

Gradient-Attention Guided Dual-Masking Synergetic Framework for Robust Text-based Person Retrieval

Tianlu Zheng^{♥*}, Yifan Zhang^{*†}, Xiang An[♣], Ziyong Feng[♣]
Kaicheng Yang^{♣†}, Qichuan Ding^{♥†}

[♥]Northeastern University ^{*†}South China University of Technology [♣]DeepGlint

zhengt2@mails.neu.edu.cn, dingqichuan@mail.neu.edu.cn, kaichengyang@deepglint.com

Abstract

Although Contrastive Language-Image Pre-training (CLIP) exhibits strong performance across diverse vision tasks, its application to person representation learning faces two critical challenges: (i) the scarcity of large-scale annotated vision-language data focused on person-centric images, and (ii) the inherent limitations of global contrastive learning, which struggles to maintain discriminative local features crucial for fine-grained matching while remaining vulnerable to noisy text tokens. This work advances CLIP for person representation learning through synergistic improvements in data curation and model architecture. First, we develop a noise-resistant data construction pipeline that leverages the in-context learning capabilities of MLLMs to automatically filter and caption web-sourced images. This yields **WebPerson**, a large-scale dataset of 5M high-quality person-centric image-text pairs. Second, we introduce the **GA-DMS** (Gradient-Attention Guided **Dual-Masking Synergetic**) framework, which improves cross-modal alignment by adaptively masking noisy textual tokens based on the gradient-attention similarity score. Additionally, we incorporate masked token prediction objectives that compel the model to predict informative text tokens, enhancing fine-grained semantic representation learning. Extensive experiments show that GA-DMS achieves state-of-the-art performance across multiple benchmarks. The data and pre-trained models are released at <https://github.com/Multimodal-Representation-Learning-MRL/GA-DMS>.

1 Introduction

The rapid advancement of large-scale vision-language pre-training (Chen et al., 2023; Yang et al., 2023a; Gu et al., 2024; Wu et al., 2023)


^{*}Equal contribution.

[†]Corresponding author.

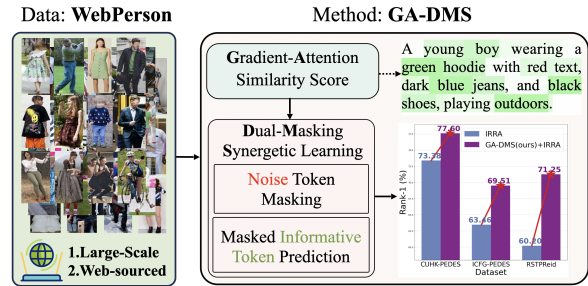
Limited data source and scale			Presence of noise and hard to learn fine-grained semantic
Dataset	Source	Size	
CUHK-PEDES	Surveillance Cameras	40K	
ICFG-PEDES	Surveillance Cameras	54K	
RSTPReid	Surveillance Cameras	20K	
MALS	Generative Model	1M+	
LUPerson-T	LUPerson	1M	
LUPerson-MLLM	LUPerson	1M	

Data Construction

Person Representation Learning



(a) Current work exhibits several deficiencies



(b) Our method for robust person representation learning

Figure 1: Current human-centric datasets are limited in diversity and scale, complicating model training due to noise interference and hindering the effective learning of fine-grained semantics.

has been driven by the unprecedented availability of web-sourced image-text pairs. As a milestone in vision-language representation learning, Contrastive Language-Image Pre-training (CLIP) (Radford et al., 2021b) employs dual encoders for visual and textual modalities and leverages a contrastive loss mechanism (Wang and Liu, 2021) to learn joint representations. Trained on 400 million noisy web-curated image-text pairs, CLIP exhibits strong zero-shot generalization and has been widely adopted for tasks including image classification (Abdelfattah et al., 2023; Peng et al., 2023; An et al., 2023), retrieval (Sain et al., 2023; Shao et al., 2023; Hu et al., 2025), and grounding (Xiao et al., 2023; An et al., 2024; Xie et al., 2025). However, CLIP shows suboptimal performance in text-based person retrieval, as evidenced by recent studies (Shao et al., 2023; Yan et al., 2023b; Li et al., 2023; Han et al., 2024; Zhao et al., 2025).

CLIP’s suboptimal performance in text-based

person retrieval stems from two key limitations. First, the scarcity and noise levels in person-centric image-text data pose significant challenges. Existing datasets such as CUHK-PEDES (Li et al., 2017), ICFG-PEDES (Ding et al., 2021), and RST-PRid (Zhu et al., 2021) are constrained in scale due to their reliance on extensive manual annotations. Although large-scale person-centric datasets like LUPerson (Fu et al., 2021) comprise approximately 200K identities and 4 million images, they lack corresponding textual descriptions. Recent efforts (Tan et al., 2024) have employed Multimodal Large Language Models (MLLMs) to address data scarcity by generating synthetic captions. However, these automatically produced annotations frequently contain inaccuracies and semantic misalignments, thereby introducing noise into the training process and requiring the implementation of corrective strategies (Zhao et al., 2024). Second, CLIP’s global contrastive learning paradigm fails to effectively capture fine-grained visual semantics crucial for distinguishing similar individuals (Yan et al., 2023b; Liu et al., 2024). This is particularly problematic as person retrieval often relies on localized attributes (e.g., clothing patterns or accessories) that require precise visual-semantic alignment.

In this work, we advance CLIP for person representation learning through synergistic improvements in data curation and model architecture (Fig.1). We initially introduce **WebPerson**, a large-scale person-centric dataset consisting of 5 million high-quality text-image pairs derived from web-sourced images. After that, we propose the **GA-DMS** (Gradient-Attention Guided Dual-Masking Synergetic) framework, which enhances cross-modal alignment by masking noisy textual tokens based on a gradient-attention similarity score. Meanwhile, we incorporate masked token prediction objectives to enforce the model to predict informative text tokens, thereby enhancing fine-grained semantic representation learning. Extensive experiments demonstrate that GA-DMS establishes new state-of-the-art performance across multiple benchmarks. The main contributions of this paper are summarized as follows:

- We **design a novel person-centric data construction pipeline** that automatically filters and annotates web-sourced images, yielding the **WebPerson** dataset with 5 million high-quality image-text pairs.
- We **propose the GA-DMS** (Gradient-

Attention Guided **Dual-Masking Synergetic**) framework to improve cross-modal alignment through gradient-attention guided noisy text token masking while enhancing fine-grained visual-semantic correspondence via masked informative token prediction objectives.

- We **conduct comprehensive experiments and demonstrate that GA-DMS achieves new state-of-the-art performance** across multiple text-based person retrieval datasets.

2 Related Works

2.1 Person Representation Learning

Early approaches to text-based person retrieval typically employ separate vision and language encoders with custom alignment losses (Zheng et al., 2020; Si et al., 2018). These methods often exhibit suboptimal modality alignment and require extensive manual annotation. The introduction of CLIP (Radford et al., 2021b) establishes a unified vision-language embedding space, significantly advancing cross-modal matching. Recent works extend CLIP with specialized modules for text-based person retrieval. IRRA (Jiang and Ye, 2023) merges visual cues into textual tokens via a cross-modal transformer and aligns global similarity distributions. MDRL (Yang et al., 2024b) designs a cross-modality global feature learning architecture to learn the global features from the two modalities and meet the demand of the task. UniPT (Shao et al., 2023) utilizes a simple vision-and-language pre-training framework to explicitly align the feature space of the image and text modalities during pre-training. However, these approaches largely ignore data noise, which critically influences cross-modal alignment in feature space. RDE (Qin et al., 2024) mitigates the adverse impact of noisy through the proposed confident consensus division and novel triplet alignment loss. ProPOT (Yan et al., 2024) transforms the identity-level matching problem into a prototype learning problem, aiming to learn identity-enriched prototypes. However, prototype aggregation compromises fine-grained semantic learning.

2.2 Person-centric Dataset

High-quality image-text paired datasets are essential for learning discriminative person representations. However, existing manually annotated datasets (e.g., CUHK-PEDES (Li et al., 2017), ICFG-PEDES (Ding et al., 2021), RSTPRid (Zhu

et al., 2021)) face severe scalability limitations due to labor-intensive annotation processes. This scalability bottleneck ultimately constrains models’ capacity to acquire diverse semantic information and learn discriminative features. Recent efforts to mitigate this issue focus on constructing large-scale datasets, such as LUPerson (Fu et al., 2021), LUPerson-T (Shao et al., 2023), LUPerson-MLLM (Tan et al., 2024), and SYNTH-PEDES (Zuo et al., 2024) demonstrate that increased data volume improves general pedestrian representation learning. Nevertheless, these datasets primarily derive from video sources, inheriting inherent scalability constraints from computationally intensive video processing pipelines. The success of multimodal large language models in cross-modal understanding (Yu et al., 2024) has inspired their application to synthetic data generation. For instance, LUPerson-MLLM (Tan et al., 2024) employs template-guided MLLMs to generate diverse textual descriptions, significantly enhancing text-to-image ReID performance. However, this approach remains limited by its dependence on existing LUPerson image collections.

3 WebPerson Dataset

3.1 Person-Centric Image Filtering

In this study, we utilize the COYO700M dataset (Byeon et al., 2022), a large-scale dataset that contains 747M image-text pairs collected from CommonCrawl, as our web-crawled images source. To filter high-quality person-centric images, we initially deploy the YOLOv11 model (Jocher and Qiu, 2024) to detect humans and extract bounding box coordinates. The specific workflow is illustrated in Fig. 2, where images are retained based on the following criteria: (i) shorter dimension exceeds 90 pixels, (ii) aspect ratio between 1:2 and 1:4, and (iii) human detection confidence above 0.85. Subsequently, YOLOv11-Pose (Jocher and Qiu, 2024) verifies pose integrity by requiring: (i) visibility of at least eight keypoints, (ii) presence of at least one hip keypoints and two head keypoints. This process yields 5 million high-quality human-centric images filtered from the COYO700M dataset.

3.2 Synthetic Caption Generation

Following the selection of 5 million high-quality human-centric images, we develop a synthetic caption generation pipeline to create diverse and precise textual descriptions. Our approach trans-

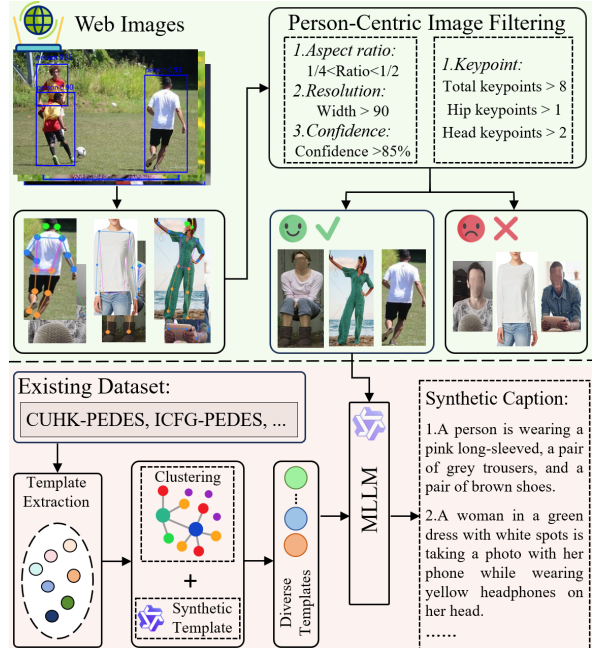


Figure 2: The details of person-centric image filtering and synthetic caption generation pipeline for constructing our WebPerson dataset.

forms existing captions from CUHK-PEDES (Li et al., 2017), ICFG-PEDES (Ding et al., 2021), and RSTPReid (Zhu et al., 2021) into structured templates using Qwen2.5-72B-Instruct (Yang et al., 2024a). The model systematically replaces fine-grained attributes (e.g., black jacket, ponytail) with standardized placeholders (e.g., [colored top], [hairstyle]).

To reduce redundancy and cluster semantically similar templates, inspired by the previous works (Yang et al., 2025; Gu et al., 2025b), we employ the OPENCLIP ViT-bigG/14 (Radford et al., 2021a) to extract text embeddings of the template texts, then we utilize the standard k -means algorithm to partition all the templates embeddings into k distinct clusters based on the nearest neighbor criterion. Within each cluster, we select the most representative template (highest cosine similarity to the centroid) along with five randomly sampled templates. To further enhance template diversity, we employ Qwen2.5-72B-Instruct (Yang et al., 2024a) to synthesize new templates from this refined set. All generated templates undergo rigorous review to eliminate biases and stereotypes, yielding a curated collection of one thousand high-quality templates. To generate diverse, high-quality captions, we leverage the in-context learning capabilities of MLLMs (Li et al., 2025; Gu et al., 2025a). Specifically, we randomly assign a template to each image and use Qwen2.5-VL-7B-Instruct and

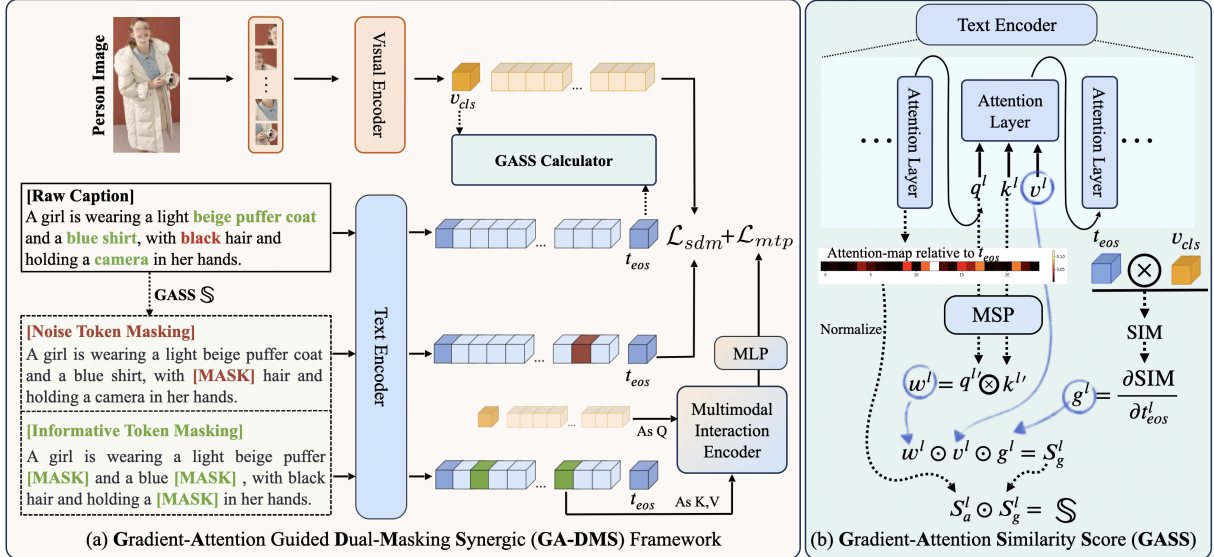


Figure 3: Overview of our proposed method. (a) The architecture of our proposed Gradient-Attention Guided Dual-Masking Synergetic (GA-DMS) framework. (b) The details of the Gradient-Attention Similarity Score (GASS).

Qwen2.5-VL-32B-Instruct (Bai et al., 2025) to produce captions that follow the given template. We adopt vLLM (Kwon et al., 2023) to accelerate large-scale inference. The details of the prompt are provided in the Appendix A.2.

4 GA-DMS Framework

This section presents our **GA-DMS** (Gradient-Attention Guided Dual-Masking Synergetic) framework (Fig. 3). In Sec. 4.1, we introduce the Gradient-Attention Similarity Score, which dynamically differentiates noise tokens and informative tokens during the training process. In Sec. 4.2, we present the dual-masking synergetic learning details and the training objective.

4.1 Gradient-Attention Similarity Score

Existing interpretability research (Selvaraju et al., 2017) on CLIP-based models has shown that intermediate layer gradients retain fine-grained image-text alignment information. Motivated by prior work (Zhao et al., 2025), we introduce a gradient-attention similarity score \mathbb{S} that quantifies each textual token’s contribution to the image-text alignment. We denote the embeddings of the text tokens and image tokens as T and V . We first calculate the global cosine similarity $\text{SIM} = T_{eos} V_{cls}^T$. The gradient importance for the l -th transformer layer’s text token T_{eos}^l is then derived as $g^l = \frac{\partial \text{SIM}}{\partial T_{eos}^l}$.

To capture fine-grained semantics, we integrate a Multi-Scale Pooling (MSP) layer within the transformer architecture. The MSP layer aggregates local contexts at multiple scales $c \in \mathcal{C}$ through

average pooling of adjacent c tokens, followed by bilinear interpolation to restore original dimensions. This process yields features enriched with multi-scale local information. The spatial importance w^l for each transformer layer l is then computed as:

$$w^l = \Phi(\text{MSP}(q_{eos}^l) \text{MSP}(k^l)^T) \quad (1)$$

where Φ is the normalization function, q_{eos}^l is the query embedding for the [eos] token at layer l , k^l represent the key embedding at layer l . The gradient-based score S_g^l of the l -th transformer layer is defined as:

$$S_g^l = g^l * w^l * v^l \quad (2)$$

where v^l is the value embedding at layer l .

Simultaneously, we compute attention-based semantic scores S_a^l for each token based on the attention maps \mathcal{M}^l from the l -th transformer layer. We denote the attention score for the [eos] token as \mathcal{W}^l , the attention-based semantic score S_a^l is computed as:

$$S_a^l = \frac{\mathcal{W}^l}{\sum_{j=1}^N \mathcal{W}_j^l} \quad (3)$$

The final gradient-attention similarity score \mathbb{S} for the effective textual tokens is defined as:

$$\mathbb{S} = \text{ReLU}\left(\frac{1}{L} \sum_{l \in L} S_g^l * S_a^l\right) \quad (4)$$

where L represents the number the final L layers of the transformer, $\mathbb{S} \in \mathbb{R}^{B \times N}$, and N is the number of tokens. This score integrates information from both gradients and attention maps to weight text tokens for masking probability computation.

4.2 Dual-Masking Synergetic Learning

4.2.1 Noise Token Masking

While Multimodal Large Language Models (MLLMs) inevitably introduce noise during large-scale data generation due to inherent hallucination effects. To mitigate this issue, we employ a noise token masking strategy to reduce the influence of noise tokens based on the gradient-attention similarity score \mathbb{S} . We calculate the masking probability for the i -th text token T_i as:

$$p(T_i) = \frac{\alpha_n}{1 + e^{-\lambda[(1-s_i)-\gamma]}} \quad (5)$$

where $s_i \in \mathbb{S}$ is the gradient-attention similarity score for the i -th token, α_n is a hyperparameter to set the upper limit of the masking probability for noise tokens. λ and γ respectively modulate the slope and midpoint of the probability distribution, thereby sharpening the differentiation between noisy and semantically relevant tokens. During training, we dynamically mask textual tokens using [mask] according to these computed probabilities.

Given the embeddings of image-text pairs $\{(v_{\text{cls}}^i, t_{\text{eos}}^i)\}_{i=1}^B$, we define the ground-truth matching distribution as $q_{i,j}$ and compute the predicted distribution as:

$$p_{i,j} = \frac{\exp(\text{sim}(v_{\text{cls}}^i, t_{\text{eos}}^j)/\tau)}{\sum_{b=1}^B \exp(\text{sim}(v_{\text{cls}}^i, t_{\text{eos}}^b)/\tau)} \quad (6)$$

where τ is a temperature parameter. Following (Jiang and Ye, 2023), we adopt the Similarity Distribution Matching (SDM) loss to align the distribution. The \mathcal{L}_{i2t} is defined as:

$$\mathcal{L}_{\text{i2t}} = \frac{1}{B} \sum_{i=1}^B \sum_{j=1}^B p_{i,j} \log \left(\frac{p_{i,j}}{q_{i,j} + \varepsilon} \right) \quad (7)$$

where ε is a small number to avoid numerical problems. We compute a symmetric loss \mathcal{L}_{t2i} by swapping $\{(v_{\text{cls}}^i, t_{\text{eos}}^i)\}$, and the SDM loss is:

$$\mathcal{L}_{\text{sdm}} = \mathcal{L}_{\text{i2t}} + \mathcal{L}_{\text{t2i}} \quad (8)$$

4.2.2 Masked Informative Token Prediction

To improve fine-grained semantic representation, we selectively mask tokens with strong image-semantic correlations and introduce a masked token prediction task to enhance local semantic learning. Similar to the Equation 5, the masking probability for the informative text tokens is defined as:

$$p(T_i) = \frac{\alpha_i}{1 + e^{-\lambda[s_i-\gamma]}} \quad (9)$$

where α_i bounds the maximum masking probability for informative tokens. For effective fine-grained visual-textual fusion during token prediction, we integrate a cross-modal interaction module (Jiang and Ye, 2023) as a decoder. This module consists of a multi-head cross-attention followed by four Transformer layers to align modalities in a shared embedding space. A final MLP layer predicts original tokens from the fused representations. Given hidden states h_i^m , $i \in \mathcal{M}$ and \mathcal{M} denotes the masked text token set, the distribution of the output token is $\mathbf{m}_i = \text{MLP}(h_i^m)$. The Masked Token Prediction (MTP) loss \mathcal{L}_{mtp} is defined as:

$$\mathcal{L}_{\text{mtp}} = -\frac{1}{|\mathcal{M}||\mathcal{V}|} \sum_{i \in \mathcal{M}} \sum_{j \in |\mathcal{V}|} y_j^i \log \frac{\exp(\mathbf{m}_j^i)}{\sum_{k=1}^{|\mathcal{V}|} \exp(\mathbf{m}_k^i)}, \quad (10)$$

where $|\mathcal{V}|$ is the size of vocabulary \mathcal{V} , and y_j is a one-hot vocabulary distribution. Finally, the total loss \mathcal{L} is define as:

$$\mathcal{L} = \mathcal{L}_{\text{sdm}} + \beta \mathcal{L}_{\text{mtp}} \quad (11)$$

where β is a loss weight.

5 Experiments

Implementation Details. Consistent with previous works (Tan et al., 2024; Jiang and Ye, 2023), we utilize the CLIP ViT-B/16 model as our backbone. Following IRR (Jiang and Ye, 2023), we incorporate a randomly initialized multimodal interaction encoder to facilitate masked token prediction. Our implementation processes 384×128 resolution images with a maximum length of $N = 77$ text sequences. We employ Adam (Kingma, 2014) as the optimizer, initialized with a learning rate of $1e-4$ and a weight decay of $4e-5$. The parameters β_1 and β_2 are set to 0.9 and 0.999, respectively. The temperature parameter τ in SDM loss is set to 0.02. We train GA-DMS for 30 epochs with a batch size of 512 on 8 NVIDIA A100 (80G) GPUs. For generating synthetic templates and captions, we utilize Qwen2.5-72B-Instruct (Yang et al., 2024a), Qwen2.5-VL-7B-Instruct, and Qwen2.5-VL-32B-Instruct (Bai et al., 2025). Additionally, vLLM (Kwon et al., 2023) is leveraged to accelerate large-scale inference. Please refer to the Appendix A.1 for more detailed hyperparameters.

Evaluation. Following previous works (Tan et al., 2024; Qin et al., 2024), we conduct a comprehensive evaluation of our method across three challenging text-to-image person retrieval datasets: CUHK-PEDES (Li et al., 2017), ICFG-PEDES (Ding et al.,

Method	Image Enc.	Text Enc.	CUHK-PEDES				ICFG-PEDES				RSTPReid			
			R1	R5	R10	mAP	R1	R5	R10	mAP	R1	R5	R10	mAP
ViTAA (Wang et al., 2020)	RN50	LSTM	55.97	75.84	83.52	-	50.98	68.79	75.78	-	-	-	-	-
SSAN (Ding et al., 2021)	RN50	LSTM	61.37	80.15	86.73	-	54.23	72.63	79.53	-	43.50	67.80	77.15	-
LBUL (Wang et al., 2022b)	RN50	BERT	64.04	82.66	87.22	-	-	-	-	-	45.55	68.2	77.85	-
SAF (Li et al., 2022b)	ViT-Base	BERT	64.13	82.62	88.4	-	-	-	-	-	-	-	-	-
TIPCB (Chen et al., 2022)	RN50	BERT	64.26	83.19	89.1	-	54.96	74.72	81.89	-	-	-	-	-
CAIBC (Wang et al., 2022a)	RN50	BERT	64.43	82.87	88.37	-	-	-	-	-	47.35	69.55	79.00	-
AXM-Net (Farooq et al., 2022)	RN50	BERT	64.44	80.52	86.77	58.70	-	-	-	-	-	-	-	-
LGUR (Shao et al., 2022)	DeiT-Small	BERT	65.25	83.12	89.00	-	59.02	75.32	81.56	-	47.95	71.85	80.25	-
IVT (Shu et al., 2022)	ViT-Base	BERT	65.69	85.93	91.15	-	56.04	73.60	80.22	-	46.70	70.00	78.80	-
LCR2S (Yan et al., 2023a)	RN50	TextCNN+BERT	67.36	84.19	89.62	59.20	57.93	76.08	82.40	38.21	54.95	76.65	84.70	40.92
UniPT (Shao et al., 2023)	ViT-Base	BERT	68.50	84.67	90.38	-	60.09	76.19	82.46	-	51.85	74.85	82.85	-

with ALBEF (Li et al., 2021) backbone:

RaSa (Bai et al., 2023)	CLIP-ViT	BERT-base	76.51	90.29	94.25	69.38	65.28	80.40	85.12	41.29	66.90	86.50	91.35	52.31
APTM (Yang et al., 2023b)	Swin-B	BERT-base	76.53	90.04	94.15	66.91	68.51	82.99	87.56	41.22	67.50	85.70	91.45	52.56

with CLIP (Radford et al., 2021b) backbone:

Han et al. (Han et al., 2021)	CLIP-RN101	CLIP-Xformer	64.08	81.73	88.19	60.08	-	-	-	-	-	-	-	-
IRRA (Jiang and Ye, 2023)	CLIP-ViT	CLIP-Xformer	73.38	89.93	93.71	66.10	63.46	80.25	85.82	38.06	60.20	81.30	88.20	47.17
FSRL (Wang et al., 2024)	CLIP-ViT	CLIP-Xformer	74.65	89.77	94.03	67.49	64.01	80.42	85.56	39.64	60.20	81.40	88.60	47.38
Propot (Yan et al., 2024)	CLIP-ViT	CLIP-Xformer	74.89	89.90	94.17	67.12	65.12	81.57	86.97	42.93	61.87	83.63	89.70	47.82
SAP-SAM (Wang et al., 2024)	CLIP-ViT	CLIP-Xformer	75.05	89.93	93.73	-	63.97	80.84	86.17	-	62.85	82.65	89.85	-
PLOT (Park et al., 2024)	CLIP-ViT	CLIP-Xformer	75.28	90.42	94.12	-	65.76	81.39	86.73	-	61.80	82.85	89.45	-
RDE (Qin et al., 2024)	CLIP-ViT	CLIP-Xformer	75.94	90.14	94.12	67.56	67.68	82.47	87.36	40.06	65.35	83.95	89.90	50.88
NAM (Tan et al., 2024)	CLIP-ViT	CLIP-Xformer	76.82	91.16	94.46	69.55	67.05	82.16	87.33	41.51	68.50	87.15	92.10	53.02
Ours (1.0 M)	CLIP-ViT	CLIP-Xformer	<u>77.02</u>	<u>91.28</u>	<u>94.58</u>	<u>69.65</u>	<u>69.07</u>	<u>83.26</u>	<u>87.64</u>	41.91	<u>70.30</u>	<u>88.00</u>	<u>92.85</u>	<u>54.89</u>
Ours (5.0 M)	CLIP-ViT	CLIP-Xformer	77.60	91.40	94.78	69.82	69.51	83.47	87.67	<u>42.30</u>	71.25	87.25	92.90	55.43

Table 1: Comparisons with state-of-the-art methods in the traditional evaluation settings. The best results are marked in **bold**, and the second-best results are underlined.

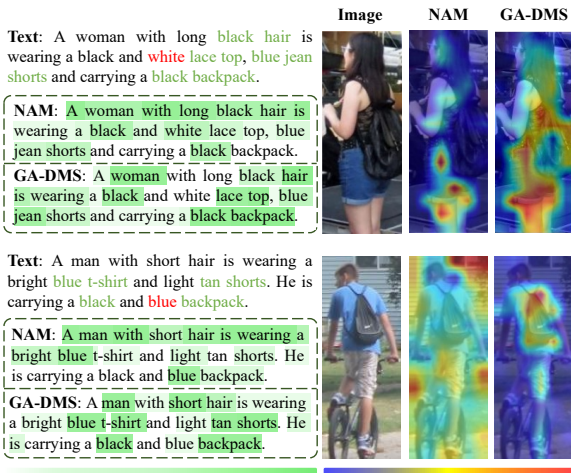


Figure 4: Visualization of token-wise weight scores and attention maps generated by NAM (Tan et al., 2024) and our GA-DMS.

2021), and RSTPReid (Zhu et al., 2021). We employ Rank- k ($k=1, 5, 10$) and mean Average Precision (mAP) as evaluation metrics for all datasets.

5.1 Comparison with Existing Methods.

We evaluate GA-DMS against state-of-the-art methods on three benchmarks: CUHK-PEDES, ICFG-PEDES, and RSTPReid. As shown in Tab. 1, our pre-trained model achieves superior performance after fine-tuning with the IRRA (Jiang and Ye, 2023), which significantly improves Rank-1 accu-

Pre-training Dataset	CUHK-PEDES		ICFG-PEDES		RSTPReid	
	R1	mAP	R1	mAP	R1	mAP
None	12.65	11.15	6.67	2.51	13.45	10.31
MALS (1.5M)	20.47	18.46	11.71	4.57	21.50	16.95
LUPerson-T (0.95M)	21.55	18.76	11.20	4.53	22.15	17.29
SYNTH-PEDES (1.0M)	57.29	51.86	57.13	31.36	42.20	32.28
LUPerson-MLLM (1.0M)	56.01	50.34	37.00	20.21	50.60	37.08
Ours (0.1 M)	58.95	52.77	38.18	19.70	47.10	36.68
Web-Person (1.0 M)	<u>66.26</u>	<u>58.54</u>	51.99	28.81	<u>55.35</u>	<u>40.57</u>
Ours (5.0 M)	68.34	60.22	<u>54.64</u>	<u>30.68</u>	57.60	42.00

Table 2: Comparisons with existing pre-training datasets in the direct transfer setting. The best results are marked in **bold**, and the second-best results are underlined.

racy and mAP on the RSTPReid dataset by 10.10% and 7.72% over the baseline of IRRA. Compared with the NAM (Tan et al., 2024), GA-DMS obtains 0.2%, 2.02%, and 1.8% improvement in Rank-1 on the CUHK-PEDES, ICFG-PEDES, and RSTPReid datasets, respectively. The primary reason for this improvement is that our proposed GA-DMS framework effectively distinguishes between noise and informative tokens based on the gradient-attention similarity score. As shown in Fig. 4, compared with NAM (Tan et al., 2024), our GA-DMS can better allocate weights to text tokens while concentrating attention on human-centric regions. This capability not only reduces the effect of noise on model training but also improves the model’s capacity to learn fine-grained semantic information. Moreover,

Pre-training Dataset	Source	Target							
		CUHK-PEDES		ICFG-PEDES		RSTPReid			
		RI	mAP	RI	mAP	RI	mAP		
None	CUHK-PEDES	73.48	66.21	43.04	22.45	52.55	39.97		
	ICFG-PEDES	33.90	31.65	63.83	38.37	47.45	36.83		
	RSTPReid	35.25	32.35	33.58	19.58	60.40	47.70		
MALS(1.5 M)	CUHK-PEDES	73.67	65.23	46.02	24.06	55.05	41.29		
	ICFG-PEDES	43.11	38.93	65.21	38.52	48.45	37.29		
	RSTPReid	44.51	39.99	40.78	25.42	64.05	50.08		
LUPerson-T(0.95 M)	CUHK-PEDES	74.28	66.52	44.83	22.72	54.25	39.26		
	ICFG-PEDES	34.66	32.51	65.33	38.45	48.30	38.51		
	RSTPReid	39.26	34.26	34.95	22.25	61.50	48.28		
SYNTH-PEDES(1.0 M)	CUHK-PEDES	74.12	65.82	57.14	32.12	55.85	40.85		
	ICFG-PEDES	60.49	54.61	66.63	39.32	49.80	37.34		
	RSTPReid	57.75	53.01	53.88	30.88	66.75	52.18		
LUPerson-MLLM(1.0 M)	CUHK-PEDES	76.59	68.06	47.17	25.41	59.35	43.76		
	ICFG-PEDES	60.75	54.42	67.18	40.27	55.65	44.05		
	RSTPReid	60.04	53.85	46.39	27.91	69.45	53.30		
Ours(0.1 M)	CUHK-PEDES	75.53	67.92	47.79	25.14	56.75	41.01		
	ICFG-PEDES	58.67	52.66	66.35	39.95	52.93	39.84		
	RSTPReid	58.49	52.50	44.41	25.98	65.90	49.28		
Ours(1.0 M)	CUHK-PEDES	<u>77.02</u>	<u>69.65</u>	<u>57.24</u>	<u>32.13</u>	61.10	45.27		
	ICFG-PEDES	<u>68.16</u>	<u>60.79</u>	<u>69.07</u>	<u>41.91</u>	59.15	44.94		
	RSTPReid	<u>68.41</u>	<u>61.28</u>	<u>56.13</u>	<u>34.64</u>	<u>70.30</u>	<u>54.89</u>		
Ours(5.0 M)	CUHK-PEDES	77.60	69.82	58.91	33.70	61.80	46.81		
	ICFG-PEDES	69.83	62.06	69.52	42.30	60.05	45.46		
	RSTPReid	69.19	62.00	57.13	35.76	71.25	55.43		

Table 3: Comparisons with existing pre-training datasets in the fine-tuning setting. The best results are marked in **bold**, and the second-best results are underlined. Gray indicates that the source and target are homologous.

upon scaling the WebPerson dataset from 1.0 M to 5.0 M, GA-DMS achieves new state-of-the-art Rank-1 accuracies of 77.6%, 69.51%, and 71.25% across three downstream datasets.

5.2 Comparison with Existing Datasets.

We conduct comprehensive comparisons between our WebPerson dataset and four existing large-scale pre-training datasets: MALS (Yang et al., 2023b), LUPerson-T (Shao et al., 2023), SYNTH-PEDES (Zuo et al., 2024), and LUPerson-MLLM (Tan et al., 2024). MALS consists of 1.5 million synthetic images generated using commercial diffusion models, with textual descriptions automatically produced by BLIP (Li et al., 2022a). LUPerson-T includes 0.95 million images, each enhanced by one of 456 templates to maximize caption diversity. SYNTH-PEDES provides 4.8 million images, each annotated with an average of 2.53 textual descriptions, generated through a hybrid architecture that combines a ResNet101-FPN (He et al., 2016) visual encoder with a GPT-2 (Radford et al., 2019) text generator for detailed person attribute modeling. Notably, LUPerson-MLLM utilizes two multimodal large language models for caption generation, supplemented by 46 ChatGPT-optimized templates obtained through iterative dialogues to enhance linguistic variation. This dataset comprises 1.0 million images, each paired with two

Masking Method		Components		CUHK-PEDES		ICFG-PEDES		RSTPReid	
CSS	GASS	SDM	MTP	RI	mAP	RI	mAP	RI	mAP
x	x	x	x	56.75	50.42	34.63	17.59	45.50	34.51
✓	x	x	✓	56.35	50.21	34.72	17.66	44.60	33.28
✓	x	✓	x	63.29	57.42	43.39	24.12	51.95	39.41
✓	x	✓	✓	62.74	57.01	42.96	23.88	50.80	38.91
x	✓	x	✓	57.29	52.28	36.24	18.96	47.90	35.97
x	✓	✓	x	<u>63.87</u>	<u>57.56</u>	<u>44.02</u>	<u>24.18</u>	<u>52.30</u>	<u>39.61</u>
x	✓	✓	✓	64.25	58.27	44.39	24.67	52.70	40.12

Table 4: Ablation on different components and masking methods. CSS: Cosine Similarity Score. GASS: Gradient-Attention Similarity Score. SDM: Similarity Distribution Matching. MTP: Masked Token Prediction.


MLLM-generated captions.

Tab. 2 presents comparative results under a direct transfer setting, where models pre-trained on WebPerson exhibit superior cross-dataset generalization across three benchmarks. Specifically, under the comparable 1M dataset, our constructed WebPerson dataset demonstrates superior performance on CUHK-PEDES and RSTPReid, and shows suboptimal performance on ICFG-PEDES. Notably, the WebPerson dataset demonstrates comparable performance to the full-scale LUPerson-MLLM even when trained on a mere 0.1M samples. These experimental results demonstrate that our proposed WebPerson dataset exhibits strong robustness and can learn person representations with enhanced transferability.


As shown in Tab. 3, we also evaluate the fine-tuning performance following LUPerson-MLLM (Tan et al., 2024), utilizing the IRRA with models pretrained on different datasets. Results indicate that WebPerson pretraining yields state-of-the-art performance across both in-domain and cross-domain scenarios. At the 1M data scale, WebPerson achieves consistent improvements over LUPerson-MLLM, with Rank-1 accuracy gains of 0.43%, 1.89%, and 0.85% on CUHK-PEDES, ICFG-PEDES, and RSTPReid respectively. The cross-domain evaluations reveal particularly significant performance enhancements, highlighting WebPerson’s exceptional representation transferability through fine-tuning.

5.3 Ablation Study

Ablation on Different Components and Masking Methods. To substantiate the efficacy of various components and the effectiveness of our proposed Gradient-Attention Similarity Score (GASS), we perform a comprehensive ablation study with a 0.5M data sample from our WebPerson dataset. As shown in Tab. 4, the integrating Masked Token Pre-



(a) Ablation on α_n for masking noise tokens



(b) Ablation on α_i for masking informative tokens

Figure 5: Ablation experiment results for α_n and α_i , which can directly influence the upper limit of the masking probability for noise and informative tokens.

diction (MTP) with GASS improves performance across all evaluation metrics, as predicting semantically rich tokens enhances fine-grained learning. The Similarity Distribution Matching (SDM) component alone enhances image-text alignment by replacing noisy tokens with learnable embeddings, achieving Rank-1 accuracy gains of 7.12%, 9.39%, and 6.8% on CUHK-PEDES, ICFG-PEDES, and RSTPReid respectively. By combining MTP with SDM, we observe enhancements across all metrics, further substantiating the efficacy of the components within our method.

When comparing Cosine Similarity Score (CSS) with Gradient-Attention Similarity Score (GASS), GASS consistently exhibits superior performance. This advantage primarily stems from GASS’s capacity to precisely weight textual tokens during training by incorporating gradient and attention information. As illustrated in Fig. 4, our method accurately allocates weights to noise textual tokens (*e.g.*, "white lace top"), thereby effectively mitigating the influence of noise on the model’s representation learning.

Ablation on α_n and α_i . In this work, our dual-masking synergistic learning method dynamically masks textual tokens according to gradient-attention similarity scores. We introduce parameters α_n and α_i to regulate the maximum masking probabilities for noise and informative tokens. Fig. 5 presents an ablation study on α_n and α_i to determine the optimal settings. For enhanced performance on three downstream datasets, we set $\alpha_n = 0.2$ and $\alpha_i = 0.3$. Additionally, our method consistently outperforms random masking baselines, confirming its effectiveness.

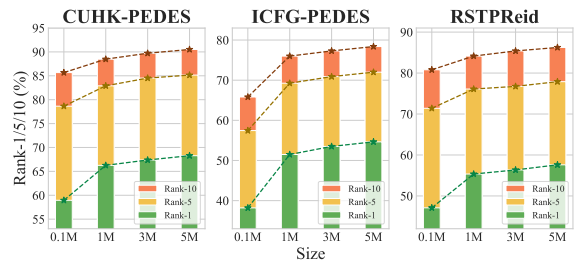


Figure 6: Data scaling analysis of WebPerson dataset. The performance of our GA-DMS method in direct transfer settings.

Data Scaling Analysis. To explore the impact of pretraining data scale on person representation learning, we systematically augmented the dataset size from 0.1M to 1M, 3M, and 5M samples for pertaining. Fig. 6 illustrates the direct transfer performance evaluation across three benchmarks at different data scales. The outcomes consistently reveal performance enhancements as the data volume increases. At the maximum scale of 5.0M samples, the model demonstrates Rank-1 accuracy improvements of 9.39%, 16.46%, and 10.50% across the three benchmarks in comparison to the 0.1M baseline, indicating a clear upward trajectory. These findings conclusively demonstrate that scaling high-quality pretraining data substantially enhances text-based person retrieval capability.

Conclusion

In this paper, we enhance CLIP for person representation learning by synergistically improving data acquisition and model architecture. First, we devise a noise-resistant data construction pipeline that leverages the in-context learning capabilities of MLLMs for automatic filtering and captioning of web-crawled images. This results in the WebPerson dataset, which comprises 5M high-quality person-centric image-text pairs. Second, we propose the GA-DMS framework, which improves cross-modal alignment by masking semantically irrelevant tokens based on a gradient-attention similarity score. Concurrently, we implement masked token prediction objectives that force the model to reconstruct informative text tokens, facilitating discriminative fine-grained feature learning for visual-semantic correspondence. Comprehensive experiments demonstrate that GA-DMS achieves state-of-the-art performance in several downstream datasets. We hope our work provides insights for the person representation learning task.

Limitations

In this work, we demonstrate the exceptional text-based person retrieval performance of the person-centric dataset constructed solely from internet images. Limited by computational resources, this paper constructs a 5M-scale WebPerson dataset, with further scaling left for community exploration.

Acknowledgement

This work was supported in part by the National Natural Science Foundation of China under Grant 62373086, by the Liaoning Province Applied Basic Research Program (2025JH2/101330131), by the Guangdong Basic and Applied Basic Research under 2023A1515140014, and by the State Key Laboratory of Robotics under Grant 2024-O14.

Ethics Statement

We abide by the ACL Code of Ethics. The data resources used in this study are publicly available.

References

- Rabab Abdelfattah, Qing Guo, Xiaoguang Li, Xiaofeng Wang, and Song Wang. 2023. Cdul: Clip-driven unsupervised learning for multi-label image classification. In *ICCV*, pages 1348–1357.
- Xiang An, Jiankang Deng, Kaicheng Yang, Jaiwei Li, Ziyong Feng, Jia Guo, Jing Yang, and Tongliang Liu. 2023. Unicom: Universal and compact representation learning for image retrieval. *arXiv preprint arXiv:2304.05884*.
- Xiang An, Kaicheng Yang, Xiangzi Dai, Ziyong Feng, and Jiankang Deng. 2024. Multi-label cluster discrimination for visual representation learning. In *ECCV*, pages 428–444. Springer.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, and 1 others. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Yang Bai, Min Cao, Daming Gao, Ziqiang Cao, Chen Chen, Zhenfeng Fan, Liqiang Nie, and Min Zhang. 2023. Rasa: Relation and sensitivity aware representation learning for text-based person search. *IJCAI*.
- Minwoo Byeon, Beomhee Park, Haecheon Kim, Sungjun Lee, Woonhyuk Baek, and Saehoon Kim. 2022. Coyo-700m: Image-text pair dataset. <https://github.com/kakaobrain/coyo-dataset>.
- Fei-Long Chen, Du-Zhen Zhang, Ming-Lun Han, Xiu-Yi Chen, Jing Shi, Shuang Xu, and Bo Xu. 2023. Vlp: A survey on vision-language pre-training. *Machine Intelligence Research*, 20(1):38–56.
- Yuhao Chen, Guoqing Zhang, Yujiang Lu, Zhenxing Wang, and Yuhui Zheng. 2022. Tipcb: A simple but effective part-based convolutional baseline for text-based person search. *Neurocomputing*.
- Zefeng Ding, Changxing Ding, Zhiyin Shao, and Dacheng Tao. 2021. Semantically self-aligned network for text-to-image part-aware person re-identification. *arXiv preprint arXiv:2107.12666*.
- Ammarah Farooq, Muhammad Awais, Josef Kittler, and Syed Safwan Khalid. 2022. Axm-net: Implicit cross-modal feature alignment for person re-identification. In *AAAI*.
- Dengpan Fu, Dongdong Chen, Jianmin Bao, Hao Yang, Lu Yuan, Lei Zhang, Houqiang Li, and Dong Chen. 2021. Unsupervised pre-training for person re-identification. In *CVPR*, pages 14750–14759.
- Dengpan Fu, Dongdong Chen, Hao Yang, Jianmin Bao, Lu Yuan, Lei Zhang, Houqiang Li, Fang Wen, and Dong Chen. 2022. Large-scale pre-training for person re-identification with noisy labels. In *CVPR*, pages 2476–2486.
- Tiancheng Gu, Kaicheng Yang, Xiang An, Ziyong Feng, Dongnan Liu, Weidong Cai, and Jiankang Deng. 2024. Rwkv-clip: A robust vision-language representation learner. *arXiv preprint arXiv:2406.06973*.
- Tiancheng Gu, Kaicheng Yang, Ziyong Feng, Xingjun Wang, Yanzhao Zhang, Dingkun Long, Yingda Chen, Weidong Cai, and Jiankang Deng. 2025a. Breaking the modality barrier: Universal embedding learning with multimodal llms. *arXiv preprint arXiv:2504.17432*.
- Tiancheng Gu, Kaicheng Yang, Chaoyi Zhang, Yin Xie, Xiang An, Ziyong Feng, Dongnan Liu, Weidong Cai, and Jiankang Deng. 2025b. Realsyn: An effective and scalable multimodal interleaved document transformation paradigm. *arXiv preprint arXiv:2502.12513*.
- Qianru Han, Xinwei He, Zhi Liu, Sannyuya Liu, Ying Zhang, and Jinhai Xiang. 2024. Clip-scgi: Synthesized caption-guided inversion for person re-identification. *arXiv preprint arXiv:2410.09382*.
- Xiao Han, Sen He, Li Zhang, and Tao Xiang. 2021. Text-based person search with limited data. *arXiv preprint arXiv:2110.10807*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *CVPR*, pages 770–778.
- Xiaoxing Hu, Kaicheng Yang, Jun Wang, Haoran Xu, Ziyong Feng, and Yupei Wang. 2025. Decoupled global-local alignment for improving compositional understanding. *arXiv preprint arXiv:2504.16801*.
- Ding Jiang and Mang Ye. 2023. Cross-modal implicit relation reasoning and aligning for text-to-image person retrieval. In *CVPR*, pages 2787–2797.

- Glenn Jocher and Jing Qiu. 2024. [Ultralytics yolo11](#).
- Diederik P Kingma. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *ACM SIGOPS*.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022a. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, pages 12888–12900. PMLR.
- Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. 2021. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705.
- Shiping Li, Min Cao, and Min Zhang. 2022b. Learning semantic-aligned feature representation for text-based person search. In *ICASSP*.
- Shuang Li, Tong Xiao, Hongsheng Li, Bolei Zhou, Dayu Yue, and Xiaogang Wang. 2017. Person search with natural language description. In *CVPR*, pages 1970–1979.
- Siyuan Li, Li Sun, and Qingli Li. 2023. Clip-reid: exploiting vision-language model for image re-identification without concrete text labels. In *AAAI*, volume 37, pages 1405–1413.
- Yanshu Li, Tian Yun, Jianjiang Yang, Pinyuan Feng, Jinfa Huang, and Ruixiang Tang. 2025. Taco: Enhancing multimodal in-context learning via task mapping-guided sequence configuration. *arXiv preprint arXiv:2505.17098*.
- Feng Liu, Minchul Kim, Zhiyuan Ren, and Xiaoming Liu. 2024. Distilling clip with dual guidance for learning discriminative human body shape representation. In *CVPR*, pages 256–266.
- Jicheol Park, Dongwon Kim, Boseung Jeong, and Suha Kwak. 2024. Plot: Text-based person search with part slot attention for corresponding part discovery. In *ECCV*. Springer.
- Fang Peng, Xiaoshan Yang, Linhui Xiao, Yaowei Wang, and Changsheng Xu. 2023. Sgva-clip: Semantic-guided visual adapting of vision-language models for few-shot image classification. *TMM*, 26:3469–3480.
- Yang Qin, Yingke Chen, Dezhong Peng, Xi Peng, Joey Tianyi Zhou, and Peng Hu. 2024. Noisy-correspondence learning for text-to-image person re-identification. In *CVPR*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, A. Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021a. Learning transferable visual models from natural language supervision. In *ICML*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, and 1 others. 2021b. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PmLR.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, and 1 others. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Aneeshan Sain, Ayan Kumar Bhunia, Pinaki Nath Chowdhury, Subhadeep Koley, Tao Xiang, and Yi-Zhe Song. 2023. Clip for all things zero-shot sketch-based image retrieval, fine-grained or not. In *CVPR*, pages 2765–2775.
- Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*, pages 618–626.
- Zhiyin Shao, Xinyu Zhang, Changxing Ding, Jian Wang, and Jingdong Wang. 2023. Unified pre-training with pseudo texts for text-to-image person re-identification. In *ICCV*, pages 11174–11184.
- Zhiyin Shao, Xinyu Zhang, Meng Fang, Zhifeng Lin, Jian Wang, and Changxing Ding. 2022. Learning granularity-unified representations for text-to-image person re-identification. In *ACMMM*.
- Xiujun Shu, Wei Wen, Haoqian Wu, Keyu Chen, Yiran Song, Ruizhi Qiao, Bo Ren, and Xiao Wang. 2022. See finer, see more: Implicit modality alignment for text-based person retrieval. In *ECCV*, pages 624–641. Springer.
- Jianlou Si, Honggang Zhang, Chun-Guang Li, Jason Kuen, Xiangfei Kong, Alex C Kot, and Gang Wang. 2018. Dual attention matching network for context-aware feature sequence based person re-identification. In *CVPR*, pages 5363–5372.
- Guanglu Song, Biao Leng, Yu Liu, Congrui Hetang, and Shaofan Cai. 2018. Region-based quality estimation network for large-scale person re-identification. In *AAAI*, volume 32.
- Wentan Tan, Changxing Ding, Jiayu Jiang, Fei Wang, Yibing Zhan, and Dapeng Tao. 2024. Harnessing the power of mllms for transferable text-to-image person reid. In *CVPR*, pages 17127–17137.
- Di Wang, Feng Yan, Yifeng Wang, Lin Zhao, Xiao Liang, Haodi Zhong, and Ronghua Zhang. 2024. Fine-grained semantics-aware representation learning for text-based person retrieval. In *ICMR*.

- Feng Wang and Huaping Liu. 2021. Understanding the behaviour of contrastive loss. In *CVPR*, pages 2495–2504.
- Zhe Wang, Zhiyuan Fang, Jun Wang, and Yezhou Yang. 2020. Vitaa: Visual-textual attributes alignment in person search by natural language. In *Computer vision—ECCV 2020: 16th European conference, glasgow, UK, August 23–28, 2020, proceedings, part XII 16*, pages 402–420. Springer.
- Zijie Wang, Aichun Zhu, Jingyi Xue, Xili Wan, Chao Liu, Tian Wang, and Yifeng Li. 2022a. Caibc: Capturing all-round information beyond color for text-based person retrieval. In *ACM MM*.
- Zijie Wang, Aichun Zhu, Jingyi Xue, Xili Wan, Chao Liu, Tian Wang, and Yifeng Li. 2022b. Look before you leap: Improving text-based person retrieval by learning a consistent cross-modal common manifold. In *ACM MM*, pages 1984–1992.
- Longhui Wei, Shiliang Zhang, Wen Gao, and Qi Tian. 2018. Person transfer gan to bridge domain gap for person re-identification. In *CVPR*, pages 79–88.
- Yu Wu, Yana Wei, Haozhe Wang, Yongfei Liu, Sibe Yang, and Xuming He. 2023. Grounded image text matching with mismatched relation reasoning. In *ICCV*, pages 2976–2987.
- Linhui Xiao, Xiaoshan Yang, Fang Peng, Ming Yan, Yaowei Wang, and Changsheng Xu. 2023. Clipvg: Self-paced curriculum adapting of clip for visual grounding. *TMM*, 26:4334–4347.
- Yin Xie, Kaicheng Yang, Xiang An, Kun Wu, Yongle Zhao, Weimo Deng, Zimin Ran, Yumeng Wang, Ziyong Feng, Roy Miles, and 1 others. 2025. Region-based cluster discrimination for visual representation learning. *arXiv preprint arXiv:2507.20025*.
- Shuanglin Yan, Neng Dong, Jun Liu, Liyan Zhang, and Jinhui Tang. 2023a. Learning comprehensive representations with richer self for text-to-image person re-identification. In *ACMMM*.
- Shuanglin Yan, Neng Dong, Liyan Zhang, and Jinhui Tang. 2023b. Clip-driven fine-grained text-image person re-identification. *TIP*, 32:6032–6046.
- Shuanglin Yan, Jun Liu, Neng Dong, Liyan Zhang, and Jinhui Tang. 2024. Prototypical prompting for text-to-image person re-identification. *arXiv preprint arXiv:2409.09427*.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, and 1 others. 2024a. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.
- Fan Yang, Wei Li, Menglong Yang, Binbin Liang, and Jianwei Zhang. 2024b. Multi-modal disordered representation learning network for description-based person search. In *AAAI*.
- Kaicheng Yang, Jiankang Deng, Xiang An, Jiawei Li, Ziyong Feng, Jia Guo, Jing Yang, and Tongliang Liu. 2023a. Alip: Adaptive language-image pre-training with synthetic caption. In *ICCV*, pages 2922–2931.
- Kaicheng Yang, Tiancheng Gu, Xiang An, Haiqiang Jiang, Xiangzi Dai, Ziyong Feng, Weidong Cai, and Jiankang Deng. 2025. Clip-cid: Efficient clip distillation via cluster-instance discrimination. In *AAAI*, volume 39, pages 21974–21982.
- Shuyu Yang, Yinan Zhou, Zhedong Zheng, Yaxiong Wang, Li Zhu, and Yujiao Wu. 2023b. Towards unified text-based person retrieval: A large-scale multi-attribute and language search benchmark. In *ACMMM*, pages 4492–4501.
- Qiyang Yu, Quan Sun, Xiaosong Zhang, Yufeng Cui, Fan Zhang, Yue Cao, Xinlong Wang, and Jingjing Liu. 2024. Capsfusion: Rethinking image-text data at scale. In *CVPR*, pages 14022–14032.
- Chenyang Zhao, Kun Wang, Janet H. Hsiao, and Antoni B. Chan. 2025. [Grad-eclip: Gradient-based visual and textual explanations for clip](#). *Preprint*, arXiv:2502.18816.
- Henry Hengyuan Zhao, Pan Zhou, and Mike Zheng Shou. 2024. Genixer: Empowering multimodal large language model as a powerful data generator. In *ECCV*, pages 129–147. Springer.
- Zhedong Zheng, Liang Zheng, Michael Garrett, Yi Yang, Mingliang Xu, and Yi-Dong Shen. 2020. Dual-path convolutional image-text embeddings with instance loss. *TOMM*, 16(2):1–23.
- Aichun Zhu, Zijie Wang, Yifeng Li, Xili Wan, Jing Jin, Tian Wang, Fangqiang Hu, and Gang Hua. 2021. Dssl: Deep surroundings-person separation learning for text-based person retrieval. In *ACMMM*, pages 209–217.
- Jialong Zuo, Jiahao Hong, Feng Zhang, Changqian Yu, Hanyu Zhou, Changxin Gao, Nong Sang, and Jingdong Wang. 2024. Plip: Language-image pre-training for person representation learning. *NIPS*, 37:45666–45702.

A Appendix

A.1 Detail Experimental Settings

We present the settings used in the training GA-DMS in Tab. 5.

Hyperparameter	Value
Temperature	0.02
Loss weight β	0.4
Multiple scales \mathcal{C}	[1,2]
Adam β_1	0.9
Adam β_2	0.999
Adam ϵ	10^{-3}
Warm-up epochs	5
Weight decay	4×10^{-5}
Batch size	512
Learning rate	10^{-4}
Learning rate scheduler	CosineAnnealingLR
Training epochs	30
GPU	8xA100(80G)

Table 5: Hyperparameters used for GA-DMS pre-training.

A.2 Detail Instruction Prompt

The prompt used to input Qwen2.5-72B-Instruct (Yang et al., 2024a) for the generation of structured templates is as follows:

First, identify the words in the title that describe pedestrian attributes, such as tops, pants, footwear, head features, accessories, age, gender, actions, etc. Then replace these words with cross-identity generic terms like ‘colored top’, ‘colored bottom’, ‘hairstyle’ etc. Complete examples are as follows:

"A man wearing a orange jersey with yellow stripes, a pair of black shorts and a pair of green shoes." → "A [man] wearing a [color top] with [color pattern], a pair of [colored bottom] and a pair of [colored shoes]."

"This lady is wearing glasses, and she has her hair in a yellow ponytail. She is wearing a striped shirt and is carrying a bag over her right shoulder." → "This [person] is wearing an [accessory], and [he/she] has a [colored hairstyle]. [He/She] is wearing a patterned top and is carrying an object over [his/her] [body part]."

"A women is wearing a light colored sweater and black pants. She has long dark hair in a pony tail. " → "A [person] is wearing a [colored top] and [colored bottom]. [He/She] has long [colored hair] in a [hairstyle]."

Do not add any extra features not included in the original description. Output only the final description without any explanation.

The prompt used for inputting Qwen2.5-VL-Instruct (Bai et al., 2025) to generate pedestrian descriptions is as follows:

"Please generate a concise caption for the pedestrian image based on the following principles:

Core Subject Focus: Only describe the dominant pedestrian elements in the frame (e.g., gender, clothing, footwear, head features, accessories, actions), focusing on the color of each part."

Description restriction: 1. Use vague color terms (e.g., dark, light) only when the color is uncertain. 2. Use generic terms like "top" or "bottom" only when the clothing type is unclear, otherwise, use specific terms like "shirt" or "shorts."

Background Suppression Rule: Do not mention background information or abstract atmospheres (e.g., cozy).

Certainty Principle: Only output visually confirmed details — omit descriptions of unclear/low-resolution areas. Invisible elements do not need be described in the sentence (e.g., items are not visible). Avoid speculative terms ("possibly", "seems", "appears to be"), do not interpret potential relationships (e.g., inferring identity or emotions), and exclude artistic style critiques (e.g., "impressionist style").

Sentence Structure Reference: "<Structured Template>", First output the most significant pedestrian elements, the sentence length is less than <sequence length> English words. Use common words and phrasing from social media or daily life, ensuring correct grammar and logic. Provide only the caption sentence without any additional output."

A.3 The Influence of Layers.

We calculate the Gradient-Attention Similarity Score (GASS) between each text token and the image using the final L layers of the text encoder. This study examines how the number of layers involved in gradient-based similarity computation influences performance. As depicted in Fig. 7, the model consistently outperforms the baseline, which lacks gradient-based masking, across all tested layer depths. Notably, employing the last 8 layers of the text encoder achieves the highest overall performance, underscoring their effectiveness in optimizing masking outcomes.

A.4 Dataset analysis

Current text-based person retrieval datasets predominantly consist of manually annotated pedestrian images from re-identification benchmarks,

Datasets	Year	#Images	#Descriptions	Data Source	#Vocabulary Size	Label Method
CUHK-PEDES (Li et al., 2017)	2017	40,206	80,412	Market, Duke, etc.	12,517	Manual
LPW (Song et al., 2018)	2018	592,438	-	Surveillance Video	-	Manual+Detector+NN
MSMT-17 (Wei et al., 2018)	2018	126,441	-	Manual Collection	-	FasterRCNN
RSTPReid (Zhu et al., 2021)	2021	20,505	41,010	MSMT-17	6,331	Manual
ICFG-PEDES (Ding et al., 2021)	2021	54,522	54,522	MSMT-17	5,848	Manual
LUPerson (Fu et al., 2021)	2021	4,180,243	-	YouTube	-	YOLOv5
LUPerson-NL (Fu et al., 2022)	2022	10,683,716	-	YouTube	-	FairMOT
MALS (Yang et al., 2023b)	2023	1,510,330	1,510,330	Automatic Synthesis	4,772	ImaginAIry
LuPerson-T (Shao et al., 2023)	2023	957,606	1,277,991	LUPerson	459	CLIP
Luperson-MLLM (Tan et al., 2024)	2024	1,020,022	2,037,239	LUPerson	39,566	MLLM
SYNTH-PEDES (Zuo et al., 2024)	2024	4,791,771	12,138,157	LUPerson-NL& LPW	8,598	SPAC
WebPerson	2025	5,002,723	10,005,446	COYO-700M	96,623	MLLM

Table 6: Statistical comparison of different datasets. WebPerson stands as the largest automatically-generated text-described person dataset, offering inherent scalability without manual annotation requirements.




Figure 7: Results of different layers to compute \mathcal{S} . The encoders contain 12 layers in total.

fundamentally limited in scale and diversity by the substantial costs of human annotation. While generative methods have shown promise for dataset augmentation, they fail to achieve the necessary scale and fidelity for practical deployment. The emergence of Multimodal Large Language Models (MLLMs) and the availability of web-scale image resources now enable a new paradigm for automated dataset construction. Our WebPerson dataset leverages novel image filtering and text generation techniques to create a comprehensive pedestrian image library with accurate textual descriptions across diverse scenarios. Compared to existing datasets, WebPerson offers three key advantages:

High-quality WebPerson surpasses existing datasets containing single-style synthetic images or low-quality surveillance footage by providing superior texture details and diverse scene variations. Our rigorous image filtering pipeline ensures exceptional visual fidelity, while the MLLM-powered text generation framework produces highly accurate and detailed descriptions. Fig. 8 showcases representative examples demonstrating precise textual characterization of pedestrian attributes.

Diversity Sourced from web data, WebPerson exhibits rich variations in images, including but not limited to scene diversity, viewpoint changes, occlusions, clothing variations, and body poses. Our caption generation strategy further ensures corresponding textual descriptions maintain sufficient diversity. This dual-modality diversity enables



Figure 8: Visualization of some examples in our WebPerson dataset.

WebPerson to serve as an effective training corpus for developing robust models that generalize well to novel and unseen data across visual tasks, language tasks, and vision-language tasks.

Large-scale As illustrated in Tab. 6, we compare the attributes of WebPerson with other prominent person datasets. WebPerson emerges as the most extensive real-world dataset, featuring high-quality image-text pairs, encompassing 5 million images and 10 million textual descriptions. Moreover, our efficient data collection and caption generation strategies enable seamless scalability in data volume.

A.5 Broader Impact

This work introduces a novel pedestrian representation learning framework that achieves fine-grained cross-modal alignment through gradient-based token-wise similarity scoring while effectively suppressing noise interference. Complementing this framework, we construct WebPerson, a large-scale human-centric dataset with diverse web-sourced image-text pairs. Together, these contributions demonstrate robust performance in human-oriented applications, including intelligent surveillance and autonomous retail systems.