q. Li, and Z. He, "Gaitpart: Temporal part-based model for gait recognition," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 14 225–14 233.
- [9] B. Lin, S. Zhang, and X. Yu, "Gait recognition via effective global-local feature representation and local temporal aggregation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 14 648–14 656.
- [10] A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever, "Improving language understanding by generative pre-training," *Preprint*, 2018.
- [11] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [12] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. Richemond, E. Buchatskaya, C. Doersch, B. Avila Pires, Z. Guo, M. Gheshlaghi Azar *et al.*, "Bootstrap your own latent-a new approach to self-supervised learning," *Advances in Neural Information Processing Systems*, vol. 33, pp. 21 271–21 284, 2020.
- [13] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 9729–9738.
- [14] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *International Conference on Machine Learning*. PMLR, 2020, pp. 1597–1607.
- [15] X. Chen and K. He, "Exploring simple siamese representation learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 15 750–15 758.
- [16] S. Abu-El-Hajja, N. Kothari, J. Lee, P. Natsev, G. Toderici, B. Varadarajan, and S. Vijayanarasimhan, "Youtube-8m: A large-scale video classification benchmark," *arXiv preprint arXiv:1609.08675*, 2016.
- [17] N. Xu, L. Yang, Y. Fan, D. Yue, Y. Liang, J. Yang, and T. Huang, "Youtube-vos: A large-scale video object segmentation benchmark," *arXiv preprint arXiv:1809.03327*, 2018.
- [18] D. Fu, D. Chen, J. Bao, H. Yang, L. Yuan, L. Zhang, H. Li, and D. Chen, "Unsupervised pre-training for person re-identification," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 14 750–14 759.
- [19] Z. Wu, Y. Xiong, S. X. Yu, and D. Lin, "Unsupervised feature learning via non-parametric instance discrimination," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [20] W. Kusakunniran, Q. Wu, H. Li, and J. Zhang, "Multiple views gait recognition using view transformation model based on optimized gait energy image," in *2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV workshops*. IEEE, 2009, pp. 1058–1064.
- [21] Y. Wang, S. Yu, Y. Wang, and T. Tan, "Gait recognition based on fusion of multi-view gait sequences," in *International Conference on Biometrics*. Springer, 2006, pp. 605–611.
- [22] R. Bodor, A. Drenner, D. Fehr, O. Masoud, and N. Papanikolopoulos, "View-independent human motion classification using image-based reconstruction," *Image and Vision Computing*, vol. 27, no. 8, pp. 1194–1206, 2009.
- [23] G. Ariyanto and M. S. Nixon, "Model-based 3d gait biometrics," in *2011 International Joint Conference on Biometrics (IJCB)*. IEEE, 2011, pp. 1–7.
- [24] R. Liao, S. Yu, W. An, and Y. Huang, "A model-based gait recognition method with body pose and human prior knowledge," *Pattern Recognition*, vol. 98, p. 107069, 2020.
- [25] T. Teepe, A. Khan, J. Gilg, F. Herzog, S. Hörmann, and G. Rigoll, "Gaitgraph: graph convolutional network for skeleton-based gait recognition," in *2021 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2021, pp. 2314–2318.
- [26] T. Teepe, J. Gilg, F. Herzog, S. Hörmann, and G. Rigoll, "Towards a deeper understanding of skeleton-based gait recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 1569–1577.
- [27] X. Li, Y. Makihara, C. Xu, and Y. Yagi, "Multi-view large population gait database with human meshes and its performance evaluation," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 4, no. 2, pp. 234–248, 2022.
- [28] Li, Xiang and Makihara, Yasushi and Xu, Chi and Yagi, Yasushi, "End-to-end model-based gait recognition using synchronized multi-view pose constraint," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 4106–4115.
- [29] X. Li, Y. Makihara, C. Xu, Y. Yagi, S. Yu, and M. Ren, "End-to-end model-based gait recognition," in *Proceedings of the Asian Conference on Computer Vision*, 2020.
- [30] C. Xu, Y. Makihara, X. Li, and Y. Yagi, "Occlusion-aware human mesh model-based gait recognition," *IEEE Transactions on Information Forensics and Security*, 2023.
- [31] W. Kusakunniran, Q. Wu, J. Zhang, Y. Ma, and H. Li, "A new view-invariant feature for cross-view gait recognition," *IEEE Transactions on Information Forensics and Security*, vol. 8, no. 10, pp. 1642–1653, 2013.
- [32] Y. Makihara, R. Sagawa, Y. Mukaigawa, T. Echigo, and Y. Yagi, "Gait recognition using a view transformation model in the frequency domain," in *European Conference on Computer Vision*. Springer, 2006, pp. 151–163.
- [33] J. Han and B. Bhanu, "Individual recognition using gait energy image," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 2, pp. 316–322, 2005.
- [34] C. Wang, J. Zhang, L. Wang, J. Pu, and X. Yuan, "Human identification using temporal information preserving gait template," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 11, pp. 2164–2176, 2011.
- [35] X. Li, Y. Makihara, C. Xu, Y. Yagi, and M. Ren, "Gait recognition via semi-supervised disentangled representation learning to identity and covariate features," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [36] K. Zhang, W. Luo, L. Ma, W. Liu, and H. Li, "Learning joint gait representation via quintuplet loss minimization," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [37] S. Hou, C. Cao, X. Liu, and Y. Huang, "Gait lateral network: Learning discriminative and compact representations for gait recognition," in *European Conference on Computer Vision*. Springer, 2020, pp. 382–398.
- [38] C. Xu, Y. Makihara, X. Li, Y. Yagi, and J. Lu, "Gait recognition from a single image using a phase-aware gait cycle reconstruction network," in *Computer Vision – ECCV 2020*, A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, Eds. Cham: Springer International Publishing, 2020, pp. 386–403.
- [39] Z. Huang, D. Xue, X. Shen, X. Tian, H. Li, J. Huang, and X.-S. Hua, "3d local convolutional neural networks for gait recognition," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 14 920–14 929.
- [40] X. Huang, D. Zhu, H. Wang, X. Wang, B. Yang, B. He, W. Liu, and B. Feng, "Context-sensitive temporal feature learning for gait recognition," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 12 909–12 918.
- [41] S. Zhang, Y. Wang, and A. Li, "Cross-view gait recognition with deep universal linear embeddings," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 9095–9104.
- [42] Z. Zhang, L. Tran, F. Liu, and X. Liu, "On learning disentangled representations for gait recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [43] C. Song, Y. Huang, Y. Huang, N. Jia, and L. Wang, "Gaitnet: An

- end-to-end network for gait based human identification," *Pattern recognition*, vol. 96, p. 106988, 2019.
- [44] J. Liang, C. Fan, S. Hou, C. Shen, Y. Huang, and S. Yu, "Gaitedge: Beyond plain end-to-end gait recognition for better practicality," in *Computer Vision – ECCV 2022*, 2022.
- [45] J. Zheng, X. Liu, W. Liu, L. He, C. Yan, and T. Mei, "Gait recognition in the wild with dense 3d representations and a benchmark," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [46] I. Kemelmacher-Shlizerman, S. M. Seitz, D. Miller, and E. Brossard, "The megaface benchmark: 1 million faces for recognition at scale," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 4873–4882.
- [47] R. Hadsell, S. Chopra, and Y. LeCun, "Dimensionality reduction by learning an invariant mapping," in *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, vol. 2. IEEE, 2006, pp. 1735–1742.
- [48] X. Chen, H. Fan, R. Girshick, and K. He, "Improved baselines with momentum contrastive learning," *arXiv preprint arXiv:2003.04297*, 2020.
- [49] M. Caron, I. Misra, J. Mairal, P. Goyal, P. Bojanowski, and A. Joulin, "Unsupervised learning of visual features by contrasting cluster assignments," *Advances in Neural Information Processing Systems*, vol. 33, pp. 9912–9924, 2020.
- [50] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [51] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020.
- [52] J. Bromley, I. Guyon, Y. LeCun, E. Säckinger, and R. Shah, "Signature verification using a 'siamese' time delay neural network," *Advances in neural information processing systems*, vol. 6, 1993.
- [53] P. Vincent, H. Larochelle, Y. Bengio, and P.-A. Manzagol, "Extracting and composing robust features with denoising autoencoders," in *Proceedings of the 25th international conference on Machine learning*, 2008, pp. 1096–1103.
- [54] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," *arXiv preprint arXiv:2111.06377*, 2021.
- [55] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arXiv preprint arXiv:1301.3781*, 2013.
- [56] A. v. d. Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.
- [57] Y. Tian, D. Krishnan, and P. Isola, "Contrastive multiview coding," in *European conference on computer vision*. Springer, 2020, pp. 776–794.
- [58] H. Murase and R. Sakai, "Moving object recognition in eigenspace representation: gait analysis and lip reading," *Pattern recognition letters*, vol. 17, no. 2, pp. 155–162, 1996.
- [59] P. S. Huang, C. J. Harris, and M. S. Nixon, "Recognising humans by gait via parametric canonical space," *Artificial Intelligence in Engineering*, vol. 13, no. 4, pp. 359–366, 1999.
- [60] H. Rao, S. Wang, X. Hu, M. Tan, Y. Guo, J. Cheng, X. Liu, and B. Hu, "A self-supervised gait encoding approach with locality-awareness for 3d skeleton based person re-identification," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [61] Y. Liu, Y. Zeng, J. Pu, H. Shan, P. He, and J. Zhang, "Selfgait: A spatiotemporal representation learning method for self-supervised gait recognition," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 2570–2574.
- [62] Y. Ge, F. Zhu, D. Chen, R. Zhao *et al.*, "Self-paced contrastive learning with hybrid memory for domain adaptive object re-id," *Advances in Neural Information Processing Systems*, vol. 33, pp. 11 309–11 321, 2020.
- [63] D. Fu, D. Chen, J. Bao, H. Yang, L. Yuan, L. Zhang, H. Li, and D. Chen, "Unsupervised pre-training for person re-identification," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 14750–14759.
- [64] Z. Zhong, L. Zheng, G. Kang, S. Li, and Y. Yang, "Random erasing data augmentation," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 34, no. 07, 2020, pp. 13 001–13 008.
- [65] C. Feichtenhofer, H. Fan, B. Xiong, R. Girshick, and K. He, "A large-scale study on unsupervised spatiotemporal representation learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 3299–3309.
- [66] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.
- [67] Y. Zhang, P. Sun, Y. Jiang, D. Yu, Z. Yuan, P. Luo, W. Liu, and X. Wang, "Bytetrack: Multi-object tracking by associating every detection box," *arXiv preprint arXiv:2110.06864*, 2021.
- [68] Y. Yuan, X. Chen, and J. Wang, "Object-contextual representations for semantic segmentation," in *European conference on computer vision*. Springer, 2020, pp. 173–190.
- [69] Y. Liu, L. Chu, G. Chen, Z. Wu, Z. Chen, B. Lai, and Y. Hao, "Paddleseg: A high-efficient development toolkit for image segmentation," 2021.
- [70] M. Gutmann and A. Hyvärinen, "Noise-contrastive estimation: A new estimation principle for unnormalized statistical models," in *Proceedings of the thirteenth international conference on artificial intelligence and statistics*. JMLR Workshop and Conference Proceedings, 2010, pp. 297–304.
- [71] R. Qian, T. Meng, B. Gong, M.-H. Yang, H. Wang, S. Belongie, and Y. Cui, "Spatiotemporal contrastive video representation learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 6964–6974.
- [72] R. Li, Y. Zhang, Z. Qiu, T. Yao, D. Liu, and T. Mei, "Motion-focused contrastive learning of video representations," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 2105–2114.
- [73] I. Dave, R. Gupta, M. N. Rizve, and M. Shah, "Tclr: Temporal contrastive learning for video representation," *Computer Vision and Image Understanding*, p. 103406, 2022.
- [74] E. Cole, X. Yang, K. Wilber, O. Mac Aodha, and S. Belongie, "When does contrastive visual representation learning work?" *arXiv preprint arXiv:2105.05837*, 2021.
- [75] S. Chopra, R. Hadsell, and Y. LeCun, "Learning a similarity metric discriminatively, with application to face verification," in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, vol. 1. IEEE, 2005, pp. 539–546.
- [76] R. Liao, C. Cao, E. B. Garcia, S. Yu, and Y. Huang, "Pose-based temporal-spatial network (ptsn) for gait recognition with carrying and clothing variations," in *Chinese conference on biometric recognition*. Springer, 2017, pp. 474–483.
- [77] Z. Wu, Y. Huang, L. Wang, X. Wang, and T. Tan, "A comprehensive study on cross-view gait based human identification with deep cnns," *IEEE transactions on pattern analysis and machine intelligence*, vol. 39, no. 2, pp. 209–226, 2016.
- [78] K. Shiraga, Y. Makihara, D. Muramatsu, T. Echigo, and Y. Yagi, "Geinet: View-invariant gait recognition using a convolutional neural network," in *2016 International Conference on Biometrics (ICB)*, June 2016, pp. 1–8.
- [79] A. Hermans, L. Beyer, and B. Leibe, "In defense of the triplet loss for person re-identification," *ArXiv*, vol. abs/1703.07737, 2017.
- [80] X. Huang, D. Zhu, H. Wang, X. Wang, B. Yang, B. He, W. Liu, and B. Feng, "Context-sensitive temporal feature learning for gait recognition," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 12 909–12 918.
- [81] Z. Huang, D. Xue, X. Shen, X. Tian, H. Li, J. Huang, and X.-S. Hua, "3d local convolutional neural networks for gait recognition," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 14 920–14 929.
- [82] C. Fan, J. Liang, C. Shen, S. Hou, Y. Huang, and S. Yu, "Opengait: Revisiting gait recognition towards better practicality," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 9707–9716.
- [83] K. Shiraga, Y. Makihara, D. Muramatsu, T. Echigo, and Y. Yagi, "Geinet: View-invariant gait recognition using a convolutional neural network," in *2016 international conference on biometrics (ICB)*. IEEE, 2016, pp. 1–8.
- [84] B. Lin, S. Zhang, and X. Yu, "Gait recognition via effective global-local feature representation and local temporal aggregation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 14 648–14 656.
- [85] H. Dou, P. Zhang, Y. Zhao, L. Dong, Z. Qin, and X. Li, "Gaitmpl: Gait recognition with memory-augmented progressive learning," *IEEE Transactions on Image Processing*, 2022.


- [86] T. Chai, A. Li, S. Zhang, Z. Li, and Y. Wang, "Lagrange motion analysis and view embeddings for improved gait recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 20 249–20 258.
- [87] M. Wang, B. Lin, X. Guo, L. Li, Z. Zhu, J. Sun, S. Zhang, Y. Liu, and X. Yu, "Gaitstrip: Gait recognition via effective strip-based feature representations and multi-level framework," in *Proceedings of the Asian Conference on Computer Vision*, 2022, pp. 536–551.
- [88] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *Advances in neural information processing systems*, vol. 27, 2014.
- [89] The European Parliament and of the Council, "General data protection regulation," 2016.



**Shiqi Yu** is currently an Associate Professor in the Department of Computer Science and Engineering, Southern University of Science and Technology, China. He received his B.E. degree in computer science and engineering from the Chu Kochen Honors College, Zhejiang University in 2002, and Ph.D. degree in pattern recognition and intelligent systems from the Institute of Automation, Chinese Academy of Sciences in 2007. He worked as an Assistant Professor and an Associate Professor in Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences from 2007 to 2010, and as an Associate Professor in Shenzhen University from 2010 to 2019. His research interests include gait recognition, face detection and computer vision.



**Chao Fan** received the B.E. and M.S. degrees from Xi'an University of Technology in 2018 and University of Science and Technology Beijing in 2021, respectively. He is currently a Ph.D. candidate with Department of Computer Science and Engineering, Southern University of Science and Technology. His research interests include gait recognition and contrastive learning.



**Saihui Hou** received the B.E. and Ph.D. degrees from University of Science and Technology of China in 2014 and 2019, respectively. He is currently an Assistant Professor with School of Artificial Intelligence, Beijing Normal University. His research interests include computer vision and machine learning. He recently focuses on gait recognition which aims to identify different people according to the walking patterns.



**Jilong Wang** received his B.S. degree from the University of Science and Technology Beijing in 2019. He is currently a Ph.D. student at the University of Science and Technology of China and studies in the Center for Research on Intelligent Perception and Computing (CRIPAC), National Laboratory of Pattern Recognition (NLPR), Institute of Automation, Chinese Academy of Sciences (CASIA). His research interests are in artificial intelligence, machine learning and computer vision.



**Yongzhen Huang** received the B.E. degree from Huazhong University of Science and Technology in 2006, and the Ph.D. degree from Institute of Automation, Chinese Academy of Sciences in 2011. He is currently an Associate Professor with School of Artificial Intelligence, Beijing Normal University. He has published one book and more than 80 papers at international journals and conferences such as TPAMI, IJCV, TIP, TSMCB, TMM, TCSVT, CVPR, ICCV, ECCV, NIPS, AAAI. His research interests include pat-

tern recognition, computer vision and machine learning.