

Language-Guided and Motion-Aware Gait Representation for Generalizable Recognition

ZhengXian Wu^{1*}, Chuanrui Zhang^{1*}, Shen’Ao Jiang^{1*}, Hangrui Xu^{1,2}, Zirui Liao¹, Luyuan Zhang¹, Li Huaqiu¹, Peng Jiao¹, Haoqian Wang^{1†}

¹The Shenzhen International Graduate School, Tsinghua University

²Hefei University of Technology

zx-wu24,zhang-cr22,jiangsa24,liao-zr24,lihq23,jiaop21@mails.tsinghua.edu.cn, hangruixu666,Chang Yu.zly@gmail.com

Abstract

Gait recognition is emerging as a promising technology and an innovative field within computer vision, with a wide range of applications in remote human identification. However, existing methods typically rely on complex architectures to directly extract features from images and apply pooling operations to obtain sequence-level representations. Such designs often lead to overfitting on static noise (e.g., clothing), while failing to effectively capture dynamic motion regions, such as the arms and legs. This bottleneck is particularly challenging in the presence of intra-class variation, where gait features of the same individual under different environmental conditions are significantly distant in the feature space.

To address the above challenges, we present a Language-guided and Motion-aware gait recognition framework, named **LMGait**. To the best of our knowledge, LMGait is the first method to introduce natural language descriptions as explicit semantic priors into the gait recognition task. In particular, we utilize designed gait-related language cues to capture key motion features in gait sequences. To improve cross-modal alignment, we propose the Motion Awareness Module (MAM), which refines the language features by adaptively adjusting various levels of semantic information to ensure better alignment with the visual representations. Furthermore, we introduce the Motion Temporal Capture Module (MTCM) to enhance the discriminative capability of gait features and improve the model’s motion tracking ability. We conducted extensive experiments across multiple datasets, and the results demonstrate the significant advantages of our proposed network. Specifically, our model achieved accuracies of 88.5%, 97.1%, and 97.5% on the CCPG, SUSTech1K, and CASIAB* datasets, respectively, achieving state-of-the-art performance.

1 Introduction

Gait recognition is a biometric technology that identifies individuals based on their posture from a distance, without requiring the subject’s cooperation (Chao et al. 2022). It has emerged as a promising modality in biometric identification. Compared with other methods, such as facial recognition (Meng et al. 2021) and fingerprint recognition (Guo,

*These authors contributed equally to this work.

†Corresponding author: wanghaoqian@tsinghua.edu.cn

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

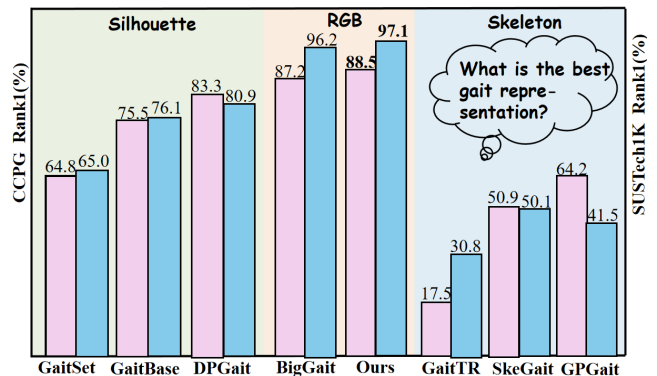


Figure 1: The Rank-1 accuracies of several networks on the CCPG (Li et al. 2023) and SUSTech1K (Shen et al. 2022) datasets are presented, involving three representations: Silhouette, Skeleton, and RGB. The results for CCPG are shown in red, while those for SUSTech1K are shown in blue.

Wu, and Tang 2018), gait recognition offers the advantages of non-contact operation and resistance to disguise.

In gait recognition tasks, the primary goal of the model is to extract highly discriminative and robust features from complex visual inputs. These gait features must effectively capture the motion patterns exhibited by individuals while maintaining consistency across a variety of challenging environments. However, the feature extraction process in real-world scenes is often hindered by several interfering factors. Specifically, variations in view angle, occlusions, dynamic backgrounds and clothing changes can introduce significant changes to the visual appearance of pedestrians. These non-gait-related factors may dominate the visual representation, leading the network to focus on static and noisy appearance cues such as clothing rather than gait-relevant features. This bias results in a significant increase in intra-class variations, where the gait features of the same individual under different conditions become dispersed in the feature space.

In recent years, an increasing number of studies have focused on effectively encoding more gait-relevant cues into feature representations. SkeletonGait++ (Fan et al. 2023b) and ParsingGait (Zheng et al. 2023) leverage skeleton and human parsing maps to guide the model to focus on dynamic

body parts rather than static noise. However, these methods are heavily dependent on precise human parsing or keypoint detection, which requires extensive manual annotations, limiting their scalability in real-world scenarios. BigGait (Ye et al. 2024) is the first to directly use raw RGB images as input, leveraging the strong feature perception ability of a pre-trained large model to extract efficient gait representation. Nevertheless, due to the absence of an explicit mechanism to suppress irrelevant information, the introduction of more gait-relevant cues also leads to increased noise. Moreover, although some methods enhance feature representation by designing architectures such as temporal-frequency separation (Ma et al. 2023), they rely on silhouette images, which contain less noise and are relatively easy to process. How to extract more gait-relevant features from the RGB input is a problem that needs to be solved.

To address the aforementioned challenges, we propose a Language-guided and Motion-aware gait perception framework, named **LMGait**. By jointly leveraging language-guided semantic priors and fine-grained temporal motion modeling, LMGait enhances dynamic gait feature extraction. As far as we know, this is the first attempt to introduce natural language descriptions as explicit semantic priors into the gait recognition task. Specifically, we design domain-specific textual descriptions that emphasize key body regions involved in gait. These semantic cues serve as informative priors and are injected into the visual feature learning pipeline to guide the model’s attention toward dynamically active body parts. To bridge the modality gap and enhance the correspondence between textual semantics and visual features, we introduce a Motion Awareness Module (MAM), which refines language representations through learnable adapters integrated into the transformer layers. MAM provides a content-adaptive refinement mechanism—it learns to emphasize gait-relevant semantics while suppressing distractive cues based on the input context. To further enhance temporal modeling, we propose a Motion Temporal Capture Module (MTCM) to comprehensively capture part-level variations throughout the walking process. This module performs pixel-wise and region-level temporal feature aggregation to learn localized and continuous motion patterns. In particular, the effectiveness of MTCM relies on prior guidance from the language branch, which facilitates the extraction of motion-related body features. We observe that applying temporal modeling directly on raw image features leads to a drop in accuracy, indicating that MTCM is effective only when features are concentrated on key motion regions. As illustrated in Fig. 5, LMGait enables the network to focus more on motion-relevant body parts, effectively reducing interference caused by clothing variations and background noise.

We evaluate LMGait on three gait benchmarks: CCPG (Li et al. 2023), SUSTech1K (Shen et al. 2022), and CASIAB* (Yu, Tan, and Tan 2006). Extensive experiments demonstrate that LMGait achieves the best results in the gait recognition task. Notably, compared to existing methods, LMGait extracts efficient gait features, enhancing its generalization ability. The main contributions are summarized as follows:

- 1) We leverage language cues as an explicit signal to guide the vision encoder to focus on key regions of motion. This approach enables our method to extract more representative gait features. As far as we know, this is the **first attempt** to introduce language-guided learning into the gait recognition task.
- 2) To better align textual and visual features in the latent space, we design a **Motion Awareness Module (MAM)** that dynamically adjusts the textual features.
- 3) We design a **Motion Temporal Capture Module (MTCM)** that captures the walking process of pedestrians from two complementary perspectives: the pixel-level and the region-level.
- 4) As illustrated in Fig. 1, our method achieves state-of-the-art performance on the CCPG, SUSTech1K, and CASIAB* datasets, with Rank-1 accuracies of 88.5%, 97.1%, and 97.5%, respectively.

2 Related Work

2.1 Gait Recognition

Gait recognition methods can be classified into two main categories: model-based and appearance-based methods. Model-based methods consider the basic physical structure of the human body and use interpretable models to represent gait characteristics (Xu et al. 2025). Common input modalities for these methods include 2D skeletons, meshes, and point clouds. For example, PoseGait (Liao et al. 2017) leverages human pose information to extract spatio-temporal features, and employs convolutional neural networks to capture higher-level spatio-temporal representations. DAGait (Wu et al. 2025) adopts a skeleton-guided data augmentation strategy, achieving improved performance in the network.

However, model-based approaches are less effective than appearance-based methods because they lack valuable gait information such as appearance. Appearance-based methods utilize human silhouettes to capture inherent spatial and temporal variations in body shape, clothing, and movement dynamics. Specifically, GaitSet (Chao et al. 2022) innovatively treats the gait sequence as a set and employs a maximum function to compress the sequence of frame-level spatial features. GaitBase (Fan et al. 2024) leverages a universal network based on ResNet and pooling layers to produce features. DeepGait (Fan et al. 2023a) achieves promising results on real-world datasets through its deep network design. Biggait (Ye et al. 2024) is the first approach to utilize RGB as input in an end-to-end framework for gait recognition.

2.2 Fine-tuning Text Encoder for Improved Visual Alignment

While CLIP and related models perform well on general datasets, their generalization often declines on domain-specific or fine-grained tasks. To address this, researchers have proposed various fine-tuning strategies to improve cross-modal alignment. One common approach is prompt tuning (Lester et al. 2021), which involves using task-specific prompts to guide the text encoder toward generating embeddings that better align with visual features. Dual

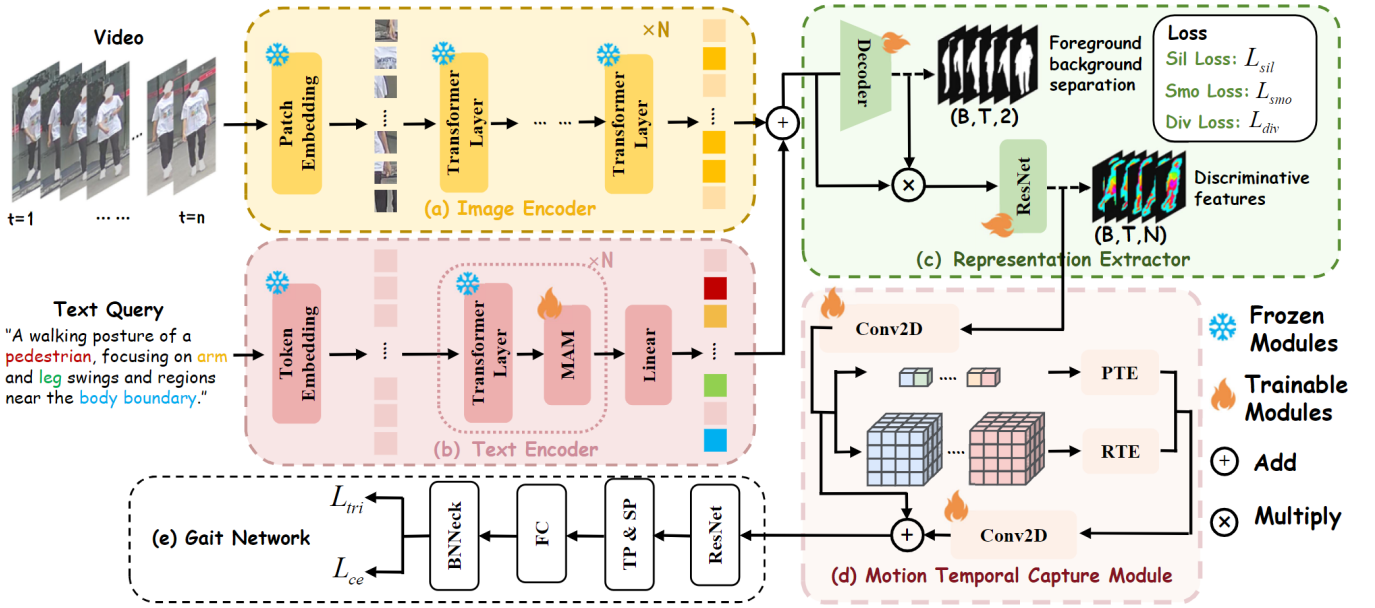


Figure 2: Pipeline of the proposed LMGait, it consists of five components. Specifically, the video input is processed through the frozen DINOv2 model for feature extraction. The text query guides the network to focus on gait-relevant regions, and it is aligned with the image feature space through the frozen CLIP text encoder and the fine-tuned MAM module. The Representation Extractor generates diverse features, while the Motion Temporal Capture Module captures posture changes during walking. Finally, the extracted features are input into the Gait Network for recognition.

Modality Prompt Tuning (DPT) (Xing et al. 2023) introduces learnable prompts to both visual and text encoders, enabling joint optimization.

In addition to prompt tuning, Low-Rank Adaptation (LoRA) (Hu et al. 2022) offers a parameter-efficient way to enhance cross-modal alignment. AdaLoRA (Zhang et al. 2023b) further improved this approach by dynamically adjusting the rank allocation during training, thereby enabling more efficient parameter usage without compromising performance. Another line of work explores adapter tuning (Houlsby et al. 2019), which adds lightweight modules to frozen pre-trained models to adjust feature distributions with minimal overhead.

3 Method

3.1 Pipeline

Gait recognition aims to accurately identify individuals in complex and dynamic environments. We propose a multi-modal gait recognition framework, **LMGait**, which integrates both visual and linguistic prior knowledge to enhance the extraction of gait-discriminative features. As illustrated in Fig. 2, the input is an RGB video sequence, denoted as $X \in \mathbb{R}^{T \times H \times W \times 3}$, where T represents the number of frames, and H and W denote the height and width of each frame. In our setting, we randomly sample t frames from each video sequence as the model input. The LMGait framework consists of five key components: (a) Image Encoder, (b) Text Encoder, (c) Representation Extractor, (d) Motion Temporal Capture Module and (e) Downstream gait recognition network. For image encoding, we employ DINOv2

as the visual feature extractor. Benefiting from its strong representation capabilities learned from large-scale pretraining, DINOv2 effectively captures the spatial distribution of human body parts. Meanwhile, we design domain-specific textual descriptions and feed them into the Text Encoder (in Section 3.2). To mitigate the modality gap and improve the alignment between linguistic cues and visual representations, we introduce a Motion Awareness Module (MAM). MAM fine-tunes the Text Encoder by adaptively refining textual representations to emphasize motion-relevant semantics and suppress distractive cues. Textual features are then fused into the visual feature map to provide semantic guidance, enabling the model to focus more effectively on motion-discriminative regions. The Representation Extractor Module (Ye et al. 2024) is designed to mitigate background interference and generate more representative feature embeddings. To capture temporal information across frames, we design the Motion Temporal Capture Module (MTCM), which learns both pixel-level and region-level motion patterns over time. Finally, the integrated multi-modal features are fed into the downstream gait recognition network to perform individual identity classification.

3.2 Text Encoder

To inject domain-specific semantic priors into the visual feature learning process, we incorporate a lightweight text encoder that transforms language prompts into informative representations, as illustrated in Fig. 3. These textual descriptions are designed to emphasize body regions typically associated with gait motion, such as arm and leg swings, offering high-level semantic cues complementary to low-level

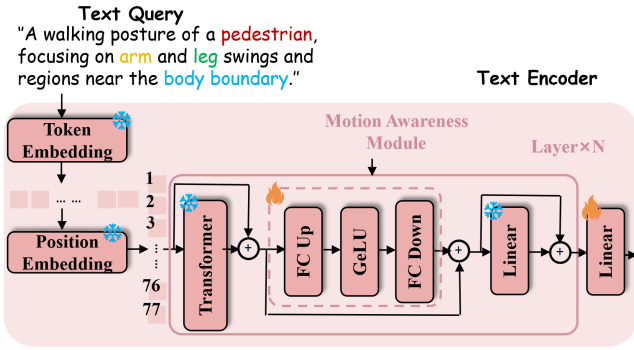


Figure 3: The architecture of the Text Encoder is designed to improve the alignment between text and image features. To achieve this, we introduce the MAM module and fine-tune it.

visual features. Formally, given an input sentence describing gait-related actions, we first tokenize the sentence using a frozen CLIP tokenizer:

$$\mathcal{T} = \{w_1, w_2, \dots, w_n\}. \quad (1)$$

Each token is mapped to an embedding vector e_i through a lookup table:

$$E = [e_1, e_2, \dots, e_n] \in \mathbb{R}^{n \times d}. \quad (2)$$

We then add fixed sinusoidal positional embeddings

$$P = [p_1, p_2, \dots, p_n] \in \mathbb{R}^{n \times d} \quad (3)$$

to incorporate sequence order, resulting in the input to the encoder:

$$H^{(0)} = E + P. \quad (4)$$

This embedded sequence is passed through a stack of L frozen Transformer layers $\{T^{(1)}, T^{(2)}, \dots, T^{(L)}\}$, each equipped with our proposed Motion Awareness Module (MAM), which adaptively adjusts the token representations:

$$H^{(l)} = \text{MAM}^{(l)} \left(T^{(l)}(H^{(l-1)}) \right), \quad l = 1, \dots, L. \quad (5)$$

Each MAM module refines the intermediate hidden states using a residual adapter structure:

$$\text{MAM}(h) = h + \phi(hW_d)W_u, \quad (6)$$

where $W_d \in \mathbb{R}^{d \times r}$ and $W_u \in \mathbb{R}^{r \times d}$ are trainable projection matrices with $r \ll d$, and $\phi(\cdot)$ is a non-linear activation function.

The design of MAM is based on the insight that the relevance of textual tokens varies with different visual inputs. For example, when the lower limbs are occluded in a video frame, semantic cues related to arm movements should become more informative. Instead of using a fixed attention mask or static prompt encoding, MAM enables the model to dynamically reweight semantic components based on their relevance to current gait scenarios. The final language representation is projected into a shared multimodal space using:

$$z_{\text{text}} = H^{(L)} \cdot W_{\text{proj}}, \quad (7)$$

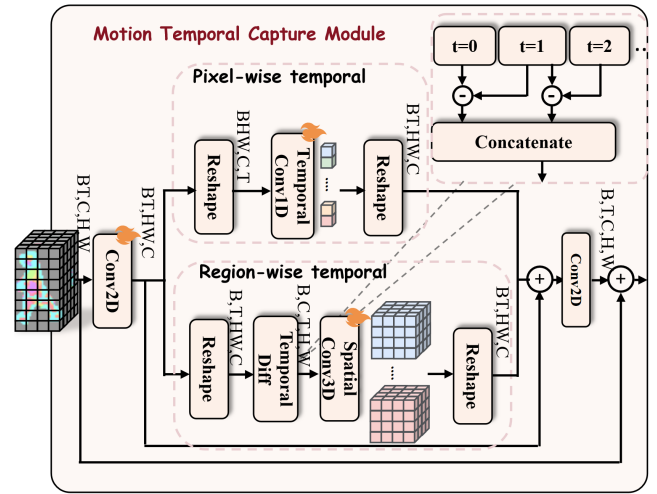


Figure 4: The architecture of the Motion Temporal Capture Module operates from two complementary perspectives: pixel-wise temporal extraction and region-wise temporal extraction. The two branches collaboratively capture gait information.

where $W_{\text{proj}} \in \mathbb{R}^{d \times d'}$ is a learnable matrix shared with the visual encoder's output space.

During training, the text feature z_{text} is fused with the DINOv2 generated visual feature map, serving as a semantic attention mask that guides the model toward dynamically active body regions. At the same time, to ensure efficient inference in real-world scenarios, the text encoder is used only during training to generate semantic modulation maps. These refined attention maps can be precomputed and cached, allowing them to be directly integrated into the visual feature stream during inference. As a result, the entire language branch can be removed at test time, enabling fast inference without compromising performance.

3.3 Motion Temporal Capture Module

Given that human walking is a continuous and temporal process, modeling gait as a time-continuous signal is more reasonable than treating each frame independently. To address this, we propose the **Motion Temporal Capture Module (MTCM)**, as illustrated in Figure 4. A key novelty of MTCM lies in its synergy with the language-guided branch. As demonstrated by our empirical analysis in Section 4.3, when RGB images are used as input, directly applying temporal modeling to unrefined visual features tends to cause the network to aggregate irrelevant or misleading cues. Benefiting from the language-guided branch that enhances attention to motion-relevant regions, MTCM achieves more effective modeling of motion trajectories. MTCM consists of two complementary components. Specifically, the **Pixel-wise Temporal Extraction (PTE)** module is designed to capture fine-grained and long-range motion patterns at individual pixel locations, while the **Region-wise Temporal Extraction (RTE)** module focuses on structured movements and posture transitions at the regional level.

Representation	Method	Testing Datasets									
		CCPG					SUSTech1K				
		CL	UP	DN	BG	Mean	NM	CL	UF	NT	Mean
		Rank-1					Rank-1				
skeleton	GaitGraph2(Teepe et al. 2021)	5.0	5.3	5.8	6.2	5.6	22.2	6.8	19.2	16.4	18.6
	Gait-TR (Zhang et al. 2023a)	15.7	18.3	18.5	17.5	17.5	33.3	21.0	34.6	23.5	30.8
	GPGait (Fu et al. 2023)	54.8	65.6	71.6	65.4	64.2	44.0	24.3	47.0	31.8	41.4
	SkeletonGait (Fan et al. 2023b)	40.4	48.5	53.0	61.7	50.9	55.0	24.7	52.0	43.9	50.1
Silhouette	GaitSet (Chao et al. 2022)	60.2	65.2	65.1	68.5	64.8	69.1	61.0	23.0	65.0	18.6
	GaitBase (Fan et al. 2024)	71.6	75.0	76.8	78.6	75.5	81.5	49.6	76.7	25.9	76.1
	DeepGaitV2 (Fan et al. 2023a)	78.6	84.8	80.7	89.2	83.3	86.5	49.2	81.9	28.0	80.9
	SkeletonGait++ (Fan et al. 2023b)	79.1	83.9	81.7	89.9	83.7	85.1	46.6	82.5	47.5	81.3
Sil+Parsing+Flow	MultiGait++ (Jin et al. 2024)	83.9	89.0	86.0	91.5	87.6	92.0	50.4	89.1	45.1	87.4
RGB	GaitEdge (Liang et al. 2022)	66.9	74.0	70.6	77.1	72.7	-	-	-	-	-
	BigGait (Ye et al. 2024)	82.6	85.9	87.1	93.1	87.2	96.1	73.3	93.2	85.3	96.2
	LMGait(ours)	84.8	87.0	88.5	93.6	88.5	96.4	79.8	93.9	87.0	97.1

Table 1: Comparison of state-of-the-art methods on the CCPG (Li et al. 2023) and SUSTech1K (Shen et al. 2022) datasets.

Pixel-wise Temporal Extraction (PTE). Let the input visual feature map sequence be $\{F^{(1)}, F^{(2)}, \dots, F^{(T)}\}$, where $F^{(t)} \in \mathbb{R}^{H \times W \times C}$ is the feature map at frame t . For a specific pixel location (i, j) , its temporal trajectory over T frames is denoted as:

$$x_{i,j} = [F_{i,j}^{(1)}, F_{i,j}^{(2)}, \dots, F_{i,j}^{(T)}] \in \mathbb{R}^{T \times C}. \quad (8)$$

To extract temporal dependencies at this location, we apply a 1D convolution along the temporal dimension:

$$y_{i,j} = \sigma(\text{Conv1D}(x_{i,j})) \in \mathbb{R}^{T' \times C'}, \quad (9)$$

where $\sigma(\cdot)$ denotes a non-linear activation function.

Region-wise Temporal Extraction (RTE). To capture spatial-temporal structures, RTE computes inter-frame feature differences that emphasize short-term motions:

$$\Delta F^{(t)} = F^{(t+1)} - F^{(t)}, \quad t = 1, \dots, T-1 \quad (10)$$

The resulting sequence $\{\Delta F^{(1)}, \dots, \Delta F^{(T-1)}\} \in \mathbb{R}^{(T-1) \times H \times W \times C}$ captures localized motion transitions. These are stacked and passed through a 3D convolutional block:

$$R = \sigma(\text{Conv3D}([\Delta F^{(1)}, \dots, \Delta F^{(T-1)}])). \quad (11)$$

The outputs of PTE and RTE are fused using channel-wise concatenation followed by a 2D convolution:

$$F_{\text{motion}} = \text{Conv2D}([\text{PTE}(F), \text{RTE}(F)]) \in \mathbb{R}^{H \times W \times C''}. \quad (12)$$

These two branches jointly enhance temporal modeling capacity. Specifically, MTCM benefits from the text-guided visual representation. These semantic priors focus the input features on motion-sensitive regions, allowing MTCM to perform effective motion aggregation.

3.4 Representation Extractor and Loss Function

To further enhance feature robustness, Representation Extractor aims to suppress background interference and generate diverse, discriminative feature representations. The decoder branch is designed to reduce the feature dimensionality by producing a two-channel output, corresponding to

background and foreground regions. The silhouette supervision is only applied during training and is not required during inference. Then, the ResNet branch is regularized by loss functions to produce smooth and discriminative features.

To achieve the best performance, we employ five loss functions: Triplet loss (L_{tri}) and cross-entropy loss (L_{ce}) are used to supervise the downstream gait recognition network.

$$L_{\text{tri}} = [D(f_i, f_k) - D(f_i, f_j) + m]_+, \quad (13)$$

$$L_{\text{ce}} = - \sum_{c=1}^C y_c \log(\hat{y}_c), \quad (14)$$

f_i , f_j , and f_k represent the feature embeddings of the anchor, positive, and negative samples, respectively. m is the predefined margin. Smoothness loss (L_{smo}) and diversity loss (L_{div}) are used to supervise the Representation Extractor module, encouraging the generation of smooth and diverse features.

$$L_{\text{smo}} = \|\text{Sobel}_x * \mathbf{F}\|_1 + \|\text{Sobel}_y * \mathbf{F}\|_1, \quad (15)$$

$$L_{\text{div}} = - \frac{1}{C} \sum_{c=1}^C \sum_{i=1}^H \sum_{j=1}^W p_{c,i,j} \log(p_{c,i,j}). \quad (16)$$

\mathbf{F} denotes the feature map generated by the ResNet branch. Sobel_x and Sobel_y are the Sobel operators along the horizontal and vertical directions, respectively. $p_{c,i,j}$ is the normalized activation at position (i, j) in channel c . Silhouette Supervision Loss (L_{rec}) is employed to supervise the generation of silhouettes. The overall loss can be formulated as:

$$\mathcal{L} = \mathcal{L}_{\text{tri}} + \mathcal{L}_{\text{ce}} + \alpha \mathcal{L}_{\text{sil}} + \beta \mathcal{L}_{\text{smo}} + \gamma \mathcal{L}_{\text{div}}. \quad (17)$$

4 Experiments

4.1 Experimental Setup

In this paper, our method is evaluated on three challenging gait datasets, **CCPG** (Li et al. 2023), **SUSTech1K** (Shen et al. 2022) and **CASIA-B** (Yu, Tan, and Tan 2006).

Representation	Method	NM	CL	Mean
Skeleton	GaitGraph2	80.3	63.8	71.8
	GaitTR	94.7	86.7	90.2
	GPgait	93.6	69.3	81.0
Silhouette	GaitSet	92.3	73.4	84.0
	GaitBase	96.5	78.0	88.7
	GaitGCI	98.4	88.5	94.5
	VPNet	98.3	90.7	94.9
RGB	BigGait	100.0	90.5	96.7
	LMGait(ours)	100.0	92.9	97.5

Table 2: Comparison of Rank-1 accuracy between state-of-the-art methods on the CASIAB* (Yu, Tan, and Tan 2006).

The **CCPG** dataset is designed for cloth-changing gait recognition. The dataset is split into 100 subjects for training and 100 for testing, with three cloth-changing scenarios during testing: full clothing change (CL-FULL), upper garment change (CL-UP), and lower garment change (CL-DN). The **CASIA-B** dataset includes 124 subjects captured from 11 different view angles, ranging from 0° to 180° . The **SUSTech1K** dataset encompasses various factors like visibility, view angles, occlusions and different environments.

The experimental implementation in this study followed the official protocols for each dataset (Fan et al. 2024), including the standard partitioning strategies for the training, gallery, and probe sets. We adopt the average Rank-1, Rank-5, and mean Average Precision (mAP) as the evaluation metrics. In the training stage, the number of frames for each sequence is set to 30 (for the CCPG dataset) and 24 frames (for the SUSTech and CASIAB datasets). The Learning Rate (LR) starts at 0.1 and is subsequently multiplied by 0.1 at the 15K, 25K, 30K, and 35K iterations. The SGD optimizer is adopted with a weight decay of $5e-4$. The loss weights α , β , and γ are 1.0, 0.01 and 5.0. We train the model for 40k iterations with batch size of (6,8). The number of text adapter and image transformer layers N is set to 12. The text encoder uses a vocabulary size of 49,408 and supports a maximum sequence length of 77 tokens. To prevent the Silhouette Supervision Loss (L_{sil}) from dominating the overall optimization, we warm up the decoder branch with 5K iterations using only L_{sil} . The decoder weights obtained from this stage are then used to initialize the decoder.

4.2 Comparison with State-Of-The-Art Methods

In this section, we compare the proposed LMGait with several popular state-of-the-art(SOTA) gait recognition methods. The experimental results on the CCPG and SUSTech1K datasets are presented in Table 1. The results on the CASIAB* dataset are shown in Table 2. LMGait achieves significantly better performance than other SOTA methods.

Taking the CCPG dataset as an example, LMGait outperforms the current skeleton-based SOTA method GP-Gait by 24.3% in Rank-1 accuracy (88.5% vs. 64.2%). For silhouette-based methods, which are the most widely used representations in gait recognition, LMGait surpasses the SOTA method SkeletonGait++ by 4.8% (88.5% vs. 83.7%). Using RGB images directly as input poses greater challenges for focusing on gait-related cues. Currently, GaitEdge

Model	Test Set				
	CASIA-B*		SUSTech1K		
	NM	BG	Clothing	Night	Overall
GaitSet	47.4	40.9	8.2	11.0	12.8
GaitBase	59.1	52.7	9.5	13.1	16.8
BigGait	77.4	71.5	43.7	44.8	56.4
Ours	81.6	75.4	45.6	48.2	59.3

Table 3: Cross-dataset evaluation is conducted by training the model on CCPG and testing on CASIA-B* and SUSTech1K with Rank-1 accuracy as the evaluation metric.

Method	CL	UP	Mean
w/o Text Branch & w/o MTCM	82.6	85.9	87.2
w/o Text Branch & w/ 3D-CNN	82.3	85.0	86.4
w/o Text Branch & w/ MTCM	82.8	84.5	86.9
w/ Text Branch & w/ MTCM	84.8	87.0	88.5

Table 4: Ablation study on the correlation between the Text Branch and the Motion Temporal Capture Module (MTCM).

and BigGait are the only two methods that operate directly on RGB inputs. Compared to BigGait, our method achieves improvements of 2.2%, 1.1%, 1.4%, and 0.5% in the CL, UP, DN, and BG scenarios, with an average gain of 1.3% (88.5% vs. 87.2%). We then conducted cross-dataset evaluation experiments by training our network on the CCPG dataset and evaluating it on the SUSTech1K and CASIA-B* datasets. The results are presented in Table 3. These datasets exhibit significant variations in terms of subjects and scenes. Compared to other methods, LMGait, by focusing on discriminative information, achieves higher accuracy in cross-dataset evaluations. We provide the complete evaluation results on all three datasets, as well as additional cross-dataset evaluation results, in the supplementary material. Based on the above analysis, we conclude that LMGait extracts better gait features and mitigates the interference from irrelevant noise.

4.3 Ablation Study and Visualization

The correlation between the Motion Temporal Capture Module (MTCM) and the Text Encoder. The MTCM and the text branch serve complementary roles in our framework. Under the explicit guidance of textual priors, the network is encouraged to focus on motion-relevant regions—such as the arms and legs—which enables MTCM to more effectively capture temporal motion patterns. As shown in Table 4, when the text branch is removed, using either MTCM or an alternative temporal module (3D-CNN) results in a performance drop compared to the baseline without any temporal modeling (second and third row vs first row). In contrast, the full model (bottom row), which incorporates both the text branch and MTCM, achieves the highest accuracy across all scenarios.

Motion Temporal Capture Module(MTCM). The MTCM module consists of two complementary branches: pixel-wise Temporal Extraction (PTE) and region-wise Temporal Extraction (RTE). These branches work collaboratively to enhance the model’s capability in capturing temporal dynamics. As shown in Table 5, removing any compo-

Method	CL	UP	DN	BG	Mean
w.o. PTE-TPConv1D	82.9	86.2	87.6	93.2	87.5
w.o. RTE Temp-Diff	83.8	86.8	88.5	93.5	88.2
w.o. RTE-SPConv3D	83.6	86.6	88.2	93.4	87.9
Ours	84.8	87.0	88.5	93.6	88.5

Table 5: The contribution of different components within the MTCM Module to Rank-1 Accuracy (on the CCPG dataset).

Method	CL	UP	DN	Mean
w.o. Text Branch & MAM	82.8	84.5	87.2	86.9
w.o. MAM + Wrong Text	81.6	83.8	86.8	86.1
w. MAM + Wrong Text	82.8	84.9	87.7	87.1
w.o. MAM + Correct Text	83.1	86.4	87.5	87.6
w. MAM + Random Text	83.8	86.6	88.2	88.1
w.MAM & Correct Text	84.8	87.0	88.5	88.5

Table 6: Effect of the Text Encoder on Rank-1 Recognition Accuracy (on the CCPG dataset).

ment of MTCM leads to a drop in overall performance.

Text Encoder. We further evaluate the contribution of the text branch to the overall performance, as summarized in Table 6. When the text branch is removed, a substantial drop in accuracy is observed (86.9% vs. 88.5%). This performance drop is caused by the MTCM module aggregating information that is irrelevant to gait. This issue becomes even more pronounced when incorrect textual inputs are provided to the model. Specifically, we constructed a misleading text query: **”Attention is given to environmental details such as the blue sky, the clean ground, and the clear windows.”** When this erroneous information is integrated into the feature maps without any refinement, the accuracy drops significantly by 2.4% (86.1% vs. 88.5%). Introducing a refinement module for these incorrect texts helps mitigate such adverse effects to some extent. When correct text descriptions are used but the refinement module is disabled, the network still benefits from the semantic priors provided by language encoder, guiding attention toward gait-relevant body. Meanwhile, we deployed the MLLM (Li et al. 2024) to generate generic descriptions. After six test runs, the average performance was lower than that achieved with our manually designed text inputs (88.1 % vs 88.5%). Detailed information can be found in our supplementary material.

Visual Encoder and Inference Efficiency. Table 7 presents a comparison between two visual encoders, CLIP and DINOv2, under the same framework. While the image branch of CLIP may be better aligned with the text branch in the shared feature space, DINOv2 is better at capturing fine-grained visual features. In the supplementary material, we also analyze that fine-tuning the visual encoder can accelerate the convergence process but leads to a drop in accuracy. In terms of **inference efficiency**, gait recognition often demands real-time performance in applications. In our framework, the text encoder is only used during training to generate semantic attention guidance related to gait. Once training is complete, the resulting attention maps can be stored and directly integrated with visual features at inference time, eliminating the need to run the text branch again. Removing

Visual Encoder	CL	UP	DN	BG	Mean
CLIP Image Encoder	68.7	75.1	72.2	79.4	73.9
ours	84.8	87.0	88.5	93.6	88.5

Table 7: The impact of different visual encoders on accuracy.

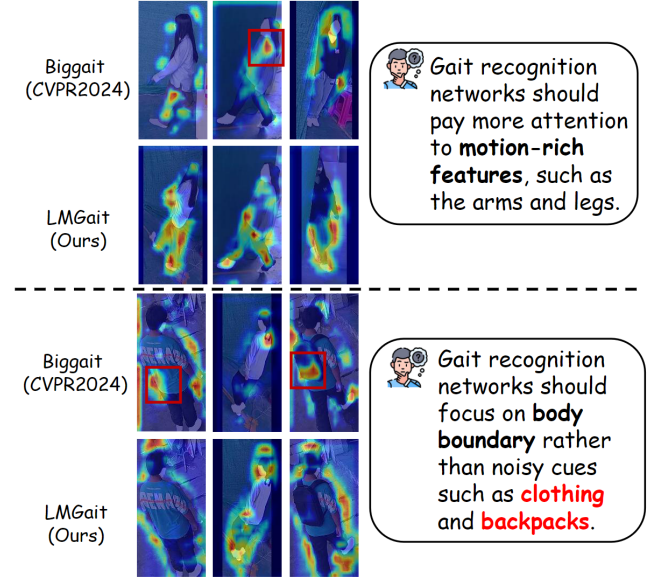


Figure 5: Visualization of attention maps from the first layer of the downstream gait recognition network.

the **loss** supervision from the Representation Extractor also leads to a performance drop. Detailed ablation studies on the visual encoder, inference efficiency, representation extractor, and loss functions can be found in our supplementary material. As shown in Figure 5, we visualize the attention maps from the first layer of the downstream gait recognition network. In contrast, our method effectively extracts more discriminative features from gait-relevant regions.

5 Conclusion

In this paper, we present a language-guided and motion-aware gait recognition framework, named LMGait. This is the first attempt to introduce natural language descriptions as explicit semantic priors into the gait recognition task. With the guidance of semantic cues, the network can extract features that are more relevant to gait. To further align the text and image spaces, we design the Motion Awareness Module (MAM) to fine-tune the textual features. Furthermore, we propose the Motion Temporal Capture Module (MTCM) to capture the changes in posture as a person walks. The text branch and the MTCM work collaboratively to focus the features on motion-related body regions. Results show LMGait outperforms previous methods on multiple datasets.

Acknowledgments

This work is supported by the Shenzhen Science and Technology Project under Grant KJZD20240903103210014.

References

- Chao, H.; Wang, K.; He, Y.; Zhang, J.; and Feng, J. 2022. GaitSet: Cross-View Gait Recognition Through Utilizing Gait As a Deep Set. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 44(07): 3467–3478.
- Fan, C.; Hou, S.; Huang, Y.; and Yu, S. 2023a. Exploring Deep Models for Practical Gait Recognition. *ArXiv*, abs/2303.03301.
- Fan, C.; Hou, S.; Liang, J.; Shen, C.; Ma, J.; Jin, D.; Huang, Y.; and Yu, S. 2024. OpenGait: A Comprehensive Benchmark Study for Gait Recognition towards Better Practicality. *ArXiv*, abs/2405.09138.
- Fan, C.; Ma, J.; Jin, D.; Shen, C.; and Yu, S. 2023b. SkeletonGait: Gait Recognition Using Skeleton Maps. *ArXiv*, abs/2311.13444.
- Fu, Y.; Meng, S.; Hou, S.; Hu, X.; and Huang, Y. 2023. GPGait: Generalized Pose-based Gait Recognition. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, 19538–19547.
- Guo, X.; Wu, F.; and Tang, X. 2018. Fingerprint Pattern Identification and Classification. In *2018 14th International Conference on Natural Computation, Fuzzy Systems and Knowledge Discovery (ICNC-FSKD)*, 1045–1050.
- Houlsby, N.; Giurgiu, A.; Jastrzebski, S.; Morrone, B.; De Laroussilhe, Q.; Gesmundo, A.; Attariyan, M.; and Gelly, S. 2019. Parameter-efficient transfer learning for NLP. In *International conference on machine learning*, 2790–2799. PMLR.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Jin, D.; Fan, C.; Chen, W.; and Yu, S. 2024. Exploring More from Multiple Gait Modalities for Human Identification. *ArXiv*, abs/2412.11495.
- Lester, B.; Al-Rfou, R.; Constant, D.; Noah; Zhang, S.; Liang, G.; Wang, P.; and Zhang, Y. 2021. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*.
- Li, B.; Zhang, Y.; Guo, D.; Zhang, R.; Li, F.; Zhang, H.; Zhang, K.; Li, Y.; Liu, Z.; and Li, C. 2024. LLaVA-OneVision: Easy Visual Task Transfer. *ArXiv*, abs/2408.03326.
- Li, W.; Hou, S.; Zhang, C.; Cao, C.; Liu, X.; Huang, Y.; and Zhao, Y. 2023. An In-Depth Exploration of Person Re-Identification and Gait Recognition in Cloth-Changing Conditions. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 13824–13833.
- Liang, J.; Fan, C.; Hou, S.; Shen, C.; Huang, Y.; and Yu, S. 2022. Gaitedge: Beyond plain end-to-end gait recognition for better practicality. In *European Conference on Computer Vision*, 375–390. Springer.
- Liao, R.; Cao, C.; Garcia, E. B.; Yu, S.; and Huang, Y. 2017. Pose-based temporal-spatial network (PTSN) for gait recognition with carrying and clothing variations. In *Chinese Conference on Biometric Recognition*, 474–483. Springer.
- Ma, K.; Fu, Y.; Zheng, D.; Cao, C.; Hu, X.; and Huang, Y. 2023. Dynamic aggregated network for gait recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22076–22085.
- Meng, Q.; Zhao, S.; Huang, Z.; and Zhou, F. 2021. Mag-Face: A Universal Representation for Face Recognition and Quality Assessment. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 14220–14229.
- Shen, C.; Fan, C.; Wu, W.; Wang, R.; Huang, G. Q.; and Yu, S. 2022. LidarGait: Benchmarking 3D Gait Recognition with Point Clouds. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 1054–1063.
- Teepe, T.; Khan, A.; Gilg, J.; Herzog, F.; Hörmann, S.; and Rigoll, G. 2021. Gaitgraph: Graph convolutional network for skeleton-based gait recognition. In *2021 IEEE international conference on image processing (ICIP)*, 2314–2318. IEEE.
- Wu, Z.; Zhang, C.; Xu, H.; Jiao, P.; and Wang, H. 2025. DAGait: Generalized Skeleton-Guided Data Alignment for Gait Recognition. *arXiv preprint arXiv:2503.18830*.
- Xing, Y.; Wu, Q.; Cheng, D.; Zhang, S.; Liang, G.; Wang, P.; and Zhang, Y. 2023. Dual modality prompt tuning for vision-language pre-trained model. *IEEE Transactions on Multimedia*, 26: 2056–2068.
- Xu, H.; Zhang, C.; Wu, Z.; Jiao, P.; and Wang, H. 2025. PSGait: Gait Recognition using Parsing Skeleton. *arXiv:2503.12047*.
- Ye, D.; Fan, C.; Ma, J.; Liu, X.; and Yu, S. 2024. BigGait: Learning Gait Representation You Want by Large Vision Models. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 200–210.
- Yu, S.; Tan, D.; and Tan, T. 2006. 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, 441–444.
- Zhang, C.; Chen, X.-P.; Han, G.-Q.; and Liu, X.-J. 2023a. Spatial transformer network on skeleton-based gait recognition. *Expert Systems*, 40(6): e13244.
- Zhang, Q.; Chen, M.; Bukharin, A.; Karampatziakis, N.; He, P.; Cheng, Y.; Chen, W.; and Zhao, T. 2023b. Adalora: Adaptive budget allocation for parameter-efficient fine-tuning. *arXiv preprint arXiv:2303.10512*.
- Zheng, J.; Liu, X.; Wang, S.; Wang, L.; Yan, C. C.; and Liu, W. 2023. Parsing is All You Need for Accurate Gait Recognition in the Wild. *Proceedings of the 31st ACM International Conference on Multimedia*.