

OpenGait: Revisiting Gait Recognition Toward Better Practicality

Chao Fan^{1,2}, Junhao Liang^{1,2}, Chuanfu Shen^{3,1}, Saihui Hou^{4,5},
Yongzhen Huang^{4,5}, Shiqi Yu^{1,2*}

¹ Department of Computer Science and Engineering, Southern University of Science and Technology

² Research Institute of Trustworthy Autonomous System, Southern University of Science and Technology

³ Department of Industrial and Manufacturing Systems Engineering, The University of Hong Kong

⁴ School of Artificial Intelligence, Beijing Normal University ⁵ WATRIX.AI

{12131100, 12132342, 11950016}@mail.sustech.edu.cn, {housaihui, huangyongzhen}@bnu.edu.cn, yusq@sustech.edu.cn

Abstract

Gait recognition is one of the most critical long-distance identification technologies and increasingly gains popularity in both research and industry communities. Despite the significant progress made in indoor datasets, much evidence shows that gait recognition techniques perform poorly in the wild. More importantly, we also find that some conclusions drawn from indoor datasets cannot be generalized to real applications. Therefore, the primary goal of this paper is to present a comprehensive benchmark study for better practicality rather than only a particular model for better performance. To this end, we first develop a flexible and efficient gait recognition codebase named OpenGait. Based on OpenGait, we deeply revisit the recent development of gait recognition by re-conducting the ablative experiments. Encouragingly, we detect some imperfect parts of certain prior works, as well as new insights. Inspired by these discoveries, we develop a structurally simple, empirically powerful, and practically robust baseline model, GaitBase. Experimentally, we comprehensively compare GaitBase with many current gait recognition methods on multiple public datasets, and the results reflect that GaitBase achieves significantly strong performance in most cases regardless of indoor or outdoor situations. Code is available at <https://github.com/ShiqiYu/OpenGait>.

1. Introduction

Gait recognition has recently gained growing interest from the vision research community. It utilizes the physiological and behavioral characteristics from walking videos to authenticate individuals's identities [34]. Compared with other biometrics, e.g., face, fingerprint, and iris, gait patterns can be captured from a distance in uncontrolled settings, without requiring any physical contact. As a walking

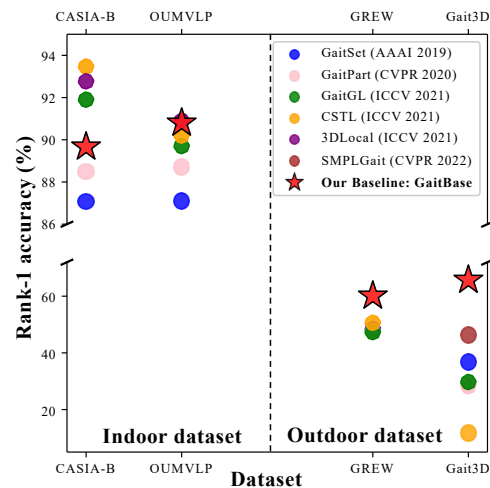


Figure 1. Performance of popular models and ours baseline on 4 major gait datasets [2, 5, 11, 12, 18, 41]. The left two are indoor datasets [32, 38], the right two are outdoor datasets [41, 43].

behavior, gait is also hard to disguise and thus promisingly robust for usual subject-related covariates, such as dressing, carrying, and standing conditions. These advantages make gait recognition suitable for public security applications, e.g., criminal investigation, and suspect tracking [35].

With the boom of deep learning, gait recognition in the laboratory [32, 38] has achieved significant progress [2, 5, 18] over the last decade. However, much evidence [41, 43] reveal that gait recognition techniques may not perform optimally in the wild. As shown in Figure 1, most existing gait recognition methods suffer an over 40% accuracy degradation when transitioning from indoor to outdoor datasets. Typically, this performance gap should be mainly caused by real-world noisy factors, such as complex occlusion, background variation, and illumination changes.

Nevertheless, our further ablative study shows that this situation is not unique, as many of the conclusions drawn

*Corresponding Author

in prior works vary with different datasets. Therefore, besides proposing an improved model for better performance, the primary objective of this paper is to present a comprehensive benchmark study to revisit gait recognition for enhanced practicality. To this end, we make efforts in the following three aspects.

Firstly, to the best of our knowledge, previous works mainly develop the models on their code repository and rely heavily on the indoor gait datasets, particularly CASIA-B [38] and OU-MVLP [32]. To accelerate the real-world applications, we appeal to pay more attention to outdoor gait datasets, such as GREW [43] and Gait3D [41]. Additionally, this paper also considers building a unified evaluation platform, which covers the various state-of-the-art methods and testing datasets, is highly desired. Accordingly, we propose a flexible and efficient gait recognition codebase with PyTorch [24] and name it **OpenGait**.

To ensure extensibility and reusability, OpenGait supports the following features: (1) **Various datasets**, *e.g.*, the indoor CASIA-B [38] and OU-MVLP [32], the outdoor GREW [43] and Gait3D [41]. (2) **State-of-the-art methods**, *e.g.*, GaitSet [2], GaitPart [5], GLN [10], GaitGL [18], SMPLGait [41], GaitEdge [16], and so on. (3) **Multiple popular frameworks**, *e.g.*, the end-to-end, multi-modality, and contrastive learning paradigms. Recently, OpenGait has been widely employed in two of the major international gait recognition competitions, *i.e.* HID [37]¹, and GREW [43]. Encouragingly, all of the top-10 winning teams at HID2022 [37] have utilized OpenGait as their codebase and extended OpenGait to develop new solutions.

Based on OpenGait, we reproduce many progressive methods [2, 5, 18], and the results have been shown in Figure 1. More importantly, we conduct a comprehensive re-evaluation of various commonly accepted conclusions by re-implementing the ablation study on recently-built outdoor gait datasets. To our surprise, we find that the MGP branch proposed by GaitSet [2], the FConv proposed by GaitPart [5], the local feature extraction branch proposed by GaitGL [18], and the SMPL branch proposed by SMPLGait [41], do not exhibit superiority on the outdoor datasets. Moreover, our thorough exploration of potential causes reveals several hidden limitations of prior gait research, such as the lack of comprehensive ablation study, outdoor dataset evaluation, and a strong backbone network.

Inspired by the above discoveries, we develop a simple yet strong baseline model, named **GaitBase**. Specifically, GaitBase is composed of several essential parts, some of which are simple and commonly utilized rather than intricately constructed modules. Even no bells or whistles, GaitBase can achieve comparable or even superior performance on indoor gait datasets. As to the datasets in the wild, GaitBase outperforms recently proposed methods and

reaches a new state-of-the-art. Furthermore, we also conduct a comprehensive study to verify that GaitBase is a structurally simple, experimentally powerful, and empirically robust baseline model for gait recognition.

In summary, this paper contributes future works from three aspects: (1) OpenGait, a unified and extensible platform, is built to facilitate the comprehensive study of gait recognition. (2) We deeply revisit the recent gait recognition developments and consequently bring many new insights for further gait recognition research. (3) We provide a structurally simple and experimentally powerful baseline model, GaitBase, which can inspire the future design of gait recognition algorithms. We hope the works in the paper can inspire researchers to develop more advanced gait recognition methods for real-world applications.

2. Related Work

According to the classical taxonomy, gait recognition methods can be roughly grouped to two categories, including the model-based and appearance-based methods.

Model-based Gait Recognition methods [15, 17, 33] tend to take the estimated underlying structure of human body as input, such as 2D/3D pose and SMPL [20] model. Specifically, PoseGait [17] used 3D body pose and human prior knowledge to overcome the changes in clothing, GaitGraph [33] introduced graph convolutional network for 2D skeleton-based gait representation learning, and HM-RGait [15] fine-tuned a pre-trained human mesh recovery network to construct the end-to-end SMPL-based model. It is also worth mentioning some model-based multi-modality frameworks, such as SMPLGait [41] that exploited the 3D geometrical information from SMPL model to enhance the gait appearance feature learning, and BiFusion [26] that integrated skeletons and silhouettes to capture the rich gait spatio-temporal features. Our OpenGait supports all the aforementioned human body models. However, while these human models are theoretically robust against noisy factors like carrying and dressing, they often struggle with low-resolution situations and, as a result, may lack practicality in some real-world scenarios.

Appearance-based Gait Recognition methods directly learn shape features from the input video, which suit the low-resolution conditions and thus attract increasing attention. With the boom of deep learning, most current appearance-based works focus on spatial feature extraction and gait temporal modeling. Specifically, GaitSet [2] innovatively regarded the gait sequence as a set and utilized a maximum function to compress the sequence of frame-level spatial features. Thanks to its simplicity and effectiveness, GaitSet has become one of the most influential gait recognition works in recent years. GaitPart [5] carefully explored the local details of input silhouette and modeled the temporal dependencies by Micro-motion Capture Module [5]. GaitGL [18] argued that the spatially global gait represen-

¹HID 2023: <https://hid2023.iapr-tc4.org>

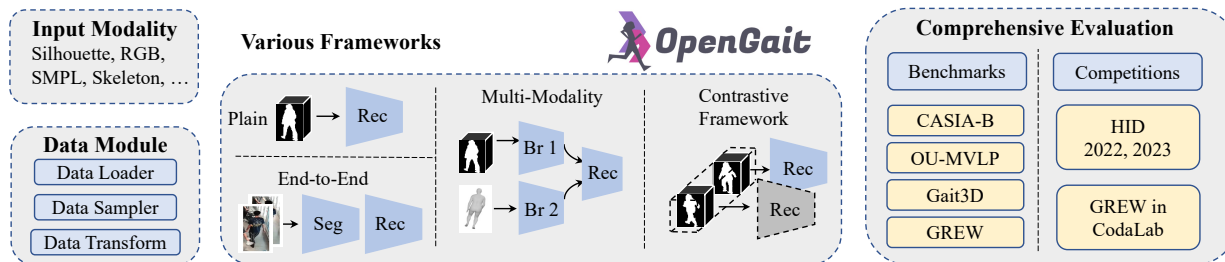


Figure 2. The design principles of proposed codebase **OpenGait**. Seg is for *segmentation*, Rec is for *recognition*, and Br is for *branch*.

tations often neglect the details, and the local region-based descriptors cannot capture the relations among neighboring parts, thus developing the global and local convolution layer [18]. CSTL [11] focused on temporal features in three scales to obtain motion representation according to the temporal contextual information. 3DLocal [12] wanted to extract the limb features through 3D local operations at adaptive scales.

In addition to the above, many other outstanding works inspire us a lot as well, *e.g.*, GaitEdge [16] designed an edge-trainable intermediate modality to build the end-to-end gait recognition framework, and GaitSSB [4] collected millions of unlabelled gait sequences and learned the general gait representation with a contrastive framework.

This paper introduces the OpenGait codebase, which is compatible with almost all of the aforementioned methods. Additionally, as research directions shift from indoor to outdoor environments, we re-analyzed many of these methods from an experimental perspective and gained new insights. Finally, we developed a simple yet powerful baseline model, GaitBase, which outperforms prior works by a significant margin on outdoor datasets.

Gait Datasets are also crucial for gait recognition research, and among them, CASIA-B [38] and OU-MVLP [32] are two of the most widely-used indoor datasets. They were proposed in 2006 and 2018, respectively, and captured by requiring subjects to walk around the laboratory, which significantly differs from the real-world scenarios. In response to the need for real-world applications, GREW [43] and Gait3D [41] were collected in the wild in 2021 and 2022, respectively. However, most existing works only verify their effectiveness on indoor datasets, which poses a high risk of vulnerability in practical usage.

Revisit Deep Gait Recognition. Recently, several survey works [7, 28, 29] have investigated almost all published papers on deep gait recognition. However, a more comprehensive and detailed analysis of these methods is still lacking. In some other fields, there are critical papers revisiting previous methods, such as recommendation systems [6], metric learning [22], and unsupervised domain adaptation [23]. OpenGait aims to address this gap in the field of gait recognition by experimentally reviewing recent works and highlighting some concerns about their robustness.

Codebase. There are many works providing infrastructure in the deep learning research community, such as a codebase for the specific research topic. For example, Amos *et al.* [1] proposed OpenFace, a face recognition library that bridges the gap between public face recognition systems and industry-leading private systems. In the field of object detection, a PyTorch toolbox called MMDetection [3] supports almost all popular detection methods, providing a convenient platform for systematic comparison. With the rapid development of gait recognition, the need for an infrastructure code platform has become increasingly prominent.

3. OpenGait

Over the past few years, numerous new frameworks and evaluation datasets have emerged for gait recognition. However, the lack of a unified and fair evaluation platform cannot be overlooked. To facilitate academic research and practical applications, this paper presents a PyTorch-based [25] toolbox, **OpenGait**, as a reasonable and dependable solution to address this issue.

3.1. Design Principles of OpenGait

As shown in Figure 2, our developed OpenGait covers the following highlight features.

Compatibility with Diverse Gait Modalities. Gait recognition has a variety of input modalities. The usual ones include silhouette images, 2D/3D skeletons, and more recently emerging modalities include SMPL parameters [15, 41] and RGB images [16, 31]. Existing open-source repositories mostly only support one of these modalities, while our OpenGait is designed to support all of these modalities.

Compatibility with Various Frameworks. Currently, more and more novel gait recognition methods have emerged, such as multi-modalities [41], end-to-end [15, 16, 31], and contrastive learning [4]. As mentioned above, most open-source methods narrowly focus on their own models, so extending to multiple frameworks is difficult. Fortunately, OpenGait supports all of the above frameworks.

Support for Various Evaluation Datasets. OpenGait is a comprehensive toolbox that includes datasets commonly used by researchers. We offer full support for indoor gait datasets such as CASIA-B and OU-MVLP, as well as newly proposed outdoor wild datasets GREW and Gait3D. OpenGait provides a range of uniquely designed functions for

these datasets, from data preprocessing to final evaluation. In addition, it is noteworthy that two of the major international competitions, HID [37] and GREW [43], are compatible with OpenGait. Many winning teams have utilized OpenGait to develop new solutions.

Support for State-of-the-arts. We reproduce many previous state-of-the-art methods, including GaitSet [2], GaitPart [5], GLN [10], GaitGL [18], Gait3D [41], SMPLGait [41], GaitEdge [16] and GaitSSB [4]. Reproduced performances are comparable to or even better than the results reported by the original papers. Rich official examples can help beginners get started with the project better and faster. In addition, this also provides an infrastructure for a more comprehensive and systematic comparison later.

3.2. Main Modules

Technically, we follow the design of most PyTorch deep learning projects and divide OpenGait into three modules, *data*, *modeling*, and *evaluation*, as shown in Figure 2.

Data module contains data loader, data sampler and data transform, which are responsible for loading, sampling, and pre-processing the input data flow respectively.

Modeling module is built on top of a base class (`BaseModel`) that pre-defines many behaviors of the deep model during training and testing phases, including optimization and inference. The four essential components of current gait recognition algorithms, namely *backbone*, *neck*, *head*, and *loss*, can be customized in this class.

Evaluation module is used to evaluate the obtained model. It is well known that different datasets have various evaluation protocols, and we unify them into OpenGait to free researchers from these tedious details.

4. Revisit Deep Gait Recognition

With the help of OpenGait, we can revisit several typical gait recognition methods comprehensively. Some insights different from those presented in the original papers have been found from our fair ablation studies.

4.1. Experimental Recheck on Previous Methods

We notice that most previous works only verify the effectiveness on the indoor gait datasets, *i.e.*, CASIA-B [38] and OU-MVLP [32], and the further ablation study is only conducted based on the CASIA-B. In this section, excluding reproducing the officially stated performance, we re-conduct some ablation studies based on the newly-built outdoor dataset, *i.e.*, Gait3D, to experimentally check the algorithm’s robustness for practical gait data.

Re-conduct Ablation Study on GaitSet. With taking the silhouettes as input, GaitSet [2] treats the gait sequence as an unordered set and uses a simple maximum pooling function along the temporal dimension, called Set Pooling (SP), to generate the set-level understanding of the entire input gait sequence. GaitSet [2] provides insights for

Table 1. The effect of MGP and Multi-scale HPP in GaitSet [2]. In the original GaitSet [2], all the obtained partial vectors will be concatenated into just one feature vector used for the final evaluation. However, following works [5, 18, 41] find this concatenation unnecessary and take the average of partial distance for evaluation. Therefore, OpenGait follows this manner and reproduces a better performance than the official results.

MGP	Multi-scale HPP	CASIA-B			Gait3D	
		NM	BG	CL	R-1	R-5
✓	×	95.9	90.3	74.2	44.3	64.7
	✓	95.8	90.4	73.2	44.3	64.4
×	×	95.3	90.5	74.0	45.8	65.1
	✓	94.5	89.1	72.3	43.7	63.8

many subsequent works thanks to its simplicity and effectiveness. However, we find that the other two important components in GaitSet [2], namely the parallel Multi-layer Global Pipeline (MGP) and pyramid-like Horizontal Pyramid Pooling (HPP) [8], do not work well enough on both the indoor CASIA-B and outdoor Gait3D datasets. Specifically, as shown in Figure 3 (a), MGP can be regarded as an additional branch that aggregates the hierarchical set-level characteristics. HPP [8] follows the fashion feature pyramid structure aiming to extract multi-scale part-level representations. As shown in Table 1, if we strip MGP or remove the multi-scale mechanism in HPP from the official GaitSet [2], the obtained model could reach the same or even better performance on both CASIA-B and Gait3D, saving over 80% training weights. This result indicates that the set-level characteristics extracted by bottom convolution blocks may be hard to benefit the final gait representation. Besides, the multi-scale mechanism in HPP also provides no extra discriminative features. The cause may be that the employed statistical pooling functions are too weak to learn additional knowledge from various-scale human body parts.

Table 2. The effect of FConv in GaitPart [5].

FConv	CASIA-B			Gait3D	
	NM	BG	CL	R-1	R-5
✓	96.2	91.5	78.7	29.2	48.6
×	95.6	88.4	76.1	36.2	57.0

Re-conduct Ablation Study on GaitPart. One of the core contributions of GaitPart [5] is to point out the importance of local details with the proposed Focal Convolution Layer (FConv). Figure 3 (b) shows the receptive field’s expansion of the top-layer neuron in the network composed of the regular and focal convolution layers. Technically, FConv splits the input feature map into several parts horizontally and then performs a regular convolution over each part separately. As shown in Table 2, we get a much higher performance (Rank-1: +7.0%) on Gait3D with changing FConv to the regular convolution layer. This phenomenon exhibits

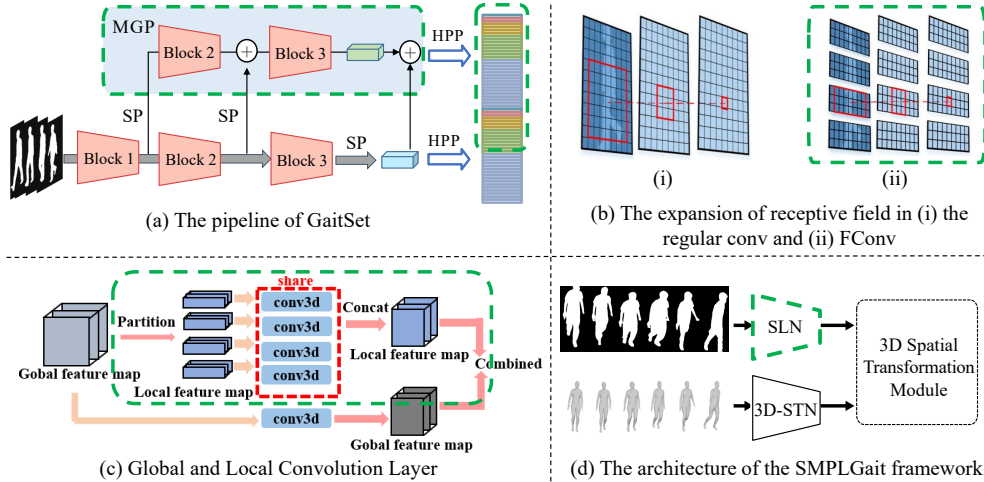


Figure 3. The key modules of four previous SOTA methods. The modules enclosed by dotted green lines are the ones we remove or replace for comparison. (a): **MGP** and multi-scale mechanism in **HPP** are removed [2]. (b): FConvs are replaced with the regular convolution layers [5]. (c): The upper local feature branch is replaced with an independent conv3d that is identical to the lower global branch [18]. (d): **SLN** branch is replaced with our stronger ResNet-like backbone [41].

Table 3. The effect of local branch in GaitGL [18].

Local Branch	CASIA-B			Gait3D	
	NM	BG	CL	R-1	R-5
✓	97.4	94.5	83.6	31.4	50.0
×	97.1	93.7	81.9	32.2	52.5

that the extraction of gait features may be seriously affected by directly splitting the feature map due to the low-quality segmentation of wild data. The recently popular shifted window mechanism [19] may address this issues. Besides, since GaitPart requires sequential frames as input, some usual real-world factors in the temporal dimension may negatively impact the final performance as well, such as the inevitable frame drop and walking speed changes.

Re-construct Ablation Study on GaitGL. GaitGL [18] argues that the spatially global-level gait representations often neglect the details of input gait frames. At the same time, the local region-based descriptors cannot capture the relations among neighboring parts, thus developing the global and local convolution layer. As shown in Fig 3 (c), the local branch can be regarded as the FConv [5] employing 3D convolution, while the global branch is a standard 3D convolution layer. Similar to GaitPart [5], as shown in Table 3, removing the local branch can achieve better performance on the outdoor Gait3D dataset.

Re-construct Ablation Study on SMPLGait. As shown in Figure 3 (d), SMPLGait [41] consists of two elaborately-designed branches, *i.e.*, silhouette (SLN) and SMPL (3D-STN) branches. They are respectively used for 2D appearance extraction and 3D knowledge learning. SMPLGait [41] takes advantage of the 3D mesh data available in Gait3D [41] and achieves a gain effect on the top of the

Table 4. The effect of SMPL branch in SMPLGait [41]. SLN presents silhouette branch used in Gait3D. ResNet-like network is the backbone of our GaitBase.

SMPL branch	Silhouette branch			
	SLN		ResNet-like Network	
	R-1	R-5	R-1	R-5
✓	46.3	64.5	55.2	75.7
×	42.9	63.9	56.5	75.2

silhouette branch. However, as shown in Table 4, our experiment demonstrates that the proposed SMPL branch not provides apparent promotion when we give the silhouette branch a strong backbone network. *e.g.*, ResNet-like network that will be built into our strong baseline.

In our view, there are three possible reasons causing the failure of the SMPL branch: a). Though the SMPL model is usually visualized as a dense mesh, its feature vector only possesses tens of dimensions that present the relatively sparse characterization of body shape and posture, making it challenging to enhance the fine-grained description of gait patterns. b). Since the SMPL model is not recognition-oriented, purposefully fine-tuning it may be more optimal than directly utilizing it to depict the subtle individual characteristics [15]. c). In the wild, estimating an accurate SMPL model that finely captures body shape and posture from a single RGB camera is still challenging. In a nutshell, introducing 3D geometrical information from the SMPL model to benefit gait representation learning is well worth further exploration.

4.2. Analysis and Discussion

Building on the findings and insights gained through our analyses, we believe that the model structures of previous gait recognition methods may not be robust enough. This is




Figure 4. The pipeline of GaitBase. DA is for *data augmentation*, ResNet9 is for *ResNet-like backbone*, TP is for *temporal pooling*, HP is for *horizontal pooling* (w/o pyramid structure [8]).

partly due to the fact that earlier research was more lab-oriented and relied on limited datasets that did not capture the complexity of outdoor environments. We provide a more detailed explanation below.

Necessity of Outdoor Evaluation. The evaluation of existing methods has primarily focused on the indoor CASIA-B [38] and OU-MVLP [32]. We argue that this approach suffers from three significant drawbacks: a) **Indoor settings.** The walking videos are captured by an all-side camera array, and the subjects are requested to follow a particular course. It makes the data collection condition obviously different from the real-world scenarios. b) **Simple background.** Simple laboratory background cannot reflect the complex background changes of wild scenes. c) **Outdated processing methods.** The raw RGB videos are processed by the outdated background subtraction algorithm.

Recently, new large-scale outdoor gait datasets, such as GREW [43] and Gait3D [41], have emerged to promote gait recognition from in-the-lab setting to in-the-wild scenario. However, despite these developments, most lately-published works [36, 39, 40] still rely excessively on indoor gait datasets. In order to advance the development and applications of gait recognition systems, we suggest the community pay more attention to outdoor datasets.

Necessity of Comprehensive Ablation Study. The ablation study is recognized as the primary means of evaluating the effectiveness of individual components within a proposed method. However, we believe that the conclusions drawn from ablative experiments performed only on CASIA-B [38] may have limited practical applicability. This is due to two primary reasons: a) CASIA-B only contains 50 subjects for evaluation. Such few testing subjects make the results vulnerable to noisy factors. b) In current applications, the pedestrian segmentation is typically achieved by deep models, rather than the outdated background subtraction algorithm used by CASIA-B.

Totally, conducting a comprehensive ablation study on various large-scale datasets can provide a more robust hyper-parameter configuration and model structure.

Necessity of A Strong Backbone. The quality of a model is largely determined by the capabilities of its backbone, and a poor backbone can unjustly overestimate the effectiveness of additional modules. The evolution of CNN net-

Table 5. Architecture of our ResNet-like backbone. k, c, b respectively denotes kernel size, channels and number of blocks.

Layer	Structure [$k \times k, c$] $\times b$
Conv0	$[3 \times 3, 64] \times 1, \text{stride} = 1$
Block1	$[3 \times 3, 64]$ $[3 \times 3, 64] \times 1, \text{stride} = 1$
Block2	$[3 \times 3, 64]$ $[3 \times 3, 128] \times 1, \text{stride} = 2$
Block3	$[3 \times 3, 128]$ $[3 \times 3, 256] \times 1, \text{stride} = 2$
Block4	$[3 \times 3, 256]$ $[3 \times 3, 512] \times 1, \text{stride} = 1$

work models has progressed from shallow to deep, resulting in the emergence of excellent backbone networks such as AlexNet [14], VGG-16 [30], and ResNet [9]. However, previous works [2, 11, 41] on gait recognition have predominantly relied on plain convolutional neural networks, consisting of several pure convolution layers. As gait recognition research advances towards applications, and larger-scale, real-world datasets [41, 43] become more accessible, it is clear that a more robust and powerful backbone network is highly desirable to achieve accurate and reliable results.

5. A Strong Baseline: GaitBase

Using the insights gained from the above analysis, this section aims to offer a simple but strong baseline model to reach the totally comparable or even better performance for both indoor and outdoor evaluations. No bells or whistles, the obtained model should be structurally simple, experimentally powerful, and empirically robust to serve as a new baseline for further research. To this end, we build a silhouette-based model, *i.e.*, GaitBase.

5.1. Pipeline

We employ several easy and widely-accepted modules to form GaitBase. As shown in Figure 4, GaitBase follows the popular set-based and part-based paradigm and takes a ResNet-like [9] network as the backbone. Specifically, the ResNet-like backbone will transform each input silhouette frame into a 3D feature map with the height, width, and channel dimensions. And then, a Temporal Pooling module will aggregate the obtained sequence of feature maps by performing the maximization along the temporal dimension, outputting a set-level understanding of the input gait sequence, *i.e.*, a 3D feature map. Next, the obtained feature map will be horizontally divided into several parts, and each part will be pooled into a feature vector. Accordingly, we get several feature vectors and further use a separate fully connected layer to map them into the metric space. Finally, we employ the widely-used BNNeck [21] to adjust feature space, and the separate triplet and cross-entropy losses are utilized to drive the training process. The overall loss function is formulated by $L = L_{triplet} + L_{ce}$.

Table 6. 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

5.2. ResNet-like Backbone

ResNet [9] is one of the most successful deep models and has been broadly used as the backbone for many vision tasks. As for gait recognition, most of the current works still take a shallow stack of several standard or customized convolution layers as the backbone. In this paper, we develop a ResNet-like network to serve as the backbone of GaitBase. As shown in Table 5, the network comprises initial convolution layer followed by a stack of four basic residual blocks. All the layers are equipped with BN [13] and ReLU activation layers which are skipped for conciseness.

5.3. Data Augmentation

In order to match the practical usage, we explore the data augmentation (DA) strategy specific to gait silhouettes that previous works have seldom considered. Our experiments show that several spatial augmentation operations, such as horizontal flipping, rotation, perspective, and affine transformations, can improve recognition accuracy. Additionally, we discover that allowing for an unfixed length of the input gait sequence during the training phase can further enhance the final performance.

5.4. Comparison with Other Methods

The driving philosophy behind the development of GaitBase is the belief that simplicity is beautiful. Accordingly, we adopt several succinct, widespread, and validated techniques to make GaitBase robust and efficient. These include treating a gait sequence as a set, using part-based models for human body description, utilizing the widely used ResNet-like network as a backbone, and leveraging the general data augmentation strategy to fit the practical usage. These settings do not require any unique designs but still reach comparable or even better performance on the existing gait datasets, regardless of indoor or outdoor ones.

6. Experiment

In this section, we mainly present a comprehensive study to show the effectiveness and robustness of GaitBase as well as its components on various gait datasets.

6.1. Datasets

Four public gait datasets are utilized, involving the most widely-used indoor datasets, CASIA-B [38] and OU-MVLP [32], and the largest gait dataset in the wild, GREW [43] and Gait3D [41]. Table 6 displays the statistics on the number of identities and sequences that are covered

Table 7. Implementation details. The batch size (q, k) indicates q subjects with k sequences per subject.

DataSet	Batch Size	MultiStep Scheduler	Steps
CASIA-B	(8, 16)	(20k, 40k, 50k)	60k
OU-MVLP	(32, 8)	(60k, 80k, 100k)	120k
GREW	(32, 4)	(80k, 120k, 150k)	180k
Gait3D	(32, 4)	(20k, 40K, 50k)	60k

in each of these datasets. The following section provides a detailed overview of the collection process for each of these datasets, emphasizing the significant differences that exist between indoor and outdoor datasets.

CASIA-B. Each subject is asked to follow a particular course with three walking conditions, *i.e.*, normal walking (NM), walking with bags (BG), and walking with coats (CL). The obtained videos were captured by an all-side camera array with 11 filming viewpoints. The gait silhouettes were generated through the old background subtraction algorithm.

OU-MVLP is one of the largest public gait datasets. However, similar to CASIA-B, each subject in OU-MVLP was required to walk in a fixed course with fixed cameras and produce the gait silhouette by the outdated background subtraction algorithm. OU-MVLP contains only one walking status for each subject and thus lacks the clothing and carrying changes, making the recognition task relatively easy.

GREW is the largest gait dataset in the wild up to date, to our 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. Besides the tens of thousands of identities, many other human attributes have been annotated, *e.g.*, 2 gender, 14 age groups, 5 carrying conditions, and 6 dressing styles. Therefore, GREW is believed to include adequate and diverse practical variations.

Gait3D is a large-scale gait dataset in the wild. It was collected in a supermarket and contains 1,090 hours of videos with $1,920 \times 1,080$ resolution and 25 FPS.

Our implementation follows official protocols, including the training/testing and gallery/probe set partition strategies. Rank-1 accuracy is used as the primary evaluation metric.

6.2. Implementation Details

Table 7 displays the main hyper-parameters of our experiments. All the source codes are based on OpenGait, which is proposed in this work. The input resolution for all datasets is set to 64×44 . We set the triplet loss margin as 0.2, the initial learning rate and the weight decay of SGD optimizer [27] as 0.1 and 0.0005, respectively. The input frame number is unfixed. Furthermore, we employ Horizontal Flipping, Random Erasing [42], Rotation, Perspective Transformation, and Affine Transformation as spatial data augmentation techniques. Their details are presented

Table 8. The performance comparison on four major datasets: rank-1 accuracy (%) are reported. CASIA-B* [16] is a re-processed version of CASIA-B. The highest two recognition results are in **bold**. DA is for *data augmentation*.

Method	Condition		Dataset								
			CASIA-B			CASIA-B*			OU-MVLP	GREW	Gait3D
	BNNek [21]	DA	NM	BG	CL	NM	BG	CL			
GaitBase (Ours)	×	×	96.7	91.9	74.9	94.3	87.9	73.5	89.6	54.7	56.9
	×	✓	97.8	93.9	77.6	96.2	91.4	77.3	88.3	57.9	62.0
	✓	×	97.3	92.9	78.0	94.5	90.0	75.9	90.8	57.7	54.7
	✓	✓	97.6	94.0	77.4	96.5	91.5	78.0	90.0	60.1	64.6
GaitSet [2]	AAAI 2019		95.8	90.0	75.4	92.3	86.1	73.4	87.1	48.4	36.7
GaitPart [5]	CVPR 2020		96.1	90.7	78.7	93.1	86.0	75.1	88.7	47.6	28.2
GaitGL [18]	ICCV 2021		97.4	94.5	83.8	94.1	90.0	81.4	89.7	47.3	29.7
CSTL [11]	ICCV 2021		98.0	95.4	87.0	-	-	-	90.2	50.6	11.7
3DLocal [12]	ICCV 2021		98.3	95.5	84.5	-	-	-	90.9	-	-
SMPLGait [41]	CVPR 2022		-	-	-	-	-	-	-	-	46.3

in the supplementary.

6.3. Comparison with Other State-of-the-Arts

We compare our GaitBase with other published state-of-the-art methods, *e.g.*, GaitSet [2], GaitPart [5], GaitGL [18], CSTL [11], 3DLocal [12], and SMPLGait [41], on the various popular datasets, *i.e.*, CASIA-B [38], CASIA-B* [16], OU-MVLP [32], GREW [43] and Gait3D [41]. Among these datasets, CASIA-B* [16] is a re-segmented version of CASIA-B.

As shown in Table 8, excluding CASIA-B, GaitBase achieves a totally competitive or even significantly better performance than other state-of-the-art methods. More importantly, GaitBase exceeds other methods by a more considerable margin on the datasets captured from practical scenarios, *i.e.*, +9.5% on GREW and +18.3% on Gait3D, richly exhibiting both the effectiveness and robustness of our strong baseline at the time of dealing the troublesome real-world noisy factors.

As for CASIA-B, our GaitBase achieves a less competitive accuracy. It illustrates that previous methods are finely designed for the CASIA-B dataset to achieve high accuracy while ignoring its bias from wild data. Additionally, our method still achieves relatively high performance on re-segmented CASIA-B*.

6.4. Ablation Study

To study the effectiveness of BNNeck and DA, we conduct additional experiments on all the mentioned datasets, as shown in Table 8.

Impact of BNNeck. It should be noted that GaitBase with BNNeck outperforms GaitBase without BNNeck on almost all datasets, which shows that BNNeck generally plays a positive role in GaitBase.

Impact of DA. The results show that with data augmentation, the performance of GaitBase on the outdoor datasets is improved (+2.4% for GREW and +9.9% for Gait3D). In comparison, the performance on the indoor datasets is degraded (-0.8% for OU-MVLP). This observation indicates

that our data augmentation strategy is more capable of suppressing over-fitting caused by the noisy variations in practical scenarios.

Impact of ResNet9. We have also tried the standard ResNet50 [9] as the backbone. Further ablative experiments show that ResNet50 does not bring significant performance improvement on Gait3D (+1.9%) but causes serious over-fitting on CASIA-B* (converges better but performs worse, *i.e.*, -5.2%, -6.8%, -9.6% on NM, BG and CL). Moreover, using ResNet50 also results in over 3× computation cost (3.939 *v.s.* 1.183 GFLOPs per silhouette image). Since the primary goal of GaitBase is to provide a practical yet powerful baseline, we finally choose the simple ResNet9.

7. Conclusion

This paper presents a codebase, OpenGait, developed for deep gait recognition. OpenGait provides a fair and easy-to-use platform for future gait recognition works, facilitating researchers to implement new ideas more efficiently. Subsequently, we implemented most state-of-the-art methods in OpenGait and compared them fairly. Some insights different from those in the original papers have been found in our fair ablation studies. Based on those findings and understandings, a simple but efficient gait recognition model, GaitBase, is proposed. GaitBase can achieve the best performance on most gait datasets, especially on datasets in the wild. So GaitBase can serve as a new baseline for future studies.

Acknowledgement. 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 Stable Support Plan Program of Shenzhen Natural Science Fund under Grant 20200925155017002, in part by the National Key Research and Development Program of China under Grant 2020AAA0140002, and in part by the Shenzhen Technology Plan Program under Grant JSGG20201103091002008, and Grant KQTD20170331093217368.

References

- [1] Brandon Amos, Bartosz Ludwiczuk, Mahadev Satyanarayanan, et al. Openface: A general-purpose face recognition library with mobile applications. *CMU School of Computer Science*, 6(2):20, 2016. **3**
- [2] Hanqing Chao, Yiwei He, Junping Zhang, and Jianfeng Feng. Gaitset: Regarding gait as a set for cross-view gait recognition. In *Proceedings of the AAAI conference on artificial intelligence*, pages 8126–8133, 2019. **1, 2, 4, 5, 6, 8**
- [3] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. **3**
- [4] Chao Fan, Saihui Hou, Jilong Wang, Yongzhen Huang, and Shiqi Yu. Learning gait representation from massive unlabelled walking videos: A benchmark. *arXiv preprint arXiv:2206.13964*, 2022. **3, 4**
- [5] Chao Fan, Yunjie Peng, Chunshui Cao, Xu Liu, Saihui Hou, Jiannan Chi, Yongzhen Huang, Qing Li, and Zhiqiang He. Gaitpart: Temporal part-based model for gait recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14225–14233, 2020. **1, 2, 4, 5, 8**
- [6] Maurizio Ferrari Dacrema, Paolo Cremonesi, and Dietmar Jannach. Are we really making much progress? a worrying analysis of recent neural recommendation approaches. In *Proceedings of the 13th ACM conference on recommender systems*, pages 101–109, 2019. **3**
- [7] Claudio Filipi Gonçalves dos Santos, Diego de Souza Oliveira, Leandro A. Passos, Rafael Gonçalves Pires, Daniel Felipe Silva Santos, Lucas Pascotti Valem, Thierry P. Moreira, Marcos Cleison S. Santana, Mateus Roder, Jo Paulo Papa, et al. Gait recognition based on deep learning: A survey. *ACM Computing Surveys (CSUR)*, 55(2):1–34, 2022. **3**
- [8] Yang Fu, Yunchao Wei, Yuqian Zhou, Honghui Shi, Gao Huang, Xinchao Wang, Zhiqiang Yao, and Thomas Huang. Horizontal pyramid matching for person re-identification. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence*, AAAI'19/IAAI'19/EAAI'19. AAAI Press, 2019. **4, 6**
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. **6, 7, 8**
- [10] Saihui Hou, Chunshui Cao, Xu Liu, and Yongzhen Huang. Gait lateral network: Learning discriminative and compact representations for gait recognition. In *European Conference on Computer Vision*, pages 382–398. Springer, 2020. **2, 4**
- [11] Xiaohu Huang, Duowang Zhu, Hao Wang, Xinggang Wang, Bo Yang, Botao He, Wenyu Liu, and Bin Feng. Context-sensitive temporal feature learning for gait recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12909–12918, 2021. **1, 3, 6, 8**
- [12] Zhen Huang, Dixiu Xue, Xu Shen, Xinmei Tian, Houqiang Li, Jianqiang Huang, and Xian-Sheng Hua. 3d local convolutional neural networks for gait recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14920–14929, 2021. **1, 3, 8**
- [13] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015. **7**
- [14] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017. **6**
- [15] Xiang Li, Yasushi Makihara, Chi Xu, Yasushi Yagi, Shiqi Yu, and Mingwu Ren. End-to-end model-based gait recognition. In *Proceedings of the Asian conference on computer vision*, 2020. **2, 3, 5**
- [16] Junhao Liang, Chao Fan, Saihui Hou, Chuanfu Shen, Yongzhen Huang, and Shiqi Yu. Gaitedge: Beyond plain end-to-end gait recognition for better practicality. *arXiv preprint arXiv:2203.03972*, 2022. **2, 3, 4, 8**
- [17] Rijun Liao, Shiqi Yu, Weizhi An, and Yongzhen Huang. A model-based gait recognition method with body pose and human prior knowledge. *Pattern Recognition*, 98:107069, 2020. **2**
- [18] Beibei Lin, Shunli Zhang, and Xin Yu. Gait recognition via effective global-local feature representation and local temporal aggregation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14648–14656, 2021. **1, 2, 3, 4, 5, 8**
- [19] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. **5**
- [20] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6):1–16, 2015. **2**
- [21] Hao Luo, Youzhi Gu, Xingyu Liao, Shenqi Lai, and Wei Jiang. Bag of tricks and a strong baseline for deep person re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 0–0, 2019. **6, 8**
- [22] Kevin Musgrave, Serge Belongie, and Ser-Nam Lim. A metric learning reality check. In *European Conference on Computer Vision*, pages 681–699. Springer, 2020. **3**
- [23] Kevin Musgrave, Serge Belongie, and Ser-Nam Lim. Unsupervised domain adaptation: A reality check. *arXiv preprint arXiv:2111.15672*, 2021. **3**
- [24] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. **2**

- [25] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. [3](#)
- [26] Yunjie Peng, Saihui Hou, Kang Ma, Yang Zhang, Yongzhen Huang, and Zhiqiang He. Learning rich features for gait recognition by integrating skeletons and silhouettes. *arXiv preprint arXiv:2110.13408*, 2021. [2](#)
- [27] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951. [7](#)
- [28] Alireza Sepas-Moghaddam and Ali Etemad. Deep gait recognition: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. [3](#)
- [29] Chuanfu Shen, Shiqi Yu, Jilong Wang, George Q Huang, and Liang Wang. A comprehensive survey on deep gait recognition: Algorithms, datasets and challenges. *arXiv preprint arXiv:2206.13732*, 2022. [3](#)
- [30] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. [6](#)
- [31] Chunfeng Song, Yongzhen Huang, Yan Huang, Ning Jia, and Liang Wang. Gaitnet: An end-to-end network for gait based human identification. *Pattern recognition*, 96:106988, 2019. [3](#)
- [32] Noriko Takemura, Yasushi Makihara, Daigo Muramatsu, Tomio Echigo, and Yasushi Yagi. Multi-view large population gait dataset and its performance evaluation for cross-view gait recognition. *IPSJ Transactions on Computer Vision and Applications*, 10(1):1–14, 2018. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [33] Torben Teepe, Ali Khan, Johannes Gilg, Fabian Herzog, Stefan Hörmann, and Gerhard Rigoll. Gaitgraph: graph convolutional network for skeleton-based gait recognition. In *2021 IEEE International Conference on Image Processing (ICIP)*, pages 2314–2318. IEEE, 2021. [2](#)
- [34] Liang Wang, Tieniu Tan, Huazhong Ning, and Weiming Hu. Silhouette analysis-based gait recognition for human identification. *IEEE transactions on pattern analysis and machine intelligence*, 25(12):1505–1518, 2003. [1](#)
- [35] Zifeng Wu, Yongzhen Huang, Liang Wang, Xiaogang Wang, and Tieniu Tan. A comprehensive study on cross-view gait based human identification with deep cnns. *IEEE transactions on pattern analysis and machine intelligence*, 39(2):209–226, 2016. [1](#)
- [36] Weilai Xiang, Hongyu Yang, Di Huang, and Yunhong Wang. Multi-view gait video synthesis. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 6783–6791, 2022. [6](#)
- [37] Shiqi Yu, Yongzhen Huang, Liang Wang, Yasushi Makihara, Shengjin Wang, Md Atiqur Rahman Ahad, and Mark Nixon. Hid 2022: The 3rd international competition on human identification at a distance. In *2022 IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–9. IEEE, 2022. [2](#), [4](#)
- [38] Shiqi Yu, Daoliang Tan, and Tieniu Tan. A framework for evaluating the effect of view angle, clothing and carrying condition on gait recognition. In *18th International Conference on Pattern Recognition (ICPR'06)*, volume 4, pages 441–444. IEEE, 2006. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [39] Weichen Yu, Hongyuan Yu, Yan Huang, and Liang Wang. Generalized inter-class loss for gait recognition. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 141–150, 2022. [6](#)
- [40] Pengyi Zhang, Huanzhang Dou, Yunlong Yu, and Xi Li. Adaptive cross-domain learning for generalizable person re-identification. In *European Conference on Computer Vision*, pages 215–232. Springer, 2022. [6](#)
- [41] Jinkai Zheng, Xinchun Liu, Wu Liu, Lingxiao He, Chenggang Yan, and Tao Mei. Gait recognition in the wild with dense 3d representations and a benchmark. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [42] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. Random erasing data augmentation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 13001–13008, 2020. [7](#)
- [43] Zheng Zhu, Xianda Guo, Tian Yang, Junjie Huang, Jiankang Deng, Guan Huang, Dalong Du, Jiwen Lu, and Jie Zhou. Gait recognition in the wild: A benchmark. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14789–14799, 2021. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)