



Learnable Pillar-based Re-ranking for Image-Text Retrieval

Leigang Qu
leigangqu@gmail.com
National University of Singapore

Meng Liu
mengliu.sdu@gmail.com
Shandong Jianzhu University

Wenjie Wang*
wenjiewang96@gmail.com
National University of Singapore

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

Liqiang Nie
nieliqiang@gmail.com
Harbin Institute of Technology
(Shenzhen)

Tat-Seng Chua
dcscts@nus.edu.sg
National University of Singapore

ABSTRACT

Image-text retrieval aims to bridge the modality gap and retrieve cross-modal content based on semantic similarities. Prior work usually focuses on the pairwise relations (*i.e.*, whether a data sample matches another) but ignores the *higher-order neighbor relations* (*i.e.*, a matching structure among multiple data samples). Re-ranking, a popular post-processing practice, has revealed the superiority of capturing neighbor relations in single-modality retrieval tasks. However, it is ineffective to directly extend existing re-ranking algorithms to image-text retrieval. In this paper, we analyze the reason from four perspectives, *i.e.*, generalization, flexibility, sparsity, and asymmetry, and propose a novel learnable pillar-based re-ranking paradigm. Concretely, we first select top-ranked intra- and inter-modal neighbors as *pillars*, and then reconstruct data samples with the neighbor relations between them and the pillars. In this way, each sample can be mapped into a multimodal pillar space only using similarities, ensuring generalization. After that, we design a neighbor-aware graph reasoning module to flexibly exploit the relations and excavate the sparse positive items within a neighborhood. We also present a structure alignment constraint to promote cross-modal collaboration and align the asymmetric modalities. On top of various base backbones, we carry out extensive experiments on two benchmark datasets, *i.e.*, Flickr30K and MS-COCO, demonstrating the effectiveness, superiority, generalization, and transferability of our proposed re-ranking paradigm.

CCS CONCEPTS

• **Information systems** → **Novelty in information retrieval; Multimedia and multimodal retrieval.**

KEYWORDS

Image-Text Matching; Cross-Modal Retrieval; Re-ranking

*Corresponding author: Wenjie Wang (wenjiewang96@gmail.com). This research is supported by the Defence Science and Technology Agency, the National Natural Science Foundation of China under Grant 62006142, and the Special Fund for distinguished professors of Shandong Jianzhu University.

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](https://permissions.acm.org).
SIGIR '23, July 23–27, 2023, Taipei, Taiwan.

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9408-6/23/07...\$15.00
<https://doi.org/10.1145/3539618.3591712>

ACM Reference Format:

Leigang Qu, Meng Liu, Wenjie Wang, Zhedong Zheng, Liqiang Nie, and Tat-Seng Chua. 2023. Learnable Pillar-based Re-ranking for Image-Text Retrieval. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '23)*, July 23–27, 2023, Taipei, Taiwan. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3539618.3591712>

1 INTRODUCTION

Image-text retrieval (ITR), as a fundamental vision-language task, has attracted substantial attention over the past decade [5, 7, 34]. It aims to bridge the heterogeneous modality gap and achieve semantical matching by bidirectional retrieval: 1) image-to-text retrieval, searching for correct descriptive texts given an image, and 2) text-to-image retrieval, finding the relevant images for a given text. By pursuing cross-modal understanding and reasoning, ITR has become essential to other visual-language applications such as visual question answering [1] and video moment retrieval [23].

Existing approaches to ITR fall into two categories: The first category, *i.e.*, *two-tower* framework [7, 19], separately maps the visual and textual samples into a joint embedding space to calculate the cross-modal similarity. In contrast, the second category, *i.e.*, *one-tower* framework [5, 34], models fine-grained interactions by the cross-attention mechanism and directly outputs the similarity. Despite impressive progress, both categories only consider the pairwise matching relations between an image and a text, as shown in Figure 1(a). They completely ignore the higher-order **neighbor relations** among multiple images and texts illustrated in Figure 1(b), and therefore suffer from suboptimal retrieval performance. As a prevailing post-processing technology to explore neighbor relations, re-ranking has revealed remarkable effectiveness in conventional single-modality retrieval tasks, such as image retrieval [16, 46], document retrieval [27, 29], and recommendation [41]. As shown by the blue line in Figure 2, the retrieval performance of image retrieval can be largely improved via query expansion re-ranking (*i.e.*, α QE [35]). This motivates us to investigate re-ranking to exploit neighbor relations in the context of ITR.

In general, existing re-ranking approaches can be divided into three main categories: Query Expansion (QE) [8, 35], Diffusion [16], and Neighbor-based methods [48]. However, they are not suitable for cross-modal ITR due to the following challenges: 1) **Architecture Generalization**. Some re-ranking methods strongly rely on specific model architectures. For example, QE can be only applied to two-tower frameworks since it requires intermediate features for expansion [35], and thus can not well serve the one-tower




Figure 1: Illustration of (a) Pairwise Relation, (b) Neighbor Relation, and (c) Pillar Space Construction. Compared with the pairwise relations, the higher-order neighbor relation is a matching structure including the query-neighbor (e.g., $q \leftrightarrow x$) and neighbor-neighbor (e.g., $z \leftrightarrow w$) relations. In a neighborhood, top-ranked neighbors are selected as pillars to construct the pillar space to represent data samples. Numbers denote similarities calculated by a base backbone.

framework. 2) **Representation Flexibility.** Existing re-ranking approaches [16, 35, 48] are designed with rigid or restricted rules from prior knowledge, making them difficult to automatically explore neighbor relations in complex multimodal data. Specifically, QE aggregates neighbors to enrich the query in a limited unidirectional way, the diffusion mechanism performs ranking on fixed manifolds, and neighbor-based methods [32, 48] heavily depend on certain pre-defined neighbor relations. 3) **Sparse Relevance.** The average ratio of positive items over the benchmark datasets of ITR (e.g., MS-COCO (5K) [20]) is much smaller than those of image retrieval (e.g., \mathcal{R} Paris6K [31]), specifically 0.02% vs. 2.3%. This makes re-ranking on ITR more challenging since noisy negative information are widely propagated among items when capturing the higher-order neighbor relations. And 4) **Modality Asymmetry.** The two asymmetry retrieval directions of ITR result in different similarity distributions and heterogeneous neighbors, yet existing re-ranking methods are designed for single-modality retrieval, thus failing to consider the asymmetry problem. As shown in Figure 2, we can observe the great performance degradation of α QE (the red line) as the aggregated neighbors increase, even though we choose a two-tower backbone to fit it. We argue that it is caused by flexibility, sparsity, and asymmetry challenges.

To handle the above challenges, we present a **Learnable Pillar-based Re-Ranking (LeaPRR)** framework for ITR. 1) First, to well fit in with one-tower and two-tower frameworks, we abandon any intermediate content features and only take the final similarities which can be calculated by any framework as the input to our re-ranking model. For the sake of informative entity (i.e., queries and items) representation, we define top-ranked multimodal neighbors as *pillars* and then reconstruct entities with the similarities between pillars and these entities, as shown in Figure 1(c). In this way, each entity can be mapped into a pillar space. Different from the conventional modality-independent content space, the pillar space is constructed with multimodal pillars to alleviate the asymmetry problem. 2) Moreover, as for the flexibility and sparsity challenges, we propose a neighbor-aware graph reasoning module based on the neighborhood centered with a given query, to flexibly explore complicated neighbor relations, including the query-neighbor and neighbor-neighbor ones. 3) To further deal with the high sparsity of positive items, we exert local and global ranking constraints to reinforce the discriminative power of refined pillar features, which




Figure 2: Retrieval performance comparison between Image Retrieval (ImR) and Image-Text Retrieval (ITR) using the classical re-ranking method α QE [35]. We use GeM (AP) [36] as the base backbone to evaluate ImR on \mathcal{R} Paris6K (Hard) [31], and VSE $^\infty$ [7] for ITR on Flickr30K [42]. The results without re-ranking are shown when the neighbor number is 0.

can adaptively absorb valuable information and resist noise from neighbors. 4) Finally, we present a structure alignment constraint by transforming contextual structures from one modality to another, further combating the asymmetry issue.

To sum up, our main contributions are as follows:

- We discuss the limitations of existing re-ranking approaches in the ITR task and present four challenges, i.e., generalization, flexibility, sparsity, and asymmetry.
- To deal with the four challenges, we reformulate re-ranking in the multi-modality field and propose a learnable pillar-based re-ranking paradigm for the first time, which is plug-and-play and applicable to all existing ITR backbones.
- Extensive experiments based on different base ITR backbones on two datasets, i.e., Flickr30K [42] and MS-COCO [20], validate the effectiveness and superiority of our model. The codes and settings are released at <https://github.com/LgQu/LeaPRR>.

2 RELATED WORK

2.1 Image-Text Retrieval

Current ITR models can be divided into two categories, i.e., *two-tower framework* and *one-tower framework*, in accordance with the manner for modality interaction. As for the two-tower framework [7, 10, 11, 33, 47], images and texts are independently mapped into a joint feature space in which the semantic similarities are calculated via cosine function or Euclidean distance. For example, Frome *et al.* [11] represented visual and textual instances in a common modality-agnostic space to assess semantic relevance. To further improve the discriminative power, Zheng *et al.* [47] argued that the commonly used ranking loss is not effective confronted with large-scale multimodality data and presented a new instance loss to exploit the intra-modal data distribution in an end-to-end manner. With the popularity of bottom-up attention [1], recent research attentions have been shifted to capturing and modeling relations between entities (e.g., objects or words) via intra-model interaction. As a representative model, VSRN [19] reasons across regions in the visual modality by using GCN and Gated Recurrent Unit network for short-long term relation modeling. Qu *et al.* [33] discussed the modality asymmetry issue and proposed a multi-view summarization architecture. Chen *et al.* [7] discovered that simple

global pooling functions can outperform the complicated models and proposed a generalized pooling operator. These two-tower models ensure satisfactory retrieval efficiency as features in each modality can be extracted in parallel and indexed in an offline way, but suffer from poor accuracy due to the coarse-grained matching.

In contrast, the one-tower framework [6, 9, 17, 34] embraces fine-grained cross-modal interactions to achieve more thorough matching between fragments (e.g., objects and words). As the pioneering work, SCAN [17] infers some latent region-word alignments by means of cross-modal interaction. Thereafter, much effort has been dedicated to mining cross-modal matching relations. For instance, Chen *et al.* [6] incorporated the iteration strategy into SCAN to perform multi-step cross-modal relation reasoning. Liu *et al.* [21] tailored a graph-based network to explore fine-grained correspondence from node level and structure level. Diao *et al.* [9] inferred multi-level semantic correlations in a similarity graph and filtered noisy alignments via the attention mechanism. Differently, Qu *et al.* [34] argued possible optimal patterns may not be explored by existing modality interaction patterns and developed a unified dynamic modality interaction modeling network for automatic routing. Despite the superior retrieval accuracy, this line of work has the problem of heavy computational overhead.

To sum up, most existing works focus on representing data samples in joint embedding space or intensively excavating fine-grained cross-modal alignments, but they ignore the efficient and practical post-processing procedure, *i.e.*, re-ranking.

2.2 Re-ranking

As an effective and practical technology, re-ranking has attracted a flurry of interest in traditional single-modality retrieval tasks, such as image retrieval [3, 8, 12, 18, 28, 44, 48] and document retrieval [4, 24–26, 43, 49]. A rich line of studies has explored **Query Expansion** by using the nearest neighbors of a query to generate a new enriched one. For instance, average query expansion (AQE) [8] was first proposed to aggregate local features by mean-pooling the top-k images in the original ranking list. Afterward, Gordo *et al.* [12] and Radenovic *et al.* [35] argued that the mean-pooling would induce much noisy information and presented AQEwD and α QE, respectively. Another thread of work is **Diffusion**-based re-ranking strategy, which depends on a neighborhood topology constructed in an offline way and utilizes it at query time for search on the manifold. For example, Iscen *et al.* [16] captured the image manifold in the region-level embedding space and proposed a fast spectral ranking method [15] to improve the computational efficiency. In addition, some research has been devoted to employing neighborhood relations, especially the **k-Reciprocal Nearest Neighbors**, for re-ranking. For example, Zhong *et al.* [48] encoded the k-reciprocal neighbors of a query into a single vector and then conducted re-ranking under the Jaccard distance. Liu *et al.* [22] built a graph based on neighborhood affinities and leveraged GCN for manifold learning. Recently, Ouyang *et al.* [28] reconstructed affinity features with a set of anchor images for contextual similarity reasoning. As one of the rare efforts to explore re-ranking for image-text retrieval, the k-reciprocal nearest neighbor searching scheme [38] was designed as post-processing to enhance retrieval performance. Despite the thrilling progresses, existing methods

depend heavily on prior knowledge, lacking sufficient flexibility to adaptively model query-neighbor and neighbor-neighbor relations. Besides, they overlook the sparsity and asymmetry issues widely existing in image-text retrieval.

3 METHODOLOGY

The overall framework of our proposed LeaPRR is depicted in Figure 3. In the following, we will first introduce the problem definition for cross-modal ITR re-ranking and then elaborate on *Pillar-based Representation* and *Neighbor-aware Graph Reasoning*. Finally, we will describe the *Learning* strategy in the pillar space.

3.1 Problem Definition

Given a query q (an image or a text) and the database in another modality, *i.e.*, $\mathcal{D} = \{d_1, \dots, d_C\}$, we first calculate the initial similarity $\phi(q, d_i)$ between the query and each item d_i in \mathcal{D} via a base ITR backbone ϕ . It can be implemented by a one- or two-tower network. The former [19] commonly treats ϕ as the cosine similarity between embeddings extracted from two branches, while the latter views ϕ as an integrated inter-modal interaction module, such as cross-attention [6, 39]. Afterward, we can attain an initial ranking list $\mathcal{R}(q, \mathcal{D})$ by sorting similarities in descending order.

The goal of ITR re-ranking is to refine the initial ranking list so those items that semantically match well with the query are ranked higher and vice versa in the refined ranking list, denoted by $\mathcal{R}^*(q, \mathcal{D})$. In this work, we focus on re-ranking top- K items in $\mathcal{R}(q, \mathcal{D})$ since users always attach more attention to top-ranked retrieval results in practical scenarios.

We consider two retrieval settings: Image-to-Text (I2T), *i.e.*, $q := I, \mathcal{D} := \mathcal{D}_T$, and Text-to-Image (T2I), *i.e.*, $q := T, \mathcal{D} := \mathcal{D}_I$. Considering the high consistency except for the retrieval direction, we will elaborate our method under the *I2T setting* without loss of generality. Besides, we will use *items* to denote all retrieved texts (or images) for a given image (or text) query, *neighbors* to refer to top-ranked items, and *entities* to represent any queries or items.

3.2 Pillar-based Representation

For two-tower I2T frameworks, we can easily access intermediate features of each modality and leverage them to enrich the query or item representation for re-ranking [13, 35]. Nevertheless, this is not possible with one-tower I2T frameworks due to the existence of cross-modal interaction modules. In the light of this, it is unrealistic to use the intermediate representations to optimize the ranking results of both two-tower and one-tower frameworks. Therefore, we resort to cross-modal similarities for image-text retrieval re-ranking. However, how to represent an image or a text and compare them only depending on similarities, becomes a problem to be solved.

As the saying goes, “birds of a feather flock together”. We can describe a person with inherent appearance characteristics, such as height, weight, skin color, face shape, and hairstyle. Meanwhile, we can also depict this person by means of the relation between him and other people. For instance, it is reasonable to believe that two people are similar in appearance if they have common family members, or similar in interest if with common idols and/or friends. This inspired us that we can not only represent an entity using its

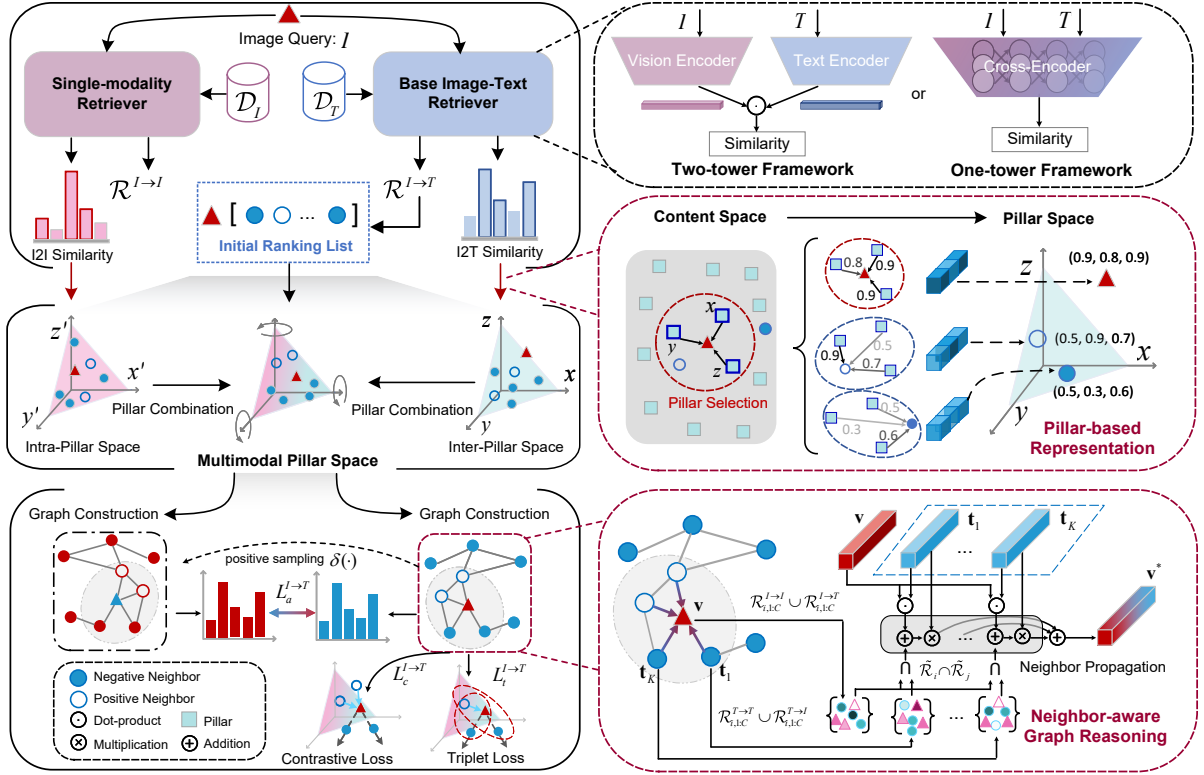


Figure 3: Schematic illustration of the proposed LeaPRR model (image-to-text direction). Given the image query and databases, the initial ranking list is obtained via the base retriever. And then the top-ranked multimodal neighbors are selected as pillars to build the multimodal pillar space. In this space, Pillar-based Representation is conducted for each data sample. Thereafter, the graph is constructed by neighbor-based and learning-based affinities and Neighbor-aware Graph Reasoning is performed. Finally, the model learns neighbor relations for re-ranking by means of three constraints.

internal content/attributes but also can do it with external relations between it and other entities.

Although others can be employed as references to represent an item, it does not mean anyone can take responsibility. As an illustration, those people that the person is unfamiliar with would not be helpful as references to discriminate him against. Therefore, the relations between an item and the selected references should not be very sparse. We will detail how to select such references and exploit them as references to represent an item, in the following.

3.2.1 Multimodal Pillar Selection. Given the initial ranking list $\mathcal{R}(I, \mathcal{D}_T) = \{T_1, \dots, T_N\}$ of I2T retrieval, we select top- L ($L < N$) texts as the above-mentioned references. In our work, these selected references are defined as *pillars*. For the sake of simplicity, we re-denote $\mathcal{R}(I, \mathcal{D}_T)$ as $\mathcal{R}^{I \rightarrow T}$, and utilize $\mathcal{R}_{1:L}^{I \rightarrow T} = \{T_1, \dots, T_L\}$ to represent the *inter-modal pillars* of the query I . To further enhance the representation ability, we propose to excavate *intra-modal pillars* from the intra-modal database \mathcal{D}_I . Concretely, we apply the image encoder ψ_I to extract features and then calculate pairwise similarities between them. After that, based on the intra-modal ranking list $\mathcal{R}(I, \mathcal{D}_I) = \{I_1, \dots, I_M\}$, $\mathcal{R}^{I \rightarrow I}$ for short, we collect a set of intra-modal pillars¹ of query I as $\mathcal{R}_{1:L}^{I \rightarrow I} = \{I_1, \dots, I_L\}$.

¹In this paper, we use the same number of intra- and inter-modal pillars for the sake of simplicity. In practice, it is acceptable to consider different numbers of pillars.

3.2.2 Pillar-based Encoding. On the basis of pillars, we encode the query image I and its top- K neighbor texts $\{T_1, \dots, T_K\}$ as follows:

$$\begin{aligned} \mathbf{v} &= [\oplus_{i=1}^L \phi(I, T_i), \oplus_{j=1}^L \pi(I, I_j)], \\ \mathbf{t}_k &= [\oplus_{i=1}^L \pi(T_k, T_i), \oplus_{j=1}^L \phi(I_j, T_k)], \end{aligned} \quad (1)$$

where \oplus denotes the concatenation operator, $\mathbf{v} \in \mathbb{R}^{2L}$ and $\mathbf{t}_k \in \mathbb{R}^{2L}$ respectively represent pillar-based representations of the query I and its k -th neighbor T_k , ϕ can be any base backbone that calculates image-text similarities, and π is the intra-modal similarity calculation model. Specifically, $\pi(I, I_j) = \cos(\psi_I(I), \psi_I(I_j))$, $\pi(T_k, T_i) = \cos(\psi_T(T_k), \psi_T(T_i))$, where ψ_I and ψ_T are image encoder and text encoder, respectively.

3.3 Neighbor-aware Graph Reasoning

One of the keys to re-ranking is to model the high-order neighbor relations in a neighborhood including a query and its neighbors. Prior approaches try to achieve it from different aspects. For instance, QE-based models [2, 35] exploit query-neighbor interaction to expand query representation by aggregating neighbors. In contrast, diffusion-based [16, 30] models perform one-way information flow by building a connected graph. Nevertheless, these models are designed based on some rules with limited flexibility. To tackle

this issue, we tailor a graph-based neighbor propagation to capture query-neighbor and neighbor-neighbor relations adaptively.

After finishing pillar encoding, we are able to represent the query and its neighbors in a pillar space. In this space, we view each item (*i.e.*, a query or a neighbor) as a node and then connect them according to their affinity scores.

3.3.1 Neighbor-based Affinity. To exploit neighbor relations from the raw feature space, we first collect top- C neighbors from both intra- and inter-modal ranking lists for each item and then merge them into one set, as,

$$\tilde{\mathcal{R}}_i = \begin{cases} \mathcal{R}_{i,1:C}^{I \rightarrow I} \cup \mathcal{R}_{i,1:C}^{I \rightarrow T}, & \text{node } i \text{ is an image,} \\ \mathcal{R}_{i,1:C}^{T \rightarrow T} \cup \mathcal{R}_{i,1:C}^{T \rightarrow I}, & \text{node } i \text{ is a text.} \end{cases} \quad (2)$$

where $\tilde{\mathcal{R}}_i$ denotes the neighbor set of the i -th node. $\mathcal{R}_{i,1:C}^{I \rightarrow T}$ refers to the set of top- C neighbors of the i -th node in the direction $I \rightarrow T$, and the same is true for other directions. Like $I \rightarrow I$, we also apply a text encoder to obtain $\mathcal{R}_{i,1:C}^{T \rightarrow T}$.

Afterward, we calculate the edge weight \mathbf{C}_{ij} between the i -th node and the j -th node depending on their common neighbors:

$$\mathbf{C}_{ij} = \frac{|\tilde{\mathcal{R}}_i \cap \tilde{\mathcal{R}}_j|}{\sum_{k=1}^{1+K} |\tilde{\mathcal{R}}_i \cap \tilde{\mathcal{R}}_k|}. \quad (3)$$

where $|\cdot|$ denotes the element number of the set. To alleviate the influence of noisy edges, we introduce a sparse factor λ and the edge weight is updated as follows,

$$\tilde{\mathbf{C}}_{ij} = \begin{cases} \mathbf{C}_{ij}, & \mathbf{C}_{ij} > \frac{\lambda}{1+K}, \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

Following the row-wise normalization of $\tilde{\mathbf{C}}$, we could obtain the neighbor-based affinity matrix, denoted as $\hat{\mathbf{A}} \in \mathbb{R}^{(1+K) \times (1+K)}$.

3.3.2 Learning-based Affinity. To endow the model with sufficient flexibility, we drive it to learn informative relations. Specifically, we first stack the pillar-based representations of the query and its neighbors to build a matrix $\mathbf{F} = [\mathbf{v}; \mathbf{t}_1; \dots; \mathbf{t}_K] \in \mathbb{R}^{(1+K) \times 2L}$. And then we calculate the affinity matrix $\hat{\mathbf{A}} \in \mathbb{R}^{(1+K) \times (1+K)}$ as,

$$\hat{\mathbf{A}} = \text{softmax}(f_Q(\mathbf{F}) \cdot f_K(\mathbf{F})^\top), \quad (5)$$

where f_Q and f_K respectively denote two fully-connected (FC) layers and $\text{softmax}(\cdot)$ is performed over each row. $\hat{\mathbf{A}}_{ij}$ represents a learning-based affinity score between node i and node j .

3.3.3 Neighbor Propagation. Combining the above two types of affinities, we derive the final affinity matrix by $\mathbf{A} = (\hat{\mathbf{A}} + \hat{\mathbf{A}})/2$ and perform the neighbor propagation in the first layer of GCN as,

$$\mathbf{F}^{(1)} = g(\mathbf{A} \cdot f_V(\mathbf{F})) + \mathbf{F}, \quad (6)$$

where f_V is an FC layer and g denotes a multi-layer perceptron for feature transformation followed by residual connection. After the first layer of GCN, we can re-compute the learning-based affinity matrix using Eqn. (5) and then perform the following propagation as Eqn. (6) does. The output of the last layer GCN, denoted as \mathbf{F}^* , is treated as the refined pillar-based representation matrix.

3.4 Learning from Pillar Space

Under pairwise supervision between images and texts, our model is optimized with the following three constraints.

3.4.1 Global Contrastive Loss. We adopt the contrastive loss [14] to enforce the query close to positives and far away from negatives as much as possible in a global perspective, which is defined as,

$$L_c^{I \rightarrow T} = -\log \frac{\sum_{i \in \mathcal{P}} \exp(s(I, T_i)/\tau)}{\sum_{j=1}^K \exp(s(I, T_j)/\tau)}, \quad (7)$$

where \mathcal{P} denotes the positive set of I , τ is a scalar temperature factor, and $s(I, T_i) = \cos(\mathbf{v}^*, \mathbf{t}_i^*)$ refers to the cosine similarity between the refined pillar-based representations of the query image I and its i -th textual neighbor T_i .

3.4.2 Local Triplet Loss. Different from the contrastive loss that forces similarities of all positive pairs essentially, the hinge-based triplet loss [37] focuses on local relative proximity and allows the variance between positives to some extent, making it less greedy. In this paper, considering negative samples in top- K neighbors are already hard, we do not conduct the hardest negative mining but turn to the hardest positive instead, which makes it complementary to the contrastive loss. Formally, we formulate the triplet loss as,

$$L_t^{I \rightarrow T} = \sum_{i \in \mathcal{N}} [\alpha - s(I, \hat{T}) + s(I, T_i)]_+, \quad (8)$$

where α is the scalar margin, \mathcal{N} represents the negative set of I , and $\mathcal{P} \cup \mathcal{N} = \mathcal{R}_{1:K}^{I \rightarrow T}$. $[x]_+ = \max(x, 0)$ is the hinge function, and \hat{T} is the hardest positive defined as $\hat{T} = \arg \min_{i \in \mathcal{P}} s(I, T_i)$.

3.4.3 Mutual Modality Alignment. Until now, we have illustrated the mechanism of the I2T submodel. Similarly, we can derive the T2I submodel by respectively assigning the query and database as the text T and D_I .

To ensemble these two submodels, a naive way is to directly average two similarities, as most of prior base backbones [17, 21, 34] do. In this way, however, two submodels are trained independently and unable to promote each other. Inspired by mutual learning [40, 45], we devise a mutual modality alignment scheme to spur two submodels to teach and promote each other together.

Taking I2T as an example, we first carry out cross-modal positive sampling for the query I and each neighbor T_i in \mathcal{D}_T , obtaining $T' = \delta(I)$ and $I'_i = \delta(T_i)$, where $\delta(\cdot)$ is the sampling function. Concretely, for a given entity, we first find its positive matches in another modality by the pairwise ground-truth information and then uniformly sample one if there are multiple positives. Afterward, we respectively calculate two similarity distributions centered with I and T' , as,

$$\begin{aligned} \mathbf{p}_I &= \text{softmax}(s(I, T_1)/\tau, \dots, s(I, T_K)/\tau), \\ \mathbf{q}_{T'} &= \text{softmax}(s(T', I'_1)/\tau, \dots, s(T', I'_K)/\tau). \end{aligned} \quad (9)$$

Finally, the alignment loss L_a is defined based on Kullback–Leibler divergence D_{KL} as,

$$L_a^{I \rightarrow T} = D_{KL}(\mathbf{p}_I || \mathbf{q}_{T'}). \quad (10)$$

Table 1: Performance comparison between the proposed LeaPRR and several baselines on the Flickr30K and MS-COCO (5K) datasets based on VSE ∞ and DIME*. The best performance for each base backbone is highlighted in bold, while the second best results are underlined. The results of the proposed method are marked with a gray background.

Method	Flickr30K Dataset							MS-COCO (5K) Dataset						
	Image-to-Text			Text-to-Image			rSum	Image-to-Text			Text-to-Image			rSum
	R@1	R@5	R@10	R@1	R@5	R@10		R@1	R@5	R@10	R@1	R@5	R@10	
• VSE ∞ [CVPR21] [7]	81.7	95.4	97.6	61.4	85.9	91.5	513.5	58.3	85.3	92.3	42.4	72.7	83.2	434.2
+ AQEwD [IJCV17] [12]	81.7	88.8	94.0	61.4	80.0	87.0	492.9	58.3	77.5	85.1	42.4	66.3	77.5	407.1
+ α QE [TPAMI18] [35]	81.7	90.9	95.7	61.4	83.4	90.7	503.8	58.3	80.6	89.3	42.4	71.0	82.1	423.7
+ \ddagger ADBAwD + AQEwD [IJCV17] [12]	79.3	88.1	92.7	37.4	69.4	79.3	446.2	55.5	75.9	83.4	29.7	58.5	71.1	374.1
+ \ddagger α DBA + α QE [TPAMI18] [35]	80.6	90.2	95.3	59.5	79.6	88.3	493.5	59.0	79.4	87.8	41.8	68.5	80.5	417.0
+ \ddagger ADBA + AQE [ICCV07] [8]	77.0	85.2	92.4	62.7	85.6	91.2	494.1	52.0	74.1	83.4	38.5	67.5	79.1	394.6
+ \ddagger ADBAwD + AQEwD [IJCV17] [12]	81.7	88.9	94.0	<u>66.4</u>	<u>87.8</u>	92.1	510.9	58.3	77.5	85.1	42.7	70.9	81.1	415.6
+ \ddagger α DBA + α QE [TPAMI18] [35]	79.0	89.4	95.4	64.8	87.9	93.3	509.8	56.0	79.6	88.3	41.6	71.3	82.4	419.2
+ DFS [CVPR17] [16]	81.7	95.4	<u>97.6</u>	61.4	86.1	91.6	513.8	58.3	85.3	<u>92.3</u>	42.4	72.7	83.2	434.3
+ FSR [CVPR18] [15]	80.0	94.8	97.2	60.9	83.5	87.6	504.0	56.8	84.5	92.2	41.8	71.0	80.2	426.5
+ KRNN [ACMMM19] [38]	<u>84.7</u>	<u>96.4</u>	98.0	64.6	87.7	<u>92.2</u>	<u>523.6</u>	<u>61.6</u>	<u>87.5</u>	93.0	<u>45.3</u>	<u>74.5</u>	<u>84.0</u>	<u>445.9</u>
+ LeaPRR (Ours)	86.2	96.6	97.6	66.6	87.9	91.9	526.7	65.0	87.8	93.0	46.8	74.6	84.1	451.5
• DIME* [SIGIR21] [34]	81.0	95.9	98.4	63.6	88.1	93.0	520.0	59.3	85.4	91.9	43.1	73.0	83.1	435.8
+ DFS [CVPR17] [16]	80.5	93.5	95.0	59.2	72.9	75.3	476.4	59.2	85.4	91.9	42.9	70.2	78.2	427.8
+ FSR [CVPR18] [15]	80.5	91.0	93.2	59.2	67.6	69.0	460.5	58.8	83.3	88.8	42.6	62.5	67.6	403.7
+ KRNN [ACMMM19] [38]	<u>83.6</u>	96.8	98.4	65.2	88.3	<u>92.9</u>	<u>525.2</u>	<u>62.9</u>	<u>86.6</u>	<u>92.5</u>	43.8	73.1	83.3	<u>442.2</u>
+ LeaPRR (Ours)	86.2	<u>96.7</u>	98.9	71.1	89.5	93.2	535.6	64.3	87.3	93.1	45.5	74.6	83.7	448.5

•: base backbones without re-ranking; \ddagger : expand database and query sequentially; \ddagger : concatenate all queries with database and expand together;
*: ensemble two single base backbones

By combining all the above three loss functions in I2T and T2I settings, we could acquire the total loss for model optimization, as,

$$L = \sum_{j \in \{I \rightarrow T, T \rightarrow I\}} \sum_{i \in \{c, t, a\}} L_i^j. \quad (11)$$

4 EXPERIMENTS

In this section, we conducted experiments on two benchmark datasets to answer the following research questions:

- **RQ1:** How does LeaPRR perform in the image-text retrieval task compared with state-of-the-art baselines?
- **RQ2:** How does each component of LeaPRR affect the retrieval performance?
- **RQ3:** How are the generalization and transferability of LeaPRR when facing different base ITR backbones and datasets?

4.1 Datasets

We validate our method on two large-scale datasets: Flickr30K [42] and MS-COCO [20].

Flickr30K [42]: It comprises 31,783 images collected from Flickr, and each image corresponds to 5 human-annotated sentences. We employ the same split as [19, 33] did, specifically, 29,783 training images, 1,000 validation images, and 1,000 testing images.

MS-COCO [20]: It includes 123,287 images and each is annotated with 5 sentences. Following the split of [17, 19, 34], we use 113,287 images for training, 5,000 images for validation, and 5,000 images for testing. In particular, we adopt the challenging evaluation setting **MS-COCO (5K)**, *i.e.*, directly testing on the full 5K images.

4.2 Experimental Settings

4.2.1 Evaluation Protocols. Following the existing work [7, 19, 34], we conduct evaluation by the standard recall metric R@K (K=1, 5, and 10) and rSum. Concretely, R@K refers to the percentage of queries for which the corrected item is retrieved within the

top-K of the ranking list and rSum means the sum of all R@K in both image-to-text and text-to-image retrieval directions for overall performance evaluation. The higher R@K and rSum are better.

4.2.2 Implementation Details. In our work, 64 intra-modal pillars and 64 inter-modal pillars are used to construct a 128-dimensional pillar space, *i.e.*, $L = 64$. Top-32 and Top-8 items in I2T and T2I settings are considered for re-ranking, respectively. In other words, $K = 32$ for I2T and $K = 8$ for T2I. Moreover, the dimension of intermediate embedding space mapped via FC layers is set to 768. The sparse factor in Eqn. (4) is set as 0.8. The number of GCN layers is set to 2 for neighbor propagation. The margin α in local triplet loss and the temperature factor τ are set as 0.2 and 1.0, respectively. In addition, we used the SGD optimizer with a momentum of 0.9 and batch size of 512. Our model is trained for 30 epochs with a learning rate of 0.01. The checkpoint with the highest rSum on the validation set is chosen for testing.

In this paper, we directly applied the visual and textual branches of the pre-trained two-tower model CAMERA[33] published in the open-source community² to calculate intra-modal similarities and collected intra-modal pillars, *i.e.*, $\mathcal{R}_{1:L}^{I \rightarrow I}$ and $\mathcal{R}_{1:L}^{T \rightarrow T}$ for all base backbones and experiments.

4.3 Performance Comparison (RQ1)

This section demonstrates the comparison with other state-of-the-art re-ranking methods, including query expansion based methods (AQE [8], AQEwD [12], α QE [35]), and them combined with database-side augmentation (DBA) [2]³, including average augmentation (ADBA), with-decay augmentation (ADBAwD), and weighted augmentation (α DBA), respectively), diffusion based methods (DFS [16], FSR [15])⁴, and reciprocal nearest neighbor based

²<https://acmmmcamera.wixsite.com/camera>.

³<https://github.com/naver/deep-image-retrieval>.

⁴<https://github.com/ducha-aiki/manifold-diffusion>.

Table 2: Ablation Study on Flickr30K regarding different pillar selection strategies. Pil. denotes Pillar.

Method	Image-to-Text		Text-to-Image	
	R@1	R@5	R@1	R@5
DIME* ^[34]	81.0	95.9	63.6	88.1
+ LeaPRR (Bottom Pil.)	78.8	94.4	64.8	87.3
+ LeaPRR (Random Pil.)	81.8	95.3	62.9	87.3
+ LeaPRR (w/o Inter-Pil.)	80.3	96.2	66.7	89.0
+ LeaPRR (w/o Intra-Pil.)	84.3	95.9	67.9	89.1
+ LeaPRR	86.2	96.7	71.1	89.5

ones (KRNN [38])⁵. These re-ranking baselines are implemented on top of two representative state-of-the-art base backbones, including a two-tower base backbones VSE ∞ [7] and a one-tower base backbone DIME*⁷ [34]. Note that we directly utilized the published pre-trained parameters to calculate image-text similarities to train and evaluate LeaPRR and other re-ranking baselines. Besides, we inherit ensemble similarities (*i.e.*, averaging two similarities calculated by two single base backbones) for challenging evaluation if available. As query expansion re-ranking baselines (*e.g.*, AQE, AQEWd, α QE, and their corresponding DBA versions) require accessing intermediate intra-modal representations, we can only applied them to two-tower base backbones.

The comparisons are summarized in Table 1. From this table, we have the following observations:

- Directly applying Query Expansion methods [12, 35] to the two-tower base backbone, *i.e.*, VSE ∞ , is not helpful. Concretely, after executing QE-based re-ranking for VSE ∞ , the performance drops, especially on R@5 and R@10. The main reasons are as follows: 1) The intrinsic sparsity of ITR (discussed in Section 1) may induce top-ranked negative neighbor information to be integrated into the query, therefore overwhelming the informative representation of the original query; and 2) the modality asymmetry makes it difficult to fuse neighbors from a different modality. Besides, performing DBA before QE (*e.g.*, $\dagger\alpha$ DBA + α QE) fails to redeem this but even makes it worsen, which is mainly attributable to the sparsity and asymmetry issues that have not been tackled.
- By combining DBA and QE at the same time for re-ranking, the performance on some metrics (*e.g.*, recall values at T2I on Flickr30K) gets improved. Among them, $\ddagger\alpha$ DBAwd + AQEWd [12] gains the better R@1 (5.0% higher than VSE ∞) and $\ddagger\alpha$ DBA + α QE [35] performs better on R@5 (2.0% gain) and R@10 (1.8% gain). The reason may be that the combination strategy relieves the asymmetry problem. In other words, the query could absorb certain informative clues from intra-modal neighbors.
- The diffusion approach DFS [16] gains marginal overall performance improvement (0.3% and 0.1% regarding rSum on Flickr30K and MS-COCO (5K), respectively) compared with the base VSE ∞ [7]. However, it fails to offer consistent improvement for DIME* [34] due to the possible more complicated manifold structure caused by cross-modal interactions in this one-tower architecture.
- The reciprocal neighbor-based method KRNN [38], acting as a strong baseline, surpasses all aforementioned baselines in terms of most metrics. This demonstrates the effectiveness of capturing and modeling reciprocal query-neighbor relations for ITR.

⁵<https://github.com/Wangt-CN/MTFN-RR-PyTorch-Code>.

⁶https://github.com/woodfrog/vse_infty.

⁷<https://sigir21.wixsite.com/dime>.

Table 3: Ablation Study on Flickr30K regarding different constraints and ensemble ways. Cont., Trip., and MMA stand for global contrastive loss, local triplet loss, and mutual modality alignment, respectively.

Method	Image-to-Text		Text-to-Image	
	R@1	R@5	R@1	R@5
DIME* ^[34]	81.0	95.9	63.6	88.1
+ LeaPRR (w/o Cont. Loss)	86.1	95.9	70.4	89.4
+ LeaPRR (w/o Trip. Loss)	85.5	96.1	69.9	88.7
+ LeaPRR (w/o MMA, only i-t)	83.0	95.0	70.1	89.3
+ LeaPRR (w/o MMA, only t-i)	86.7	97.8	66.9	89.2
+ LeaPRR (w/o MMA, i-t & t-i)	86.0	96.8	68.7	89.5
+ LeaPRR	86.2	96.7	71.1	89.5

- LeaPRR outperforms all the baselines over most metrics on top of two base backbones, *e.g.*, it gains 15.9% and 12.7% improvement regarding rSum for DIME* [34] on Flickr30K and MS-COCO (5K), respectively. The remarkable and consistent improvement is attributed to 1) the excellent **flexibility** endowed by pillar-based representation and neighbor-aware graph reasoning, and 2) the settlement of the **sparsity** and **asymmetry** challenges powered by multimodal pillars, local-global constraints, and the mutual modality alignment.

4.4 In-depth Analysis (RQ2 & RQ3)

4.4.1 Ablation Study. On top of the challenging state-of-the-art base backbone DIME*, We conduct a series of experiments on Flickr30K to verify the effectiveness of each component of LeaPRR and explore their impacts as follows.

- **Impact of Pillar-based Represen.** Considering our proposed model performs a series of learning in the pillar space, the quality of this space plays a key role in the overall performance. In this section, we hence first justified the effectiveness of our multimodal pillar selection strategy by designing the following variants: 1) **Bottom Pil.**, collecting L items located in the tail of the intra- and inter-modal ranking list to constitute the pillar set; 2) **Random Pil.**, randomly and uniformly sampling L items from the intra- and inter-modal databases; 3) **w/o Inter-Pil.**, removing inter-modal pillars; and 4) **w/o Intra-Pil.**, excluding intra-modal pillars.

From Table 2, we observe that the first two variants Bottom Pil. and Random Pil. get performance degradation as compared to the full model, indicating that inferior pillars would deteriorate retrieval performance. Nevertheless, the performance achieved by these two variants is still comparable with the base backbone DIME* [34], demonstrating that our LeaPRR can tolerate bad pillars and avoid performance collapse. Besides, the improvements achieved by the full model LeaPRR compared with two single-pillar variants w/o Intra-Pil and w/o Inter-Pil. verifies the effectiveness of the top- L pillars selected from the intra- or inter-modal database. The remarkable performance improvement achieved by DIME*+LeaPRR strongly attests to the effectiveness of our proposed selection strategy and the complementary intra-modal and inter-modal pillars.

- **Impact of Neighbor-aware Graph Reasoning.** To assess how the neighbor-aware graph module in Section 3.3 contributes to the performance, we design three variants to compare with the base backbone DIME* and LeaPRR by removing the neighbor-based affinity, the learning-based affinity, and both. As displayed in Figure 4(c), the full model achieves the best I2T and T2I retrieval performance compared with other variants, which demonstrates

the effectiveness of the neighbor-aware graph reasoning module. Besides, we can observe **w/o NA** performs better than **w/o LA**, showing learning to estimate edge weights and construct the graph in the pillar space is more important, because it endows LeaPRR with flexibility. Finally, by comparing the two retrieval directions, more performance improvement by contrasting the full model and **w/o LA & NA** is achieved in I2T. One of the reasons is that the positive distribution is sparser in T2I than I2T. From this perspective, it verifies the existence of the sparsity challenge in Section 1.

• **Impact of Learning Objectives.** To thoroughly analyze the effectiveness of the three learning objectives presented in Section 3.4, we carried out ablation experiments by eliminating each of them. Concretely, we separately remove the global contrastive objective (**w/o Cont. Loss**), the local triplet objective (**w/o Trip. Loss**), and mutual modality alignment (**w/o MMA**). Besides, as to MMA, we further explored the contribution of each branch by building the following variants: 1) **w/o MMA, only i-t**, 2) **w/o MMA, only t-i**, and 3) **w/o MMA, i-t & t-i** (*i.e.*, separately training two single models and then combining them).

By analyzing the comparison results reported in Table 3, we have the following observations: 1) The contrastive objective and the triplet objective are able to enhance the learning of LeaPRR from different perspectives and collaborate with each other to improve the performance. Considering most of the raw content information has been abandoned in the pillar space, sufficient optimization driven by these two constraints is crucial; 2) Compared with the contrastive one, the triplet objective seems to contribute more, which reflects the importance of focusing on local hard structures (*e.g.*, relations among the query, hard positives, and hard negatives); 3) Two single submodels still gain obvious performance improvement in terms of all recall metrics strongly validate the effectiveness of our proposed re-ranking framework; 4) The MMA module plays a critical role in combining two single submodels, aligning cross-modal local structures, and mutual enhancement; And 5) considerable overall performance (rSum) improvement can be gained, though some of metrics (*e.g.*, R@5 in I2T) may sacrifice in ensemble models.

4.4.2 Parameter Sensitivity. To further delve into how LeaPRR performs with regard to pillars and neighbors, we conducted extensive experiments using varying numbers of pillars and neighbors.

• **Impact of the Pillar Number (L).** As shown in Figure 4(a), we evaluated and compared the retrieval performance by selecting various numbers of top-ranked pillars with a fixed exponential interval. From the compared results, we found that saturation points exist in all metrics, roughly at $L = 32$ or 64, except for I2T R@5. Too few pillars are inadequate to represent data samples, while too many pillars have a higher risk to introduce noisy pillars. From this knowable, an appropriate number of pillars is beneficial to make pillar representations rich and discriminative. More importantly, almost all the settings improve the base backbone from different depths, further indicating the effectiveness and robustness.

• **Impact of the Neighbor Size (K).** To investigate the sensitivity of LeaPRR to neighbors, we evaluated its performance under different numbers of neighbors, *i.e.*, K , in each submodel and reported the corresponding results in Figure 4(b).

These results show that: 1) The overall performance, *i.e.*, rSum, improves with the increasing neighbors at the beginning. When

Table 4: Quantitative results on the Flickr30K and MS-COCO (5K) datasets based on five base ITR architectures for generalization evaluation. The performance improvement compared with base backbones achieved by the proposed LeaPRR is marked in green.

Method	Image-to-Text		Text-to-Image	
	R@1	R@5	R@1	R@5
(Flickr30K)				
SCAN* [BCCV18] [17]	69.0	89.9	47.8	77.7
+ KRNN [ACMMM19] [38]	69.6	92.4	52.0	79.3
+ LeaPRR (Ours)	75.1 ↑6.1	93.4 ↑3.5	57.8 ↑10.0	79.3 ↑1.6
VSRN* [BCCV19] [19]	71.5	90.8	54.8	80.9
+ KRNN [ACMMM19] [38]	76.0	94.1	55.9	81.6
+ LeaPRR (Ours)	78.8 ↑7.3	93.6 ↑2.8	61.0 ↑6.2	82.5 ↑1.6
CAMERA* [ACMMM20] [33]	78.0	95.1	60.3	85.9
+ KRNN [ACMMM19] [38]	81.5	96.4	63.2	87.5
+ LeaPRR (Ours)	84.9 ↑6.9	96.8 ↑1.7	64.6 ↑4.3	87.4 ↑1.5
VSE∞ [CVPR21] [7]	81.7	95.4	61.4	85.9
+ KRNN [ACMMM19] [38]	84.7	96.4	64.6	87.7
+ LeaPRR (Ours)	86.2 ↑4.5	96.6 ↑1.2	66.6 ↑5.2	87.9 ↑2.0
DIME* [SIGIR21] [34]	81.0	95.9	63.6	88.1
+ KRNN [ACMMM19] [38]	83.6	96.8	65.2	88.3
+ LeaPRR (Ours)	86.2 ↑5.2	96.7 ↑0.8	71.1 ↑7.5	89.5 ↑1.4
(MS-COCO (5K))				
SCAN* [BCCV18] [17]	47.3	78.2	35.3	65.9
+ KRNN [ACMMM19] [38]	53.4	82.1	38.2	68.4
+ LeaPRR (Ours)	58.1 ↑10.8	84.2 ↑6.0	41.3 ↑6.0	69.4 ↑3.5
VSRN* [BCCV19] [19]	53.0	81.1	40.5	70.6
+ KRNN [ACMMM19] [38]	56.3	84.1	41.1	71.0
+ LeaPRR (Ours)	60.0 ↑7.0	84.8 ↑3.7	43.0 ↑2.5	71.7 ↑1.1
CAMERA* [ACMMM20] [33]	55.1	82.9	40.5	71.7
+ KRNN [ACMMM19] [38]	57.9	85.2	43.0	73.6
+ LeaPRR (Ours)	61.5 ↑6.4	85.9 ↑3.0	44.1 ↑3.6	73.0 ↑1.3
VSE∞ [CVPR21] [7]	58.3	85.3	42.4	72.7
+ KRNN [ACMMM19] [38]	61.6	87.5	45.3	74.5
+ LeaPRR (Ours)	65.0 ↑6.7	87.8 ↑2.5	46.8 ↑4.4	74.6 ↑1.9
DIME* [SIGIR21] [34]	59.3	85.4	43.1	73.0
+ KRNN [ACMMM19] [38]	62.9	86.6	43.8	73.1
+ LeaPRR (Ours)	64.3 ↑5.0	87.3 ↑1.9	45.5 ↑2.4	74.6 ↑1.6

*: ensemble two single base backbones

the number increases to a range of values for each submodel, the performance reaches the saturation point and then decreases. This indicates that exploiting a proper number of neighbors can enhance performance and increase the probability that bottom-ranked positives are recalled. And 2) compared with the *i-t* submodel, the *t-i* one gets saturated and starts to rapidly deteriorate earlier. It may attribute to the intrinsic ratio of positives to negatives in existing datasets. As described in Section 4.1, each image in both datasets corresponds to five sentences, *i.e.* more negative samples exist in the graph and the top-K neighbor set for the *t-i* submodel than for the *i-t* one. More importantly, regardless of the number of neighbors, all the settings manage to achieve overall performance improvement over the base backbone, which also verifies the effectiveness and robustness of LeaPRR.

4.4.3 Generalization Analysis. As discussed in Section 1, generalization is essential for re-ranking. To verify whether LeaPRR can perform well for different base ITR architectures, we plug it into five representative ones: two one-tower (SCAN* and DIME*) and three two-tower (VSRN*, CAMERA*, and VSE∞) architectures and compare it with KRNN [38]. Results on Flickr30K and MS-COCO (5K) are reported in Table 4. The superior performance, especially the R@1 performance, on both datasets validates the strong generalization capability of our method and the superiority over the current state-of-the-art KRNN, regardless of base backbones.

4.4.4 Transferability Analysis. The transferability includes backbone transferability and dataset transferability. These two reflect the zero-shot ability across different backbones and datasets, respectively. The bulk of previous efforts on ITR focuses on representing

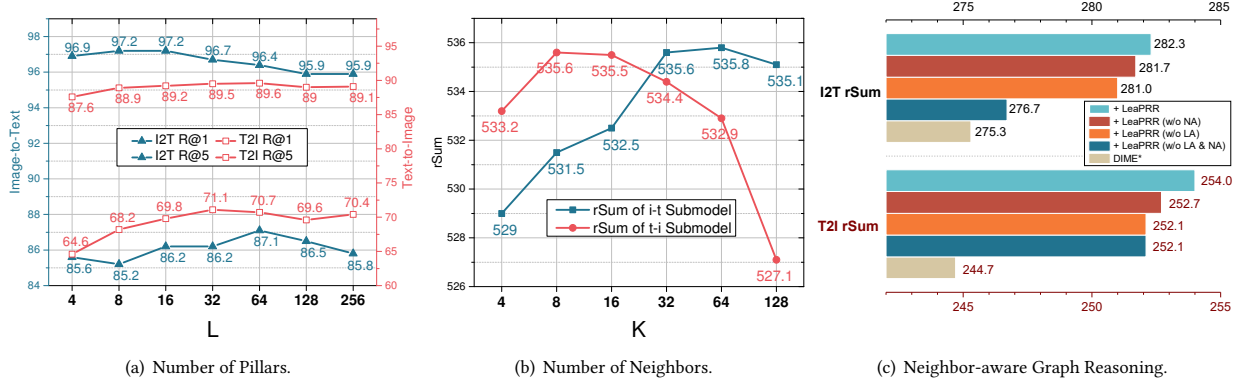


Figure 4: Impact of (a) the pillar number L , (b) the neighbor number K , and (c) the neighbor-aware graph reasoning module on retrieval performance. Evaluation performed on the Flickr30K dataset and the base backbone DIME* [34]. NA and LA in (c) denote neighbor-based affinity and learning-based affinity, respectively.




Figure 5: Results on Flickr30K based on five base ITR models for transferability evaluation. LeaPRR are trained on similarities calculated by source base backbones in the training set, and then are evaluated on similarities of the test set calculated by target base backbones. Δ rSum means the improvement achieved by LeaPRR compared with corresponding base backbones after being transferred.

and aligning visual and textual content. Due to the existence of a large discrepancy between content spaces, these methods suffer from limited transferability. In contrast, our proposed LeaPRR goes beyond raw content information and embraces pillars to exploit relations between entities for representation. We believe such pillar-based representations are more robust and transferable.

• **Backbone Transferability.** To investigate the backbone transferability, we train LeaPRR using similarities from a source base backbone and evaluated it on similarities of another target one. As shown by the impressive improvement on the overall metric (rSum) in Figure 5, LeaPRR exhibits superior transferability across almost all base ITR architectures. It is attributed to the robust pillar-based representation which is consistent across various architectures.

• **Dataset Transferability.** To further delve into the transferability across datasets, we also carried out corresponding experiments, the results of which are reported in Table 5. The comparison and

Table 5: Dataset transferability results achieved by LeaPRR on top of DIME* [34]. The transferability results and the overall performance improvement on rSum compared with base backbones are marked in bold and green, respectively.

Target \ Source	Flickr30K	MS-COCO
Base	520.0	438.8
Flickr30K	535.9 ↑15.9	446.8 ↑11.0
MS-COCO	529.1 ↑9.1	448.5 ↑12.7

performance improvement over the base backbone illustrates the outstanding dataset transferability of LeaPRR. It is worth noting that the improvement for Flickr30K \rightarrow MS-COCO is highly comparable to that of learning on MS-COCO, specifically 11.0 vs. 12.7, although the scale of Flickr30K is only a quarter of that of MS-COCO. We ascribe the strong transferability to the extensive soundness of the pillar-based representation, *i.e.*, representing an entity with its top-ranked neighbors, as discussed in Section 3.2.

5 CONCLUSION AND FUTURE WORK

In this paper, we proposed to model the higher-order neighbor relations for image-text retrieval in the re-ranking paradigm. To this end, we first delved into the limitations of prior re-ranking methods in the context of multi-modality and proposed four challenges for image-text retrieval re-ranking. To tackle the four issues, we then reformulated multi-modal re-ranking, constructed a new pillar space with top-ranked neighbors, and proposed a learning-based framework, which can fit in with one-tower and two-tower frameworks and flexibly explore complex neighbor relations among entities. Comprehensive experiments demonstrate the effectiveness, superiority, generalization, and transferability of the proposed method.

In the future, to further promote re-ranking performance, we plan to introduce abundant multimodal external knowledge for the modeling of the higher-order neighbor relations. Besides, inspired by the proposed re-ranking paradigm, we plan to explore the transferability between different downstream retrieval tasks from a broader perspective.

REFERENCES

- [1] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-Up and Top-Down Attention for Image Captioning and Visual Question Answering. In *CVPR*. IEEE, 6077–6086.
- [2] Relja Arandjelović and Andrew Zisserman. 2012. Three Things Everyone should Know to Improve Object Retrieval. In *CVPR*. IEEE, 2911–2918.
- [3] Song Bai and Xiang Bai. 2016. Sparse Contextual Activation for Efficient Visual Re-ranking. *TIP* 25, 3 (2016), 1056–1069.
- [4] Fei Cai, Shangsong Liang, and Maarten De Rijke. 2014. Personalized Document Re-ranking based on Bayesian Probabilistic Matrix Factorization. In *SIGIR*. ACM, 835–838.
- [5] Micael Carvalho, Rémi Cadène, David Picard, Laure Soulier, Nicolas Thome, and Matthieu Cord. 2018. Cross-modal Retrieval in the Cooking Context: Learning Semantic Text-Image Embeddings. In *SIGIR*. ACM, 35–44.
- [6] Hui Chen, Guiguang Ding, Xudong Liu, Zijia Lin, Ji Liu, and Jungong Han. 2020. IMRAM: Iterative Matching with Recurrent Attention Memory for Cross-Modal Image-Text Retrieval. In *CVPR*. IEEE, 12655–12663.
- [7] Jiacheng Chen, Hexiang Hu, Hao Wu, Yuning Jiang, and Changhu Wang. 2021. Learning the Best Pooling Strategy for Visual Semantic Embedding. In *CVPR*. IEEE, 15789–15798.
- [8] Ondrej Chum, James Philbin, Josef Sivic, Michael Isard, and Andrew Zisserman. 2007. Total Recall: Automatic Query Expansion with a Generative Feature Model for Object Retrieval. In *ICCV*. IEEE, 1–8.
- [9] Haiwen Diao, Ying Zhang, Lin Ma, and Huchuan Lu. 2021. Similarity Reasoning and Filtration for Image-Text Matching. In *AAAI*. AAAI Press, 1218–1226.
- [10] Fartash Faghri, David J Fleet, Jamie Ryan Kiros, and Sanja Fidler. 2018. VSE++: Improving Visual-Semantic Embeddings with Hard Negatives. In *BMVC*. BMVA Press, 1–13.
- [11] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc'Aurelio Ranzato, and Tomas Mikolov. 2013. DeViSE: A Deep Visual-Semantic Embedding Model. In *NeurIPS*. Curran Associates, Inc., 2121–2129.
- [12] Albert Gordo, Jon Almazan, Jerome Revaud, and Diane Larlus. 2017. End-to-end Learning of Deep Visual Representations for Image Retrieval. *IJCV* 124, 2 (2017), 237–254.
- [13] Albert Gordo, Filip Radenovic, and Tamara Berg. 2020. Attention-based Query Expansion Learning. In *ECCV*. Springer, 172–188.
- [14] Raia Hadsell, Sumit Chopra, and Yann LeCun. 2006. Dimensionality Reduction by Learning an Invariant Mapping. In *CVPR*. IEEE, 1735–1742.
- [15] Ahmet Iscen, Yannis Avrithis, Giorgos Tolias, Teddy Furon, and Ondrej Chum. 2018. Fast Spectral Ranking for Similarity Search. In *CVPR*. IEEE, 7632–7641.
- [16] Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, Teddy Furon, and Ondrej Chum. 2017. Efficient Diffusion on Region Manifolds: Recovering Small Objects with Compact CNN Representations. In *CVPR*. IEEE, 2077–2086.
- [17] Kuang-Huei Lee, Xi Chen, Gang Hua, Houdong Hu, and Xiaodong He. 2018. Stacked Cross Attention for Image-Text Matching. In *ECCV*. Springer, 201–216.
- [18] Jianjun Lei, Tianyi Qin, Bo Peng, Wanqing Li, Zhaoqing Pan, Haifeng Shen, and Sam Kwong. 2022. Reducing Background Induced Domain Shift for Adaptive Person Re-Identification. *TII* (2022), 1–12.
- [19] Kunpeng Li, Yulun Zhang, Kai Li, Yuanyuan Li, and Yun Fu. 2019. Visual Semantic Reasoning for Image-Text Matching. In *ICCV*. IEEE, 4654–4662.
- [20] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft COCO: Common Objects in Context. In *ECCV*. Springer, 740–755.
- [21] Chunxiao Liu, Zhedong Mao, Tianzhu Zhang, Hongtao Xie, Bin Wang, and Yongdong Zhang. 2020. Graph Structured Network for Image-Text Matching. In *CVPR*. IEEE, 10921–10930.
- [22] Chundi Liu, Guangwei Yu, Maksims Volkovs, Cheng Chang, Himanshu Rai, Junwei Ma, and Satya Krishna Gorti. 2019. Guided Similarity Separation for Image Retrieval. In *NeurIPS*. Curran Associates, Inc., 1–12.
- [23] Meng Liu, Xiang Wang, Liqiang Nie, Xiangnan He, Baoquan Chen, and Tat-Seng Chua. 2018. Attentive Moment Retrieval in Videos. In *SIGIR*. ACM, 15–24.
- [24] Sean MacAvaney, Franco Maria Nardini, Raffaele Perego, Nicola Tonello, Nazli Goharian, and Ophir Frieder. 2020. Efficient Document Re-ranking for Transformers by Precomputing Term Representations. In *SIGIR*. ACM, 49–58.
- [25] Sean MacAvaney, Franco Maria Nardini, Raffaele Perego, Nicola Tonello, Nazli Goharian, and Ophir Frieder. 2020. Training curricula for open domain answer re-ranking. In *SIGIR*. ACM, 529–538.
- [26] Sean MacAvaney, Andrew Yates, Kai Hui, and Ophir Frieder. 2019. Content-based Weak Supervision for Ad-hoc Re-ranking. In *SIGIR*. ACM, 993–996.
- [27] Yoshitomo Matsubara, Thuy Vu, and Alessandro Moschitti. 2020. Reranking for Efficient Transformer-based Answer Selection. In *SIGIR*. ACM, 1577–1580.
- [28] Jianbo Ouyang, Hui Wu, Min Wang, Wengang Zhou, and Houqiang Li. 2021. Contextual Similarity Aggregation with Self-attention for Visual Re-ranking. In *NeurIPS*. Curran Associates, Inc., 3135–3148.
- [29] Liang Pang, Jun Xu, Qingyao Ai, Yanyan Lan, Xueqi Cheng, and Jirong Wen. 2020. Setrank: Learning a Permutation-invariant Ranking Model for Information Retrieval. In *SIGIR*. ACM, 499–508.
- [30] Shanmin Pang, Jin Ma, Jianru Xue, Jihua Zhu, and Vicente Ordonez. 2018. Deep Feature Aggregation and Image Re-ranking with Heat Diffusion for Image Retrieval. *TMM* 21, 6 (2018), 1513–1523.
- [31] James Philbin, Ondrej Chum, Michael Isard, Josef Sivic, and Andrew Zisserman. 2008. Lost in Quantization: Improving Particular Object Retrieval in Large Scale Image Databases. In *CVPR*. IEEE, 1–8.
- [32] Danfeng Qin, Stephan Gammeter, Lukas Bossard, Till Quack, and Luc Van Gool. 2011. Hello Neighbor: Accurate Object Retrieval with K-reciprocal Nearest Neighbors. In *CVPR*. IEEE, 777–784.
- [33] Leigang Qu, Meng Liu, Da Cao, Liqiang Nie, and Qi Tian. 2020. Context-Aware Multi-View Summarization Network for Image-Text Matching. In *MM*. ACM, 1047–1055.
- [34] Leigang Qu, Meng Liu, Jianlong Wu, Zan Gao, and Liqiang Nie. 2021. Dynamic Modality Interaction Modeling for Image-Text Retrieval. In *SIGIR*. ACM, 1104–1113.
- [35] Filip Radenović, Giorgos Tolias, and Ondrej Chum. 2018. Fine-tuning CNN Image Retrieval with no Human Annotation. *TPAMI* 41, 7 (2018), 1655–1668.
- [36] Jerome Revaud, Jon Almazán, Rafael S Rezende, and Cesar Roberto de Souza. 2019. Learning with Average Precision: Training Image Retrieval with a Listwise Loss. In *ICCV*. IEEE, 5107–5116.
- [37] Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. Facenet: A Unified Embedding for Face Recognition and Clustering. In *CVPR*. IEEE, 815–823.
- [38] Tan Wang, Xing Xu, Yang Yang, Alan Hanjalic, Heng Tao Shen, and Jingkuan Song. 2019. Matching Images and Text with Multi-modal Tensor Fusion and Re-ranking. In *Proceedings of the ACM International Conference on Multimedia*. ACM, 12–20.
- [39] Xi Wei, Tianzhu Zhang, Yan Li, Yongdong Zhang, and Feng Wu. 2020. Multi-Modality Cross Attention Network for Image and Sentence Matching. In *CVPR*. IEEE, 10941–10950.
- [40] Haokun Wen, Xuemeng Song, Xin Yang, Yibing Zhan, and Liqiang Nie. 2021. Comprehensive Linguistic-visual Composition Network for Image Retrieval. In *SIGIR*. ACM, 1369–1378.
- [41] Yunjia Xi, Weiwen Liu, Jieming Zhu, Xilong Zhao, Xinyi Dai, Ruiming Tang, Weinan Zhang, Rui Zhang, and Yong Yu. 2022. Multi-Level Interaction Reranking with User Behavior History. In *SIGIR*. ACM, 1336–1346.
- [42] Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. From Image Descriptions to Visual Denotations: New Similarity Metrics for Semantic Inference over Event Descriptions. *TACL* 2 (2014), 67–78.
- [43] George Zerveas, Navid Rekasas, Daniel Cohen, and Carsten Eickhoff. 2022. Mitigating Bias in Search Results Through Contextual Document Reranking and Neutrality Regularization. In *SIGIR*. ACM, 2532–2538.
- [44] Xuanmeng Zhang, Minyue Jiang, Zhedong Zheng, Xiao Tan, Errui Ding, and Yi Yang. 2020. Understanding Image Retrieval Re-ranking: A Graph Neural Network Perspective. *arXiv:2012.07620* (2020).
- [45] Ying Zhang and Huchuan Lu. 2018. Deep Cross-Modal Projection Learning for Image-Text Matching. In *ECCV*. Springer, 686–701.
- [46] Zhedong Zheng, Tao Ruan, Yunchao Wei, Yi Yang, and Tao Mei. 2020. VehicleNet: Learning Robust Visual Representation for Vehicle Re-identification. *TMM* 23, 1520–9210 (2020), 2683–2693.
- [47] Zhedong Zheng, Liang Zheng, Michael Garrett, Yi Yang, Mingliang Xu, and Yi-Dong Shen. 2020. Dual-Path Convolutional Image-Text Embeddings with Instance Loss. *TOMM* 16, 2 (2020), 1–23.
- [48] Zhun Zhong, Liang Zheng, Donglin Cao, and Shaozi Li. 2017. Re-ranking Person Re-identification with K-reciprocal Encoding. In *CVPR*. IEEE, 1318–1327.
- [49] Shengyao Zhuang and Guido Zuccon. 2021. TILDE: Term Independent Likelihood Model for Passage Re-ranking. In *SIGIR*. ACM, 1483–1492.