

Text-Based Person Search with Limited Data

Xiao Han^{1,2}

xiao.han@surrey.ac.uk

Sen He^{1,2}

sen.he@surrey.ac.uk

Li Zhang³

lizhangfd@fudan.edu.cn

Tao Xiang^{1,2}

t.xiang@surrey.ac.uk

¹ CVSSP

University of Surrey

² iFlyTek-Surrey Joint Research Centre
on Artificial Intelligence

³ School of Data Science
Fudan University

Abstract

Text-based person search (TBPS) aims at retrieving a target person from an image gallery with a descriptive text query. Solving such a fine-grained cross-modal retrieval task is challenging, which is further hampered by the lack of large-scale datasets. In this paper, we present a framework with two novel components to handle the problems brought by limited data. Firstly, to fully utilize the existing small-scale benchmarking datasets for more discriminative feature learning, we introduce a cross-modal momentum contrastive learning framework to enrich the training data for a given mini-batch. Secondly, we propose to transfer knowledge learned from existing coarse-grained large-scale datasets containing image-text pairs from drastically different problem domains to compensate for the lack of TBPS training data. A transfer learning method is designed so that useful information can be transferred despite the large domain gap. Armed with these components, our method achieves new state of the art on the CUHK-PEDES dataset with significant improvements over the prior art in terms of Rank-1 and mAP. Our code is available at <https://github.com/BrandonHanx/TextReID>.

1 Introduction

Text-based person search (TBPS) [20] is the problem of retrieving a target person from an image gallery with a descriptive text query. It is more flexible compared to image-based person search when the query image is difficult to obtain. It has thus gained increasing attention in the research community [23, 37, 38]. TBPS has various potential applications such as video surveillance and personal photo album search.

Despite the existing efforts, TBPS is still far from being solved. One of the reasons is that it is intrinsically challenging as a fine-grained cross-modal retrieval task, where all images belong to the same category, *i.e.*, pedestrian. This contrasts with the more widely studied generic image-text retrieval task [4, 21, 32, 33]. The fine-grained nature dictates that more discriminative features must be learned to distinguish visual cues and textual attributes. This

is thus the focus of existing TBPS methods. Specifically, prior works [1, 5, 11, 17, 30] typically use a two-stream architecture for fast inference, where both streams are initialized from backbones pre-trained on large-scale unimodal data, *e.g.*, ResNet [14] and BERT [9]. For the purpose of learning more discriminative features, many methods [1, 11, 17, 36, 40] also take advantage of multi-scale learning, where feature maps with different receptive fields and word/phrase/sentence embeddings are used for the visual and textual stream, respectively.

TBPS is also faced with a second challenge which has been largely ignored, that is, the lack of training data. Collecting a large-scale TBPS dataset and annotating it with highly fine-grained text descriptions is tedious and expensive. As a result, most existing TBPS dataset is orders of magnitude smaller than those coarse-grained generic image-text pair datasets [4, 21]. Having only limited data has clear negative effects on a TBPS model’s ability to learn discriminative cross-modal features for fine-grained retrieval.

Existing methods [1, 5, 11, 17, 30] are ill-equipped to address this limited data problem. More specifically, most of their learning objectives require the training data to be organized into positive and negative pairs. However, previous works construct negative pairs merely from a certain mini-batch, which does not make full use of the available TBPS data. Pre-training on larger image-text pair datasets is also an obvious option to compensate for the lack of training data. Nevertheless, the visual and textual streams in previous works are initialized from models that are separately pre-trained on unimodal data rather than image-text pairs. Information useful for cross-domain matching is thus not exploited. Cross-modal pre-training may have been attempted. However, as shown in this work, without a careful design, a naive pre-training then fine-tuning strategy would lead to negative transfer.

To overcome the learning with limited data problem, in this work, we propose a framework with two novel components for TBPS. Firstly, to fully utilize the existing small-scale benchmarking datasets for more discriminative feature learning, we introduce a cross-modal momentum contrastive learning (or CM-MoCo) framework to enrich the training data for a given mini-batch. CM-MoCo decouples the number of negative pairs with the mini-batch size to obtain more negative cross-modal counterparts for each image or description. To implement such a framework, in addition to the two gradient-updated encoders (query encoders), we introduce another two momentum-updated encoders (key encoders) for two modalities and maintain three different queues to store visual features, textual features, and identities from previous batches. Further, a contrastive loss is formulated in a cross-modal manner, which treats the features from query encoders, key encoders and queues as anchors, positive samples and negative samples, respectively. Secondly, a cross-modal transfer learning method is proposed to benefit from large-scale coarse-grained image-text pair datasets. Instead of the commonly used pre-training + fine-tuning strategy, we propose to freeze the text encoder of the pre-trained model to embed each word and then adopt one bidirectional GRU layer (Bi-GRU) [7] to contextualize words. Empirically, this transfer learning strategy can effectively prevent the negative transfer suffered by the naive full model transfer strategy (See Table 2 for experimental results).

The main contributions of this paper are: (1) A novel cross-modal momentum contrastive learning framework is proposed to better utilize the existing small-scale TBPS datasets. (2) To effectively transfer the knowledge learned from large-scale generic image-text pairs, we propose to perform cross-modal pre-training, but for the text modality, only word embedding is transferred. (3) Extensive experiments are conducted to show that our proposed framework outperforms existing methods on CUHK-PEDES [20] by large margins.

2 Related work

2.1 Text-based person search

Li *et al.* [20] first propose TBPS with a challenging dataset CUHK-PEDES and a baseline built upon a recurrent neural network with gated neural attention. Following this, many single-scale methods are proposed for better investigating the intra- and inter-modal fine-grained differences with the help of instance loss [41], cross-modal projection loss [39], adversarial loss [30] and cross-modal knowledge adaption [5].

Besides, some multi-scale methods are proposed to learn the semantic relevance between specific image regions and phrases/words in descriptions. Many works implement such architecture by making local image features attend to corresponding noun phrases and words through a variety of attention mechanisms [10, 11, 26, 35, 40]. Additionally, some works adopt side information to help align two modalities, *e.g.*, pose information [17], semantic segmentation maps [36] and attribute labels [1]. Most of these multi-scale architectures merely use global features during inference, because calculating the similarity between local features increases both the inference time and the offline features storage space.

For fast inference speed and less memory consumption, neither multi-scale architecture nor side information is involved in our method. Instead of manually designing more complicated network architectures or collecting more side information, from a more general perspective, we focus on a much more practical and under-studied problem in TBPS, *i.e.*, the scarcity of data. Being orthogonal to all previous methods, our method can be easily extended or integrated. Many solid experiments show our method is comparable and even better with other more complex methods.

2.2 Contrastive learning

The basic idea of contrastive learning is to map the original data into a latent feature space where the similarities between positive/negative pairs are maximized/minimized [13]. The instance discrimination is the most prevalent pretext task, whose positive pairs consist of two augmented views of the same instance, and the other pairs are defined to be negative. MoCo [15] and SimCLR [2] suggest that large quantities of data pairs are crucial to the performance of contrastive learning. Most recently, BYOL [12] and SimSiam [3] prove that negative pairs are unnecessary and the invariant observation of the same concept matters.

In this work, to fully exploit the available annotated dataset, we apply momentum contrastive learning to TBPS. There are two differences between our work and the classic instance discrimination contrastive learning framework: (1) Our task is identity-level rather than instance-level, because each identity has more than one image and description in the dataset. (2) In our task, the similarity is measured in a cross-modal manner.

2.3 Vision-language pre-training

With the advent of Transformer [34] and BERT [9], there has been a surging interest in applying self-supervised learning to multimodal tasks. This is usually done by pre-training on large-scale generic image-text pairs and then fine-tuning on downstream tasks. ViLBERT [24] and LXMERT [33] introduce the two-stream architecture, where two Transformers are applied to images and text independently followed by another Transformer for cross-modal fusion. In addition to that, many works [4, 18, 21, 31] adopt the single-stream architecture and achieve much better performance. In such an architecture, a single Transformer is applied to both images and text.

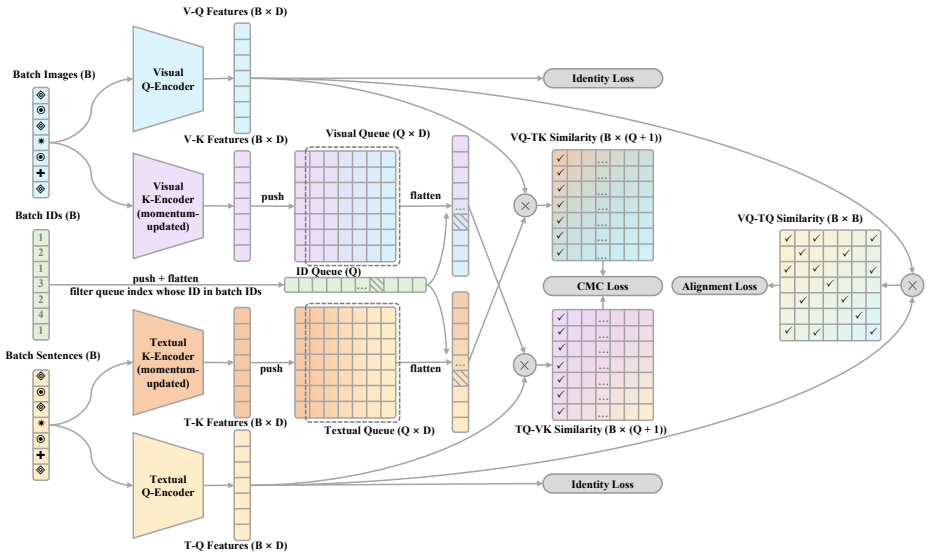


Figure 1: Illustrative diagram of our proposed architecture. The two Q-Encoders are gradient-updated while two K-Encoders are initialized from Q-Encoders and momentum-updated. The cross-modal contrastive (CMC) loss, alignment loss and identity loss are employed during training. The whole model is trained in an end-to-end manner and only two Q-Encoders are used for inference.

Although single stream models have achieved great success, its crucial component, cross-modal attention between two modalities, triggers the inevitable latency and significant computation during training and inference. To tackle this problem, CLIP [29] and BriVL [16] utilise larger datasets, larger batch size, and contrastive learning on the basis of two-stream architecture. LightningDOT [32] adopts a faster two-stream model as the main inference model and another stronger single-stream model as the re-ranker, achieving a satisfactory balance between accuracy and efficiency.

In this work, we incorporate pre-trained two-stream models into TBPS. We study how to effectively transfer the knowledge pre-trained on large-scale coarse-grained image-text pairs for fine-grained TBPS in spite of the big domain gap between them.

3 Methodology

Given a text query t , the goal of TBPS is to retrieve an image v that best matches the content in t from a gallery. The retrieval is successful if t and v share the same identity.

Our proposed framework is illustrated in Figure 1. It consists of two query encoders f_q^V and f_q^T (Visual and Textual Q-Encoder) along with two key encoders f_k^V , f_k^T (Visual and Textual K-Encoder), parameterized by θ_q^V , θ_q^T , θ_k^V , θ_k^T , respectively. During training, both the query and key encoders are used to process the input from its own modality. The outputs of the key encoders are pushed into queues which are used to construct negative pairs for contrastive learning. During inference, only two query encoders are used for feature extraction. The retrieval is done by first computing the cosine similarity between the query feature and the offline extracted features of all candidates in the gallery, and then selecting the candidate that has the highest similarity score.

In the following sections, we will first introduce our proposed cross-modal momentum contrastive learning pipeline (CM-MoCo) in Section 3.1, and then explain how we effectively transfer the knowledge learned from large-scale generic image-text pairs in Section 3.2.

3.1 Learning from limited TBPS data

Cross-modal momentum contrastive learning. One of the limitations to learn more discriminative features in previous work is caused by the limited negative pairs during the training stage. Note that MoCo [15] provides a mechanism of building dynamic queues decoupled with batch size, which makes it possible to learn from more negative samples beyond a certain batch. Inspired by this, we propose cross-modal momentum contrastive learning to make the best use of current TBPS data.

Concretely, given a batch of person images $V = \{v_1, \dots, v_B\}$, a batch of descriptions $T = \{t_1, \dots, t_B\}$ and their identities $ID = \{id_1, \dots, id_B\}$, we feed V and T into their corresponding query encoder and key encoder to obtain their normalized features:

$$\mathbf{V}^q = f_q^V(V), \mathbf{V}^k = f_k^V(V), \mathbf{T}^q = f_q^T(T), \mathbf{T}^k = f_k^T(T), \quad (1)$$

where $\mathbf{V}^q, \mathbf{V}^k, \mathbf{T}^q$, and $\mathbf{T}^k \in \mathbb{R}^{B \times D}$. $\mathbf{V}^k, \mathbf{T}^k$ and ID will be pushed into three queues, *i.e.*, visual queue (Q^V), textual queue (Q^T) and identity queue (Q^{ID}) for negative pair construction.

To learn discriminative cross-modal features, for each image query v_i , we define the cross-modal contrastive loss among its query feature \mathbf{V}_i^q (as anchor), its corresponding textual key feature \mathbf{T}_i^k (as positive key) and the keys stored in textual queue \tilde{Q}^T (as negative keys), where $\tilde{Q}^T = \{\mathbf{n} | \mathbf{n} \in Q^T \wedge Q^{ID}(\mathbf{n}) \notin ID\}$ indicating its identities are not in the current batch. In the meanwhile, we also apply the cross-modal contrastive learning in a symmetrical way when regarding each description t_i as a query. The overall cross-modal contrastive loss \mathcal{L}_{cmc} is computed as Equation 2, where τ_c denotes the tuneable temperature.

$$\mathcal{L}_{cmc} = - \sum_{i=1}^B \log \left[\frac{e^{(\mathbf{V}_i^q \mathbf{T}_i^k / \tau_c)}}{e^{(\mathbf{V}_i^q \mathbf{T}_i^k / \tau_c)} + \sum_{\mathbf{n} \in \tilde{Q}^T} e^{(\mathbf{V}_i^q \mathbf{n} / \tau_c)}} \right] - \sum_{i=1}^B \log \left[\frac{e^{(\mathbf{T}_i^q \mathbf{V}_i^k / \tau_c)}}{e^{(\mathbf{T}_i^q \mathbf{V}_i^k / \tau_c)} + \sum_{\mathbf{n} \in \tilde{Q}^V} e^{(\mathbf{T}_i^q \mathbf{n} / \tau_c)}} \right]. \quad (2)$$

After calculating the cross-modal contrastive loss, two query encoders are updated by the back propagation gradients. Following MoCo [15], the parameters of two key encoders, θ_k^V and θ_k^T , are updated by the rule given in Equation 3, where m is a momentum parameter.

$$\theta_k^V = m \cdot \theta_k^V + (1 - m) \cdot \theta_q^V, \quad \theta_k^T = m \cdot \theta_k^T + (1 - m) \cdot \theta_q^T. \quad (3)$$

With our proposed \mathcal{L}_{cmc} , each query sample is compared with a large number of negative key samples. It thus allows the model to learn more discriminative features for TBPS. To better understand CM-MoCo, please refer to the pseudocode in the supplementary material.

Joint training. Following previous works, we also incorporate widely used alignment loss \mathcal{L}_{align} [36] and identity loss \mathcal{L}_{id} [25, 41] into our end-to-end training pipeline (See details in the supplementary material). The overall loss \mathcal{L} is the summation of the three losses:

$$\mathcal{L} = \mathcal{L}_{cmc} + \mathcal{L}_{align} + \mathcal{L}_{id}. \quad (4)$$

Post-processing. In the inference stage, we adopt the cross-modal k -reciprocal rerank algorithm [11] to further improve the performance. The pair-wise rerank similarity is calculated by Jaccard Distance of k -nearest unimodal neighbors and k -nearest cross-modal neighbors, and then added to the original cosine similarity (See details in the supplementary material).

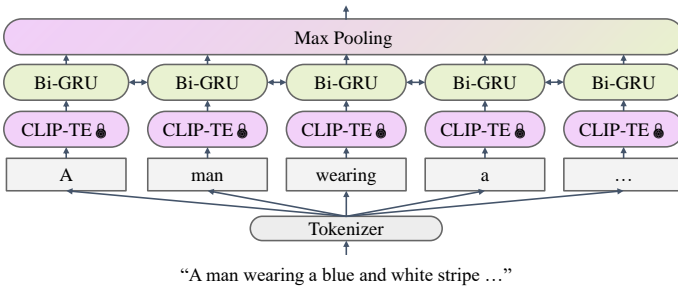


Figure 2: Illustrative diagram of our textual stream. We first tokenize a sentence into words, and then independently feed these words into a pre-trained frozen text backbone (*e.g.*, CLIP text encoder) to get word-type embeddings. On top of that, a Bi-GRU with a max pooling layer is used to contextualize all word embeddings.

3.2 Transferring knowledge from generic image-text pairs

A conventional way for TBPS is to initialize our visual and textual encoders with backbones separately pre-trained on unimodal data, *e.g.*, ResNet50 [14] pre-trained on ImageNet [8] and BERT [31] pre-trained on large corpora. However, this initialization brings a significant heterogeneous gap which is difficult to be bridged with the current limited data of TBPS. To tackle this issue, a straightforward way is to initialize our encoders pre-trained on large-scale generic image-text pairs, *e.g.*, MSCOCO captions [22], Flickr30k [28] and WIT [29], and then fine-tune the whole model for TBPS.

Unexpectedly, we empirically find this intuitive transfer strategy yields poor performance. This negative transfer is likely to be caused by the domain gap between the TBPS domain and that of the generic datasets. Such domain gap especially exists in the textual side, even under the unimodal pre-training scenario. Due to the highly fine-grained specialty of TBPS, the description sentences in TBPS are much longer than those in the generic data, and every word matters. However, text backbones pre-trained on generic data are likely to get more coarse-grained information and neglect some detailed words, thus negative transfer happens (See more discussions in the supplementary material). Nonetheless, we believe pre-training on generic data can still offer us more meaningful embeddings for each word because of large-scale contrastive learning.

To address this problem, we propose a transfer learning strategy with three alterations for the text stream while leaving the visual stream unchanged. Concretely, as illustrated in Figure 2, we take CLIP Text Encoder (CLIP-TE) [29] pre-trained on WIT as an example to demonstrate our alterations. Firstly, one of our alterations is feeding the whole sentence into CLIP-TE in a word-by-word manner to obtain word-type embeddings. For each word, its word-type embedding is represented by the [EOS] token from the last layer of CLIP-TE. Secondly, CLIP-TE is frozen in the whole training stage. In the actual implementation, an offline dictionary storing all word embeddings is computed in advance, and the frozen CLIP-TE is removed from the training process. Thirdly, to compensate the lack of sentence-level semantic information, we append a Bi-GRU [7] followed by max pooling to contextualize all word-type embeddings in a sentence. Empirically, as shown in Table 2, the model with our proposed transfer learning strategy yields a significant performance boost. This strategy allows the textual stream to effectively transfer the knowledge learned from large-scale generic image-text pairs.

4 Experiments

4.1 Experimental setup

Dataset. We conduct experiments on the CUHK-PEDES dataset [20], which is currently the only benchmark for TBPS. It contains 40,206 images of 13,003 different people, where each image has two descriptive sentences annotated by different people. As per standard, the dataset is split into 11,003 identities with 34,054 images in the training set, 1,000 identities with 3,074 images in the test set, and the remaining for the validation set. The average length of all sentences is 23 and the vocabulary size is 9408.

Implementation details. During training, we use random horizontally flipping, random crop with padding, and random erasing [25] as image data augmentation methods. Following previous works [11, 36, 40], the stride of the last block in the ResNet are set to 1 to increase the resolution of the final feature map. The feature dimension D for both modalities is set to 256. All images are resized to 384×128 . Our model is trained with Adam optimizer [19] for 80 epochs with an initial learning rate 1×10^{-4} , which is decayed by a factor 0.1 at the 40th epoch and 70th epoch, respectively. At the beginning, we spend 5 warm-up epochs linearly increasing the learning rate from 1×10^{-5} to 1×10^{-4} . Each mini-batch has 128 image-text pairs with 4 images/sentences for each identity. τ_c in \mathcal{L}_{cmc} is set to 0.07 and the momentum m is set to 0.999. Following NAFS [11], the number of nearest neighbors k used in rerank is set to 5. All experiments are conducted on one V100 GPU with Pytorch [27].

Evaluation protocol. As per standard, we evaluate our model in a bi-directional manner, where the performance is measured by Rank-K ($K=1, 5, 10$). Specifically, given a text/image query, Rank-K reports the percentage of successful searches among all searches, where each successful search retrieves at least one corresponding person correctly among the top K results. In addition, for a comprehensive evaluation, we also adopt the mean Average Precision (mAP) of all queries [10, 35, 40] as another retrieval criterion. Empirically, Rank-K can reflect models' accuracy on the first few retrieval results while mAP puts more emphasis on the order of the entire retrieval sequence predicted by the models.

4.2 Main results

We compare our method with most published works on TBPS. We integrate the proposed CM-MoCo and transfer the knowledge from generic data to finalize our model. For a comprehensive comparison, we instantiate our method with both CLIP ResNet50 and CLIP ResNet101. Please notice that we cannot reach 100% fair comparisons with all other methods, because the details of the implementations vary. For example, in the pursuit of a more general and efficient structure, being designed in a single-scale architecture with only 256 feature dimensions leads our method into a natural inferior position.

As shown in Table 1, our models' performances are comparable and even better with other complicated methods in terms of all metrics no matter whether rerank post-processing is used. Specifically, for the most important metric Rank-1, our model (ResNet101) gains approximately 9%/1% absolute improvement over the previous single-/multi-scale state-of-the-art method. For the image to text retrieval, our method outperforms all others by a large margin, indicating our method's superiority in aligning two modalities. The highest mAP also demonstrates the entire retrieval sequence predicted by our models has a top-quality order. In addition to the satisfactory performance, our method also has the merits over the training efficiency, retrieval speed and offline features storage against previous state-of-the-art methods (See detailed comparisons in the supplementary material).

Method	Arch.	Dim.	Text to Image w/o Rerank				Image to Text w/o Rerank			
			Rank-1	Rank-5	Rank-10	mAP	Rank-1	Rank-5	Rank-10	mAP
GNA-RNN [20]	S	512	19.05	-	53.64	-	-	-	-	-
Dual Path [41]	S	2048	44.40	66.26	75.07	-	-	-	-	-
CMPM/C ⁺ [39]	S	512	49.37	71.69	79.27	-	60.96	84.42	90.83	-
MIA [26]	M	1024	53.10	75.00	82.90	-	-	-	-	-
PMA [17]	M	1024	54.12	75.45	82.97	-	-	-	-	-
TIMAM [†] [30]	S	512	54.51	77.56	84.78	-	67.40	88.65	93.91	-
CKMA [5]	S	512	54.69	73.65	81.86	-	-	-	-	-
ViTAA [‡] [36]	M	256	54.92	75.18	82.90	51.60	65.71	88.68	93.75	45.75
CMAAM [1]	M	512	56.68	77.18	84.86	-	-	-	-	-
HGAN [†] [40]	M	512	59.00	79.49	86.62	-	71.16	90.05	95.06	-
NAFS (G) [‡] [11]	M	768	59.36	79.13	86.00	54.07	71.89	90.99	95.28	50.16
MGEL [35]	M	512	60.27	80.01	86.74	-	71.87	91.38	95.42	-
AXM-Net [10]	M	512	61.90	79.41	85.75	57.38	-	-	-	-
TIPCB [‡] [6]	M	2048	63.63	82.82	89.01	56.78	73.55	92.26	96.03	51.78
Ours (ResNet50)	S	256	61.65	80.98	86.78	58.29	75.96	93.40	96.55	55.05
Ours (ResNet101)	S	256	64.08	81.73	88.19	60.08	78.99	95.02	97.17	56.78

Method	Arch.	Dim.	Text to Image w/ Rerank				Image to Text w/ Rerank			
			Rank-1	Rank-5	Rank-10	mAP	Rank-1	Rank-5	Rank-10	mAP
ViTAA [‡] [36]	M	256	54.92	74.77	82.49	52.60	66.17	88.61	93.56	46.39
NAFS (G) [‡] [11]	M	768	59.62	78.90	85.72	55.02	72.67	90.92	95.12	50.92
TIPCB [‡] [6]	M	2048	63.37	81.56	87.57	60.02	74.04	92.06	95.61	53.78
Ours (ResNet50)	S	256	61.94	80.52	86.45	59.45	76.26	93.46	96.58	55.67
Ours (ResNet101)	S	256	64.40	81.27	87.96	61.19	78.99	95.02	97.23	57.31

Table 1: Comparisons with previous methods on the CUHK-PEDES. Only global features are used during inference for our reproduced NAFS [11]. "Arch."/"Dim." is the abbreviation for architecture/feature dimension. S/M stands for the methods designed in single-/multi-scale architecture, and all single-scale methods are highlighted with gray background. [†] stands for the results from HGAN [40]. [‡] stands for the results reproduced with public codes/checkpoints released by their authors. Overall 1st/2nd best in red/blue.

4.3 Ablation studies

We evaluate the effectiveness of our framework by introducing our proposed CM-MoCo in Section 3.1 and transfer learning from cross-modal pre-training in Section 3.2. We adopt a baseline model trained without our proposed \mathcal{L}_{cmc} and transfer learning, where both visual and textual backbones are initialized with the models separately pre-trained on unimodal data, *i.e.*, ResNet101 and BERT. To illustrate the power of cross-modal pre-training with generic image-text pairs in TBPS, in addition to adopting CLIP pre-trained on the unreleased huge-scale WIT [29] (about 400 million pairs), we pre-train another model on a relative smaller dataset, MSCOCO captions [22] (about 0.57 million pairs), to reach a comprehensive comparison (See more details for pre-training settings in the supplementary material).

We can draw the following conclusions from the Table 2: (1) Without our transfer strategy, no matter whether and how many generic image-text pairs are used for pre-training, directly fine-tuning the pre-trained model results in a negative transfer. This transfer strategy can effectively alleviate the domain gap coming from the textual side. More specifically, it leads to at least 6.66% performance improvement. (2) When the domain gap is well resolved, cross-modal pre-training yields better performance than unimodal pre-training. This is even more significant when the image-text pairs dataset for pre-training is scaled up, *i.e.*, from MSCOCO to WIT. (3) The proposed CM-MoCo yields consistent improvement for all models. It further boosts the performance with 1.5% improvement in average.

Visual Backbone	Textual Backbone	Paired Data	Transfer Strategy	CM-MoCo	Rank-1	Rank-5	Rank-10
ResNet101	BERT	-			51.95	71.30	79.78
ResNet101	BERT	-	✓		58.24	78.64	85.62
ResNet101	BERT	-	✓	✓	59.62	79.13	86.08
ResNet101	BERT	MSCOCO [22]			53.30	72.77	80.02
ResNet101	BERT	MSCOCO [22]	✓		59.96	79.74	86.83
ResNet101	BERT	MSCOCO [22]	✓	✓	60.79	79.58	87.35
CLIP ResNet101	CLIP-TE	WIT [29]			0.15	0.76	1.23
CLIP ResNet101	CLIP-TE	WIT [29]	✓		62.52	80.57	87.12
CLIP ResNet101	CLIP-TE	WIT [29]	✓	✓	64.08	81.73	88.19

Table 2: Ablation experimental results for proposed components. Only the results without rerank for text-to-image task are reported. Paired data denotes the dataset used for pre-training. Transfer strategy denotes the proposed strategy mentioned in Section 3.2.

Embed. Type	GRU	Fixed	Rank-1	Rank-5	Rank-10	CM-MoCo	Queue Size	Rank-1	Rank-5	Rank-10
Contextualized			0.15	0.76	1.23		0	62.52	80.57	87.12
Contextualized	✓		57.23	75.39	82.39	✓	1024	63.52	81.77	88.52
Contextualized	✓	✓	61.39	79.73	86.84	✓	2048	64.08	81.73	88.19
Word-type	✓	✓	62.52	80.57	87.12	✓	4096	63.06	81.50	87.69

Table 3: Results for different text encoders. Table 4: Results for different queue sizes.

Transfer learning strategy. In this part, we validate the design of our transfer learning strategy proposed in Section 3.2. There are three proposed alterations, *i.e.*, GRU, fixed textual encoder and non-contextualized word embedding. Table 3 demonstrates the effectiveness of each alteration. It is clear that appending a GRU to the pre-trained textual encoder directly allows the textual stream to learn fine-grained textual information on the basis of the previously learned coarse-grained one. Further, fixed textual encoder and non-contextualized word embedding boost the performance with about 4% and 1%, respectively. This confirms our conjecture that large-scale cross-modal pre-training can provide more meaningful word embeddings and image features, while the domain gap problem is needed to be addressed carefully. In summary, this strategy allows us to effectively transfer the knowledge from cross-modal pre-training.

Queue sizes. We also evaluate the effect of the queue sizes that is used to store keys and construct negative pairs for CM-MoCo. From the results reported in Table 4, it is obvious that a large queue size (1024 and 2048) in CM-MoCo improves the performance. However, further increasing this queue size yields worse performance. This phenomenon is probably due to the limited data in the TBPS dataset. A queue size which is too large for the entire dataset may store too many obsolete keys, which will mislead the learning direction. It is therefore detrimental for further cross-modal contrastive learning.

4.4 Qualitative results

Figure 3 visualizes some typical retrieval results (See more successful results and failure cases in the supplementary material). Concretely, Figure 3(a) is a retrieval result with a detailed text query, our method successfully retrieves the target image with a high similarity score. Figure 3(b) and 3(c) show the results of two different text queries for the same target image. The second query describes the lady's coat and shoes as "black clothes" and "black shoes", while the third query is more fine-grained with "high-level black coat" and "tall black high heels" and thus gets a much better result. We can draw two conclusions from this figure: (1) Our method appears to be very effective in distinguishing fine-grained details in a given text query. (2) When the text query has ambiguity, our method can still give reasonable results. (3) The result sequence retrieved by our method conforms to human intuition.



(a) He is wearing grey pants, a dark long sleeved sweater, and a light collared shirt underneath. He is carrying a black backpack on two shoulders.



(b) The woman with dark hair has all black clothes and shoes with a white handbag.



(c) The woman is wearing a high-level black coat and tall black high heels on her feet.

Figure 3: Typical retrieval results. The image with red box is the correct matching. The number on top of each image represents the predicted similarity with given text query.

5 Conclusion

In this paper, we propose a novel method for text-based person search (TBPS). Our model learns more discriminative features by using the proposed cross-modal momentum contrastive learning strategy, and effectively transfers the knowledge learned from generic image-text pairs to compensate the data-scarce problem. This is demonstrated by the fact that our approach clearly outperforms prior art on CUHK-PEDES, often by a big margin.

Acknowledgments

Xiao Han appreciates Freda Shi for her helpful discussion and Kecheng Zheng for sharing details of his implementation.

References

- [1] Surbhi Aggarwal, R. Venkatesh Babu, and Anirban Chakraborty. Text-based person search via attribute-aided matching. In *WACV*, 2020.
- [2] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020.
- [3] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *CVPR*, 2021.
- [4] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *ECCV*, 2020.
- [5] Yucheng Chen, Rui Huang, Hong Chang, Chuanqi Tan, Tao Xue, and Bingpeng Ma. Cross-modal knowledge adaptation for language-based person search. *IEEE TIP*, 2021.
- [6] Yuhao Chen, Guoqing Zhang, Yujiang Lu, Zhenxing Wang, Yuhui Zheng, and Ruili Wang. Tipcb: A simple but effective part-based convolutional baseline for text-based person search. *arXiv preprint*, 2021.
- [7] Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *EMNLP*, 2014.
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.
- [9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, 2019.
- [10] Ammarah Farooq, Muhammad Awais, Josef Kittler, and Syed Safwan Khalid. Axmnet: Cross-modal context sharing attention network for person re-id. *arXiv preprint*, 2021.
- [11] 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*, 2021.
- [12] Jean-Bastien Grill, Florian Strub, Florent Althé, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Ávila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap your own latent - A new approach to self-supervised learning. In *NeurIPS*, 2020.
- [13] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *CVPR*, 2006.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.

- [15] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020.
- [16] Yuqi Huo, Manli Zhang, Guangzhen Liu, Haoyu Lu, Yizhao Gao, Guoxing Yang, Jingyuan Wen, Heng Zhang, Baogui Xu, Weihao Zheng, *et al.* Wenlan: Bridging vision and language by large-scale multi-modal pre-training. *arXiv preprint*, 2021.
- [17] 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.
- [18] Wonjae Kim, Bokyoung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *ICML*, 2021.
- [19] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [20] Shuang Li, Tong Xiao, Hongsheng Li, Bolei Zhou, Dayu Yue, and Xiaogang Wang. Person search with natural language description. In *CVPR*, 2017.
- [21] Xiujuan Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, *et al.* Oscar: Object-semantics aligned pre-training for vision-language tasks. In *ECCV*, 2020.
- [22] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014.
- [23] Xiangtan Lin, Pengzhen Ren, Yun Xiao, Xiaojun Chang, and Alex Hauptmann. Person search challenges and solutions: A survey. In *IJCAI*, 2021.
- [24] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vlbart: Pretraining task-agnostic violinguistic representations for vision-and-language tasks. In *NeurIPS*, 2019.
- [25] Hao Luo, Wei Jiang, Youzhi Gu, Fuxu Liu, Xingyu Liao, Shenqi Lai, and Jianyang Gu. A strong baseline and batch normalization neck for deep person re-identification. *IEEE TMM*, 2019.
- [26] Kai Niu, Yan Huang, Wanli Ouyang, and Liang Wang. Improving description-based person re-identification by multi-granularity image-text alignments. *IEEE TIP*, 2020.
- [27] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, *et al.* Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019.
- [28] Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *ICCV*, 2015.
- [29] 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.

- [30] Nikolaos Sarafianos, Xiang Xu, and Ioannis A. Kakadiaris. Adversarial representation learning for text-to-image matching. In *ICCV*, 2019.
- [31] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. VL-BERT: pre-training of generic visual-linguistic representations. In *ICLR*, 2020.
- [32] Siqi Sun, Yen-Chun Chen, Linjie Li, Shuohang Wang, Yuwei Fang, and Jingjing Liu. Lightningdot: Pre-training visual-semantic embeddings for real-time image-text retrieval. In *NAACL-HLT*, 2021.
- [33] Hao Tan and Mohit Bansal. LXMERT: Learning cross-modality encoder representations from transformers. In *EMNLP-IJCNLP*, 2019.
- [34] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [35] Chengji Wang, Zhiming Luo, Yaojin Lin, and Shaozi Li. Text-based person search via multi-granularity embedding learning. In *IJCAI*, 2021.
- [36] Zhe Wang, Zhiyuan Fang, Jun Wang, and Yezhou Yang. Vitaa: Visual-textual attributes alignment in person search by natural language. In *ECCV*, 2020.
- [37] Zheng Wang, Zhixiang Wang, Yinqiang Zheng, Yang Wu, Wenjun Zeng, and Shin’ichi Satoh. Beyond intra-modality: A survey of heterogeneous person re-identification. In *IJCAI*, 2020.
- [38] Mang Ye, Jianbing Shen, Gaojie Lin, Tao Xiang, Ling Shao, and Steven CH Hoi. Deep learning for person re-identification: A survey and outlook. *IEEE TPAMI*, 2021.
- [39] Ying Zhang and Huchuan Lu. Deep cross-modal projection learning for image-text matching. In *ECCV*, 2018.
- [40] Kecheng Zheng, Wu Liu, Jiawei Liu, Zheng-Jun Zha, and Tao Mei. Hierarchical gumbel attention network for text-based person search. In *ACM MM*, 2020.
- [41] 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.