

Towards Unified Text-based Person Retrieval: A Large-scale Multi-Attribute and Language Search Benchmark

Shuyu Yang
Xi'an Jiaotong University
Xi'an, China
ysy653@stu.xjtu.edu.cn

Yaxiong Wang*
Hefei University of Technology
Hefei, China
wangyx15@stu.xjtu.edu.cn

Yinan Zhou
Xi'an Jiaotong University
Xi'an, China
zyn13572297710@stu.xjtu.edu.cn

Li Zhu*
Xi'an Jiaotong University
Xi'an, China
zhuli@mail.xjtu.edu.cn

Zhedong Zheng
National University of Singapore
Singapore
zdzheng@nus.edu.sg

Yujiao Wu
Peng Cheng Laboratory
Shenzhen, China
yujiaowu111@gmail.com

ABSTRACT

In this paper, we introduce a large Multi-Attribute and Language Search dataset for text-based person retrieval, called MALS, and explore the feasibility of performing pre-training on both attribute recognition and image-text matching tasks in one stone. In particular, MALS contains 1,510,330 image-text pairs, which is about 37.5× larger than prevailing CUHK-PEDES, and all images are annotated with 27 attributes. Considering the privacy concerns and annotation costs, we leverage the off-the-shelf diffusion models to generate the dataset. To verify the feasibility of learning from the generated data, we develop a new joint Attribute Prompt Learning and Text Matching Learning (APTML) framework, considering the shared knowledge between attribute and text. As the name implies, APTML contains an attribute prompt learning stream and a text matching learning stream. (1) The attribute prompt learning leverages the attribute prompts for image-attribute alignment, which enhances the text matching learning. (2) The text matching learning facilitates the representation learning on fine-grained details, and in turn, boosts the attribute prompt learning. Extensive experiments validate the effectiveness of the pre-training on MALS, achieving state-of-the-art retrieval performance via APTML on three challenging real-world benchmarks. In particular, APTML achieves a consistent improvement of +6.96%, +7.68%, and +16.95% Recall@1 accuracy on CUHK-PEDES, ICFG-PEDES, and RSTPReid datasets by a clear margin, respectively. The dataset, model, and code are available at <https://github.com/Shuyu-XJTU/APTM>.

CCS CONCEPTS

• Information systems → Multimedia databases; Multimedia and multimodal retrieval; • Computing methodologies → Multi-task learning.

*Corresponding authors.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '23, October 29–November 3, 2023, Ottawa, ON, Canada

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0108-5/23/10...\$15.00

<https://doi.org/10.1145/3581783.3611709>

KEYWORDS

Text-based Person Retrieval; Image-Text Alignment; Attribute Prompt Learning; Synthetic Data; Multi-Attribute Recognition

ACM Reference Format:

Shuyu Yang, Yinan Zhou, Zhedong Zheng, Yaxiong Wang, Li Zhu, and Yujiao Wu. 2023. Towards Unified Text-based Person Retrieval: A Large-scale Multi-Attribute and Language Search Benchmark. In *Proceedings of the 31st ACM International Conference on Multimedia (MM '23)*, October 29–November 3, 2023, Ottawa, ON, Canada. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3581783.3611709>



Figure 1: Selected image-text pairs from our MALS (top) and CUHK-PEDES (bottom). We could observe that the visual gap between synthetic data and real one is relatively small. In MALS, image-text pairs match almost as well as manual annotation, although there are some flaws occasionally. It is worth noting that images in MALS are high-fidelity with rich and diverse variations in terms of pose, appearance, background, etc. (Best viewed when zooming in.)

1 INTRODUCTION

Given the pedestrian description, text-based person retrieval aims to locate the person of interest from a large pool of candidates [28]. Compared to conventional image-based person retrieval [31, 68, 74], text-based person retrieval provides an intuitive way to form queries. Such techniques can be widely applied to promote public safety, such as locating lost children in large areas like airports. However, as a type of cross-modal learning task, text-based person retrieval harvests little benefits from large-scale cross-modal pretraining. The reasons stem from two aspects: 1) **Lack of Data**.

Due to privacy concerns, we usually can not collect enough data for the current data-hungry deeply-learned models. 2) **Lack of High-quality Annotation.** The language annotation process is also tedious and inevitably introduces annotator biases. As a result, the sentences are usually quite short, which can not comprehensively describe the characteristic of the target person [12, 29].

In response to these problems, we propose to construct a synthetic image-text dataset, borrowing the power of the off-the-shelf diffusion models and the image caption model. In this way, we could generate unlimited images and acquire high-quality annotations automatically. Furthermore, to make the synthetic data beneficial for real-world language-based person retrieval, there are still two challenges that need to be addressed: **(1) Realism of synthetic image-text pairs.** The visual disparity between synthetic and real-world image-text pairs constitutes a major challenge in the construction of a meaningful text-pedestrian benchmark. For text inputs, we utilize descriptions derived from real-world text-based person data to guide the diffusion models. Therefore, the generated images closely resemble those found in the real world. We further apply a post-processing mechanism as a supplementary step (See Section 3.) to further refine the synthetic images and rectify any remaining discrepancies. **(2) Diversity of annotations (sentences & attributes).** To generate a large-scale cross-modal dataset, the human-annotated description will inevitably be used multiple times, resulting in poor text diversity. To handle this limitation, we employ an off-the-shelf caption generation model to augment the descriptions for each synthetic image. Besides, we propose an automatic attribute extraction mechanism that mines the key attributes from the descriptions to further enrich the annotations.

In this way, we collect a new large-scale cross-modal dataset, *i.e.*, **Multi-Attribute and Language Search** dataset for person retrieval (MALS) with rich annotations. It is worth noting that while diffusion models have been recently studied for data augmentation [2, 45, 48], these works mainly focus on coarse-grained category recognition benchmarks such as ImageNet [11] and EuroSAT [19]. Differently, person retrieval requires a more detailed representation since the variations among individuals are comparatively small. Therefore, the MALS dataset focuses on providing fine-grained details, which is crucial for text-based person retrieval tasks. Furthermore, extensive experiments verify that the knowledge learned from MALS is also scalable to real-world applications in terms of both text-based person retrieval and pedestrian attribute recognition tasks.

To verify the value of the collected dataset, we introduce an **Attribute Prompt Learning and Text Matching Learning (APT)** framework for text-based person retrieval. As shown in Figure 2, the proposed APTM comprises three modules, the image encoder, text encoder, and cross encoder. We utilize text to acquire attribute annotation by the proposed **Explicit Matching (EM)** and **Implicit Extension (IE)** mechanism, and further map attributes to a set of **Attribute Prompts**. **Image-text contrastive learning (ITC)** and **image-attribute contrastive learning (IAC)** act on the embeddings of feature encoders, while **image-text matching (ITM)**, **image-attribute matching (IAM)**, **masked language modeling (MLM)**, and **masked attribute prompt modeling (MAM)** are imposed on the respective predictions from the cross encoder. The above constraints are jointly optimized during pre-training to learn an effective model. In summary, we highlight the contributions of this paper as follows:

- We observe that data scarcity largely compromises text-based person retrieval. Therefore, we introduce a new large-scale multi-attribute and language search benchmark, called MALS. Compared with the existing datasets, such as CUHK-PEDES, our benchmark contains about $37.5\times$ images with rich attribute annotations. (See Table 1.)
- Based on MALS, we also introduce a new joint Attribute Prompt Learning and Text Matching Learning (APT) framework, to facilitate the representation learning. As the name implies, we explicitly leverage both the attribute recognition task and the text-based person retrieval task to regularize the model training. The two tasks are complementary and benefit each other.
- The proposed approach achieves a competitive recall rate on three challenging real-world benchmarks including CUHK-PEDES, ICFG-PEDES, and RSTPReid. Besides, we observe that the text matching task facilitates attribute recognition as well. Fine-tuning APTM on PA-100K, *i.e.*, a prevalent pedestrian attribute recognition dataset, we obtain competitive performance 82.58% mA.

2 RELATED WORK

Language-based Person Search. Text-to-image person retrieval is more challenging than general cross-modal retrieval tasks because of its fine-grained nature. Existing efforts can be classified as cross-modal attention-based [29, 46, 49, 62] approaches or cross-modal attention-free approaches [9, 12, 61, 76] depending on the alignment strategy. To align representations from both modalities in a shared feature space, the cross-modal attention-free approaches build various model structures or objective functions [76]. In contrast, cross-modal attention-based approaches require pair-wise inputs and encourage building cross-modal correspondences between regions and words or regions and phrases with more interactions between modalities. It is worth noting that both strategies have their advantages as well as disadvantages. In general, cross-modal attention-free techniques are more efficient. More specifically, their complexity is $O(M+N)$ for M gallery and N queries. The complexity of cross-modal attention-based approaches, in comparison, rises to $O(MN)$ due to the pair-wise inputs. Nonetheless, these techniques typically result in noticeably superior retrieval performance. It is because cross-modal attention-based approaches reduce modality gaps more effectively with more cross-modality communication in an early stage. In this paper, we leverage cross-modal attention-free features to quickly find the candidates and then deploy the attention-based module to refine the final ranking score.

Attribute-based Person Re-identification. Attribute-based person re-identification [16, 30, 32, 33, 39, 56, 70] aims to identify individuals across different cameras or time periods based on their attributes, such as clothing color, gender, height, *etc.*, rather than relying solely on visual appearance. One of the earliest works on pedestrian attributes is by Lin *et al.* [32], who propose a framework for person re-identification using color, texture, and contour clues. In particular, Lin *et al.* extract discriminative features from each pedestrian image and train several attribute classifiers. Following this work, Han *et al.* [16] further propose to fuse part features with attribute attentions, while He *et al.* [18] study to jointly train multiple attribute classifiers in a coherent manner. In contrast to the fixed horizontal splitting, attribute localization is also studied in [47]. To encourage the interaction between attributes, both

Datasets	MALS (Ours)	CUHK-PEDES [29]	ICFG-PEDES [12]	RSTPReid [79]	Lin <i>et al.</i> [32]	PA-100K [35]
#Images	1,510,330	40,206	54,522	20,505	32,668	100,000
Data Source	Automatic Synthesis	Market [73] & Duke [43], <i>etc.</i>	MSMT-17 [65]	MSMT-17 [65]	Market [73] & Duke [43]	Manual Collection
#Avg Texts/Image	1	2	1	2	-	-
#Avg Text Length	26.96	23.54	37.2	25.8	-	-
Surrounding	Indoor & Outdoor	Indoor & Outdoor	Indoor & Outdoor	Indoor & Outdoor	Outdoor	Outdoor
Resolution	531 × 208	246 × 90	378 × 142	546 × 219	128 × 64	225 × 85
Annotation	Sentence & Attribute	Sentence	Sentence	Sentence	Attribute	Attribute
#Attribute	27	-	-	-	27	26

Table 1: Comparison between MALS and other real-world datasets for text-based person retrieval and pedestrian attribute recognition. Current datasets typically collect images from existing person re-ID datasets and manually provide corresponding natural language descriptions or attribute annotations. In contrast, MALS leverages generative models to generate a large-scale dataset including 1.5M image-text pairs. For each benchmark, the table shows the number of images, data source, the average texts per image, average text length, the main surrounding and the average resolution of images, types of annotations as well as the number of attributes.

Attribute Category	Name	Label
gender	gender	female(0), male(1)
age	age	young(0), teenager(1), adult(2), old(3)
length hair	hair	short hair(0), long hair(1)
wearing hat	hat	yes(0), no(1)
carrying backpack	backpack	yes(0), no(1)
carrying handbag	handbag	yes(0), no(1)
carrying bag	bag	yes(0), no(1)
sleeve length	sleeve	long sleeve(0), short sleeve(1)
length of lower-body	length_lower	long lower body clothing(0), short(1)
type of lower-body	type_lower	dress(0), pants(1)
color of upper-body	black, white, red, purple, yellow, blue, green, gray	(0), (1), (2), (3), (4), (5), (6), (7)
color of lower-body	black, white, purple, yellow, blue, green, pink, gray, brown	(0), (1), (2), (3), (4), (5), (6), (7), (8)

Table 2: Attribute space consists of 27 attributes. Here we show the attribute category, the name in the annotation file, and the available label choices.

Nguyen *et al.* [40] and Tang *et al.* [54] propose to build a graph representation of the attributes for each person, where nodes represent attribute embeddings and edges represent correlations between them. In addition to traditional attributes, Wang *et al.* [58] leverages both appearance and personality traits to learn representations of both visual appearance and personality traits and combine them for re-identification. Finally, there are several works that make attributes more robust against occlusions or pose variations. For instance, Jing *et al.* [24] propose a multi-modal framework that fuses attribute-based features with pose-based features to enhance re-identification accuracy under challenging conditions. In this paper, we also leverage robust attribute learning to facilitate text-based person retrieval. We find that attribute learning is complementary to image-text matching, and vice versa.

3 BENCHMARK

Existing text-based person retrieval datasets [12, 29, 79] typically collect pedestrian images from existing person re-identification datasets and manually annotate corresponding text descriptions. However, such practice greatly limits the scale and diversity due to annotation costs and privacy concerns, as shown in Table 1. The great success of recent diffusion models [3, 20, 44] inspires us to collect pedestrian images from the synthetic domain. There are two primary advantages: (1) Comparing to 3D Game Engine [51, 59, 67]

or Generative Adversarial Networks (GANs) [23, 55, 75, 77], diffusion models have shown a strong and stable ability to synthesize images with high authenticity to text, significantly reducing the gap between synthetic and real data. (2) Using synthetic pedestrian images also circumvents privacy concerns. The construction of our benchmark consists of the following steps:

Image-text Pair Generation. We utilize the off-the-shelf diffusion model, ImaginAlry [13] which could generate new pedestrian images. To make the generated samples reasonable as well as close to the real-world pedestrian images, we employ the textual descriptions of the CUHK-PEDES [29] dataset and the ICFG-PEDES [12] dataset as prompts. We feed the prompts into ImaginAlry and collect the corresponding synthetic images, resulting in a pair of aligned samples. To ensure the generation of high-quality full-body pedestrian images with controlled variability, we set the image size as 576×384 and adjust the random seed to get the high-quality samples. By randomizing the noise during inference, massive and diverse pedestrian images are collected.

Post-Processing. Due to the lack of fine-grained and controllable generation capabilities of the text-to-image generation model, many generated images cannot meet the requirement of training the pedestrian retrieval networks. Two main issues stand: (1) the low-quality images, including grayscale and blur images. To overcome this weakness, we simply sort images by file size and delete images whose size is smaller than $24k$ to filter out blurred images. Then we compute the mean variance of the difference between the 3 channels of every image and remove images whose mean variance is less than a presetting threshold. (2) the noisy images, *e.g.*, multiple persons in one image, only part of a person, and no person. To remedy this issue, we apply OpenPose [4, 5, 50, 66] to detect human key points and filter out the undesired person images. We also leverage the detected key points as a tight bounding box to re-crop the samples. With the above steps, we acquire the final pedestrian images.

Caption Calibration. The prompts used to generate images are the straightforward choice to serve as the text descriptions. However, this fashion would result in poor diversity of the textural descriptions, since multiple images usually share the same text. To cope with this problem, we leverage the cross-modal model, BLIP [27] to produce more fitting captions for every synthetic image and form the final image-text pairs.

Attribute Annotation. The associated attributes often highlight the key characteristics of both image and text samples, and many works of text-based person retrieval indicate the potential of attribute for performance improvement [9, 12, 49]. Inspired by this, we further augment our MALS with the attribute annotation, so that a more informative and comprehensive benchmark can be constructed. Considering the cost of manual annotation, we obtain the attribute annotation in an automatic manner. We first define the attribute space in the same way as Market-1501 Attribute [32], and then propose two mechanisms to obtain attributes, Explicit Matching (EM) and Implicit Extension (IE). EM deploys the correspondence of specific attributes based on keywords in the text, such as the word "man" corresponding to the attribute "gender: male". IE assigns corresponding attribute candidates based on distinctive features that are not mentioned in the text, such as allocating samples that do not mention "hat" in their descriptions to the attribute "hat: no". Finally, 27 different types of attributes are collected, as shown in Table 2.

MALS Benchmark. Following the above steps, a high-fidelity, diverse and large-scale benchmark for the text-based person retrieval task is built. As shown in Figure 1, we can observe the quality of visual images and textual sentences are comparable with CUHK-PEDES. Compared with existing text-based person retrieval datasets in Table 1, MALS has the following advantages:

- **High-fidelity Images:** Compared with the images with poor lighting and blur texture, collected from surveillance cameras, images of MALS are of higher quality benefiting from the ability of the diffusion model (see Figure 1), which means that the synthetic images are more visually appealing and realistic.
- **Diversity:** MALS contains a wide range of variations in the images, including but not limited to variations in background, viewpoint, occlusion, clothing, and body pose. Thanks to our caption calibration step, the associated textual descriptions are also diverse enough. Therefore, MALS can support us to train robust models that generalize well to new and unseen data in vision tasks, language tasks, and vision-language tasks.
- **Fewer Privacy Concerns:** Unlike several traditional benchmarks of text-based person retrieval capturing images without consent, the samples of our MALS are all synthetic images generated by the off-the-self stable diffusion model, which avoids ethical and legal issues.
- **Large-scale Pairs:** MALS contains 1.5M image-text pairs (see Table 1), while existing datasets usually provide no more than 100k of aligned image-text. This magnitude of the dataset enables a comprehensive pre-training study.
- **Rich Annotations:** Each image-text pair in MALS is annotated with appropriate attribute labels, indicating that MALS is not only effective for text-image matching and attribute prompt learning, but also explores the feasibility of pre-training for both attribute recognition and image-text matching in one stone.

4 METHOD

We leverage MALS as a pre-training dataset and devise a new simple joint Attribute Prompt Learning and Text Matching Learning (APTM) framework, as shown in Figure 2. The overall pipeline is typically divided into two steps, *i.e.*, pre-training and fine-tuning.

During pre-training, we perform Attribute Prompt Learning (APL) and Text Matching Learning (TML) to learn the common knowledge of text-based person retrieval and pedestrian attribute recognition. In the second step, the parameters are further optimized toward a specific downstream task. In this section, we elaborate on the details of the pre-training stage, as we mainly study the benefits of our MALS for pre-training.

4.1 APTM Architecture

As shown in Figure 2, APTM is a multi-task framework, containing one image-attribute stream and one image-text stream with weight-shared encoders and MLP-based headers. In particular, the framework comprises three encoders, *i.e.*, Image Encoder (E_I), Text Encoder (E_T), Cross Encoder (E_C), and two MLPs-based headers. Before pre-training, we utilize text to acquire attribute annotation by Explicit Matching and Implicit Extension mechanism and then map attributes to a set of Attribute Prompts as one of the inputs of the image-attribute stream. During pre-training, the image-text stream and the image-attribute stream are jointly trained. We deploy Random Mask to generate masked text and masked attribute prompts, and then the Image Encoder maps the image into embedding V and the Text Encoder extracts different text representations by encoding Text, Masked Text, Attribute Prompts, and Masked Attribute Prompts separately, denoted as L , \hat{L} , L_A and \hat{L}_A , respectively. In the task of ITC and ITM, V is paired with L , while in the context of IAC and IAM, V is paired with L_A . Further, V is also fed into cross encoder with \hat{L} or \hat{L}_A for MLM or MAM task.

Image Encoder. Without loss of generality, we deploy Swin Transformer (Swin-B) [36] as Image Encoder (E_I). Given an image (I) of resolution of 384×128 , we split it into N^I non-overlapping patches with a patch size of 32×32 , where $N^I = 48$. Then, these patches are linearly embedded and passed into the transformer layers of E_I , yielding a set of high-dimensional embeddings V , the [CLS] embedding v^{cls} is taken as the representation of the entire image.

Text Encoder. Following existing works [49], we intuitively employ BERT [25] as Text Encoder (E_T) for a fair comparison. Specifically, the text (T) is first tokenized as $N^T + 1$ tokens and fed into the first 6 layers of BERT. The output text embeddings $\{l^{cls}, l^1, l^2, \dots, l^{N^T}\}$ is denoted as L , where l^i ($i \in [1, N^T]$) represents the embedding of the i_{th} text token. The embedding of the [CLS] token, *i.e.*, l^{cls} is treated as the whole text representation.

Cross Encoder. The cross encoder is to fuse the image and text representations to perform the prediction tasks. Specifically, we adopt the last 6 layers of BERT as Cross Encoder (E_C). As shown in Figure 2, the image and text embeddings are fed into E_C and fused by the cross attention mechanism to capture their semantic relationship. Finally, the joint representation can be obtained: $C = \{c^{cls}, c^1, c^2, \dots, c^{N^T}\}$.

4.2 Attribute Prompt Learning

Motivations. Attributes often emphasize crucial characteristics of pedestrian images, such as gender and hair, which are vital for performing cross-modal alignment and distinguishing between candidates. Moreover, synthetic and real descriptions exhibit a considerable overlap in attribute keywords, leading us to believe that accentuating the similar attribute space can also alleviate the

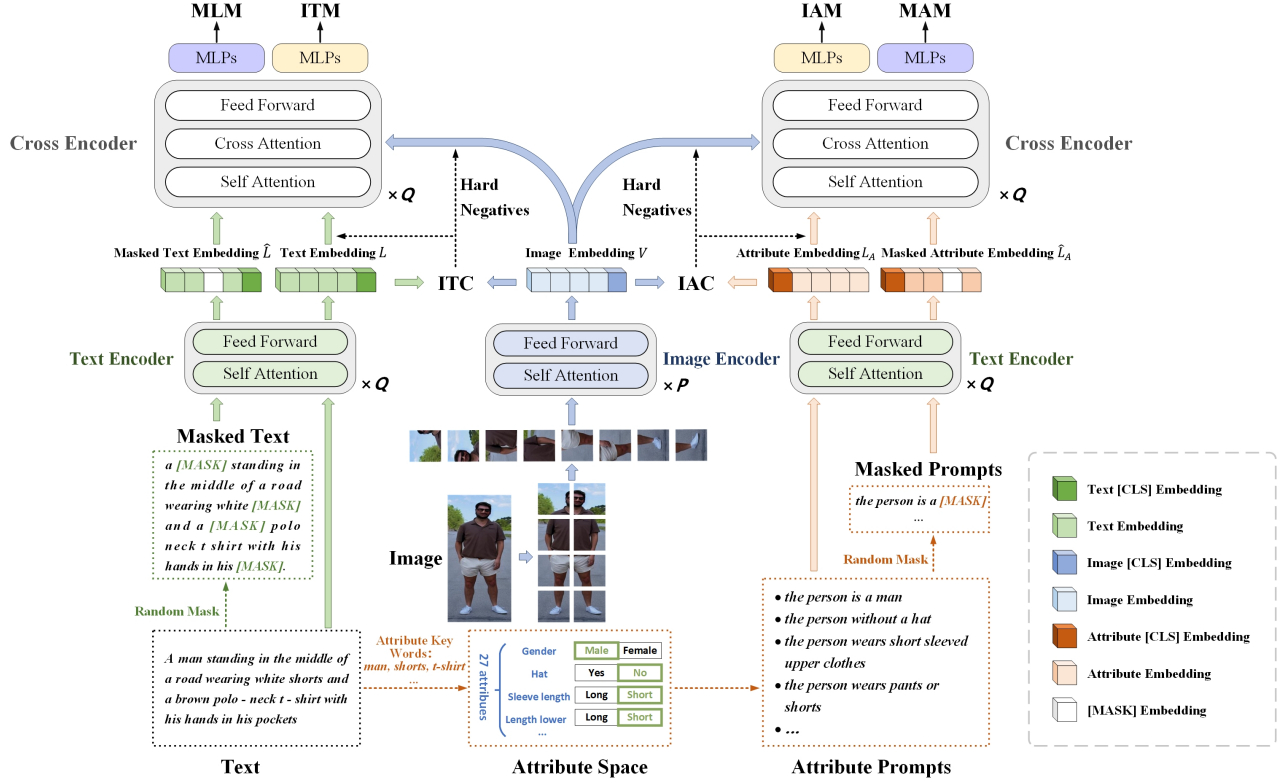


Figure 2: Overview of the proposed Attribute Prompt Learning and Text Matching Learning (APT) framework for pre-training on MALS. APTM framework contains one image-attribute stream and one image-text stream with weight-shared encoders. In particular, the framework comprises three encoders, i.e., Image Encoder (E_I), Text Encoder (E_T), Cross Encoder (E_C), and two MLPs-based headers. The Image Encoder and Text Encoder are to produce the embeddings of the image and text, respectively, while the cross encoder seeks to fuse the image and text embeddings for the subsequent predictions.

domain gap. To better leverage the attribute information for image-attribute alignment, we have opted not to rely on conventional classifier-based multi-attribute learning methods. Instead, we convert attribute labels into attribute prompts with prompt templates, as illustrated in Figure 2. We then align the attribute prompts with the corresponding image, which forms the fundamental basis of our Attribute Prompt Learning. Drawing inspiration from cross-modal Learning, we utilize Image-Attribute Contrastive Learning (IAC), Image-Attribute Matching (IAM), and Masked Attribute Language Modeling (MAM) to effectively align images with their attributes. **Image-Attribute Contrastive Learning (IAC)** concentrates on mastering the ability to differentiate between positive and negative pairs. Given a set of attribute texts $\{T_a^k\}^{2|A|}$ in a mini-batch, $k \in [1, 2|A|]$, where A is the attribute set of 27 binary attributes. For an image I , if any of its attribute labels correspond with the attribute set, we consider the corresponding attribute text and I as a matched (image, attribute prompt) pair. If not, they are considered unmatched. As exemplified in Figure 2, "the person is a man" is a matched attribute prompt of the image while "the person is a woman" is not. We denote the set of all matched (image, attribute prompt) pairs in a mini-batch as B_a . The matching score between an image I and its paired attribute prompt T_a is estimated as follows:

$$S_{i2a}(I) = \frac{\exp(s(F_I, F_{T_a})/\tau)}{\exp(s(F_I, F_{T_a})/\tau) + \exp(s(F_I, F_{\tilde{T}_a})/\tau)}, \quad (1)$$

where \tilde{T}_a is the opposite attribute prompt of T_a , which is constructed by replacing the true attribute as the false one, e.g., $man \Rightarrow woman$, τ is a learnable temperature parameter, F_I and F_{T_a} are the mapped features of their respective [CLS] embedding by two different FCs, $s(\cdot, \cdot)$ is the cosine similarity. Finally, the formulation of the IAC loss is presented below:

$$\mathcal{L}_{iac} = -\frac{1}{|B_a|} \sum_{(I, T_a) \in B_a} \log S_{i2a}(I). \quad (2)$$

Image-Attribute Matching Learning (IAM) aims to predict whether the input image and attribute prompt are matched. In particular, IAM is specified as a binary classification problem to facilitate the image-attribute alignment: the positive sample is the paired image-attribute prompt, while the unpaired is the negative one. Mathematically, assume $|B|$ images are sampled in a mini-batch, 5 attribute prompts are randomly constructed to form $5|B|$ (image, attribute prompt) pairs, denote as \tilde{B}_a . Subsequently, the image-attribute prompt tuples are passed through the Cross Encoder to get the [CLS] embedding c^{cls} , their matching score is given by an MLP with Sigmoid activation: $p^{\text{match}}(I, T_a) = \text{Sigmoid}(\text{MLP}(c^{cls}))$, the IAM loss is defined as:

$$\mathcal{L}_{iam} = -\frac{1}{|\tilde{B}_a|} \sum_{(I, T_a) \in \tilde{B}_a} (y_a^{\text{match}} \log p^{\text{match}}(I, T_a) + (1 - y_a^{\text{match}})(1 - \log p^{\text{match}}(I, T_a))), \quad (3)$$

where y_a^{match} is 1 if (I, T_a) is matched, 0 otherwise.

Masked Attribute Language Modeling (MAM) seeks to predict the masked words using the matched (image, attribute prompt) as a clue. To this end, we first adopt the following strategies to randomly mask the $2|A|$ attribute prompts: 1) mask out the text tokens with a probability of 25%; Among the masked tokens, 2) 10% and 80% is replaced with random tokens and the special token [MASK], respectively; 3) 10% remain unchanged. Then, given an image-attribute prompt pair (I, T_a) in B_a , we obtain corresponding masked attribute prompt \hat{T}_a following the aforementioned strategies. Then, (I, \hat{T}_a) is input into encoders to get the output of E_C : $\hat{C} = \{\hat{c}^{\text{cls}}, \hat{c}^1, \hat{c}^2, \dots, \hat{c}^{N^T}\}$. If \hat{t}_a^j is the masked token in \hat{T}_a , $j \in [1, N^T]$, its prediction probability is given by an MLP with Softmax activation: $p_j^{\text{mask}}(I, \hat{T}_a) = \text{Softmax}(\text{MLP}(\hat{c}^j))$. Finally, the MAM loss is defined as follows:

$$\mathcal{L}_{\text{mam}} = \mathbb{E}_{\hat{t}_a^j \sim \hat{T}_a; (I, \hat{T}_a) \sim \hat{B}_a} H(y_j^{\text{mask}}, p_j^{\text{mask}}(I, \hat{T}_a)), \quad (4)$$

where y_j^{mask} is a one-hot distribution in which the ground-truth token \hat{t}_a^j has the probability of one, and \hat{B}_a is the $2|A|$ (image, masked attribute prompt) pairs of the mini-batch. The overall APL loss is: $\mathcal{L}_{\text{APL}} = \frac{1}{3}(\mathcal{L}_{\text{iac}} + \mathcal{L}_{\text{iam}} + \mathcal{L}_{\text{mam}})$. To prevent overfitting, label smoothing is further employed. Typically, we apply a random noise perturbation to y_a^{match} , which remedies overconfident predictions. **Why APL Works Better.** In comparison to the Classification-based Multi-Attribute Learning (CMAL) approaches, APL has three distinct advantages: 1) Explicit emphasis on the attributions. Naïve classification-based practices implicitly highlight the key attributes through a classification procedure, whereas APL explicitly constructs the attribute prompt, which leads to more effective learning than implicit classification procedures. 2) More informative inputs. APL introduces information-rich inputs by constructing supplementary attribute prompts, providing richer information for cross-modal alignment learning. In contrast, traditional CMAL only utilizes a classification loss and introduces no auxiliary information. 3) Greater flexibility for framework augmentation. Thanks to the constructed attribute prompts, APL enables powerful cross-modal learning objectives such as image-text contrastive learning (ITC), Image-text matching (ITM), and masked language modeling (MLM) to be equipped after modification to be attribute-oriented, resulting in increased potential for performance improvement. During experiments, APL outperforms several naïve CMAL variants, which well verifies the superiority of APL.

4.3 Text matching Learning

As a type of cross-modal retrieval problem, the core of text-based person retrieval is to align the text query and the image candidates. Hence, we also incorporate the tasks of Image-Text Contrastive Learning (ITC), Image-Text Matching Learning (ITM), and Masked Language Modeling (MLM) to impose the alignment constraints.

Image-Text Contrastive Learning (ITC) focuses on learning to differentiate between positive and negative pairs. In our case, it is intuitive to treat the paired image-text (I, T) as the positive sample, while the unmatched image-text is the negative pair. Formally, we randomly sample $|B|$ pairs of images and text in each mini-batch. Similar to Eq. 1, given a matched pair (I, T) , we initially extract their respective representations F_I and F_T . The matching score is then estimated as follows:

$$S_{\text{I2t}}(I) = \frac{\exp(s(F_I, F_T)/\tau)}{\sum_{i=1}^{|B|} \exp(s(F_I, F_{T_i})/\tau)}, \quad (5)$$

Similarly, given the text, the matching score of the paired image $S_{\text{T2i}}(T)$ can be calculated. Finally, the ITC loss is formulated as:

$$\mathcal{L}_{\text{itc}} = -\frac{1}{2|B|} \sum_{(I, T) \in B} (\log S_{\text{I2t}}(I) + \log S_{\text{T2i}}(T)), \quad (6)$$

where B is the data set of the mini-batch.

Image-Text Matching Learning (ITM) targets to predict whether the input image and the text are matched, analogous to IAM. Nevertheless, randomly sampling an unpaired item (text or image) is overly facile for the classification. Therefore, we employ a hard example mining strategy. For each text in a mini-batch, we sample its hard negative image according to the similarity of $S_{\text{I2t}}(T)$, *i.e.*, pick the unpaired image whose similarity is the highest as the hard negative. We also sample one hard negative text for each image in a similar manner. Finally, $|B|$ positive image-text pairs and $2|B|$ negative pairs, denoted as \hat{B} , will pass through the Cross Encoder and one MLP with Sigmoid activation. Following these steps, the ITM loss can be calculated similarly as described in Eq. 3.

Masked Language Modeling (MLM) endeavors to predict the masked words using the image and text clue. Given an image-text pair (I, T) in B , we obtain corresponding masked text \hat{T} following the same masking strategies as MAM. Subsequently, (I, \hat{T}) are passed through the encoders to obtain embeddings. The MLM loss \mathcal{L}_{mlm} is analogously imposed following Eq. 4. Given the above optimization objectives, the full pre-training loss is formulated as: $\mathcal{L} = \mathcal{L}_{\text{itc}} + \mathcal{L}_{\text{itm}} + \mathcal{L}_{\text{mlm}} + \beta \mathcal{L}_{\text{APL}}$, where β denotes the APL loss weight, and we empirically set 0.8.

5 EXPERIMENT

5.1 Experimental Setup

Datasets. We evaluate our approach on three public text-based person retrieval datasets and one pedestrian attribute dataset, *i.e.*, **CUHK-PEDES** [29], **RSTPReid** [79], **ICFG-PEDES** [12] and **PA-100K** [35]. In particular, CUHK-PEDES [29] includes 80,440 description phrases and 40,206 photos of 13,003 people, while RSTPReid [79] comprises 20,505 images of 4,101 people and is created by compiling MSMT17 [65] data. ICFG-PEDES [12] is also incubated from MSMT17 and has 54,522 images of 4,102 individuals. To make a fair comparison, the data splits of three text-based person retrieval datasets keep the same as the previous works [14, 69]. PA-100K [35] is constructed by 100,000 pedestrian images from 598 real outdoor scenes. Every image is labeled by 26 attributes, and the standard splits are adopted for performance evaluation [35].

Implementation Details. We pre-train APTM with Pytorch on 4 NVIDIA A100 GPUs for 32 epochs, and the mini-batch size is 150. We adopt the AdamW [38] optimizer with a weight decay of 0.01. The learning rate is decayed from $1e^{-4}$ to $1e^{-5}$ following a linear schedule, after a warm-up schedule beginning at $1e^{-5}$ in the first 2,600 steps. Every image input is resized to 384×128 . Random horizontal flipping, RandAugment [10] and random erasing [78] are employed for image augmentation. APTM takes text with no more than 56 tokens as input. During pre-training, the Image Encoder is initialized with Swin Transformer_{base} [37], while the Text Encoder and Cross Encoder are initialized by the first and the last 6 layers of

Method	R1	R5	R10	mAP
CNN-RNN [42]	8.07	-	32.47	-
GNA-RNN [29]	19.05	-	53.64	-
PWM-ATH [7]	27.14	49.45	61.02	-
GLA [6]	43.58	66.93	76.2	-
Dual Path [76]	44.40	66.26	75.07	-
CMPM+CMPC [71]	49.37	-	79.21	-
MIA [41]	53.10	75.00	82.90	-
A-GANet [34]	53.14	74.03	81.95	-
VITAA [60]	55.97	75.84	83.52	51.60
IMG-Net [63]	56.48	76.89	85.01	-
CMAAM [1]	56.68	77.18	84.86	-
HGAN [72]	59.00	79.49	86.62	-
NAFS [15]	59.94	79.86	86.70	54.07
DSSL [79]	59.98	80.41	87.56	-
MGEL [57]	60.27	80.01	86.74	-
SSAN [12]	61.37	80.15	86.73	-
NAFS [15]	61.50	81.19	87.51	-
TBPS [17]	61.65	80.98	86.78	-
TIPCB [9]	63.63	82.82	89.01	-
LBUL [62]	64.04	82.66	87.22	-
CAIBC [61]	64.43	82.87	88.37	-
AXM-Net [14]	64.44	80.52	86.77	58.73
SRFC [52]	64.88	83.02	88.56	-
LGUR [46]	65.25	83.12	89.00	-
CFine [69]	69.57	85.93	91.15	-
TextReid (RN50) [17]	61.65	80.98	86.78	58.29
TextReid (30K) [17]	61.81	80.39	86.90	-
IVT (Baseline) [49]	55.75	75.68	84.13	-
IVT (4M) [49]	65.59	83.11	89.21	-
IVT (MALS) [49]	66.10	83.79	89.46	-
Baseline	66.44	84.92	90.76	59.19
APTM (Ours)	76.53	90.04	94.15	66.91

Table 3: Performance Comparison on CUHK-PEDES.

BERT_{base} [25], respectively. Therefore, there are 214.5M trainable parameters in APTM. After pre-training, the model is fine-tuned on the downstream datasets for 30 epochs. The learning rate is set as $1e^{-4}$ and is warmed up in the first 3 epochs. Then we apply a linear scheduler to gradually decay the learning rate.

5.2 Comparison with Existing Methods

We adapt APTM to downstream text-based person retrieval tasks and pedestrian attribute recognition tasks. Following previous practices [49, 69], we report Recall@1,5,10, and mAP for text-based person retrieval to compare the results. For the attribute recognition task, accuracy (Acc), precision (Prec), recall rate (Rec), and F1 value (F1) are adopted to evaluate the performance.

Text-based Person Retrieval. We evaluate APTM on CUHK-PEDES, RSTPReid, and ICFG-PEDES datasets and optimize ITC, ITM, and MLM loss during finetuning. Besides image data augmentation mentioned in pretraining, we adopt EDA[64] for text data augmentation and set the mini-batch size as 120. In reference, for each text query, we first compute its cosine similarity with all images and take the top-128 image candidates. Then we calculate the matching probability between the text query and every selected image candidate for ranking. The proposed method has achieved the SOTA recall rate on all three datasets. Specifically, our model has surpassed 6.96% recall@1 rate on CUHK-PEDES, compared to the second-best method 69.57% (See Table 3). Similarly, as shown in Table 4 and Table 5, we could observe that our method arrives at

Method	R1	R5	R10	mAP
DSSL [79]	32.43	55.08	63.19	-
LBUL [62]	45.55	68.20	77.85	-
IVT [49]	46.70	70.00	78.80	-
CAIBC [61]	47.35	69.55	79.00	-
CFine [69]	50.55	72.50	81.60	-
Baseline	47.20	70.85	80.00	36.36
APTM (Ours)	67.50	85.70	91.45	52.56

Table 4: Performance Comparison on RSTPReid.

Method	R1	R5	R10	mAP
Dual Path [76]	38.99	59.44	68.41	-
CMPM+CMPC [71]	43.51	65.44	74.26	-
MIA [41]	46.49	67.14	75.18	-
SCAN [26]	50.05	69.65	77.21	-
VITAA [60]	50.98	68.79	75.78	-
SSAN [12]	54.23	72.63	79.53	-
IVT [49]	56.04	73.60	80.22	-
LGUR [46]	59.02	75.32	81.56	-
CFine [69]	60.83	76.55	82.42	-
Baseline	57.49	75.84	82.60	32.41
APTM (Ours)	68.51	82.99	87.56	41.22


Table 5: Performance Comparison on ICFG-PEDES.

Method	mA	Acc	Prec	Rec	F1
HP-net [35]	74.21	72.19	82.97	82.09	82.53
strongBaseline [22]	79.38	78.56	89.41	84.78	86.55
ALM [53]	80.68	77.08	84.21	88.84	86.46
RethinkPAR [21]	81.61	79.45	87.66	87.59	87.62
Baseline (Image only)	71.68	54.51	60.47	83.60	70.18
Baseline (wo MAM)	80.43	79.91	89.26	86.49	87.85
Baseline	81.49	79.89	88.59	87.09	87.83
APTM (Ours)	82.58	80.17	88.31	87.84	88.07

Table 6: Performance Comparison on PA-100K. Baseline (Image only) denotes only finetuning the Image Encoder, while Baseline (wo MAM) removes MAM loss. Baseline refers to training the model without pretraining on MALS.

67.50% and 68.51% R1 on RSTPReid and ICFG-PEDES, respectively. Furthermore, we compare two traditional methods, TextReID [17] and IVT [49], with APTM based on MALS. We use ResNet50 as the visual backbone of TextReID and use 30k image-text pairs from MALS for TextReID pretraining since TextReID uses instance loss. IVT also explores the pretraining on 4M general image-text pairs, while MALS only consists of 1.5M image-text pairs. As shown in Table 3, after pretraining on MALS, the performance of all three methods has improved.

Pedestrian Attribute Recognition. Pedestrian attribute recognition aims at mining the attributes of target people when given a pedestrian image. We apply the attribute prompt learning part of APTM to predict the attributes of images from PA-100K. Similar to MALS, we construct the attribute prompts for PA-100K and finetune the model. During inference, we compute the matching probability between every image and every pair of attribute prompts for ranking. An attribute prompt with a higher matching probability means the image is more relevant to the corresponding attribute. Our method achieves competitive results as shown in Table 6. Baseline refers to training the model without pretraining on MALS. Baseline (Image only) denotes only finetuning the Image Encoder and its following MLPs, which are used to predict the attribute label. Baseline (wo MAM) does not optimize MAM loss. The results among the three baselines and APTM indicates the rationality of our APTM.



(a) Ablation study on the pre-training data number.

Method	L_{TML}	L_{V1}	L_{V2}	L_{APL}	R1
M1	✓				69.62
M2	✓	✓			69.46
M3	✓		✓		69.64
APTМ	✓			✓	71.38

(b) Ablation study on attribute-oriented objectives.

Figure 3: Ablation Study on the pre-training data scale and the optimization objectives in our APTM. (a) We apply 0M, 0.3M, 0.6M, 0.9M, 1.2M, and 1.5M data pairs to pre-train, and then report the fine-tuned recall rate on three datasets respectively. We can observe that the performance is consistently improved as the data scale increases. (b) L_{TML} refers to the sum of ITC loss, ITM loss, and MLM loss, and L_{APL} denotes APL loss. L_{V1} and L_{V2} are losses of two naïve CMAL variants.

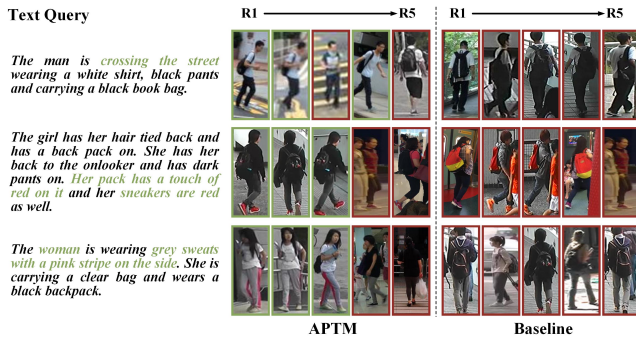


Figure 4: Qualitative text-to-image retrieval results of APTM and baseline, placing in descending order from right to left based on similarity. The green boxes indicate the correct matches, and the images in the red boxes are the wrong matches. The green texts highlight the details that our results successfully match.

APTМ obtain 0.97% improvement on mA compared with the results of RethinkPAR [21]. A recent work SOLIDER [8] reports 86.37% mA. Since SOLIDER [8] adopts a more powerful backbone, we do not include the result of SOLIDER for a fair comparison.

5.3 Ablation Study

Effectiveness of Pre-Training. Table 3 compares the performance on the CUHK-PEDES dataset, where the "Baseline" means the APTM without pre-training. We can observe that our baseline, reaching 66.44%, 84.92%, and 90.76% on R1, R5, and R10, respectively, is a competitive method. Pre-training APTM on MALS leads to improvement of 10.09%, 5.12%, and 3.39% on Recall@1, 5 and 10. Similar results could be observed on RSTPReid and ICFG-PEDES, reported in Table 4 and Table 5. To intuitively show the benefits of pre-training, three qualitative results of APTM and baseline are shown in Figure 4, indicating the superiority of APTM.

The Impact of Pre-training Scale. In generic vision and language pre-training tasks, the scale of the training dataset usually plays an important role. A larger amount of pre-training data often means better performance. To thoroughly study the effectiveness of MALS, we further explore the impact of the data scale during pre-training. Specifically, we respectively adopt 0, 0.3M, 0.6M, 0.9M, 1.2M and,

1.5M data of MALS to pre-train 32 epochs and then evaluate the finetuned performance on CUHK-PEDES, ICFG-PEDES and RSTPReid. The results are compared in Figure 3a, as the data scale increases, the recall tends to improve as well. From 0 to 0.3M, fine-tuning performance on three datasets increases noticeably, while from 0.3M to 1.5M, the rate of improvement gradually diminishes.

Effectiveness of APL Loss. We also conduct an ablation study to investigate how to leverage the attribute annotation, as shown in Table 3b. All compared model variants are pre-trained on 0.03M data of MALS and then finetuned on CUHK-PEDES. We adopt Recall@1 as an evaluation measure. First, we evaluate the effectiveness of APL loss, *i.e.*, M1, APTM. The results show that pretraining without APL loss hurts performance. Furthermore, we replace APL with several naïve CMAL variants separately, and report its finetuning performance on CUHK-PEDES in Table 3b: (1)Method V1: Image embedding and Text Embedding are used to give the prediction of attributes by mapping the Embedding to low-dimensional features, separately. The BCE Loss is adopted as the objective function. (2)Method V2: Use the joint representation of the image-text pair to predict attributes by mapping the Embedding to low-dimensional features. The attribute classification loss is BCE Loss, too. In Table 3b, compared with M2 and M3, APL outperforms both of them, which verifies the superiority of APTM.

6 CONCLUSION

We introduce MALS, a new large-scale benchmark for multi-attribute recognition and language-based person search. Our benchmark comprises 1,510,330 image-text pairs with rich attribute annotations, which is about 37.5 times larger than widely-used CUHK-PEDES. Extensive experiments verify that pretraining on MALS is scalable to real-world scenarios. To regularize the model training, we propose to jointly learn from the two complementary tasks, *i.e.*, text-based person retrieval and pedestrian attribute recognition. On three public benchmarks, including CUHK-PEDES, ICFG-PEDES, and RSTPReid, our approach has achieved a competitive recall rate. We hope our work could contribute to the community with a new viewpoint on unified text-based person retrieval.

Acknowledgments This work is supported by the National Key Research and Development Project, China (No. 2019YFB2102501).

REFERENCES

- [1] Surbhi Aggarwal, Venkatesh Babu Radhakrishnan, and Anirban Chakraborty. 2020. Text-based person search via attribute-aided matching. In *WACV*. 2617–2625.
- [2] Shekoofeh Azizi, Simon Kornblith, Chitwan Saharia, Mohammad Norouzi, and David J Fleet. 2023. Synthetic Data from Diffusion Models Improves ImageNet Classification. *arXiv preprint arXiv:2304.08466* (2023).
- [3] Tim Brooks, Aleksander Holynski, and Alexei A Efros. 2023. Instructpix2pix: Learning to follow image editing instructions. In *CVPR*. 18392–18402.
- [4] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh. 2019. OpenPose: Realtime Multi-Person 2D Pose Estimation using Part Affinity Fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2019).
- [5] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. 2017. Realtime multi-person 2d pose estimation using part affinity fields. In *CVPR*. 7291–7299.
- [6] Dapeng Chen, Hongsheng Li, Xihui Liu, Yantao Shen, Jing Shao, Zejian Yuan, and Xiaogang Wang. 2018. Improving deep visual representation for person re-identification by global and local image-language association. In *ECCV*. 54–70.
- [7] Tianlang Chen, Chenliang Xu, and Jiebo Luo. 2018. Improving text-based person search by spatial matching and adaptive threshold. In *WACV*. 1879–1887.
- [8] Weihua Chen, Xianzhe Xu, Jian Jia, Hao Luo, Yaohua Wang, Fan Wang, Rong Jin, and Xiuyu Sun. 2023. Beyond Appearance: a Semantic Controllable Self-Supervised Learning Framework for Human-Centric Visual Tasks. In *CVPR*. 15050–15061.
- [9] 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* 494 (2022), 171–181.
- [10] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. 2020. Randaugment: Practical automated data augmentation with a reduced search space. In *CVPR workshop*. 702–703.
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In *CVPR*. 248–255. <https://doi.org/10.1109/CVPR.2009.5206848>
- [12] 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* (2021).
- [13] Bryce Drennan. 2022. *imaginAIry*. <https://github.com/brycedrennan/imaginAIry>. Accessed: 2022-05-04.
- [14] Ammarah Farooq, Muhammad Awais, Josef Kittler, and Syed Safwan Khalid. 2022. AXM-Net: Implicit Cross-Modal Feature Alignment for Person Re-identification. In *AAAI*, Vol. 36. 4477–4485.
- [15] Chenyang Gao, Guanyu Cai, Xinyang Jiang, Feng Zheng, Jun Zhang, Yifei Gong, Pai Peng, Xiaowei Guo, and Xing Sun. 2021. Contextual non-local alignment over full-scale representation for text-based person search. *arXiv preprint arXiv:2101.03036* (2021).
- [16] Kai Han, Jianyuan Guo, Chao Zhang, and Mingjian Zhu. 2018. Attribute-aware attention model for fine-grained representation learning. In *ACM MM*. 2040–2048.
- [17] Xiao Han, Sen He, Li Zhang, and Tao Xiang. 2021. Text-Based Person Search with Limited Data. In *BMVC*.
- [18] Keke He, Zhanxiong Wang, Yanwei Fu, Rui Feng, Yu-Gang Jiang, and Xiangyang Xue. 2017. Adaptively weighted multi-task deep network for person attribute classification. In *ACM MM*. 1636–1644.
- [19] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. 2019. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* (2019).
- [20] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. 2022. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626* (2022).
- [21] Jian Jia, Houjing Huang, Xiaotang Chen, and Kaiqi Huang. 2021. Rethinking of pedestrian attribute recognition: A reliable evaluation under zero-shot pedestrian identity setting. *arXiv preprint arXiv:2107.03576* (2021).
- [22] Jian Jia, Houjing Huang, Wenjie Yang, Xiaotang Chen, and Kaiqi Huang. 2020. Rethinking of pedestrian attribute recognition: Realistic datasets with efficient method. *arXiv preprint arXiv:2005.11909* (2020).
- [23] Yiqi Jiang, Weihua Chen, Xiuyu Sun, Xiaoyu Shi, Fan Wang, and Hao Li. 2021. Exploring the quality of gan generated images for person re-identification. In *ACM MM*. 4146–4155.
- [24] Ya Jing, Chenyang Si, Junbo Wang, Wei Wang, Liang Wang, and Tieniu Tan. 2020. Pose-guided multi-granularity attention network for text-based person search. In *AAAI*, Vol. 34. 11189–11196.
- [25] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, Vol. 1. 2.
- [26] Kuang-Huei Lee, Xi Chen, Gang Hua, Houdong Hu, and Xiaodong He. 2018. Stacked cross attention for image-text matching. In *ECCV*. 201–216.
- [27] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. BLIP: Bootstrapping Language-Image Pre-training for Unified Vision-Language Understanding and Generation. In *ICML*.
- [28] Shuang Li, Tong Xiao, Hongsheng Li, Wei Yang, and Xiaogang Wang. 2017. Identity-aware textual-visual matching with latent co-attention. In *ICCV*. 1890–1899.
- [29] Shuang Li, Tong Xiao, Hongsheng Li, Bolei Zhou, Dayu Yue, and Xiaogang Wang. 2017. Person search with natural language description. In *CVPR*. 1970–1979.
- [30] Shuzhao Li, Huimin Yu, and Roland Hu. 2020. Attributes-aided part detection and refinement for person re-identification. *Pattern Recognition* 97 (2020), 107016.
- [31] Shengcai Liao, Yang Hu, Xiangyu Zhu, and Stan Z Li. 2015. Person re-identification by local maximal occurrence representation and metric learning. In *CVPR*. 2197–2206.
- [32] Yutian Lin, Liang Zheng, Zhedong Zheng, Yu Wu, Zhilan Hu, Chenggang Yan, and Yi Yang. 2019. Improving person re-identification by attribute and identity learning. *Pattern recognition* 95 (2019), 151–161.
- [33] Hefei Ling, Ziyang Wang, Ping Li, Yuxuan Shi, Jiazong Chen, and Fuhao Zou. 2019. Improving person re-identification by multi-task learning. *Neurocomputing* 347 (2019), 109–118.
- [34] Jiawei Liu, Zheng-Jun Zha, Richang Hong, Meng Wang, and Yongdong Zhang. 2019. Deep adversarial graph attention convolution network for text-based person search. In *ACM MM*. 665–673.
- [35] Xihui Liu, Haiyu Zhao, Maoqing Tian, Lu Sheng, Jing Shao, Shuai Yi, Junjie Yan, and Xiaogang Wang. 2017. Hydraplus-net: Attentive deep features for pedestrian analysis. In *ICCV*. 350–359.
- [36] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*. 10012–10022.
- [37] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. 2021. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. In *ICCV*.
- [38] Ilya Loshchilov and Frank Hutter. 2018. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations (ICLR)*.
- [39] Jinghao Luo, Yaohua Liu, Changxin Gao, and Nong Sang. 2019. Learning what and where from attributes to improve person re-identification. In *ICIP*. IEEE, 165–169.
- [40] Binh X Nguyen, Binh D Nguyen, Tuong Do, Erman Tjiputra, Quang D Tran, and Anh Nguyen. 2021. Graph-based person signature for person re-identifications. In *CVPR*. 3492–3501.
- [41] Kai Niu, Yan Huang, Wanli Ouyang, and Liang Wang. 2020. Improving description-based person re-identification by multi-granularity image-text alignments. *IEEE Transactions on Image Processing (TIP)* 29 (2020), 5542–5556.
- [42] Scott Reed, Zeynep Akata, Honglak Lee, and Bernt Schiele. 2016. Learning deep representations of fine-grained visual descriptions. In *CVPR*. 49–58.
- [43] Ergys Ristani, Francesco Solera, Roger Zou, Rita Cucchiara, and Carlo Tomasi. 2016. Performance measures and a data set for multi-target, multi-camera tracking. In *ECCV workshop*. Springer, 17–35.
- [44] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *CVPR*. 10684–10695.
- [45] Mert Bulent Sariyildiz, Karteek Alahari, Diane Larlus, and Yannis Kalantidis. 2023. Fake it till you make it: Learning transferable representations from synthetic ImageNet clones. In *CVPR*.
- [46] 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 *ACM MM*. 5566–5574.
- [47] Yuxuan Shi, Zhen Wei, Hefei Ling, Ziyang Wang, Jialie Shen, and Ping Li. 2020. Person retrieval in surveillance videos via deep attribute mining and reasoning. *IEEE Transactions on Multimedia* 23 (2020), 4376–4387.
- [48] Jordan Shipard, Arnold Wiliem, Kien Nguyen Thanh, Wei Xiang, and Clinton Fookes. 2023. Diversity is Definitely Needed: Improving Model-Agnostic Zero-shot Classification via Stable Diffusion. [arXiv:2302.03298 \[cs.CV\]](https://arxiv.org/abs/2302.03298)
- [49] Xiujun Shu, Wei Wen, Haoqian Wu, Keyu Chen, Yiran Song, Ruizhi Qiao, Bo Ren, and Xiao Wang. 2023. See finer, see more: Implicit modality alignment for text-based person retrieval. In *ECCV workshop*.
- [50] Tomas Simon, Hanbyul Joo, Iain Matthews, and Yaser Sheikh. 2017. Hand Key-point Detection in Single Images using Multiview Bootstrapping. In *CVPR*.
- [51] Xiaoxiao Sun and Liang Zheng. 2019. Dissecting person re-identification from the viewpoint of viewpoint. In *CVPR*. 608–617.
- [52] Wei Suo, Mengyang Sun, Kai Niu, Yiqi Gao, Peng Wang, Yanning Zhang, and Qi Wu. 2022. A Simple and Robust Correlation Filtering Method for Text-Based Person Search. In *ECCV*. Springer, 726–742.
- [53] Chufeng Tang, Lu Sheng, Zhaoxiang Zhang, and Xiaolin Hu. 2019. Improving pedestrian attribute recognition with weakly-supervised multi-scale attribute-specific localization. In *ICCV*. 4997–5006.
- [54] Geyu Tang, Xingyu Gao, and Zhenyu Chen. 2022. Learning semantic representation on visual attribute graph for person re-identification and beyond. *ACM Transactions on Multimedia Computing, Communications and Applications* (2022).

- [55] Hao Tang, Dan Xu, Gaowen Liu, Wei Wang, Nicu Sebe, and Yan Yan. 2019. Cycle in cycle generative adversarial networks for keypoint-guided image generation. In *ACM MM*. 2052–2060.
- [56] Chiat-Pin Tay, Sharmili Roy, and Kim-Hui Yap. 2019. Aanet: Attribute attention network for person re-identifications. In *CVPR*. 7134–7143.
- [57] Chengji Wang, Zhiming Luo, Yaojin Lin, and Shaozi Li. 2021. Text-based person search via multi-granularity embedding learning. In *IJCAI*. 1068–1074.
- [58] Jingya Wang, Xiadian Zhu, Shaogang Gong, and Wei Li. 2018. Transferable joint attribute-identity deep learning for unsupervised person re-identification. In *CVPR*. 2275–2284.
- [59] Yanan Wang, Shengcai Liao, and Ling Shao. 2020. Surpassing real-world source training data: Random 3d characters for generalizable person re-identification. In *ACM MM*. 3422–3430.
- [60] Zhe Wang, Zhiyuan Fang, Jun Wang, and Yezhou Yang. 2020. Vitaa: Visual-textual attributes alignment in person search by natural language. In *ECCV*. 402–420.
- [61] Zijie Wang, Aichun Zhu, Jingyi Xue, Xili Wan, Chao Liu, Tian Wang, and Yifeng Li. 2022. CAIBC: Capturing All-round Information Beyond Color for Text-based Person Retrieval. In *ACM MM*. 5314–5322.
- [62] Zijie Wang, Aichun Zhu, Jingyi Xue, Xili Wan, Chao Liu, Tian Wang, and Yifeng Li. 2022. Look Before You Leap: Improving Text-based Person Retrieval by Learning A Consistent Cross-modal Common Manifold. In *ACM MM*. 1984–1992.
- [63] Zijie Wang, Aichun Zhu, Zhe Zheng, Jing Jin, Zhouxin Xue, and Gang Hua. 2020. IMG-Net: inner-cross-modal attentional multigranular network for description-based person re-identification. *Journal of Electronic Imaging (JEI)* 29, 4 (2020), 043028.
- [64] Jason Wei and Kai Zou. 2019. EDA: Easy Data Augmentation Techniques for Boosting Performance on Text Classification Tasks. In *EMNLP-IJCNLP*. 6382–6388.
- [65] Longhui Wei, Shiliang Zhang, Wen Gao, and Qi Tian. 2018. Person transfer gan to bridge domain gap for person re-identification. In *CVPR*. 79–88.
- [66] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. 2016. Convolutional pose machines. In *CVPR*.
- [67] Suncheng Xiang, Dahong Qian, Mengyuan Guan, Binjie Yan, Ting Liu, Yuzhuo Fu, and Guanjie You. 2021. Less is more: Learning from synthetic data with fine-grained attributes for person re-identification. *ACM Transactions on Multimedia Computing, Communications and Applications* (2021).
- [68] Tong Xiao, Shuang Li, Bochao Wang, Liang Lin, and Xiaogang Wang. 2016. End-to-end deep learning for person search. *arXiv preprint arXiv:1604.01850* 2, 2 (2016), 4.
- [69] Shuanglin Yan, Neng Dong, Liyan Zhang, and Jinhui Tang. 2022. CLIP-Driven Fine-grained Text-Image Person Re-identification. *arXiv preprint arXiv:2210.10276* (2022).
- [70] Yan Zhang, Xusheng Gu, Jun Tang, Ke Cheng, and Shoubiao Tan. 2019. Part-based attribute-aware network for person re-identification. *IEEE Access* 7 (2019), 53585–53595.
- [71] Ying Zhang and Huchuan Lu. 2018. Deep cross-modal projection learning for image-text matching. In *ECCV*. 686–701.
- [72] Kecheng Zheng, Wu Liu, Jiawei Liu, Zheng-Jun Zha, and Tao Mei. 2020. Hierarchical Gumbel Attention Network for Text-based Person Search. In *ACM MM*.
- [73] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. 2015. Scalable Person Re-Identification: A Benchmark. In *ICCV*.
- [74] Wei-Shi Zheng, Shaogang Gong, and Tao Xiang. 2011. Person re-identification by probabilistic relative distance comparison. In *CVPR*. IEEE, 649–656.
- [75] Zhedong Zheng, Xiaodong Yang, Zhiding Yu, Liang Zheng, Yi Yang, and Jan Kautz. 2019. Joint discriminative and generative learning for person re-identification. In *CVPR*.
- [76] Zhedong Zheng, Liang Zheng, Michael Garrett, Yi Yang, Mingliang Xu, and Yi-Dong Shen. 2020. Dual-path convolutional image-text embeddings with instance loss. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* 16, 2 (2020), 1–23.
- [77] Zhedong Zheng, Liang Zheng, and Yi Yang. 2017. Unlabeled samples generated by gan improve the person re-identification baseline in vitro. In *ICCV*. 3754–3762.
- [78] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. 2020. Random erasing data augmentation. In *AAAI*, Vol. 34. 13001–13008.
- [79] 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 *ACM MM*. 209–217.