

# Harnessing the Power of MLLMs for Transferable Text-to-Image Person ReID

Wentao Tan<sup>1,3</sup> Changxing Ding<sup>1,2\*</sup> Jiayu Jiang<sup>1</sup> Fei Wang<sup>1,3</sup> Yibing Zhan<sup>3</sup> Dapeng Tao<sup>4,5</sup>

<sup>1</sup>South China University of Technology <sup>2</sup>Pazhou Lab, Guangzhou <sup>3</sup>JD Explore Academy, Beijing

<sup>4</sup>Yunnan University <sup>5</sup>Yunnan United Vision Technology Co., Ltd., Kunming

{ftwentaotan, 202320111494, ft\_feiw}@mail.scut.edu.cn, chxding@scut.edu.cn

zhanyibing@jd.com, dapeng.tao@gmail.com

<https://github.com/WentaoTan/MLLM4Text-ReID>

## Abstract

Text-to-image person re-identification (ReID) retrieves pedestrian images according to textual descriptions. Manually annotating textual descriptions is time-consuming, restricting the scale of existing datasets and therefore the generalization ability of ReID models. As a result, we study the transferable text-to-image ReID problem, where we train a model on our proposed large-scale database and directly deploy it to various datasets for evaluation. We obtain substantial training data via Multi-modal Large Language Models (MLLMs). Moreover, we identify and address two key challenges in utilizing the obtained textual descriptions. First, an MLLM tends to generate descriptions with similar structures, causing the model to overfit specific sentence patterns. Thus, we propose a novel method that uses MLLMs to caption images according to various templates. These templates are obtained using a multi-turn dialogue with a Large Language Model (LLM). Therefore, we can build a large-scale dataset with diverse textual descriptions. Second, an MLLM may produce incorrect descriptions. Hence, we introduce a novel method that automatically identifies words in a description that do not correspond with the image. This method is based on the similarity between one text and all patch token embeddings in the image. Then, we mask these words with a larger probability in the subsequent training epoch, alleviating the impact of noisy textual descriptions. The experimental results demonstrate that our methods significantly boost the direct transfer text-to-image ReID performance. Benefiting from the pre-trained model weights, we also achieve state-of-the-art performance in the traditional evaluation settings.

## 1. Introduction

Text-to-image person re-identification (ReID) [13–16, 24, 41, 42, 44, 50, 52, 61] is a task that retrieves pedestrian

\*Corresponding author

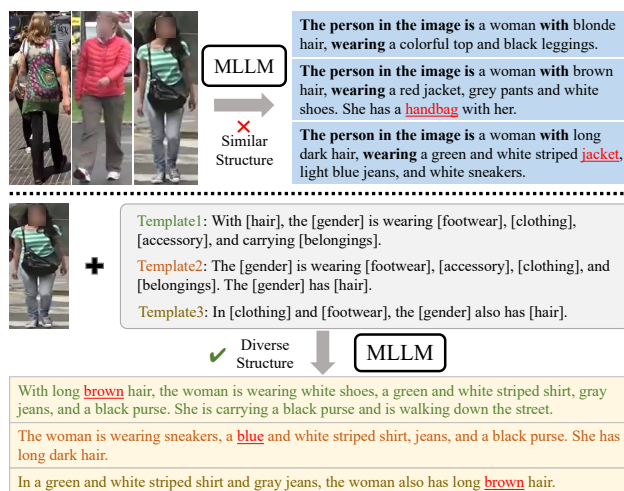


Figure 1. Illustration of textual descriptions generated by an MLLM (*i.e.*, Qwen [3]). (Top) The description patterns are similar for different images. (Bottom) Our proposed Template-based Diversity Enhancement (TDE) method significantly enhances the description pattern diversity. It is worth noting that some errors are present in the generated descriptions shown in this figure.

images according to textual descriptions. It is a powerful tool when probe images of the target person are unavailable and only textual descriptions exist. It has various potential applications, including video surveillance [6], social media analysis [31], and crowd management [18]. However, it remains challenging mainly because annotating textual descriptions for pedestrian images is time-consuming [64]. Consequently, existing datasets [15, 31, 70] for text-to-image person ReID are usually small, resulting in insufficient deep model training.

Previous studies on text-to-image ReID usually assumed that training and testing data are drawn from the same domain. They proposed novel model architectures [4, 15, 38, 41, 57, 58], loss functions [60, 66, 69], and pre-training strategies [42, 64] to improve model performance for each database. However, researchers have recently discovered

that the cross-dataset generalization ability of their approaches is significantly low [41], limiting real-world applications. Since annotating textual descriptions is time-consuming, collecting training data for each target domain is infeasible. Therefore, training a model that can be directly deployed to various target domains is necessary.

Accordingly, we study the transferable text-to-image ReID problem. The term “transferable” is derived from the seminal work CLIP [38], which refers to a large-scale pre-trained model’s capacity that directly applies its knowledge to other domains or tasks without fine-tuning on labeled data. Due to the rapid advancements in multi-modal large language models (MLLMs) [3, 8, 29, 65], we utilize them to generate textual descriptions automatically and employ them to replace traditional manual annotations. Specifically, we utilize the large-scale LUPerson dataset [17] as the image source and generate textual descriptions using MLLMs. The obtained image-text pairs are utilized to train a model directly evaluated in existing text-to-image ReID databases. However, to improve the model’s transfer ability, two essential challenges must be addressed: (1) guiding MLLMs to generate diverse textual descriptions for a single image and (2) reducing the impact of the noise in the synthesized textual descriptions.

First, MLLMs tend to generate descriptions with similar sentence structures, as shown in Fig. 1. This causes the text-to-image ReID model to overfit specific sentence patterns, reducing the model’s ability to generalize to various human description styles encountered in real-world applications. To address this issue, we propose a Template-based Diversity Enhancement (TDE) method that instructs MLLMs to conduct image captioning according to given description templates. Obtaining these templates with minimal effort involves performing multi-turn dialogues with ChatGPT [37] and prompting it to generate diverse templates. Then, we randomly integrate one of these templates into the MLLM’s captioning instruction, resulting in vivid descriptions with varied sentence structures. This approach significantly enhances textual description diversity.

Second, although MLLMs are highly effective, the generated descriptions still contain errors. This implies that certain words in a textual description may not match the paired image. Thus, we propose a novel Noise-aware Masking (NAM) method to address this problem. Specifically, we compute the similarities between each text token and all image tokens in the paired image for a specific textual description. The similarity scores between the unmatched word and image tokens are usually low. Hence, we identify potentially incorrect words and mask them with a large probability in the next training epoch before they are fed into a text encoder. Furthermore, NAM and Masked Language Modeling (MLM) are similar but have two key differences: (1) MLM masks all tokens with equal probab-

ity, while NAM masks them based on their noise level. (2) MLM applies cross-entropy loss to predict the masked tokens, whereas NAM focuses on masking words without predicting potentially noisy words. In the experimentation section, we demonstrate NAM’s ability to effectively alleviate the impact of noisy textual descriptions.

To the best of our knowledge, this is the first study focusing on the transferable text-to-image ReID problem by harnessing the power of MLLMs. We innovatively generate diverse textual descriptions and minimize the impact of the noise contained in these descriptions. The experimental results show that our method performs excellently on three popular benchmarks in both direct transfer and traditional evaluation settings.

## 2. Related Works

**Text-to-Image Re-Identification.** Existing approaches for this task improve model performance from three perspectives: model backbone [4, 24], feature alignment strategies [24, 41, 66], and pre-training [42, 64].

The first method category improves the model backbone. Early approaches adopted the VGG model [9, 31] and LSTM [35, 62, 66] as image and text encoders, respectively. These encoders gradually evolve into ResNet-50 [15, 16, 22, 52] and BERT [12, 32, 40, 43, 70] models. Moreover, the CLIP [21, 38] and ALBEF-based encoders [4, 28, 64] have recently become popular. Notably, the CLIP model contains jointly pre-trained image and text encoders. Thus, its cross-modal alignment capabilities are advantageous and have proven more effective than the individually pre-trained encoders [42]. Moreover, the ALBEF model [28] performs interaction between visual and textual features, which improves the feature representation capacity but brings in significant computational cost.

The second category of methods enhances feature alignment strategies. Previous methods aligned an image’s holistic features with its textual description [1, 43, 49, 51, 56, 59, 66]. Subsequent approaches [10, 19, 25, 26, 33, 36, 45, 53, 54] focused on aligning the image-text pair’s local features to suit the fine-grained retrieval nature of text-to-image ReID. These approaches can be divided into explicit and implicit alignment methods. Explicit methods [15, 52] extract the visual- and textual-part features and then compute the alignment loss between them. Implicit methods can also align local features [16, 41, 64]. For example, Jiang et al. [24] applied MLM to text tokens and then predicted the masked tokens using image token features. This indirectly realizes local feature alignment between the image patch and noun phrase representations.

Since existing databases are small, two recent studies explored pre-training for text-to-image ReID. Shao *et al.* [42] utilized the CLIP model to predict the attributes of a pedestrian image. Then, they inserted these attributes into man-

usually defined description templates. As a result, they obtained a large number of pre-training data. Similarly, Yang *et al.* [64] utilized the text descriptions from the CUHK-PEDES [31] and ICFG-PEDES [15] datasets to synthesize images using a diffusion model [39]. Then, they used the BLIP model [29] to caption these images and obtain a large-scale pre-training dataset. However, these two studies targeted at pre-training and did not investigate the direct transfer setting where no target domain data is available for fine-tuning. Moreover, they overlooked the noise or diversity issues generated in the obtained textual descriptions.

The above methods achieve excellent in-domain performance; however, their cross-dataset performance is usually significantly low [41]. This paper explores the transferable text-to-image ReID task with minimal manual operations. Also, we address the challenges in textual descriptions generated by MLLMs.

**Multi-modal Large Language Models.** Multi-modal Large Language Models (MLLMs) [34, 46, 47, 65, 71] are built on Large Language Models (LLMs) [5, 11, 63, 67, 68] and incorporate textual and non-textual information as input [3, 8, 20]. This paper only considers MLLMs that use both texts and images as input signals. The input text (*i.e.*, the “instruction” or “prompt”) describes the tasks assigned to MLLMs to understand the image’s content. Regarding MLLM architecture, most studies [30, 34, 48] first map the image patch and text token embeddings into a shared feature space and then perform decoding using a LLM. Some methods [2] improve the interaction and alignment strategies between the image and text tokens during decoding, facilitating more stable training [27].

In this paper, we utilize MLLMs to eliminate the need to manually annotate textual descriptions. We also explore strategies to address the diversity and noise issues in the obtained textual descriptions, facilitating the development of a transferable text-to-image ReID model.

### 3. Methods

The overview of our solution to the transferable text-to-image ReID problem is illustrated in Fig. 2. Section 3.1 addresses diversity issues associated with textual descriptions generated by MLLMs. Section 3.2 discusses the reduction of noise impact in the descriptions. And section 3.3 outlines the loss function utilized for model optimization.

#### 3.1. Generating Diverse Descriptions

Manually annotating textual descriptions for pedestrian images is time-consuming and hardly scalable. Fortunately, MLLMs have advanced rapidly and provide effective image captioning. Therefore, we decide to utilize MLLMs to create large-scale text annotations for training a model with excellent transfer capacity.

**Instruction Design.** We adopt the LUPerson database [17] as the image source because it holds a significant amount of images that were captured in diverse environments. A technical aspect of using MLLMs lies in designing an effective instruction, which usually depends on user experience. We solve this problem using a multi-turn dialogue with ChatGPT [37], and this process is detailed in the supplementary material. The resulting instruction is as follows:

*“Write a description about the overall appearance of the person in the image, including the attributes: clothing, shoes, hairstyle, gender and belongings. If any attribute is not visible, you can ignore it. Do not imagine any contents that are not in the image.”*

This is considered a *static instruction* as it is fixed for all images. In this paper, the textual descriptions generated using the *static instruction* are denoted as *static texts* or  $T^s$ .

**Diversity Enhancement.** An MLLM generates textual descriptions with similar sentence patterns for different images using the *static instruction*, as illustrated in Fig. 1. This causes the text-to-image ReID model to overfit these sentence patterns, limiting its generalization to real-world descriptions. We attempt to improve the *static instruction*, but the obtained sentence patterns remained limited. Although using more MLLMs can bring in multiple sentence patterns, these patterns are still far from diverse.

Again, we resort to ChatGPT to solve this problem. Specifically, we propose a Template-based Diversity Enhancement (TDE) method. First, we generate two descriptions for each of a set of images using two MLLMs [3, 8] according to the *static instruction*. Then, we feed these descriptions to ChatGPT to capture their sentence patterns (*i.e.*, description templates). With the guidance of these templates, we instruct ChatGPT to create more templates. Finally, it produces 46 templates after multi-turn dialogues, which are detailed in the supplementary material. We randomly select one of the templates and insert it into the *static instruction*, obtaining a *dynamic instruction* as follows:

*“Generate a description about the overall appearance of the person, including clothing, shoes, hairstyle, gender, and belongings, in a style similar to the template: ‘{template}’. If some requirements in the template are not visible, you can ignore them. Do not imagine any contents that are not in the image.”*

The ‘{template}’ is replaceable. Furthermore, the textual descriptions generated according to the *dynamic instruction* are referred to as *dynamic texts* ( $T^d$ ). As illustrated in Fig. 1, MLLMs can follow the sentence patterns specified in the templates, significantly enhancing the diversity of the obtained textual descriptions.

**Dataset Description.** We utilize the publicly available Qwen [3] and Shikra [8] models in this paper. By harnessing the power of the two MLLMs, we obtain the large-

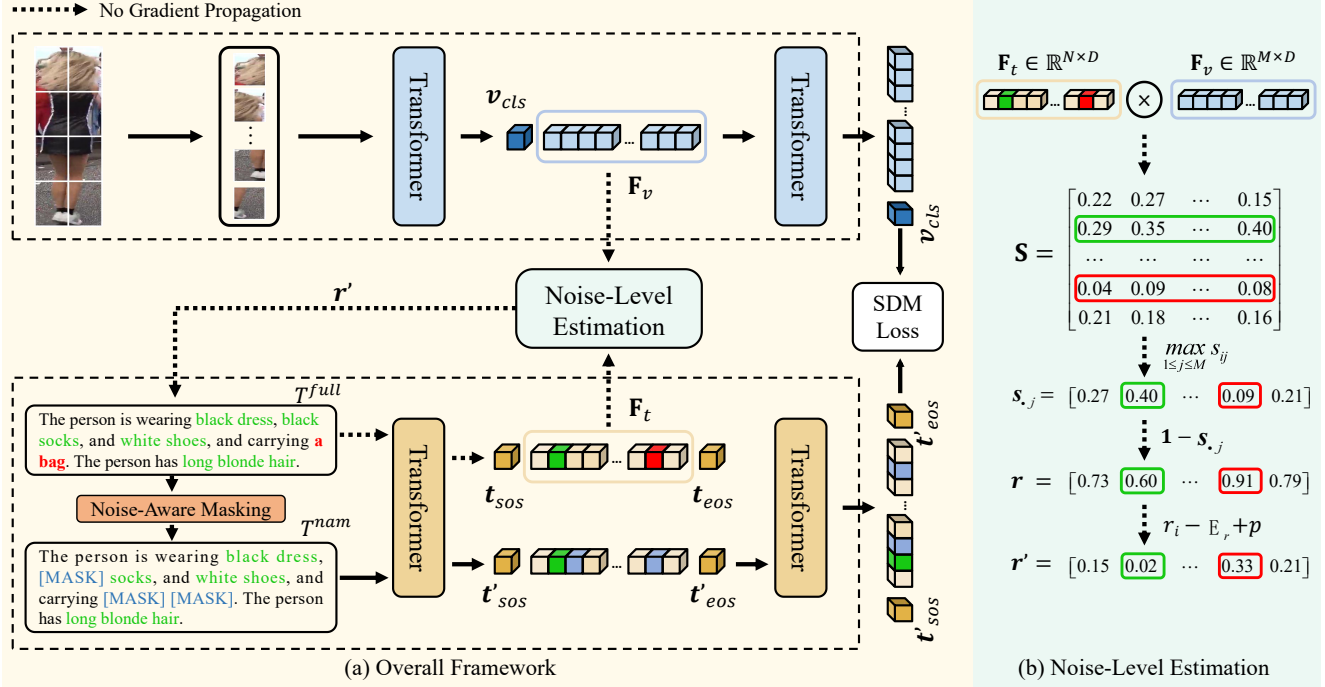


Figure 2. Overview of our framework. We adopt the CLIP-ViT/B-16 model as the backbone. Our framework uses one pedestrian image, the original textual description  $T^{full}$ , and a masked textual description  $T^{nam}$  as input during training.  $T^{nam}$  is obtained by applying NAM to  $T^{full}$ . To perform NAM, we first compute the similarity matrix  $\mathbf{S}$  between the text tokens  $\mathbf{F}_t$  of  $T^{full}$  and the image tokens  $\mathbf{F}_v$  according to their embeddings at the  $l$ -th layer of the encoders. Then, we estimate the probability of each text token’s noisiness according to the similarity between its embedding and the image token embeddings. The similarity distribution matching (SDM) loss is computed between the global visual feature  $v_{cls}$  of the pedestrian image and the global textual feature  $t'_{eos}$  of  $T^{nam}$ . The model’s optimization quality is enhanced by masking noisy words in  $T^{full}$ . (Best viewed in color.)

scale LUPerson-MLLM dataset. This dataset comprises 1.0 million images, and each image has four captions,  $T_{qwen}^s$ ,  $T_{shikra}^s$ ,  $T_{qwen}^d$ , and  $T_{shikra}^d$ . The first and the last two captions are generated according to the static and dynamic instructions, respectively. We reserve the  $T^s$  for each image as we observe that its description is usually complementary to that of  $T^d$ . In the following section, we will train the model with LUPerson-MLLM. For simplicity, we refer all the above MLLM-generated descriptions as  $T^{full}$ .

### 3.2. Noise-Aware Masking

Although MLLMs are powerful, they cannot describe images very precisely. As depicted in Fig. 1 and Fig. 2, a few words do not match the described image in the obtained textual descriptions. Existing methods [23, 29] usually discard the noisy descriptions, losing the other valuable information contained in the matched words. Accordingly, we propose a novel noise-aware masking (NAM) method that identifies noisy text tokens and fully uses the matched text tokens for model training.

**Image Encoder.** An image is divided into  $M$  non-overlapped patches. These image tokens are concatenated with the [CLS] token and are fed into the image encoder.

Then, the [CLS] token embedding at the last image encoder layer is used as the global image feature, denoted as  $v_{cls} \in \mathbb{R}^d$ . The feature dimension is represented by  $d$ .

**Text Encoder.** We tokenize each textual description  $T^{full}$  into a sequence of  $N$  tokens. The  $N$  of each sentence varies according to its length. The token sequence is bracketed with [SOS] and [EOS] to represent the start and the end of the sequence. Meanwhile, we examine each text token’s noise level in  $T^{full}$ , which is computed and stored in the previous training epoch. These values are used to perform NAM on  $T^{full}$  to obtain  $T^{nam}$ . After that,  $T^{full}$  and  $T^{nam}$  are fed into the text encoder independently. At the final text encoder layer, the global feature  $t'_{eos}$  of  $T^{nam}$  is utilized to calculate loss.  $T^{full}$  is only used for NAM, which means it is not used for loss computation.

**Noise-Aware Masking.** We utilize the image and text encoders’ token embeddings in the  $l$ -th layers for the noise-level estimation of  $T^{full}$ . These embeddings are denoted as  $\mathbf{F}_v = [v_1^l, \dots, v_M^l]$  and  $\mathbf{F}_t = [t_1^l, \dots, t_N^l]$ , respectively, where  $v_j^l \in \mathbb{R}^d$  and  $t_j^l \in \mathbb{R}^d$ .

Furthermore, we calculate the token-wise similarity between a single text-image pair as follows:

$$\mathbf{S} = \mathbf{F}_t^T \mathbf{F}_v, \quad (1)$$

where  $\mathbf{S} \in \mathbb{R}^{N \times M}$  is a similarity matrix and  $s_{ij}$  represents the cosine similarity between the  $i$ -th text token embedding and the  $j$ -th image token embedding. If one text token does not match the image, the similarity scores between this token’s embedding and those of all the image tokens will be consistently be low. Therefore, the noise level of the  $i$ -th text token in  $T^{full}$  can be estimated via:

$$r_i = 1 - \left( \max_{1 \leq j \leq M} s_{ij} \right). \quad (2)$$

By applying Eq.(2) to each row of  $\mathbf{S}$ , we obtain a vector  $\mathbf{r} = [r_1, \dots, r_N]$  that records the noise-level of all text tokens.

Moreover, NAM applies the masking operation to all the text tokens in  $T^{full}$  with different probabilities, which can be determined based on the noise-level values recorded in  $\mathbf{r}$ . However, in the initial training stage, the values of elements in  $\mathbf{r}$  may be high. This results in excessive masking of important tokens and hinders learning. To resolve this issue, we modify the expectation value of all  $\mathbf{r}$  elements into a constant number as described below:

$$\mathbb{E}_r = \frac{1}{N} \sum_{i=1}^N r_i, \quad (3)$$

$$\mathbf{r}' = [r_1 - \mathbb{E}_r + p, \dots, r_N - \mathbb{E}_r + p], \quad (4)$$

where  $p$  is the average masking ratio. We utilize the  $\mathbf{r}'$  values as the final probability that a text token might be masked. We include the pseudo code and visualization of NAM in the supplementary materials.

**Discussion.** Computing  $\mathbf{r}'$  and then applying NAM to obtain  $T^{nam}$  in each iteration requires two forward passes. This additional time cost cannot be overlooked in large-scale training. In contrast, our strategy computes  $\mathbf{r}'$  for the next training epoch, which requires only one forward pass for each iteration. Furthermore, we initialize the  $\mathbf{r}'$  values with the constant  $p$  in the first training epoch.

### 3.3. Optimization

Following [24], we adopt the similarity distribution matching (SDM) loss to optimize our model. Given a mini-batch of  $B$  matched image-text pairs  $\{(\mathbf{v}_{cls}^i, \mathbf{t}_{eos}^i)\}_i^B$ , we first establish the matching relationship between each image and text (*i.e.*,  $\{(\mathbf{v}_{cls}^i, \mathbf{t}_{eos}^j), y_{i,j}\} (1 \leq i, j \leq B)$ ), where  $y_{i,j} = 1$  and  $y_{i,j} = 0$  denote a positive and a negative image-text pair, respectively. Then, we calculate the ground truth matching distribution  $\mathbf{q}_i$  for the  $i$ -th image, where its  $j$ -th element is  $q_{i,j} = y_{i,j} / \sum_{b=1}^B y_{i,b}$ . Finally, we align the predicted probability distribution  $\mathbf{p}_i$  with  $\mathbf{q}_i$  as follows:

$$\mathcal{L}_{i2t} = \frac{1}{B} \sum_{i=1}^B KL(\mathbf{p}_i || \mathbf{q}_i) = \frac{1}{B} \sum_{i=1}^B \sum_{j=1}^B p_{i,j} \log\left(\frac{p_{i,j}}{q_{i,j} + \epsilon}\right), \quad (5)$$

where  $\epsilon$  is a small number to avoid numerical problems and

$$p_{i,j} = \frac{\exp(\text{sim}(\mathbf{v}_{cls}^i, \mathbf{t}_{eos}^j)/\tau)}{\sum_{b=1}^B \exp(\text{sim}(\mathbf{v}_{cls}^i, \mathbf{t}_{eos}^b)/\tau)}. \quad (6)$$

$\text{sim}(\mathbf{u}, \mathbf{v}) = \mathbf{u}^\top \mathbf{v} / \|\mathbf{u}\| \|\mathbf{v}\|$  denotes the cosine similarity between  $\mathbf{u}$  and  $\mathbf{v}$ ,  $\tau$  is a temperature coefficient.

The SDM loss from text to image  $\mathcal{L}_{t2i}$  can be computed by exchanging the position of  $\mathbf{v}_{cls}$  and  $\mathbf{t}_{eos}$  in Eq. (5) and Eq. (6). Finally, the complete SDM loss is computed as follows:

$$\mathcal{L}_{sdm} = \mathcal{L}_{i2t} + \mathcal{L}_{t2i}. \quad (7)$$

It is worth noting that since we randomly sample images from the large-scale LUPerson database, we assume that each image in a sampled batch has a unique identity.

## 4. Experiments

### 4.1. Datasets and Settings

**CUHK-PEDES.** CUHK-PEDES [31] is a pioneer dataset in the text-to-image ReID field. Each image in this dataset has two textual descriptions. The training set comprises data on 11,003 identities, including 34,054 images and 68,108 textual descriptions. In contrast, the testing set contains 3,074 images and 6,156 textual descriptions from 1,000 identities.

**ICFG-PEDES.** ICFG-PEDES [15] contains of 54,522 images from 4,102 identities. Each image has one textual description. The training set consists of 34,674 image-text pairs corresponding to 3,102 identities, while the testing set comprises 19,848 image-text pairs from the remaining 1,000 identities.

**RSTPReid.** RSTPReid [70] includes 20,505 images captured by 15 cameras from 4,101 identities. Each identity has five images captured with different cameras and each image has two textual descriptions. According to the official data division, the training set incorporates data from 3,701 identities, while both the validation and testing sets include data from 200 identities, respectively.

**LUPerson.** LUPerson [17] contains 4,180,243 pedestrian images sampled from 46,260 online videos, covering a variety of scenes and view points. The images are from over 200K pedestrians.

**Evaluation Metrics.** Like existing works [4, 24, 42, 64], we adopt the popular Rank- $k$  accuracy ( $k=1,5,10$ ) and mean Average Precision (mAP) as the evaluation metrics for the three databases. Moreover, we consider the following two evaluation settings.

**Direct Transfer Setting.** For this setting, the model is only trained on the LUPerson-MLLM dataset, and the above three benchmarks are tested immediately. This setting directly evaluates the quality of our dataset and the effectiveness of the proposed methods (*i.e.*, TDE and NAM).

Table 1. Ablation study on each key component in the direct transfer setting. ‘CLIP’ refers to directly using the original CLIP encoders provided in [38].

Method	$T_{qwen}^s$	$T_{shikra}^s$	$T_{qwen}^d$	$T_{shikra}^d$	NAM	CUHK-PEDES			ICFG-PEDES			RSTPReID		
						R1	R5	mAP	R1	R5	mAP	R1	R5	mAP
CLIP						12.65	27.16	11.15	6.67	17.91	2.51	13.45	33.85	10.31
Static Text	✓					37.65	57.86	33.40	23.78	42.77	11.18	36.30	60.60	26.25
		✓				39.70	62.60	36.09	19.02	35.63	9.67	36.90	62.65	28.33
	✓	✓				46.00	66.82	41.27	26.74	44.22	13.23	41.10	66.95	30.21
Dynamic Text			✓			40.72	62.36	37.21	24.16	41.24	11.32	38.65	64.70	28.81
				✓		43.63	65.46	39.08	22.07	39.57	11.35	38.80	63.45	28.60
			✓	✓		48.86	69.41	44.09	28.43	46.37	14.23	44.25	66.15	32.99
TDE	✓	✓	✓	✓		50.32	71.36	45.74	29.12	47.96	15.13	45.70	70.75	33.23
NAM	✓	✓	✓	✓	✓	52.64	71.62	46.48	32.61	50.79	16.48	47.75	70.75	34.73

**Fine-tuning Setting.** In this setting, we first pre-train our model on the LUPerson-MLLM dataset and then fine-tune it on each of the three benchmarks respectively.

## 4.2. Implementation Details

Similar to previous studies [7, 24], we adopt CLIP-VIT-B/16 [38] as the image encoder and a 12-layer transformer as our text encoder. The input image resolution is resized to  $384 \times 128$  pixels. Additionally, we apply random horizontal flipping, random cropping, and random erasing as data augmentation for the input images. Each textual description is first tokenized, with a maximum length of 77 tokens (including the [SOS] and [EOS] tokens). The hyper-parameter  $p$  is set to 0.15 and the temperature coefficient  $\tau$  in Eq. (6) is set to 0.02. The model is trained using the Adam optimizer with a learning rate of  $1e-5$  and cosine learning rate decay strategy. We train each model on 8 TITAN-V GPUs, with 64 images per GPU. The training process lasts for 30 epochs. The versions of the mentioned LLM/MLLMs are ChatGPT-3.5-Turbo, Qwen-VL-Chat-7B, and Shikra-7B.

## 4.3. Ablation Study

We randomly sample 0.1 million images from our LUPerson-MLLM database to accelerate the ablation study on the direct transfer evaluation setting. Then, we increase the amount of training images to 1.0 million to enhance the transfer ability of our text-to-image ReID models.

**Effectiveness of TDE.** The experiments in Table 1 show that *dynamic instruction* is better than *static instruction*. For example, the model using only  $T_{qwen}^d$  outperforms that the one using  $T_{qwen}^s$  by about 3% in Rank-1 performance on the CUHK-PEDES database. On the same database and evaluation metric, the model that uses only  $T_{shikra}^d$  outperforms the one using  $T_{shikra}^s$  by about 4%. These experimental results indicate that enhancing sentence pattern diversity improves the transfer ability of ReID models. Therefore, we use the four descriptions for each image in the subsequent experiments. It is worth noting that none of the above experiments employ NAM. Instead, they mask every text token with an equal probability of  $p$ .

**Effectiveness of NAM.** MLLM-generated textual descriptions often contain noise, which is harmful for model training. Replacing the equal masking strategy with our NAM method improves our model’s Rank-1 performance by 2.32%, 3.49%, and 2.05% on the three databases, respectively. These improvements are even higher than the benefits of combining dynamic and static texts (*i.e.*, 1.46%, 0.69%, and 1.45%). These experimental results demonstrate that NAM identifies the noisy words in the text and effectively reduces their impact. NAM allows the model to accurately align visual and textual features, thereby enhancing the direct transfer text-image ReID performance.

**The Layer where NAM Computes S.**  $S$  contains pairwise similarity scores between features in  $F_v$  and  $F_t$ . This experiment investigates the optimal layer for obtaining  $F_v$  and  $F_t$ . The results are plotted in Fig. 3. We observe that the model’s performance consistently improves regardless of the layer used to provide  $F_v$  and  $F_t$ . We also notice that the adopted encoders’ 10-th layer yields the best overall performance. Compared to the last encoder layer, the 10-th layer may offer more fine-grained information, facilitating more accurate similarity computation between token pairs.

**The Overall Masking Ratio for NAM.** Our NAM method masks different text tokens with unequal probabilities, but it maintains an overall probability of  $p$ . In this experiment, we explore the optimal  $p$  value. To demonstrate NAM’s advantages, we also include the results of the masking tokens with equal probabilities (referred to as “EM”). As shown in Table 4, NAM consistently outperforms EM with various  $p$  values. The optimal value of  $p$  is about 0.15.

**Combination of NAM and MLM.** MLM requires the model to predict the masked text tokens. It has proven effective and is widely applied in NLP models. Recent text-to-image ReID studies [24] confirm that MLM loss is beneficial when the textual descriptions are manually annotated. However, our NAM doesn’t predict the masked tokens as the textual descriptions generated by MLLMs may be noisy. Table 2 shows that applying MLM loss to NAM is harmful, indicating the MLLM description noise is a crucial issue.

**The Data Size Impact.** The dataset size is essential to




Figure 3. Results of different layers for NAM to compute  $S$ . The encoders contain 12 layers in total. Best viewed with zoom-in.




Figure 4. Results of different overall masking ratios  $p$  for NAM. ‘EM’ represents masking all text tokens with the same probability  $p$ . Best viewed with zoom-in.

Table 2. Results of the combination of NAM and the MLM loss.

Method	CUHK-PEDES		ICFG-PEDES		RSTPReid	
	R1	mAP	R1	mAP	R1	mAP
EM	50.32	45.74	29.12	15.13	45.70	33.23
NAM	52.64	46.48	32.61	16.48	47.75	34.73
NAM w/ MLM loss	48.79	43.86	27.36	14.16	44.45	33.07

Table 3. Comparisons with existing pre-training datasets in the direct transfer setting.

Pretrain 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 [64] (1.5 M)	19.36	18.62	7.93	3.52	22.85	17.11
LUPerson-T [42] (0.95 M)	21.88	19.96	11.46	4.56	22.40	17.08
Ours (0.1 M)	52.64	46.48	32.61	16.53	47.75	34.73
Ours (1.0 M)	57.61	51.44	38.36	20.43	51.50	37.34

training. More pre-trained data improves the performance. We investigate the effect of training data size on the direct transfer ReID performance and summarize the results in Fig. 5. It is evident that the model’s direct transfer performance steadily improves as the data amount increases. Finally, compared with the model using only 0.1 million training images, the Rank-1 performance of the model using 1.0 million training images is significantly promoted by 5.75% on the challenging ICFG-PEDES database, indicating that our approach can scale to large-scale database.

#### 4.4. Comparisons with State-of-the-Art Methods

**Comparisons with Other Pre-training Datasets.** MALS [64] and LUPerson-T [42] are two pre-training datasets in the field of text-to-image ReID. MALS [64] contains 1.5 M images, with textual descriptions obtained using the BLIP model [29]. However, it does not address the diversity and




Figure 5. Training data size’s impact on our methods’ direct transfer ReID performance. ‘0 M’ refers to directly using the original CLIP encoders.

Table 4. Comparisons with existing pre-training datasets in the fine-tuning setting.

Init Parameters	Source	Target					
		CUHK-PEDES		ICFG-PEDES		RSTPReid	
		R1	mAP	R1	mAP	R1	mAP
CLIP [38]	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
	PSTPReid	35.25	32.35	33.58	19.58	60.40	47.70
MALS [64] (1.5 M)	CUHK-PEDES	74.05	66.57	44.53	22.66	53.55	39.17
	ICFG-PEDES	40.38	36.83	64.37	38.85	49.00	38.20
	PSTPReid	38.40	34.47	34.11	20.82	61.90	48.08
LuPerson-T [42] (0.95 M)	CUHK-PEDES	74.37	66.60	44.30	22.67	53.75	38.98
	ICFG-PEDES	35.07	32.47	64.50	38.22	48.05	38.21
	PSTPReid	38.29	34.43	35.81	21.62	62.20	48.33
Ours (0.1 M)	CUHK-PEDES	74.64	67.44	46.19	24.08	56.15	40.84
	ICFG-PEDES	56.70	51.23	65.30	39.90	52.60	39.76
	PSTPReid	56.69	51.40	42.70	25.69	64.05	49.27
Ours (1.0 M)	CUHK-PEDES	76.82	69.55	49.38	26.92	59.60	44.70
	ICFG-PEDES	61.20	55.60	67.05	41.51	54.80	42.56
	PSTPReid	62.99	57.20	48.44	30.03	68.50	53.02

noise issues in the obtained descriptions. LUPerson-T [42] contains 0.95 M images that were also sampled from the LUPerson database [42]. It utilizes the CLIP model to predict pedestrian attributes and inserts them into manually defined templates as textual descriptions. We utilize the three databases to train the CLIP-ViT/B-16 model, incorporating the SDM loss. Finally, we evaluate the model’s performance in both direct transfer and fine-tuning settings.

Comparisons on the direct transfer setting are summarized in Table 3. It is shown that the model trained on the LUPerson-MLLM dataset achieves significantly better performance, even when we only sample 0.1 M images. This is because TDE enables diverse description generation. Moreover, NAM efficiently alleviates the impact of noise in textual descriptions. Combining both techniques results in a model that exhibits exceptional transfer abilities. In comparison, neither [64] nor [42] consider the noise problem in their obtained textual descriptions.

Table 4 displays the model comparisons in the fine-tuning setting. In this experiment, we adopt the IRRA method [24] in the fine-tuning stage and initialize its parameters with each of the above three pre-trained models, respectively. The fine-tuned models are evaluated on both

Table 5. Comparisons with state-of-the-art methods in the traditional evaluation settings.

Method	Image Enc.	Text Enc.	CUHK-PEDES				ICFG-PEDES				RSTPReid			
			R1	R5	R10	mAP	R1	R5	R10	mAP	R1	R5	R10	mAP
CMPM/C [66]	RN50	LSTM	49.37	-	79.27	-	43.51	65.44	74.26	-	-	-	-	-
VITAA [52]	RN50	LSTM	55.97	75.84	83.52	-	50.98	68.79	75.78	-	-	-	-	-
DSSL [69]	RN50	BERT	59.98	80.41	87.56	-	-	-	-	-	32.43	55.08	63.19	-
SSAN [15]	RN50	LSTM	61.37	80.15	86.73	-	54.23	72.63	79.53	-	43.50	67.80	77.15	-
LapsCore [55]	RN50	BERT	63.40	-	87.80	-	-	-	-	-	-	-	-	-
LBUL [54]	RN50	BERT	64.04	82.66	87.22	-	-	-	-	-	45.55	68.2	77.85	-
SAF [32]	ViT-Base	BERT	64.13	82.62	88.4	-	-	-	-	-	-	-	-	-
TIPCB [10]	RN50	BERT	64.26	83.19	89.1	-	54.96	74.72	81.89	-	-	-	-	-
CAIBC [53]	RN50	BERT	64.43	82.87	88.37	-	-	-	-	-	47.35	69.55	79.00	-
AXM-Net [16]	RN50	BERT	64.44	80.52	86.77	58.70	-	-	-	-	-	-	-	-
LGUR [41]	DeiT-Small	BERT	65.25	83.12	89.00	-	59.02	75.32	81.56	-	47.95	71.85	80.25	-
IVT [43]	ViT-Base	BERT	65.69	85.93	91.15	-	56.04	73.60	80.22	-	46.70	70.00	78.80	-
LCR <sup>2</sup> S [61]	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 [42]	ViT-Base	BERT	68.50	84.67	90.38	-	60.09	76.19	82.46	-	51.85	74.85	82.85	-
<i>with CLIP [38] backbone:</i>														
Han et al. [21]	CLIP-RN101	CLIP-Xformer	64.08	81.73	88.19	60.08	-	-	-	-	-	-	-	-
IRRA [24]	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
MALS [64] + IRRA	CLIP-ViT	CLIP-Xformer	74.05	89.48	93.64	66.57	64.37	80.75	86.12	38.85	61.90	80.60	89.30	48.08
LUPerson-T [42] + IRRA	CLIP-ViT	CLIP-Xformer	74.37	89.51	93.97	66.60	64.50	80.24	85.74	38.22	62.20	83.30	89.75	48.33
Ours (1.0 M) + IRRA	CLIP-ViT	CLIP-Xformer	<b>76.82</b>	<b>91.16</b>	<b>94.46</b>	<b>69.55</b>	<b>67.05</b>	<b>82.16</b>	<b>87.33</b>	<b>41.51</b>	<b>68.50</b>	<b>87.15</b>	<b>92.10</b>	<b>53.02</b>
<i>with ALBEF [28] backbone:</i>														
RaSa [4]	CLIP-ViT	BERT-base	76.51	90.29	94.25	<b>69.38</b>	65.28	80.40	85.12	41.29	66.90	86.50	91.35	52.31
APTM [64]	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
Ours (1.0 M) + APTM	Swin-B	BERT-base	<b>78.13</b>	<b>91.19</b>	<b>94.50</b>	68.75	<b>69.37</b>	<b>83.55</b>	<b>88.18</b>	<b>42.42</b>	<b>69.95</b>	<b>87.35</b>	<b>92.30</b>	<b>54.17</b>

in-domain and cross-domain text-to-image ReID scenarios. According to the results in Table 4, two conclusions can be derived. First, compared with the CLIP model [38], pre-training using the three pre-training datasets exhibits performance promotion for in-domain and cross-domain tasks. Second, pre-training using LUPerson-MLLM exhibits the most remarkable performance promotion. For example, in the ICFG-PEDES  $\rightarrow$  CUHK-PEDES setting, LUPerson-MLLM outperforms the other two models by 20.82% and 26.13% in Rank-1 accuracy, respectively. These experimental results further validate the effectiveness of our methods.

**Comparisons in the Traditional Evaluation Settings.** Comparisons with state-of-the-art approaches are summarized in Table 5. We observe that our method achieves the best performance. With our pre-trained model parameters, the Rank-1 accuracy and mAP of IRRA are improved by 8.30% and 5.85% on the RSTPReid database, respectively. Besides, pre-training with our LUPerson-MLLM dataset is more effective than with the MALS and LUPerson-T datasets. This is because we effectively resolve the diversity and noise issues in the MLLM descriptions, facilitating more robust and discriminative feature learning.

## 5. Conclusion and Limitations

This paper explores the challenging transferable text-to-image ReID problem by harnessing the image captioning capability of MLLMs. We acknowledge diversity and noise as critical issues in utilizing the obtained textual descriptions. To address these two problems, we introduce the Template-based Diversity Enhancement (TDE) method to

encourage diverse description generation and construct a large-scale dataset named LUPerson-MLLM. In addition, we proposed the NAM method to mitigate the impact of noisy textual descriptions. Extensive experiments demonstrate that TDE and NAM significantly improve the model’s transfer power. However, these methods have limitations: the effectiveness of TDE is limited by the number of sentence templates; NAM may occasionally fail to mask noisy tokens. In the future, we aim to explore more powerful methods to address diversity and noise issues in MLLM-generated descriptions.

**Broader Impacts.** TDE addresses fixed sentence patterns generated by MLLMs, inspiring effective instruction design to harness MLLMs’ capabilities. Meanwhile, NAM tackles text noise generated by MLLMs, facilitating wider MLLM adoption for practical real-world problems.

**Acknowledgement.** This work was partially supported by the Major Science and Technology Innovation 2030 “New Generation Artificial Intelligence” key project (No. 2021ZD0111700), the National Natural Science Foundation of China under Grants 62076101 and 62172354, the Guangdong Basic and Applied Basic Research Foundation under Grant 2023A1515010007, the Guangdong Provincial Key Laboratory of Human Digital Twin under Grant 2022B1212010004, and the Yunnan Provincial Major Science and Technology Special Plan Projects under Grant 202202AD080003. We also gratefully acknowledge the support and resources provided by the Yunnan Key Laboratory of Media Convergence, the CAAI Huawei MindSpore Open Fund and the TCL Young Scholars Program.



## References

- [1] Surbhi Aggarwal, Venkatesh Babu Radhakrishnan, and Anirban Chakraborty. Text-based person search via attribute-aided matching. In *WACV*, 2020. 2
- [2] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *NeurIPS*, 2022. 3
- [3] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 2023. 1, 2, 3
- [4] Yang Bai, Min Cao, Daming Gao, Ziqiang Cao, Chen Chen, Zhenfeng Fan, Liqiang Nie, and Min Zhang. Rasa: Relation and sensitivity aware representation learning for text-based person search. *IJCAI*, 2023. 1, 2, 5, 8
- [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *NeurIPS*, 2020. 3
- [6] Maryam Bukhari, Sadaf Yasmin, Sheneela Naz, Muazzam Maqsood, Jehyeok Rew, and Seungmin Rho. Language and vision based person re-identification for surveillance systems using deep learning with lip layers. *Image and Vision Computing*, 2023. 1
- [7] Cuiqun Chen, Mang Ye, and Ding Jiang. Towards modality-agnostic person re-identification with descriptive query. In *CVPR*, 2023. 6
- [8] Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. Shikra: Unleashing multi-modal llm’s referential dialogue magic. *arXiv preprint arXiv:2306.15195*, 2023. 2, 3
- [9] Tianlang Chen, Chenliang Xu, and Jiebo Luo. Improving text-based person search by spatial matching and adaptive threshold. In *WACV*, 2018. 2
- [10] Yuhao Chen, Guoqing Zhang, Yujiang Lu, Zhenxing Wang, and Yuhui Zheng. Tipcb: A simple but effective part-based convolutional baseline for text-based person search. *Neuro-computing*, 2022. 2, 8
- [11] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022. 3
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2
- [13] Changxing Ding and Dacheng Tao. Trunk-branch ensemble convolutional neural networks for video-based face recognition. *IEEE TPAMI*, 2018. 1
- [14] Changxing Ding, Kan Wang, Pengfei Wang, and Dacheng Tao. Multi-task learning with coarse priors for robust part-aware person re-identification. *IEEE TPAMI*, 2022.
- [15] Zefeng Ding, Changxing Ding, Zhiyin Shao, and Dacheng Tao. Semantically self-aligned network for text-to-image part-aware person re-identification. *arXiv preprint arXiv:2107.12666*, 2021. 1, 2, 3, 5, 8
- [16] Ammarah Farooq, Muhammad Awais, Josef Kittler, and Syed Safwan Khalid. Axm-net: Implicit cross-modal feature alignment for person re-identification. In *AAAI*, 2022. 1, 2, 8
- [17] Dengpan Fu, Dongdong Chen, Jianmin Bao, Hao Yang, Lu Yuan, Lei Zhang, Houqiang Li, and Dong Chen. Unsupervised pre-training for person re-identification. In *CVPR*, 2021. 2, 3, 5
- [18] Hiren Galiyawala and Mehul S Raval. Person retrieval in surveillance using textual query: a review. *Multimedia Tools and Applications*, 2021. 1
- [19] Chenyang Gao, Guanyu Cai, Xinyang Jiang, Feng Zheng, Jun Zhang, Yifei Gong, Pai Peng, Xiaowei Guo, and Xing Sun. Contextual non-local alignment over full-scale representation for text-based person search. *arXiv preprint arXiv:2101.03036*, 2021. 2
- [20] Jiaming Han, Renrui Zhang, Wenqi Shao, Peng Gao, Peng Xu, Han Xiao, Kaipeng Zhang, Chris Liu, Song Wen, Ziyu Guo, et al. Imagebind-llm: Multi-modality instruction tuning. *arXiv preprint arXiv:2309.03905*, 2023. 3
- [21] Xiao Han, Sen He, Li Zhang, and Tao Xiang. Text-based person search with limited data. *arXiv preprint arXiv:2110.10807*, 2021. 2, 8
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 2
- [23] Ruifei He, Shuyang Sun, Xin Yu, Chuhui Xue, Wenqing Zhang, Philip Torr, Song Bai, and Xiaojuan Qi. Is synthetic data from generative models ready for image recognition? In *ICLR*, 2023. 4
- [24] Ding Jiang and Mang Ye. Cross-modal implicit relation reasoning and aligning for text-to-image person retrieval. In *CVPR*, 2023. 1, 2, 5, 6, 7, 8
- [25] Ya Jing, Chenyang Si, Junbo Wang, Wei Wang, Liang Wang, and Tieniu Tan. Pose-guided multi-granularity attention network for text-based person search. In *AAAI*, 2020. 2
- [26] Kuang-Huei Lee, Xi Chen, Gang Hua, Houdong Hu, and Xiaodong He. Stacked cross attention for image-text matching. In *ECCV*, 2018. 2
- [27] Chunyuan Li, Zhe Gan, Zhengyuan Yang, Jianwei Yang, Linjie Li, Lijuan Wang, and Jianfeng Gao. Multimodal foundation models: From specialists to general-purpose assistants. *arXiv preprint arXiv:2309.10020*, 2023. 3
- [28] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *NeurIPS*, 2021. 2, 8
- [29] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, 2022. 2, 3, 4, 7
- [30] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with

- frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. 3
- [31] Shuang Li, Tong Xiao, Hongsheng Li, Bolei Zhou, Dayu Yue, and Xiaogang Wang. Person search with natural language description. In *CVPR*, 2017. 1, 2, 3, 5
- [32] Shiping Li, Min Cao, and Min Zhang. Learning semantic-aligned feature representation for text-based person search. In *ICASSP*, 2022. 2, 8
- [33] Zechao Li, Hao Tang, Zhimao Peng, Guo-Jun Qi, and Jinhui Tang. Knowledge-guided semantic transfer network for few-shot image recognition. *TNNLS*. 2
- [34] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023. 3
- [35] Long Short-Term Memory. Long short-term memory. *Neural computation*, 2010. 2
- [36] Kai Niu, Yan Huang, Wanli Ouyang, and Liang Wang. Improving description-based person re-identification by multi-granularity image-text alignments. *TIP*, 2020. 2
- [37] OpenAI. Chatgpt. <https://openai.com/blog/chatgpt/>, 2022. 2, 3
- [38] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 1, 2, 6, 7, 8
- [39] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 3
- [40] Nikolaos Sarafianos, Xiang Xu, and Ioannis A Kakadiaris. Adversarial representation learning for text-to-image matching. In *ICCV*, 2019. 2
- [41] Zhiyin Shao, Xinyu Zhang, Meng Fang, Zhifeng Lin, Jian Wang, and Changxing Ding. Learning granularity-unified representations for text-to-image person re-identification. In *ACM MM*, 2022. 1, 2, 3, 8
- [42] Zhiyin Shao, Xinyu Zhang, Changxing Ding, Jian Wang, and Jingdong Wang. Unified pre-training with pseudo texts for text-to-image person re-identification. In *ICCV*, 2023. 1, 2, 5, 7, 8
- [43] Xiujun Shu, Wei Wen, Haoqian Wu, Keyu Chen, Yiran Song, Ruizhi Qiao, Bo Ren, and Xiao Wang. See finer, see more: Implicit modality alignment for text-based person retrieval. In *ECCV*, 2022. 2, 8
- [44] Wentao Tan, Changxing Ding, Pengfei Wang, Mingming Gong, and Kui Jia. Style interleaved learning for generalizable person re-identification. *IEEE TMM*, 2024. 1
- [45] Hao Tang, Chengcheng Yuan, Zechao Li, and Jinhui Tang. Learning attention-guided pyramidal features for few-shot fine-grained recognition. *Pattern Recognition*. 2
- [46] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 3
- [47] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. 3
- [48] Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu, and Lijuan Wang. Git: A generative image-to-text transformer for vision and language. *TMLR*, 2022. 3
- [49] Liwei Wang, Yin Li, and Svetlana Lazebnik. Learning deep structure-preserving image-text embeddings. In *CVPR*, 2016. 2
- [50] Pengfei Wang, Changxing Ding, Wentao Tan, Mingming Gong, Kui Jia, and Dacheng Tao. Uncertainty-aware clustering for unsupervised domain adaptive object re-identification. *IEEE TMM*, 2023. 1
- [51] Yuyu Wang, Chunjuan Bo, Dong Wang, Shuang Wang, Yunwei Qi, and Huchuan Lu. Language person search with mutually connected classification loss. In *ICASSP*. IEEE, 2019. 2
- [52] Zhe Wang, Zhiyuan Fang, Jun Wang, and Yezhou Yang. Vitaa: Visual-textual attributes alignment in person search by natural language. In *ECCV*, 2020. 1, 2, 8
- [53] Zijie Wang, Aichun Zhu, Jingyi Xue, Xili Wan, Chao Liu, Tian Wang, and Yifeng Li. Caibc: Capturing all-round information beyond color for text-based person retrieval. In *ACM MM*, 2022. 2, 8
- [54] Zijie Wang, Aichun Zhu, Jingyi Xue, Xili Wan, Chao Liu, Tian Wang, and Yifeng Li. Look before you leap: Improving text-based person retrieval by learning a consistent cross-modal common manifold. In *ACM MM*, 2022. 2, 8
- [55] Yushuang Wu, Zizheng Yan, Xiaoguang Han, Guanbin Li, Changqing Zou, and Shuguang Cui. Lapscore: language-guided person search via color reasoning. In *ICCV*, 2021. 8
- [56] Ziqiang Wu, Bingpeng Ma, Hong Chang, and Shiguang Shan. Refined knowledge transfer for language-based person search. *TMM*, 2023. 2
- [57] Yi Xie, Jianqing Zhu, Huanqiang Zeng, Canhui Cai, and Lixin Zheng. Learning matching behavior differences for compressing vehicle re-identification models. In *VCIP*, 2020. 1
- [58] Yi Xie, Fei Shen, Jianqing Zhu, and Huanqiang Zeng. View-point robust knowledge distillation for accelerating vehicle re-identification. *EURASIP J ADV SIG PR*, 2021. 1
- [59] Yi Xie, Hanxiao Wu, Fei Shen, Jianqing Zhu, and Huanqiang Zeng. Object re-identification using teacher-like and light students. In *BMVC*, 2021. 2
- [60] Yi Xie, Huaidong Zhang, Xuemiao Xu, Jianqing Zhu, and Shengfeng He. Towards a smaller student: Capacity dynamic distillation for efficient image retrieval. In *CVPR*, 2023. 1
- [61] Shuanglin Yan, Neng Dong, Jun Liu, Liyan Zhang, and Jinhui Tang. Learning comprehensive representations with richer self for text-to-image person re-identification. In *ACM MM*, 2023. 1, 8
- [62] Shuanglin Yan, Hao Tang, Liyan Zhang, and Jinhui Tang. Image-specific information suppression and implicit local alignment for text-based person search. *TNNLS*, 2023. 2

- [63] Aiyuan Yang, Bin Xiao, Bingning Wang, Borong Zhang, Chao Yin, Chenxu Lv, Da Pan, Dian Wang, Dong Yan, Fan Yang, et al. Baichuan 2: Open large-scale language models. *arXiv preprint arXiv:2309.10305*, 2023. 3
- [64] Shuyu Yang, Yinan Zhou, Zhedong Zheng, Yaxiong Wang, Li Zhu, and Yujiao Wu. Towards unified text-based person retrieval: A large-scale multi-attribute and language search benchmark. In *ACM MM*, 2023. 1, 2, 3, 5, 7, 8
- [65] Shukang Yin, Chaoyou Fu, Sirui Zhao, Ke Li, Xing Sun, Tong Xu, and Enhong Chen. A survey on multimodal large language models. *arXiv preprint arXiv:2306.13549*, 2023. 2, 3
- [66] Ying Zhang and Huchuan Lu. Deep cross-modal projection learning for image-text matching. In *ECCV*, 2018. 1, 2, 8
- [67] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 2023. 3
- [68] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *NeurIPS*, 2023. 3
- [69] Zhedong Zheng, Liang Zheng, Michael Garrett, Yi Yang, Mingliang Xu, and Yi-Dong Shen. Dual-path convolutional image-text embeddings with instance loss. *ACM TOMM*, 2020. 1, 8
- [70] Aichun Zhu, Zijie Wang, Yifeng Li, Xili Wan, Jing Jin, Tian Wang, Fangqiang Hu, and Gang Hua. Dssl: Deep surroundings-person separation learning for text-based person retrieval. In *ACM MM*, 2021. 1, 2, 5
- [71] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023. 3