LIDAR GAIT: Benchmarking 3D Gait Recognition with Point Clouds

Chuanfu Shen[1,2] Chao Fan[2] Wei Wu[2] Rui Wang[2] George Q. Huang[1] Shiqi Yu[2]

[1]The University of Hong Kong   [2]Southern Univeristy of Science and Technology

noahshen@connect.hku.hk gqhuang@hku.hk yusq@sustech.edu.cn
{12131100,12032501,12232385}@mail.sustech.edu.cn

Abstract

Video-based gait recognition has achieved impressive results in constrained scenarios. However, visual cameras neglect human 3D structure information, which limits the feasibility of gait recognition in the 3D wild world. In this work, instead of extracting gait features from images, we explore precise 3D gait features from point clouds and propose a simple yet efficient 3D gait recognition framework, termed multi-view projection network (MVPNet). MVPNet first projects point clouds into multiple depth maps from different perspectives, and then fuse depth images together, to learn the compact representation with 3D geometry information. Due to the lack of point cloud datasets, we build the first large-scale Lidar-based gait recognition dataset, LIDAR GAIT, collected by a Lidar sensor and an RGB camera mounted on a robot. The dataset contains 25,279 sequences from 1,050 subjects and covers many different variations, including visibility, views, occlusions, clothing, carrying, and scenes. Extensive experiments show that, (1) 3D structure information serves as a significant feature for gait recognition. (2) MVPNet not only competes with five representative point-based methods, but it also outperforms existing camera-based methods by large margins. (3) The Lidar sensor is superior to the RGB camera for gait recognition in the wild. LIDAR GAIT dataset and MVPNet code will be publicly available.

1. Introduction

Gait is an essential biometric, which has the unique advantage of human identification at a distance without physical contact. Gait empowers many real-world applications such as human retrieval, forensic identification, and serving robots. Recently, great progress has been made to promote gait recognition from in-the-lab setting [20, 41, 49] to in-the-wild scenario [18, 51, 55, 57]. Despite these studies have made significant contributions to recent advances [5, 7, 12, 28, 54], two inherent problems still remain: (1) *lack of 3D geometry information*, and (2) *poor feasibility*
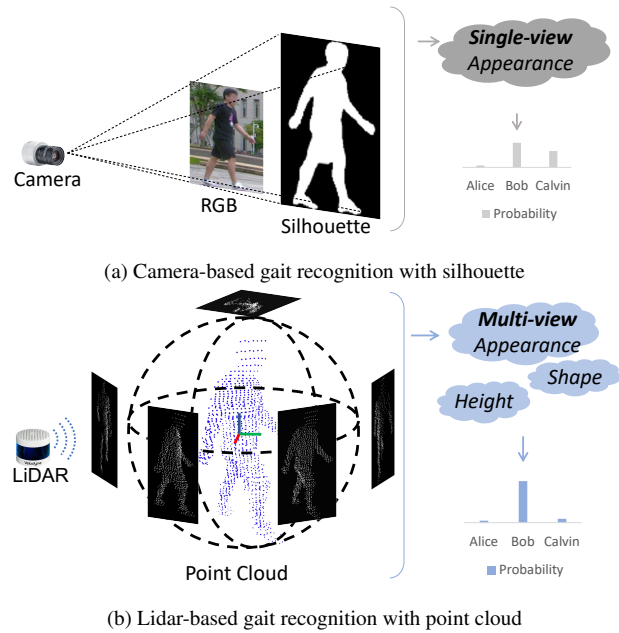

(a) Camera-based gait recognition with silhouette

(b) Lidar-based gait recognition with point cloud

Figure 1. Illustration of **(a)** camera-based and **(b)** Lidar-based gait recognition. Camera-based gait recognition commonly uses silhouettes to learn shape information from a single view. Lidar-based gait recognition can use 3D structure, shape, and scale information to identify a subject.

in the real-world scenario.

Existing camera-based methods [19, 52] are counterintuitive to human nature. When recognizing a subject [55, 56], humans consider not only the 2D appearance characteristics, but also 3D geometry structure information like height, shape, and viewpoints. Differently, camera-based gait recognition methods either capture 2D representations from a single viewpoint as shown in Fig. 1a, or exploit 3D representations from estimated 3D pose/mesh models which is usually imprecise in various challenging conditions of low resolution, poor illumination, untrained posture, etc. Fortunately, 3D sensors provide precise 3D perception like human nature, *e.g.* recognizing a subject from multiple views as illustrated in Fig. 1b.

1

Existing gait recognition research is mostly based on publicly camera-based datasets, neglecting the fact that walking videos are not guaranteed in unconstrained scenes. Many factors in the real scenes like poor visibility and complex background [6], cause that gait recognition suffers from accurate human detection and segmentation [25, 30].

With visual ambiguity and the abovementioned limitations, we believe that cameras are far behind the requirements of practical gait recognition. Considering the remarkable success of 3D applications [6, 15, 32], it is highly desirable to study 3D gait recognition with a new modality, Lidar, to significantly provide 3D structure information and precisely percept humans in the wild. As Lidar sensors have been commonly used in robotics navigation [10, 31] and autonomous driving [15, 32], it also motivates us to endow gait recognition in a future scenario, where facilitating Lidar-based gait recognition on robots.

In this paper, we build the first large-scale Lidar-based dataset to facilitate the research of gait recognition with 3D point clouds, named LIDAR GAIT. Specifically, the dataset is captured in the wild, and collected by a Velodyne VLS128 Lidar sensor and an RGB camera mounted on a mobile robot. Compared to existing datasets listed in Tab. 1, our LIDAR GAIT dataset has the following distinctive attributes: (1) **Precision**. The LIDAR GAIT dataset provides 3D point clouds as gait representations with high precision and density, providing precise and robust 3D structure information for recognition. (2) **Scalability**. The dataset captures 25,279 sequences from 1,050 subjects, which is scalable for statistical evaluation. (3) **Diversity**. The dataset not only contains realistic challenges including illumination, occlusion, dressing, carrying, and *et al.*, but it also provides detailed annotations. Therefore, the LIDAR GAIT dataset helps the community to study the impact of diverse challenging factors. (4) **Multimodality**. Our dataset captured data streams from a Lidar and a camera at the same time, endowing the exploration of sensor fusion for robust gait recognition.

Since point clouds formats differently with pixels in images and point-based gait recognition has been barely investigated, we explore five cutting-edge methods from point-based object classification. Consequently, we observe that all implemented point-based methods perform second-optimal to the methods using silhouettes from the camera, encouraging us to investigate the reason. From analysis, we think the main reason is that those representative point-based methods are specifically designed for object classification. These object classification methods mainly focus on global context modeling and neglects local structure information.

To address this issue, we propose a simple yet effective baseline method, named multi-view projection network (MVPNet). The proposed MVPNet first projects 3D point clouds into depth images from three views. Then it utilizes convolutional networks to efficiently extract spatial features with local connectivity from projection, rather than other point-based methods learning global context from sparse point clouds. Extensive experiments show that (1) 3D structural information significantly contributes to performance improvement, and our proposed MVPNet is effective to maintain 3D structure for gait recognition. (2) Equipped with a Lidar sensor, point-based gait recognition performs stably well on cross-view challenges, demonstrating the effectiveness of Lidar for practical application. (3) There is a lot of room for the exploration of fine-grained feature extraction from sparse point cloud data.

To summarize, our main contributions are as follows: (1) We carry out one of the first studies of 3D gait recognition with point clouds, bringing precise perception and 3D geometry of humans for better practicality in real-world scenarios. (2) We build the first large-scale Lidar-based gait recognition benchmark, with various annotations ranging from occlusions, viewpoint, carrying, clothing, distance, and scenes. (3) We propose a novel point cloud gait recognition framework, named MVPNet, outperforming camera-based methods by a large margin.

2. Related Work

Gait Recognition. According to the used representations, gait recognition can be generally divided into 2D and 3D representations-based methods.

The majority of 2D representations-based methods study gait characteristics directly from images, termed appearance-based [5,8,12,27,37,38,41,45] methods, which have made surprising high performance based on silhouettes [17, 24, 25] together with other gait templates [4, 17, 44]. The alternative approaches learn human structure [22, 26, 43] and dynamics [43] as gait representations, but they are heavily constrained by model-based estimation models. 3D representation methods are generally extracted by sensors [13, 18] or estimation models [24, 26]. The commonly used 3D sensors such as Kinect, provide 3D structured data, but they only facilitate in an indoor environment [13]. Meanwhile, multi-cameras reconstruction [2] and 3D estimation models [24,26,43,52,55] provide considerable 3D geometry, but the performance is far behind the requirements of real-world applications as reported in [57].

Gait Recognition Benchmark. There are three types of publicly available datasets: in-the-lab [20, 41, 49], synthetic [9], and in-the-wild datasets [18, 33, 55, 57]. The in-the-lab datasets [20, 41, 45, 49], represented by CASIA series [42,45,49] and OU-ISIR series [20, 41], advance the investigation of the feasibility of gait recognition. The recent synthetic datasets [9] are to overcome the difficulty in data acquisition and annotation of gait, providing more synthetic data with a variety of annotations but introduc-

Table 1. Comparison of publicly available datasets for gait recognition.

Dataset	Year	Subject #	Seq #	View #	Data Type	3D structure	Multimodal	In-the-wild
CASIA-B [49]	2006	124	13,640	11	RGB, Silhouettes	✗	✗	✗
CASIA-C [42]	2006	153	1,530	1	Infrared, Silhouettes	✗	✗	✓
KY4D [21]	2010	42	168	16	Silhouettes, RGB, 3D Volumetrics	✓	✗	✗
TUM-GAID [18]	2012	305	3,370	1	Audio, Video, Depth	✓	✓	✓
SZTAKI-LGA [3]	2016	28	11	1	3D Point Cloud	✓	✗	✓
OU-MVLP [41]	2018	10,307	288,596	14	Silhouettes	✗	✗	✗
FVG [51]	2019	226	2856	3	RGB	✗	✗	✓
PCG [48]	2020	30	60	1	3D Point Cloud	✓	✗	✗
GREW [57]	2021	26,345	128,671	882	Silhouettes, 2D/3D Skeleton, Flow	✗	✗	✗
Gait3D [55]	2022	4000	25309	39	Silhouettes, 2D/3D Skeleton, 3D Mesh	✓	✗	✓
OUMVLP-Mesh [23]	2022	10,307	288,596	14	3D Mesh	✓	✗	✗
LIDAR GAIT	-	**1,050**	**25,279**	**12**	**RGB, Silhouettes, 3D Point Cloud**	✓	✓	✓

ing cross-domain issue [25] at the same time. The in-the-wild datasets [33, 55, 57] are to promote gait recognition research in the unconstrained environment. The recent works [1, 3, 48] based on Lidar sensor are closely related to our work, while the main concern is that the existing datasets include at most 30 subjects, which cannot guarantee statistically reliable performance evaluation of Lidar-based gait recognition. Because of insufficient 3D representations for data-driven gait recognition, as shown in Tab. 1, a dataset with accurate 3D representations is essential.

Point Cloud and 3D Object Classification. 3D scanners (*e.g.* Lidar and depth camera) project to the targets and then generate point cloud sets. Each point represents a data point in Cartesian coordinates (X, Y, Z). Point cloud data is sparsely distributed, remaining a significant challenge in modeling correlation and geometry. 3D object classification explore projection-based [16,39,47], point-wise [35,36,53], and graph-wise models [46, 47] to capture discriminative feature on point cloud data for object classification. In this paper, we select many representative models of 3D object classification and compare them with our proposed method to comprehensively study 3D gait recognition with point clouds.

3. The LIDAR GAIT Benchmark

3.1. Overview of LIDAR GAIT Dataset

The LIDAR GAIT dataset is constructed by a 128-beam Lidar scanner and a monocular camera mounted at a fixed position on a robot. The LIDAR GAIT dataset includes 1,050 identities, 25,279 sequences, 762,896 point-cloud frames, and 3,044,499 silhouettes in total. Each sequence consists of continuous frames in RGB and points. Moreover, every frame of two modalities is timestamped, serving LIDAR GAIT as a synchronized multimodal dataset. We obtained informed consent from all participants, and we protect personal privacy by blurring faces. To the end, LIDAR GAIT annotates various variances unincluded in existing in-the-wild datasets [33,55,57].
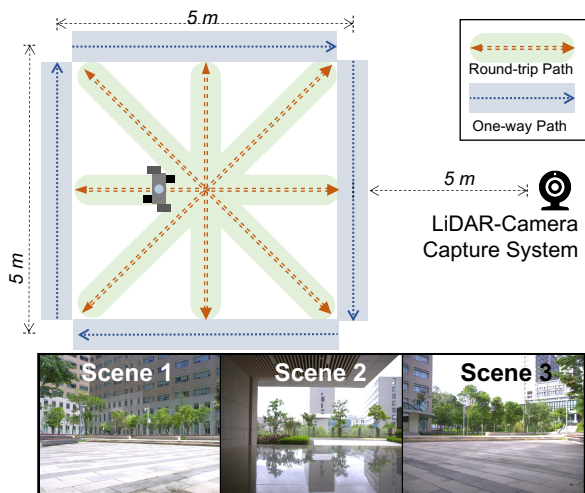

Figure 2. Data acquisition setup. Each participant is first instructed to normally walk along four round-trip paths and four one-way paths, then walk again with a random variance along the same paths.

3.2. Dataset Construction

3.2.1 Data Collection

The data was collected in July 2022 in three scenes on a university campus using an industrial camera and a Lidar scanner. The camera captured RGB imagery streams at 1,280 x 980 resolution and 30 frames per second (FPS) while the point-cloud streams were recorded at 10 FPS. We synchronized Lidar and the camera to the GPS clock and timestamped each frame, which can help collaborate two modalities for a robust gait recognition system. As shown in Figure 2, each subject first walks in normal condition along the *one-way paths* and the *round-trip paths*, then walks again but with a random attribute, for example, riding a bike as shown in Fig. 3a. In total, 48 gait sequences (= [4 × 2 (round-trip) + 4 × 1(one-way)] × 2 (twice) × 2 (modality)) can be captured for each subject. The dataset is captured in the multiple scenes at the campus and includes a variety of in-the-wild variants such as changing lighting, walking in crowds, occlusions, carrying, and complex backgrounds.

3

(a) Samples of 12 views in two conditions for a subject.

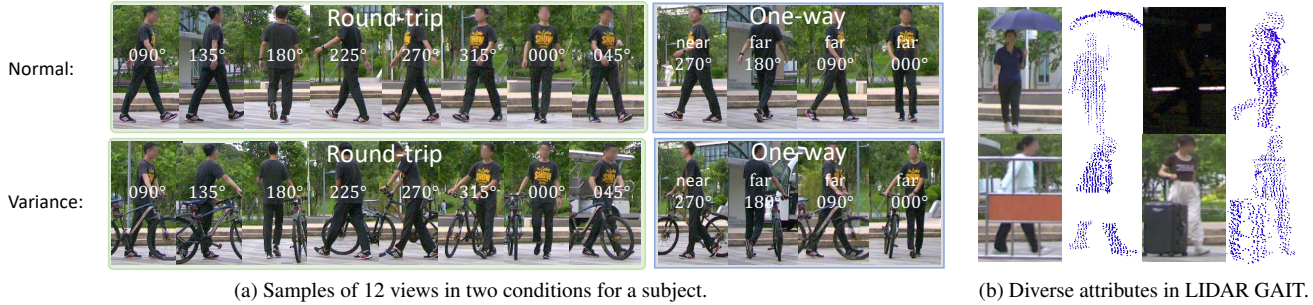(b) Diverse attributes in LIDAR GAIT.

Figure 3. (a) Each participant walks normally (top row), followed by walking with a random variance (bike for this subject) as shown in the bottom row. (b) LIDAR GAIT collects data in both point cloud and RGB modality with diverse realistic variances.

3.2.2 Variances

As illustrated in Fig. 3, the proposed dataset concludes the significant challenges: (1) **View.** The subjects walk along four green round-trip paths, resulting in eight viewpoints ranging from $0°, 45°, 90°, 135°, 180°, 235°, 270°, 315°$. The four blue one-way paths generate four extra viewpoints: far-$0°$, far-$90°$, far-$180°$, near-$270°$. (2) **Occlusion.** The portions of the body are obstructed by occlusions as shown in Fig. 3b. (3) **Illumniation.** More than 70 identities are recorded at night. (4) **Carrying.** The dataset conclude various carrying objects including umbrellas, carton, tote bag, cellphone, clothes, chessboard, kids, balls, etc. (5) **Clothing and uniform.** Our dataset includes participants wearing different coats but also contains subjects wearing the same uniform. (6) **Others.** There are other variants in the dataset such as riding a bike, playing ball, walking with kids, etc. More samples of variants can be found in **the supplementary material**.

3.2.3 Annotations

The continuous data streams are first manually segmented into sequences according to the predefined trajectories in Fig. 2. Then, each sequence is labeled by aforementioned variances such as view, occlusion, and illumination. In the end, camera-based sequences and Lidar-based sequences are separately processed to obtain gait representations.

Camera-based Data. The human boxes and trajectories are automatically generated by YOLOX [14] and Byte-Track [50], respectively. Since the LIDAR GAIT dataset is recorded in a public area, which resulted in multiple people in some frames, and a wrong trajectory when participants overlapped with the passers. We manually select the of-the-interest boxes as the ground truth in such cases. In the end, PaddleSeg [29] is performed to obtain silhouettes in size of 128×128. However, detection and segmentation perform well only under good lighting conditions but fail to deliver good performance at night.

Lidar-based Data. It is relatively easy to obtain human point clouds because there is only one person walking in the experimental area at a time. The point cloud sets of pedestrians are obtained after applying noise removal and ground removal on each frame. To protect the privacy of uninvolved passers, we only release the area range to $[-5, -12m]$ for the X axis, $[-3m, 3m]$ for the Y axis, and $[-2m, 3m]$ for the Z axis.

3.3. Dataset Statistics and Evaluation Metrics

The statistics about the distributions of sequence, frame, and attribute are detailed in Fig. 4. From Fig. 4a, we observe that the sequence scale of the two modalities is equal, but RGB modality has three times of frame number of Lidar modality. From Fig. 4b, the average pixel number of a silhouette is 3,200, while the average points number of a point cloud frame is 800. The rate of pixels vs points is approximately 1:1 and 4:1 when silhouettes are in the size of 64×64 and 128×128, respectively. To the end, the distribution of attributes is shown in Fig. 4c, demonstrating the variety of the LIDAR GAIT dataset.

To establish a more challenging and realistic setting, the LIDAR GAIT dataset is evaluated under open-set setting [34, 41], where train and test set splits are without sample intersection. The LIDAR GAIT dataset is separated into three splits: a train set with 250 identities and 5,988 sequences, a valid set with 250 identities and 5981 sequences, and a test set with 550 identities and 13,310 sequences. During the test, the sequences in normal conditions sets as gallery sets, and the sequences in variant conditions are taken as probe sets. The evaluation protocol follows the cross-view recognition setting as commonly used in CASIA-B [49] and OUMVLP [41], where probe sets of the same view calculate the similarity to gallery sets of each view. To evaluate the impact of attributes, the probe sets are grouped into many subsets according to the attributes, then perform cross-view retrieval task. The prevailing Rank-1 accuracy is adopted as the evaluation metric.
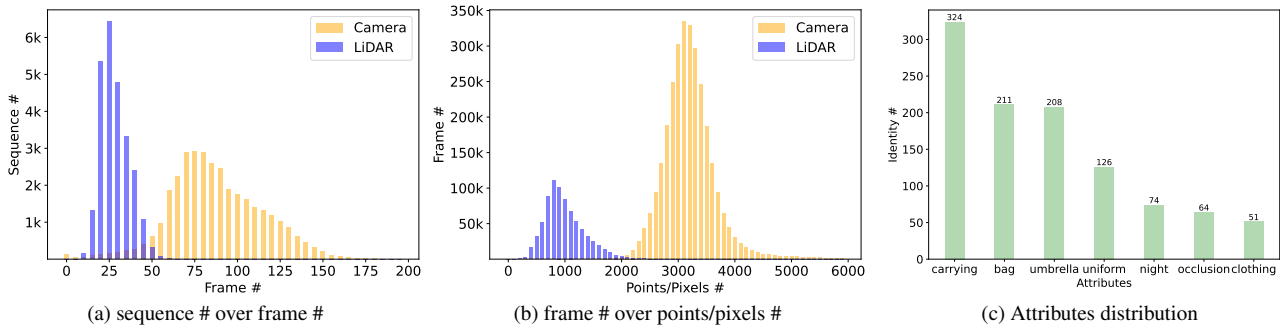
Figure 4. Statistics about LIDAR GAIT dataset. Lidar modality and RGB modality are represented in blue and yellow, respectively. It shows that LIDAR GAIT dataset is scalable, multimodal, and diverse for the study of 3D gait recognition. Best viewed in color.

4. Gait Recognition with Point Clouds

4.1. Problem Setting

In this section, we introduce the problem setting of 3D gait recognition with point clouds. Given a point cloud dataset $\mathcal{P} = \{\mathcal{P}_i^j | i = 1, 2..., N; j = 1, 2, ..., m_i\}$ with N identities and m_i sequence for each identity y_i. Each point cloud sequence $\mathcal{P}_i^j \in \mathbb{R}^{T \times N \times C}$ is with T frames and N points for each frame, where C is the number of feature channels. Our goal is to learn a network $N_\theta(\cdot)$ that can produce the feature embedding $\mathcal{F}_i^j$ to represent the associated identity y_i.

In this paper, we propose the MVPNet, as shown in Fig. 5, to tackle the 3D gait recognition task, formulated as:

$$\mathcal{F}_i^j = N_\theta(\mathcal{G}(\mathcal{D}(\mathcal{P}_i^j))) \tag{1}$$

where $\mathcal{D}$ is the essential data augmentation to avoid overfitting and improve the network generalization ability. The function $\mathcal{G}$ operates on point clouds and generates multi-views depth images from multiple virtual cameras. The feature extractor N_θ is composed of three components. 1) a feature encoder $\mathcal{E}$ that captures spatially local connectivity from projected multi-view images. 2) a multi-view feature fusion layer $\mathcal{S}$ that combine the learned 2D descriptors from multiple point-views into a compact 3D representation. 3) a temporal aggregation network $\mathcal{T}$ that models dynamical conjunction along sequential input, which can be formulated as:

$$N_\theta(\cdot) = \mathcal{T}(\mathcal{S}(\mathcal{E}(\cdot))) \tag{2}$$

After point clouds are transformed into depth images, N_θ is utilized to extract compact embedding as the final representations for recognition. With the spatial representations effectively exploited from multiple viewpoints, we combine the learned 2D descriptors from multiple point-views into a compact spatial-temporal representation, which contains informative 3D geometry. In the subsections below, we will go into further detail about the proposed approach.

4.2. Multi-view Representations Generation

Point Cloud. Point cloud is a set of sparse points in Euclidean space, where each point represents a data point in Cartesian coordinates (X, Y, Z). Lidar will generate points with additional information like intensity, but intensity reflects the fabric of clothing. To avoid adding gait-irrelated features, we only reserve the 3D coordinates of points as input features. Therefore, the input data for our MVPNet is a sequential point cloud set $\mathcal{P}_i^j \in \mathbb{R}^{T \times N \times 3}$. For more details about the impact of the number of T and N, the experiments are presented in **the supplementary material**.

Data Augmentation. Point cloud augmentation plays a significant role in 3D object classification [35,36] from two aspects: 1) data augmentations enlarge the dataset to prevent overfitting from the limited availability of the dataset, 2) and act as the regularizer to learn more representative feature by randomly perturbing points. To fairly evaluate MVPNet with other point-based methods, we adopt data augmentation [16,35,53], formulated as: $\mathcal{P}_i^{j'} = \mathcal{D}(\mathcal{P}_i^j)$ where $\mathcal{P}_i^{j'}$ is the augmented point clouds.

Multi-view Representation. As point cloud maintains 3D geometric information, Humans also recognize an object in a similar way, which constructs the 3D perception of the target from multiple perspectives. The similarity motivates us to conduct MVPNet for information aggregation from multiple perspectives. Here we acquire multi-view representations from three virtual cameras, *i.e.* range-view, side-view, and bird's-eye-view (BEV) as illustrated in Fig. 5. The side-view and BEV cameras are orthogonally placed at a distance of 1.2 units from the pedestrian center [16]. The range-view depth images are with high density and local connectivity, while the other two orthogonal cameras obtain depth images with sparse connectivity. The depth images can be formulated as: $\mathcal{G}(\mathcal{P}_i^j) = \{I_{i,r}^j, I_{i,s}^j, I_{i,b}^j\}$, where $I_{i,r}^j, I_{i,s}^j, I_{i,b}^j$ represent depth sequences at range-view, side-view, and BEV respectively.
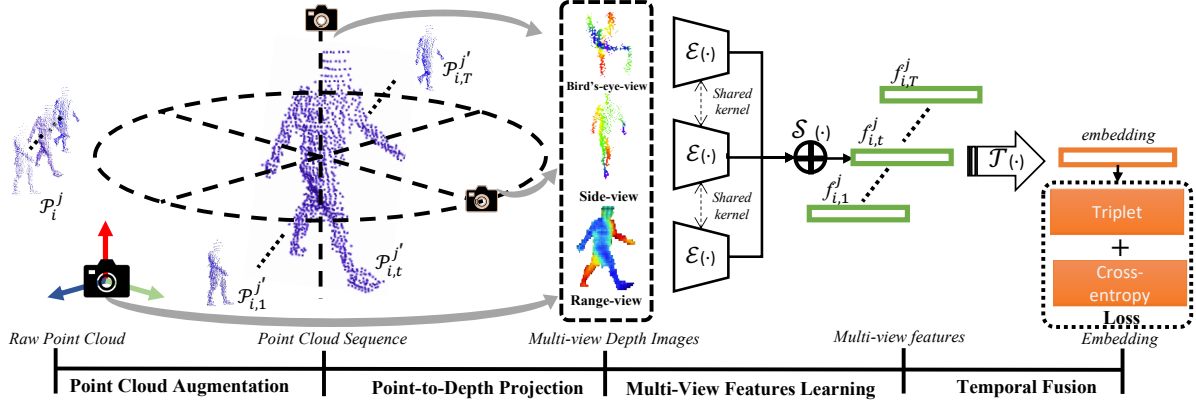
Figure 5. The framework of MVPNet for 3D gait recognition with point clouds. MVPNet receives a sequence of point clouds, then extract and combines representations from range-view, side-view, and bird's-eye-view.

4.3. Multi-view Representation Learning

Feature Extractor. The commonly used backbone [11] is adopted as the feature encoder $\mathcal{E}$. For all three views, feature encoder $\mathcal{E}$ with shared kernels is applied to learn multi-view representations $f_{i,mv}^j \in \mathbb{R}^{T \times D \times 3}$, where D is the number of features channel. The multi-view representations $f_{i,mv}^j$ consists of three representations $f_{i,r}^j, f_{i,s}^j, f_{i,b}^j$ learned from range-view, side-view, and BEV respectively.

Multi-View Feature Fusion Module is designed to combine learned multiple representations from different perspectives into a dense representation $F_i^j \in \mathbb{R}^{T \times D_{fusion}}$, where D_{fusion} refers to the dimension of the fused multiview features. We study various fusion strategies and present more details and comparisons in **the supplementary material**.

Temporal Feature Fusion aggregates sequential dense multi-view representation $F_i^j \in \mathbb{R}^{T \times D_{fusion}}$ into sequencelevel embedding $\mathcal{F}_i^j \in \mathbb{R}^{D_{fusion}}$, many temporal modeling methods exist in the literature [5, 52, 54]. Set Pooling is adopted for both modalities to fuse dynamic cues to fairly compare with the camera-based method.

4.4. Traning and Inference

In this work, the model is trained with a combined loss function, which is formulated as:

$$L = \alpha L_{tri} + \beta L_{ce} \qquad (3)$$

where L_{tri} is BA^+ triplet loss [5], L_{ce} is the cross-entropy loss [28]. α and β are the weighted hyperparameters.

During inference, the whole sequence of point clouds is used, then data augmentation $\mathcal{D}$ only applies normalization on the input points. Euclidean distance is used to measure the similarity of each probe-gallery pair and calculate Rank 1 recognition accuracy.

5. Experiments

5.1. Experimental Setup

The LIDAR GAIT dataset is used to evaluate all the experiments. We follow the dataset split and evaluation metric as detailed in Sec. 3.3. To evaluate the performance of different modalities in the wild, we select the most representative methods for each modality and compare our proposed MVPNet with these methods.

5.1.1 Comparative Methods

Lidar-based Methods. We implement five commonly used approaches in point cloud classification including PointNet [35], PointNet++ [36], PointTransformer [53], DGCNN [46], and SimpleView [16]. Among them, the methods [35, 36, 53] are point-wise models. DGCNN [46] is a representative graph-wise model, and SimpleView [16] is chosen as a representative projection-based method. *Implementation details*. During training, the frame number, and point number are set to 10 and 512, respectively. All the points and the whole sequence are used for evaluation. For all point-based methods including MVPNet, we apply identical data augmentation including normalization, random scaling, and random shifting during training. Triplet loss with batch size ($p = 32, k = 8$) is applied for all models, where p and k donate the number of identities and sequences per identity, respectively. SGD [40] with a learning rate of 0.1 optimizes all point-based methods for 30,000 iterations.

Camera-based Methods. To evaluate the performance of the camera-based modality, we implement four cuttingedge methods: GaitSet [5], OpenGait [11], GaitPart [12], and GaitGL [28]. *Implementation details*. Since LIDAR GAIT has the equivalent scale of the training set to CASIA-B, the network parametric setting is identical to the config-

6

Table 2. Evaluation of the impact of different attributes. Rank-1 accuracy (%) on the *valid + test* set is reported.

Modality	Method Publication	Model	Mean	Bag	Clothing	Probe Seuqence Umberalla	Uniform	Occlusion	Night
Camera	AAAI2019	GaitSet [5]	64.43	68.33	35.57	64.78	63.53	65.21	23.33
	CVPR2019	GaitPart [12]	64.18	67.65	35.91	63.08	63.62	65.28	20.18
	GITHUB	OpenGait [11]	65.85	70.01	38.04	65.96	64.96	67.05	23.82
	ICCV2021	GaitGL [28]	76.09	78.98	51.23	**76.92**	76.38	77.68	21.58
Lidar	CVPR2017	PointNet [35]	23.63	30.00	19.57	14.52	21.40	44.96	21.17
	NIPS2017	PointNet++ [36]	32.19	32.72	23.64	27.68	27.49	41.53	33.79
	TOG2019	DGCNN [46]	36.42	40.42	30.37	27.78	32.61	58.74	39.82
	ICML2021	SimpleView [16]	50.75	56.98	40.99	28.54	49.62	78.99	52.77
	ICCV2021	PointTransformer [53]	38.24	40.93	30.79	29.16	35.67	55.10	37.99
	-	MVPNet (Range-view)	79.24	80.68	56.65	60.29	73.20	92.32	83.19
	-	**MVPNet (Ours)**	82.41	83.31	61.91	64.11	**77.40**	**94.81**	**85.53**

uration for CASIA-B. The frame number in the training is set to 30, and the whole sequence is used in the test evaluation. The silhouettes resolution is in the size of 64×64. The optimizer, weight decay, and initial learning rate (LR) are set to Adam, 0.0005, and 0.0001, respectively. LR is multiplied by 0.1 at the 30,000th and 60,000th iterations, and all camera-based models train with 70,000 iterations in total.

5.1.2 Implementation Details of MVPNet

MVPNet is a projection-based method, which transfers point clouds into images and utilizes CNNs to extract appearance features. The commonly used backbone [11, 55] is utilized in MVPNet as a feature encoder. The range-view images are size-normalized by [20], and other rendered images use point-level normalization [35]. All the projection-based images from point clouds are with a resolution of 64×64. The feature extractor is identical to [11]. The optimizer and batch size is the same with silhouette-based methods. The weighted parameters for triplet loss and cross-entropy loss are set to 1 and 0.1, respectively. The code for all the experiments is written on codebase [11].

5.2. Comparative Results

To compare with other point-based methods and understand what factors really impact gait recognition in the wild, we detailed conducted experiments on various realistic factors including illumination, distance, carrying, clothing, and occlusion. We separate all the sequences with the same variance into multiple subsets. For example, to evaluate the impact of illumination, the probe sets only contain sequences with bad illumination conditions, and the gallery sets are all the sequences in the test set with normal conditions. Th cross-view recognition on each subset is evaluated, and the cross-view accuracy matrix is like the results in Fig. 6. We only report the average of the accuracy matrix in Tab. 2, and we have the following observations: (1) MVPNet shows its superiority to both other point-based method and camera-
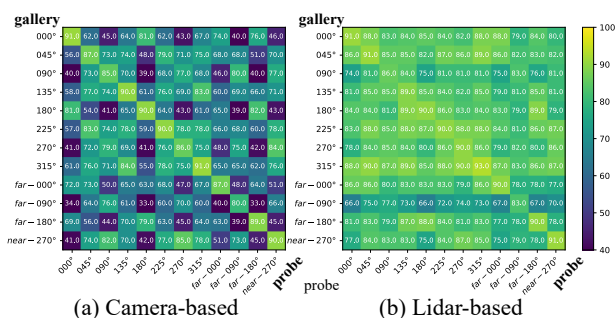

(a) Camera-based   (b) Lidar-based

Figure 6. Comparison between Lidar and Camera for gait recognition. We report rank-1 accuracy (%) on cross-view protocol. Best viewed in color and pdf.

based methods, which is mainly beneficial by integration with 3D geometry and multi-view appearance information. (2) MVPNet achieves state-of-the-art results in all conditions except the umbrella subset. It is mainly caused that umbrellas are erased after segmentation on RGB images, while the umbrellas are kept in point sets. (3) The methods using silhouettes make a poor performance at night. Point-based methods provide more promising results, and MVPNet outperforms others by a large margin. (4) All silhouette-learning models [5, 11, 12, 28] achieve higher accuracy than point-based models in point cloud classification. This concludes that it is necessary to specifically design point-based models for 3D gait recognition.

Cross-view Gait Recognition. We conduct a detailed comparison of two modalities of cross-view gait recognition in Fig. 6. Note that we utilize identical feature encoders for two modalities to make ablative results, and we can make the following observations: (1) the distance from subjects to sensors indeed impacts the performance for both two modalities. (2) Camera-based method achieves poor performance when query sets are at views of $0°, 90°, 180°$ (see purple pixel in Fig. 6a). The same phenomenon can be found on CASIA-B [5] and OUMVLP [41]. However,
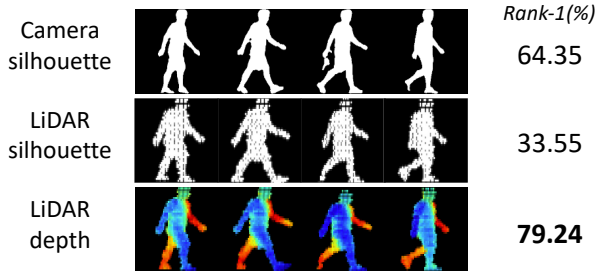
Figure 7. Ablation study on the effectiveness of depth information for performance. Best viewed in color.

Lidar-based methods can perform stably in cross-view settings, validating the effectiveness of 3D structure for cross-view gait recognition.

5.3. Ablation Study

Effectiveness of 3D Geometry Information. To evaluate the effectiveness of depth information for gait recognition, we make a comparison among four types of data as input: (1) Camera silhouettes: the camera-based silhouettes are obtained by segmentation results of RGB images. (2) Lidar silhouettes: Lidar silhouettes are obtained by range-view projection of point cloud sets. (3) Lidar depth: the depth information is added. From Fig. 7, we can observe that: (1) When depth information is not included, the performance of Lidar silhouettes is much lower than the accuracy of camera silhouettes. This is because the camera has a much higher resolution than Lidar, so the camera can catch the target with more details. (2) Though Lidar generates point clouds in sparse space, the depth of information makes a magnificent improvement to the accuracy. The integration of depth information can improve rank-1 accuracy from 33.55% to 79.24%, validating the necessity and effectiveness of 3D information for gait recognition.

Effectiveness of Each View. From Fig. 8, we have the following observations: (1) The comparison between front-view and range-view, indicates that the convolutions can take more advantage from range-view depth maps. It is because the closer view makes the pixels within projected depth images sparser. (2) The side-view depth maps obtained from point cloud projection are reliable and effective, which provide a meaningful cues from another perspective. (3) The accuracy on bird-eye views are relatively low, but BEV images provide interesting evidence that gait recognition is the potential to be achieved at the bird-eye view.

6. Discussion

Ethical Discussion. All the subjects involved in the dataset signed a written consent to agree that their data can be collected, processed, used, and shared for research purposes. The dataset can be distributed only for non-commercial re-
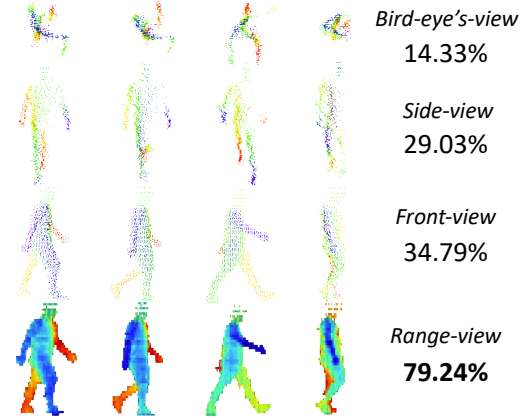

Figure 8. Performance comparison between different views. Best viewed in color.

search purposes with the case-by-case dataset access application. The human faces of each RBG image are blurred to protect sensitive privacy. The recorded data can only be used for 20 years. After this date, all data will be deleted and not allowed to be used.

Potential Negative Societal Impacts. Point cloud gait recognition will endow robotics with the advanced perception of humans, which can benefit healthcare and social security. The main concern is to prevent biometrics security leaks.

7. Conclusions

In this paper, we introduce the Lidar sensor to provide reliable anthropometric parameters for the human body, and to perceive pedestrians in unconstrained scenes. First, we proposed a novel multi-view projection network for point cloud gait recognition, named MVPNet, to exploit 3D human geometry from multi-view representations. Moreover, we build the first large-scale multimodal 3D point cloud gait recognition dataset, termed LIDAR GAIT, to facilitate the research of gait recognition with point cloud data. LIDAR GAIT contains 25,279 sequences with 1,050 subjects and covers various visibility, views, occlusions, clothing, carry, and scenes. Lastly, our proposed method achieves remarkable results on the LIDAR GAIT dataset, showing the superior of LIDAR and the effectiveness of MVPNet.

Future Work. MVPNet has obtained remarkable results in various scenarios, yet it performs not well enough if subjects carry an umbrella. The reason should be that the umbrellas are wrongly included in the point cloud. Better performance can be achieved if the umbrellas are removed from the point clouds. Besides, MVPNet only takes one modality as input currently. LIDAR GAIT dataset is a multimodal dataset with synchronized RGB images and point clouds. Much better results should be achieved if the two modalities are fused.

A. Ethical Discussion

The LIDAR GAIT dataset has been reviewed by Institutional Review Board (IRB), and IRB also approves the use of the dataset. Our data acquisition has been authorized by each subject involved in the dataset. We have applied several processes to protect sensitive privacy, including blurring faces, limiting the duration of data usage, and requiring distribution agreements. Besides, all the data collectors obtained qualification certificates of ethics training.

B. Ablative Results for Multi-view Fusion

We study five ways to fuse multiple features $f_{mv} \in \mathbb{R}^{3 \times C}$ from three views into a multi-view feature:

- **Concatenation.** The given multiple features are naively concatenated together. Then the fused features are in the shape of $\mathcal{F} \in \mathbb{R}^{3 \times C}$.

- **Weighted Fusion.** The multiple features from three views are fused by a fully connected network (FCN) along channel dimension. The final features are in the shape of $\mathcal{F} \in \mathbb{R}^{C}$.

- **Mean Pooling.** Statistically averaging multiple features along view dimension. All the final features of statistical pooling are in the shape of $\mathcal{F} \in \mathbb{R}^{C}$.

- **Max Pooling.** Statistically selecting maximum features along view dimension.

- **Min Pooling.** Statistically selecting minimum features along view dimension.

We show the overall rank-1 recognition accuracy evaluated in cross-view protocol over the validation and test set in Tab. 3. It can be seen that: (1) All fusion strategies, excluding min pooling, consistently improve performance, showing the effectiveness of multi-view features for performance improvement. (2) Though the mean pooling and weighted fusion show comparative performance, we choose statistically average pooling as the optimal fusion strategy because channel-wise requires more parameters.

Table 3. Ablative results of different fusion strategies.

	Fusion	Rank-1	Rank-5	Parameters #
Range-view	-	79.24	92.98	3.254M
Side-view	-	29.03	54.85	3.252M
BEV-view	-	14.33	34.26	3.252M
MVPNet	Concat	80.12	92.91	7.415M
MVPNet	Wighted	82.32	94.10	5.351M
MVPNet	Max	80.74	93.28	3.254M
MVPNet	Min	79.25	92.58	3.254M
MVPNet	Mean	**82.41**	**94.14**	3.254M

C. Effect of Sequence Length and Point Number per frame

We study the effect of sequence length and point number per frame during the training process. The sequence length and point number are set to fixed values for the training process. During the inference process, the whole sequence is used, and the point number per frame is the same as the point number in the training process.

It can be seen that: (1) MVPNet performs better when it trains with sequence length less than one gait cycle (10 frames in LIDAR GAIT) as shown in Fig. 9. (2) The point number per frame does not impact recognition performance. It is because MVPNet relies on range-view depth images, and the projection-based methods are less sensitive to point numbers than other point-based methods such as DGCNN.
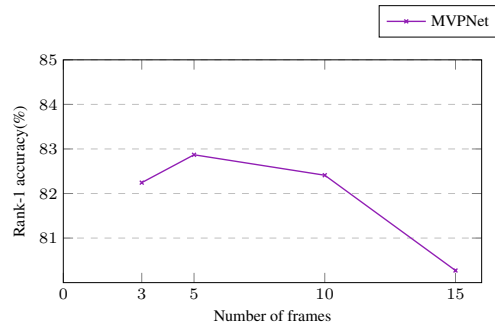

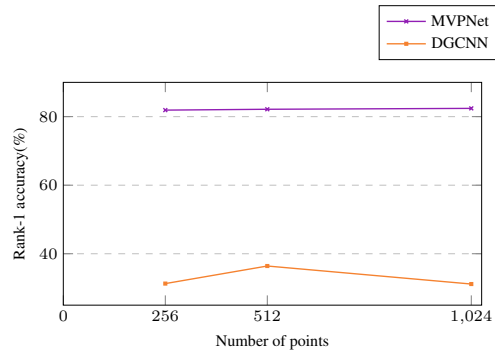Figure 9. The performance of MVPNet training with different sequence length.

Figure 10. The performance of MVPNet and DGCNN training with different point numbers.

We have also studied the impact of sequence length on the inference process. Compared to the camera-based method using silhouettes, we can observe that: (1) Both camera-based and Lidar-based models obtain better performance with the increasing frames of the sequences. (2) When only given one frame for each sequence of probe and gallery, the Lidar-based method can surprisingly achieve
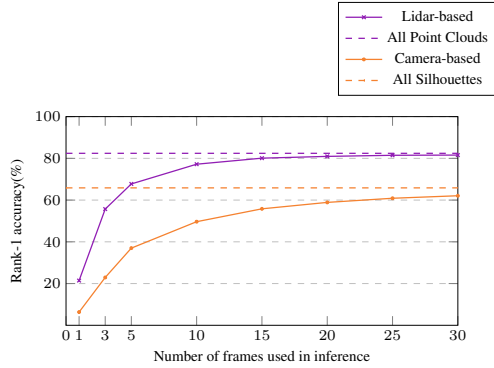
Figure 11. The performance comparison between Lidar and camera on used frame number in inference.

21.44 % rank-1 recognition accuracy and 47.31 % rank-5 recognition accuracy, showing the effectiveness of Lidar-based gait recognition in the few-shot setting.

D. Qualitative results

To analyze the performance gap between our MVPNet and representative PointNet, We visualize the feature distribution on the LIDAR GAIT dataset. We can observe that MVPNet can capture features with clear discrimination. As shown in Fig. 12b, MVPNet prominently learns the inter-class margin and makes the intra-class distribution more compact. However, the representative point-wise model, PointNet, can only obtain global features as shown in Fig. 12a. PointNet captures features with less discrimination, and its intra-class features distribute sparsely.

E. Exemplar Sequences of LIDAR GAIT Dataset

To demonstrate the necessity of the LIDAR GAIT dataset, We show several exemplar sequences of the LIDAR GAIT dataset under normal, occlusion, and poor illumination conditions in Fig. 13 - 15.

Fig. 13 shows that Lidar provides informative geometry as significant cues that extend gait recognition from 2D to 3D space. The most considerable advantage of Lidar for gait recognition is that it allows for perspective from another viewpoint, as shown in the bottom row of Fig. 13.

When the pedestrians are occluded, as shown in Fig. 14, the silhouettes obtained by segmentation methods are typical with lower quality. The conventional segmentation methods are based on 2D cameras, but humans live in 3D space, making it difficult to separate the off-the-interest pedestrian from 2D space. Lidar with precise 3D information can obtain high-quality gait representation under the condition of occlusion.

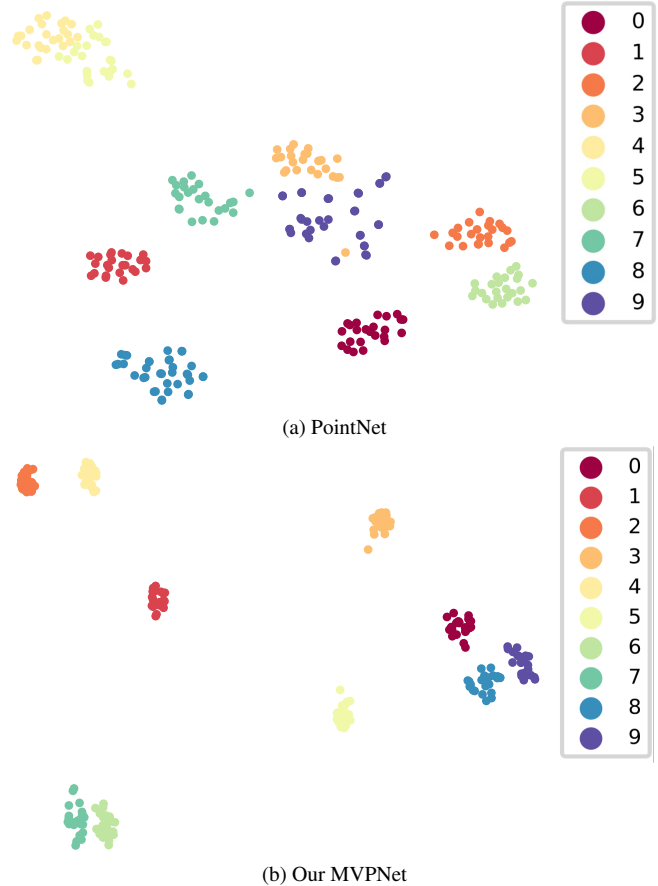When the pedestrians are occluded, as shown in Fig. 14,

(a) PointNet

(b) Our MVPNet

Figure 12. Feature distributions are visulized by t-SNE.

the silhouettes obtained by segmentation methods are typical with lower quality. The conventional segmentation methods are based on 2D cameras, but humans live in 3D space, making it difficult to separate the of-the-interest pedestrian from 2D space. With precise 3D information, Lidar can obtain high-quality gait representation under occlusion.

In Fig. 16, we show gait representations in existing in-the-wild datasets, GREW and Gait3D. We can observe that failure gait representations commonly exist because of various factors in the real world. Therefore, it is necessary to investigate a new way to obtain robust gait representation in such complex scenarios.

References

[1] Jeongho Ahn, Kazuto Nakashima, Koki Yoshino, Yumi Iwashita, and Ryo Kurazume. 2v-gait: Gait recognition using 3d lidar robust to changes in walking direction and measurement distance. In *2022 IEEE/SICE International Symposium on System Integration (SII)*, pages 602–607, 2022. 3

Figure 13. Exemplar sequences of LIDAR GAIT dataset under normal conditions. Eight frames in three modalities are visualized. The top three rows show three gait representations in RGB images, silhouettes, and point clouds. The bottom four rows represent gait representations in RGB images, silhouettes, front-view point clouds, and side-view point clouds. It shows Lidar provides informative 3D geometry. (Best viewed in color.)

[2] Gunawan Ariyanto and Mark S Nixon. Model-based 3d gait biometrics. In *2011 international joint conference on biometrics (IJCB)*, pages 1–7, 2011. 2

[3] Csaba Benedek, Bence Gálai, Balázs Nagy, and Zsolt Jankó.

Lidar-based gait analysis and activity recognition in a 4d surveillance system. *IEEE Transactions on Circuits and Systems for Video Technology*, 28(1):101–113, 2016. 3

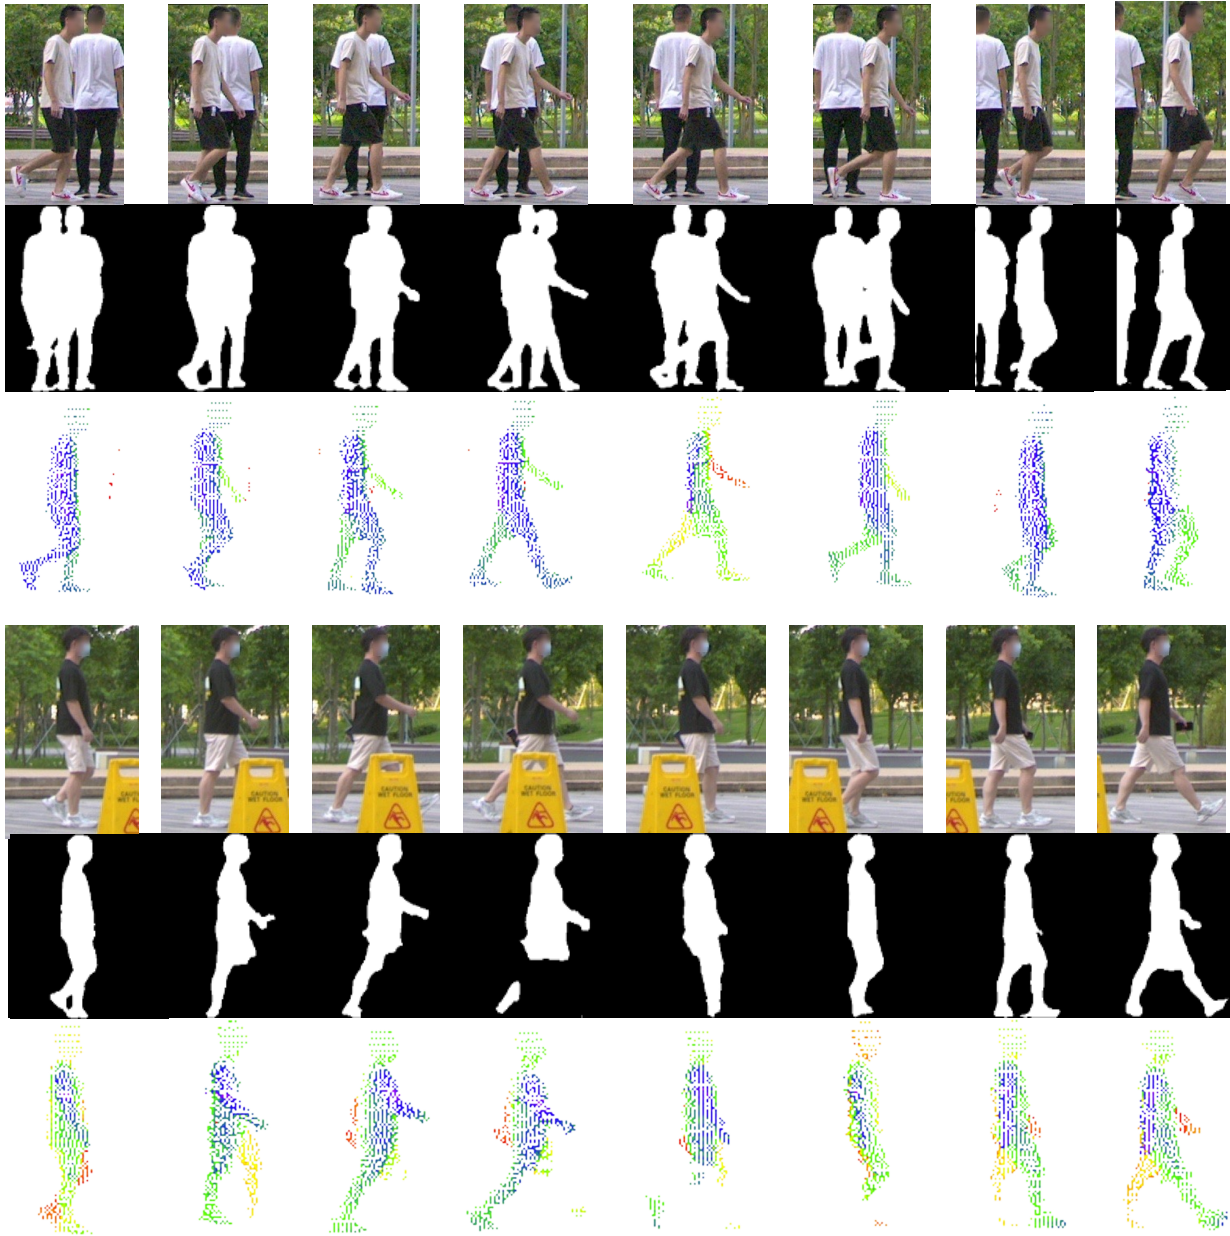[4] A.F. Bobick and J.W. Davis. The recognition of hu-

Figure 14. Exemplar sequences of LIDAR GAIT dataset under occlusions. The top three rows show gait representations when another subject overlaps the of-the-interest pedestrian. The bottom three rows show gait representations occluded by a static obstruction. It indicates that Lidar can provide robust gait representations under occlusion conditions. (Best viewed in color.)

man movement using temporal templates. *IEEE TPAMI*, 23(3):257–267, 2001. 2

[5] Hanqing Chao, Yiwei He, Junping Zhang, and Jianfeng Feng. Gaitset: Regarding gait as a set for cross-view gait recognition. In *AAAI*, pages 8126–8133, 2019. 1, 2, 6, 7

[6] Peishan Cong, Xinge Zhu, Feng Qiao, Yiming Ren, Xidong Peng, Yuenan Hou, Lan Xu, Ruigang Yang, Dinesh Manocha, and Yuexin Ma. Stcrowd: A multimodal dataset for pedestrian perception in crowded scenes. In *CVPR*, pages

19608–19617, 2022. 2

[7] Huanzhang Dou, Pengyi Zhang, and Wei Su. Metagait: Learning to learn an omni sample adaptive representation for gait recognition. In *ECCV*, 2022. 1

[8] Huanzhang Dou, Pengyi Zhang, Wei Su, Yunlong Yu, and Xi Li. Metagait: Learning to learn an omni sample adaptive representation for gait recognition. In *ECCV*, pages 357–374, 2022. 2

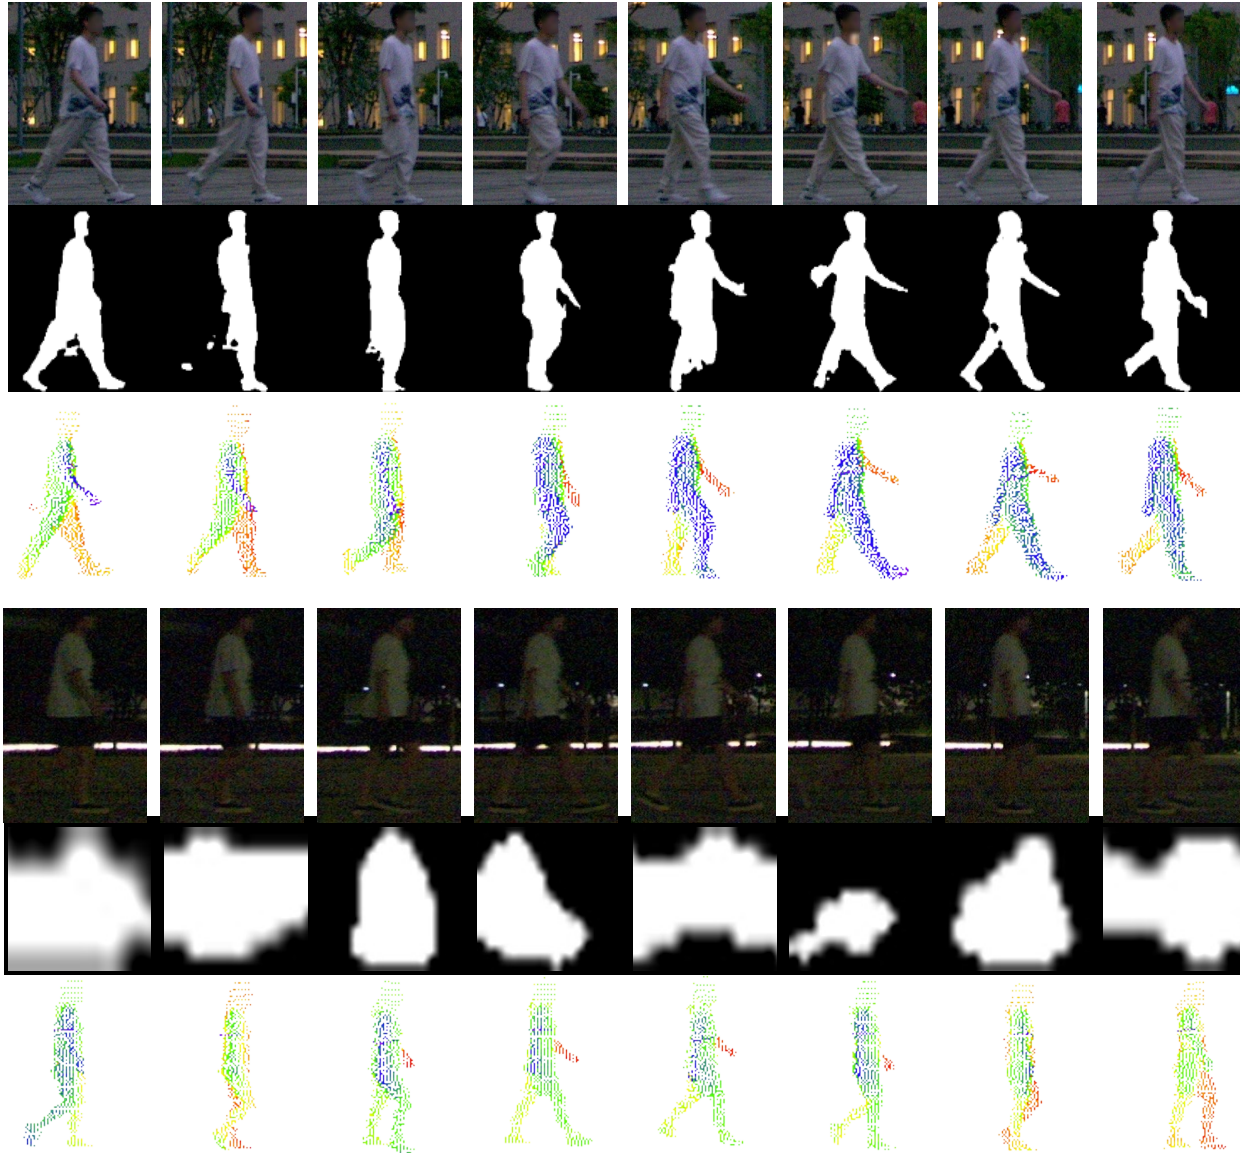[9] Huanzhang Dou, Wenhu Zhang, Pengyi Zhang, Yuhan Zhao,

Figure 15. Exemplar sequences of LIDAR GAIT dataset under poor illumination. When illumination is extremely low, human segmentation is barely performed. In contrast, Lidar provides robust gait representation with point clouds regardless of lighting. (Best viewed in color.)

Songyuan Li, Zequn Qin, Fei Wu, Lin Dong, and Xi Li. Versatilegait: A large-scale synthetic gait dataset with fine-grainedattributes and complicated scenarios. *arXiv preprint arXiv:2101.01394*, 2021. 2

[10] Mahsa Ehsanpour, Fatemeh Saleh, Silvio Savarese, Ian Reid, and Hamid Rezatofighi. Jrdb-act: A large-scale dataset for spatio-temporal action, social group and activity detection. In *CVPR*, pages 20983–20992, 2022. 2

[11] Chao Fan, Junhao Liang, Chuanfu Shen, Saihui Hou, Yongzhen Huang, and Shiqi Yu. Opengait: Revisiting gait recognition toward better practicality. *arXiv preprint arXiv:2211.06597*, 2022. 6, 7

[12] 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 *CVPR*, pages 14225–14233, 2020. 1, 2, 6, 7

[13] Péter Fankhauser, Michael Bloesch, Diego Rodriguez, Ralf Kaestner, Marco Hutter, and Roland Siegwart. Kinect v2 for mobile robot navigation: Evaluation and modeling. In *2015 International Conference on Advanced Robotics (ICAR)*, pages 388–394, 2015. 2

[14] Zheng Ge, Songtao Liu, Feng Wang, Zeming Li, and Jian Sun. Yolox: Exceeding yolo series in 2021. *arXiv preprint arXiv:2107.08430*, 2021. 4

[15] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The Inter-*
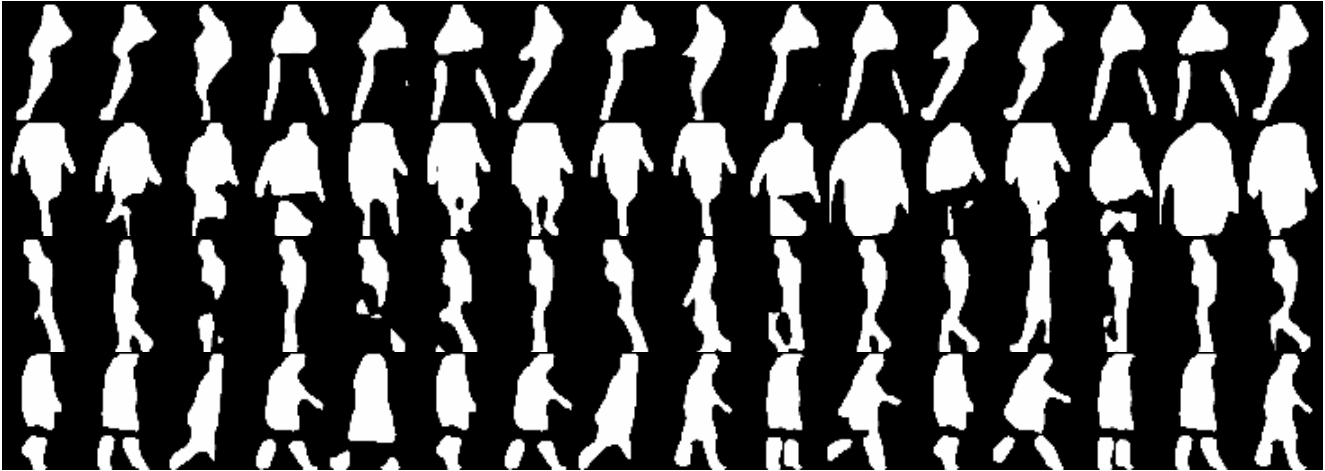
(a) GREW

(b) Gait3D

Figure 16. Failure silhouettes in the existing in-the-wild datasets. The existing in-the-wild datasets face the issues that segmentation methods fail to provide gait representations with high quality by the effect of many real-world factors.

national Journal of Robotics Research, 32(11):1231–1237, 2013. 2

[16] Ankit Goyal, Hei Law, Bowei Liu, Alejandro Newell, and Jia Deng. Revisiting point cloud shape classification with a simple and effective baseline. In *ICML*, pages 3809–3820, 2021. 3, 5, 6, 7

[17] Jinguang Han and Bir Bhanu. Individual recognition using gait energy image. *IEEE TPAMI*, 28(2):316–322, 2005. 2

[18] Martin Hofmann, Jürgen Geiger, Sebastian Bachmann, Björn Schuller, and Gerhard Rigoll. The tum gait from audio, image and depth (gaid) database: Multimodal recognition of subjects and traits. *Journal of Visual Communication and Image Representation*, 25(1):195–206, 2014. 1, 2, 3

[19] 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 *ICCV*, pages 12909–12918, 2021. 1

[20] H. Iwama, M. Okumura, Y. Makihara, and Y. Yagi. The ou-isir gait database comprising the large population dataset and performance evaluation of gait recognition. *IEEE TIFS*, 7, Issue 5:1511–1521, 2012. 1, 2, 7

[21] Y. Iwashita, R. Baba, K. Ogawara, and R. Kurazume. Person identification from spatio-temporal 3d gait. In *Int. Conf. Emerging Security Technologies (EST)*, 2010. 3

[22] Xiang Li, Yasushi Makihara, Chi Xu, and Yasushi Yagi. End-to-end model-based gait recognition using synchronized multi-view pose constraint. In *ICCVW*, pages 4106–4115, 2021. 2

[23] Xiang Li, Yasushi Makihara, Chi Xu, and Yasushi Yagi. Multi-view large population gait database with human meshes and its performance evaluation. *IEEE TBIOM*, 4(2):234–248, 2022. 3

[24] Xiang Li, Yasushi Makihara, Chi Xu, Yasushi Yagi, Shiqi Yu, and Mingwu Ren. End-to-end model-based gait recognition. In *ACCV*, 2020. 2

[25] 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

[26] Rijun Liao, Shiqi Yu, Weizhi An, and Yongzhen Huang. A model-based gait recognition method with body pose and human prior knowledge. *PR*, 98:107069, 2020. 2

[27] Beibei Lin, Shunli Zhang, and Feng Bao. Gait recognition with multiple-temporal-scale 3d convolutional neural network. In *ACM MM*, pages 3054–3062, 2020. 2

[28] Beibei Lin, Shunli Zhang, and Xin Yu. Gait recognition via effective global-local feature representation and local temporal aggregation. In *ICCV*, pages 14648–14656, 2021. 1, 6, 7

[29] Yi Liu, Lutao Chu, Guowei Chen, Zewu Wu, Zeyu Chen, Baohua Lai, and Yuying Hao. Paddleseg: A high-efficient development toolkit for image segmentation. *arXiv preprint arXiv:2101.06175*, 2021. 4

[30] Zhijian Liu, Haotian Tang, Alexander Amini, Xingyu Yang, Huizi Mao, Daniela Rus, and Song Han. Bevfusion: Multi-task multi-sensor fusion with unified bird's-eye view representation. *arXiv*, 2022. 2

[31] Roberto Martin-Martin, Mihir Patel, Hamid Rezatofighi, Abhijeet Shenoi, JunYoung Gwak, Eric Frankel, Amir Sadeghian, and Silvio Savarese. Jrdb: A dataset and benchmark of egocentric robot visual perception of humans in built environments. *PAMI*, 2021. 2

[32] Jieru Mei, Alex Zihao Zhu, Xinchen Yan, Hang Yan, Siyuan Qiao, Liang-Chieh Chen, and Henrik Kretzschmar. Waymo open dataset: Panoramic video panoptic segmentation. In *ECCV*, pages 53–72, 2022. 2

[33] Zihao Mu, Francisco M Castro, Manuel J Marín-Jiménez, Nicolás Guil, Yan-Ran Li, and Shiqi Yu. Resgait: The real-scene gait dataset. In *2021 IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–8, 2021. 2, 3

[34] Mark S Nixon, Tieniu Tan, and Rama Chellappa. *Human identification based on gait*, volume 4. Springer Science & Business Media, 2010. 4

[35] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, pages 652–660, 2017. 3, 5, 6, 7

[36] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *NeurIPS*, 30, 2017. 3, 5, 6, 7

[37] Chuanfu Shen, Beibei Lin, Shunli Zhang, George Q Huang, Shiqi Yu, and Xin Yu. Gait recognition with mask-based regularization. *arXiv preprint arXiv:2203.04038*, 2022. 2

[38] 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. 2

[39] Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *ICCV*, pages 945–953, 2015. 3

[40] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum in deep learning. In *ICML*, pages 1139–1147, 2013. 6

[41] 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 Trans. on Computer Vision and Applications*, 10(4):1–14, 2018. 1, 2, 3, 4, 7

[42] Daoliang Tan, Kaiqi Huang, Shiqi Yu, and Tieniu Tan. Efficient night gait recognition based on template matching. In *ICPR*, pages 1000–1003, 2006. 2, 3

[43] Torben Teepe, Ali Khan, Johannes Gilg, Fabian Herzog, Stefan Hörmann, and Gerhard Rigoll. GaitGraph: Graph convolutional network for skeleton-based gait recognition. In *ICIP*, pages 2314–2318, 2021. 2

[44] Chen Wang, Junping Zhang, Jian Pu, Xiaoru Yuan, and Liang Wang. Chrono-gait image: A novel temporal template for gait recognition. In *ECCV*, pages 257–270, 2010. 2

[45] Liang Wang, Tieniu Tan, Huazhong Ning, and Weiming Hu. Silhouette analysis-based gait recognition for human identification. *IEEE TPAMI*, 25(12):1505–1518, 2003. 2

[46] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *Acm Transactions On Graphics (tog)*, 38(5):1–12, 2019. 3, 6, 7

[47] Xin Wei, Ruixuan Yu, and Jian Sun. View-gcn: View-based graph convolutional network for 3d shape analysis. In *CVPR*, pages 1850–1859, 2020. 3

[48] Hiroyuki Yamada, Jeongho Ahn, Oscar Martinez Mozos, Yumi Iwashita, and Ryo Kurazume. Gait-based person identification using 3d lidar and long short-term memory deep networks. *Advanced Robotics*, 34(18):1201–1211, 2020. 3

[49] Shiqi Yu, Daoliang Tan, and Tieniu Tan. A framework for evaluating the effect of view angle, clothing and carrying condition on gait recognition. In *ICPR*, pages 441–444, 2006. 1, 2, 3, 4

[50] Yifu Zhang, Peize Sun, Yi Jiang, Dongdong Yu, Fucheng Weng, Zehuan Yuan, Ping Luo, Wenyu Liu, and Xinggang Wang. Bytetrack: Multi-object tracking by associating every detection box. In *ECCV*, pages 1–21, 2022. 4

[51] Ziyuan Zhang, Luan Tran, Feng Liu, and Xiaoming Liu. On learning disentangled representations for gait recognition. *IEEE TPAMI*, 2020. 1, 3

[52] Ziyuan Zhang, Luan Tran, Xi Yin, Yousef Atoum, Xiaoming Liu, Jian Wan, and Nanxin Wang. Gait recognition via disentangled representation learning. In *CVPR*, pages 4710–4719, 2019. 1, 2, 6

[53] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16259–16268, 2021. 3, 5, 6, 7

[54] Jinkai Zheng, Xinchen Liu, Xiaoyan Gu, Yaoqi Sun, Chuang Gan, Jiyong Zhang, Wu Liu, and Chenggang Yan. Gait recognition in the wild with multi-hop temporal switch. In *ACMMM*, pages 6136–6145, 2022. 1, 6

[55] Jinkai Zheng, Xinchen Liu, Wu Liu, Lingxiao He, Chenggang Yan, and Tao Mei. Gait recognition in the wild with dense 3d representations and a benchmark. In *ICCV*, pages 20228–20237, 2022. 1, 2, 3, 7

[56] Zhedong Zheng, Xiaohan Wang, Nenggan Zheng, and Yi Yang. Parameter-efficient person re-identification in the 3d space. *IEEE TNNLS*, 2022. 1

[57] 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 *CVPR*, pages 14789–14799, 2021. 1, 2, 3