

Learning Token-Based Representation for Image Retrieval

Hui Wu¹, Min Wang^{2*}, Wengang Zhou^{1,2*}, Yang Hu¹, Houqiang Li^{1,2}

¹ CAS Key Laboratory of GIPAS, University of Science and Technology of China

² Institute of Artificial Intelligence, Hefei Comprehensive National Science Center
wh241300@mail.ustc.edu.cn, wangmin@iai.ustc.edu.cn, {zhwg, eeyhu, lihq}@ustc.edu.cn

Abstract

In image retrieval, deep local features learned in a data-driven manner have been demonstrated effective to improve retrieval performance. To realize efficient retrieval on large image database, some approaches quantize deep local features with a large codebook and match images with aggregated match kernel. However, the complexity of these approaches is non-trivial with large memory footprint, which limits their capability to jointly perform feature learning and aggregation. To generate compact global representations while maintaining regional matching capability, we propose a unified framework to jointly learn local feature representation and aggregation. In our framework, we first extract deep local features using CNNs. Then, we design a tokenizer module to aggregate them into a few visual tokens, each corresponding to a specific visual pattern. This helps to remove background noise, and capture more discriminative regions in the image. Next, a refinement block is introduced to enhance the visual tokens with self-attention and cross-attention. Finally, different visual tokens are concatenated to generate a compact global representation. The whole framework is trained end-to-end with image-level labels. Extensive experiments are conducted to evaluate our approach, which outperforms the state-of-the-art methods on the Revisited Oxford and Paris datasets.

Introduction

Given a large image corpus, image retrieval aims to efficiently find target images similar to a given query. It is challenging due to various situations observed in large-scale dataset, *e.g.*, occlusions, background clutter, and dramatic viewpoint changes. In this task, image representation, which describes the content of images to measure their similarities, plays a crucial role. With the introduction of deep learning into computer vision, significant progress (Babenko et al. 2014; Gordo et al. 2017; Cao, Araujo, and Sim 2020; Xu et al. 2018; Noh et al. 2017) has been witnessed in learning image representation for image retrieval in a data-driven paradigm. Generally, there are two main types of representation for image retrieval. One is global feature, which maps an image to a compact vector, while the other is local feature, where an image is described with hundreds of short vectors.

*Corresponding Author

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

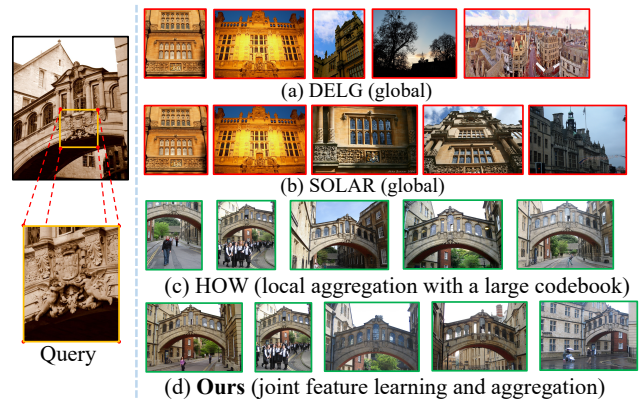


Figure 1: Top-5 retrieval results of different methods, including DELG (Cao, Araujo, and Sim 2020), SOLAR (Ng et al. 2020), HOW (Tolias, Jenicek, and Chum 2020) and ours. Query image is on the left (black outline) with a target object (orange box), and the right are the top-ranking images for the query. Our approach achieves similar results as HOW, which use large visual codebook to aggregate local features, with *lower memory and latency*. Green solid outline: positive images for the query; red solid outline: negative results.

In global feature based image retrieval (Radenović, Tolias, and Chum 2018; Babenko and Lempitsky 2015), although the representation is compact, it usually lacks capability to retrieve target images with only partial match. As shown in Fig. 1 (a) and (b), when the query image occupies only a small region in the target images, global features tend to return false positive examples, which are somewhat similar but do not indicate the same instance as the query image.

Recently, many studies have demonstrated the effectiveness of combining deep local features (Tolias, Jenicek, and Chum 2020; Noh et al. 2017; Teichmann et al. 2019) with traditional ASMK (Tolias, Avrithis, and Jégou 2013) aggregation method in dealing with background clutter and occlusion. In those approaches, the framework usually consists of two stages: feature extraction and feature aggregation, where the former extracts discriminative local features, which are further aggregated by the latter for the efficient retrieval. However, they require offline clustering and coding proce-

dures, which lead to a considerable complexity of the whole framework with a high memory footprint and long retrieval latency. Besides, it is difficult to jointly learn local features and aggregation due to the involvement of large visual codebook and hard assignment in quantization.

Some existing works such as NetVLAD (Arandjelovic et al. 2016) try to learn local features and aggregation simultaneously. They aggregate the feature maps output by CNNs into compact global features with a learnable VLAD layer. Specifically, they discard the original features and adopt the sum of residual vectors of each visual word as the representation of an image. However, considering the large variation and diversity of content in different images, these visual words are too coarse-grained for the features of a particular image. This leads to insufficient discriminative capability of the residual vectors, which further hinders the performance of the aggregated image representation.

To address the above issues, we propose a unified framework to jointly learn and aggregate deep local features. We treat the feature map output by CNNs as original deep local features. To obtain compact image representations while preserving the regional matching capability, we propose a tokenizer to adaptively divide the local features into groups with spatial attention. These local features are further aggregated to form the corresponding visual tokens. Intuitively, the attention mechanism ensures that each visual token corresponds to some visual pattern and these patterns are aligned across images. Furthermore, a refinement block is introduced to enhance the obtained visual tokens with self-attention and cross-attention. Finally, the updated attention maps are used to aggregate original local features for enhancing the existing visual tokens. The whole framework is trained end-to-end with only image-level labels.

Compared with the previous methods, there are two advantages in our approach. First, by expressing an image with a few visual tokens, each corresponding to some visual pattern, we implicitly achieve local pattern alignment with the aggregated global representation. As shown in Fig. 1 (d), our approach performs well in the presence of background clutter and occlusion. Secondly, the global representation obtained by aggregation is compact with a small memory footprint. These facilitate effective and efficient semantic content matching between images. We conduct comprehensive experiments on the Revisited Oxford and Paris datasets, which are further mixed with one million distractors. Ablation studies demonstrate the effectiveness of the tokenizer and the refinement block. Our approach surpasses the state-of-the-art methods by a considerable margin.

Related Work

In this section, we briefly review the related work including local feature and global feature based image retrieval.

Local feature. Traditionally local features (Lowe 2004; Bay, Tuytelaars, and Van Gool 2006) are extracted using hand-crafted detectors and descriptors. They are first organized in bag-of-words (Sivic and Zisserman 2003; Zhou et al. 2010) and further enhanced by spatial validation (Philbin et al. 2007), hamming embedding (Jégou, Douze, and Schmid 2008) and query expansion (Chum

et al. 2007). Recently, tremendous advances (Mishchuk et al. 2017; Tolias, Jenicek, and Chum 2020; Dmytro Mishkin 2018; Noh et al. 2017; Tian et al. 2019; Cao, Araujo, and Sim 2020) have been made to learn local features suitable for image retrieval in a data-driven manner. Among these approaches, the state-of-the-art approach is HOW (Tolias, Jenicek, and Chum 2020), which uses attention learning to distinguish deep local features with image-level annotations. During testing, it combines the obtained local features with the traditional ASMK (Tolias, Avrithis, and Jégou 2013) aggregation method. However, HOW cannot jointly learn feature representation and aggregation due to the very large codebook and the hard assignment during the quantization process. Moreover, its complexity is considerable with a high memory footprint. Our method uses a few visual tokens to effectively represent image. The feature representation and aggregation are jointly learned.

Global feature. Compact global features reduce memory footprint and expedite the retrieval process. They simplify image retrieval to a nearest neighbor search and extend the previous query expansion (Chum et al. 2007) to an efficient exploration of the entire nearest neighbor graph of the dataset by diffusion (Fan et al. 2019). Before deep learning, they are mainly developed by aggregating hand-crafted local features, *e.g.*, VLAD (Jégou et al. 2011), Fisher vectors (Perronnin et al. 2010), ASMK (Tolias, Avrithis, and Jégou 2013). Recently, global features are obtained simply by performing the pooling operation on the feature map of CNNs. Many pooling methods have been explored, *e.g.*, max-pooling (MAC) (Tolias, Sicre, and Jégou 2016), sum-pooling (SPoC) (Babenko and Lempitsky 2015), weighted-sum-pooling (CroW) (Kalantidis, Mellina, and Osindero 2016), regional-max-pooling (R-MAC) (Tolias, Sicre, and Jégou 2016), generalized mean-pooling (GeM) (Radenović, Tolias, and Chum 2018), and so on. These networks are trained using ranking (Radenović, Tolias, and Chum 2018; Revaud et al. 2019) or classification losses (Deng et al. 2019). Differently, our method tokenizes the feature map into several visual tokens, enhances the visual tokens using the refinement block, concatenates different visual tokens and performs dimension reduction. Through these steps, our method generates a compact global representation while maintaining the regional matching capability.

Methodology

An overview of our framework is shown in Fig. 2. Given an image, we first obtain the original deep local features $F \in \mathcal{R}^{C \times H \times W}$ through a CNN backbone. These local features are obtained with limited receptive fields covering part of the input image. Thus, we follow (Ng et al. 2020) to apply the Local Feature Self-Attention (LFSA) operation on F to obtain context-aware local features $F^c \in \mathcal{R}^{C \times H \times W}$. Next, we divide them into L groups with spatial attention mechanism, and the local features of each group are aggregated to form a visual token $t \in \mathcal{R}^C$. We denote the set of obtained visual tokens as $T = [t^{(1)}, t^{(2)}, \dots, t^{(L)}] \in \mathcal{R}^{L \times C}$. Furthermore, we introduce a refinement block to update the obtained visual tokens T based on the previous local features F^c . Finally, all the visual tokens are concatenated and

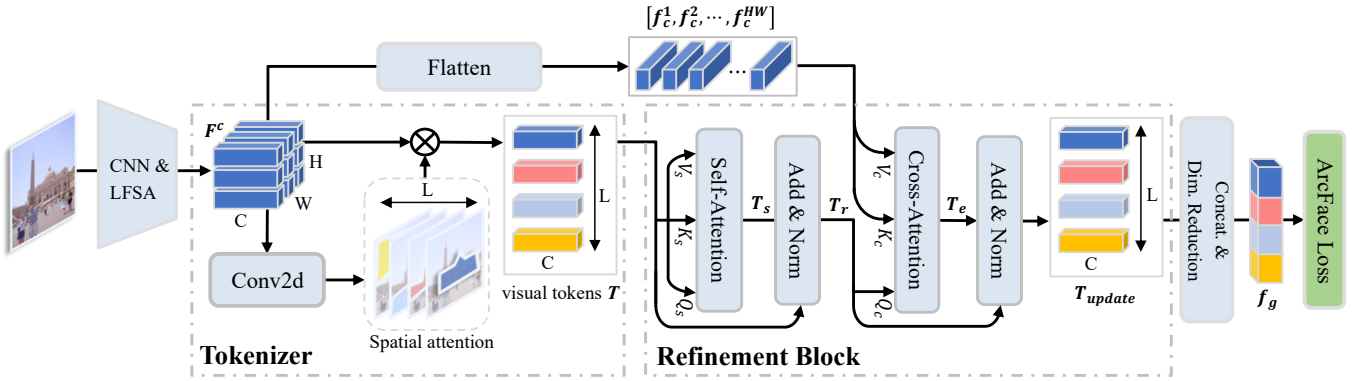


Figure 2: An overview of our framework. Given an image, we first use a CNN and a Local Feature Self-Attention (LFSA) module to extract local features F^c . Then, they are tokenized into L visual tokens with spatial attention. Further, a refinement block is introduced to enhance the obtained visual tokens with self-attention and cross-attention. Finally, we concatenate all the visual tokens to form a compact global representation f_g and reduce its dimension.

we reduce its dimension to form the final global descriptor f_g . ArcFace margin loss is used to train the whole network.

Tokenizer

To effectively cope with the challenging conditions observed in large datasets, such as noisy backgrounds, occlusions, *etc.*, image representation is expected to find patch-level matches between images. A typical pipeline to tackle these challenges consists of local descriptor extraction, quantization with a large visual codebook created usually by k -means and descriptor aggregation into a single embedding. However, due to the offline clustering and hard assignment of local features, it is difficult to optimize feature learning and aggregation simultaneously, which further limits the discriminative power of the image representation. To alleviate this problem, we here use spatial attention to extract the desired visual tokens. By training, the attention module can adaptively discover discriminative visual patterns.

For the set of local features F^c , we generate L attention maps $\mathbf{A} = \{\mathbf{a}^{(1)}, \mathbf{a}^{(2)}, \dots, \mathbf{a}^{(L)}\}$, which are implemented by L 1×1 convolutional layers. We denote the parameters of the convolution layers as $\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_L] \in \mathcal{R}^{L \times C}$, and the attention maps are calculated as

$$\mathbf{a}_{h,w}^{(i)} = \frac{\exp(\mathbf{w}_i \cdot \mathbf{F}_{h,w}^c)}{\sum_{l=1}^L \exp(\mathbf{w}_l \cdot \mathbf{F}_{h,w}^c)}. \quad (1)$$

Then, the visual tokens $\mathbf{T} = [\mathbf{t}^{(1)}, \mathbf{t}^{(2)}, \dots, \mathbf{t}^{(L)}]$ are computed as

$$\mathbf{t}^{(i)} = \frac{1}{\gamma(\mathbf{a}^{(i)})} \sum_{h \in H, w \in W} \mathbf{a}_{h,w}^{(i)} \mathbf{F}_{h,w}^c, \quad (2)$$

where $\gamma(\mathbf{a}^{(i)}) = \sum_{h \in H, w \in W} \mathbf{a}_{h,w}^{(i)}$ and $\mathbf{t}^{(i)} \in \mathcal{R}^C$.

Relation to GMM. Tokenization aggregates local features into visual tokens, which helps to capture discriminative visual patterns, leading to a more general and robust image representation. The visual tokens represent several specific

region patterns in an image, which share the similar rationale of learning a Gaussian Mixture Model (GMM) over the original local features of an image. GMM is a probabilistic model that assumes all the data points are generated from a mixture of a finite number of Gaussian distributions with unknown mean vectors and data variances. Formally,

$$p(\mathbf{f}_i | z = j) = \frac{1}{(2\pi\sigma_j^2)^{\frac{C}{2}}} \exp\left(-\frac{1}{2\sigma_j^2} \|\mathbf{f}_i - \mathbf{c}_j\|^2\right), \quad (3)$$

$$p(\mathbf{f}_i) = \sum_{j=1}^N p(z = j) p(\mathbf{f}_i | z = j).$$

Here, N is the number of Gaussian mixtures, \mathbf{f}_i is the local feature of an image with dimension C . z is the latent cluster assignment variable, \mathbf{c}_j and σ_j correspond to the mean vector and variance of the j -th Gaussian distribution, respectively. The Expectation-Maximization algorithm is commonly used to solve this problem. Iteratively, it estimates for each point a probability of being generated by each component of the model and updates the mean vectors:

$$\mathbf{c}_j = \frac{\sum_{i=1}^M p(z = j | \mathbf{f}_i) \mathbf{f}_i}{\sum_{i=1}^M p(z = j | \mathbf{f}_i)}, \quad j = 1, 2, \dots, N, \quad (4)$$

$$p(z = j | \mathbf{f}_i) = \frac{p(z = j) \cdot \exp\left(-\frac{\|\mathbf{f}_i - \mathbf{c}_j\|^2}{2\sigma_j^2}\right)}{\sum_{l=1}^N p(z = l) \cdot \exp\left(-\frac{\|\mathbf{f}_i - \mathbf{c}_l\|^2}{2\sigma_l^2}\right)},$$

where M is the total feature number. In GMM, $p(z = j | \mathbf{f}_i)$ represents the posterior probability of a local feature $\mathbf{f}_i \in \mathcal{R}^C$ being assigned to the j -th cluster.

In our approach, considering that $\mathbf{w}_i \cdot \mathbf{F}_{h,w}^c = \frac{1}{2} \|\mathbf{F}_{h,w}^c\|^2 + \frac{1}{2} \|\mathbf{w}_i\|^2 - \frac{1}{2} \|\mathbf{F}_{h,w}^c - \mathbf{w}_i\|^2$, Eq. (1) can be reformulated as

$$\mathbf{a}_{h,w}^{(i)} = \frac{\phi(\|\mathbf{w}_i\|^2) \cdot \exp\left(-\frac{1}{2} \|\mathbf{F}_{h,w}^c - \mathbf{w}_i\|^2\right)}{\sum_{l=1}^L \phi(\|\mathbf{w}_l\|^2) \cdot \exp\left(-\frac{1}{2} \|\mathbf{F}_{h,w}^c - \mathbf{w}_l\|^2\right)}, \quad (5)$$

where $\phi(\|\mathbf{w}_i\|^2) = \exp(\frac{1}{2}\|\mathbf{w}_i\|^2) / \sum_{j=1}^L \exp(\frac{1}{2}\|\mathbf{w}_j\|^2)$. We set σ in Eq. (3) as 1 and $p(z = j)$ as $\phi(\|\mathbf{w}_i\|^2)$. $\mathbf{a}_{h,w}^{(i)}$ can be interpreted as a soft cluster assignment of local features $\mathbf{F}_{h,w}^c$ to i -th visual pattern, which is the same as the meaning of $p(z = j|\mathbf{f}_i)$. \mathbf{w}_i is the mean vector corresponding to the i -th visual pattern, whose L_2 norm is proportional to the probability of occurrence of that visual pattern. Since \mathbf{W} is the network parameter that learned from the image data distribution, it enables the alignment of the corresponding visual patterns across different images. Further, as shown in Eq. (2), the visual token $\mathbf{t}^{(i)}$ obtained by weighted aggregation are equivalent to the updated mean vectors \mathbf{c}_j in GMM.

Refinement Block

The two major components of our refinement block are the relationship modeling and visual token enhancement. The former allows the propagation of information between tokens, which helps to produce more robust features, while the latter is utilized to relocate the visual patterns in the original local features and extract the corresponding features for enhancing the existing visual tokens \mathbf{T} .

Relationship modeling. During tokenization, different attention maps are used separately. It excludes any relative contribution of each visual token to the other tokens. Thus, we employ the self-attention mechanism (Vaswani et al. 2017) to model the relationship between different visual tokens $\mathbf{t}^{(i)}$, generating a set of relation-aware visual tokens $\mathbf{T}_r = [\mathbf{t}_r^{(1)}, \mathbf{t}_r^{(2)}, \dots, \mathbf{t}_r^{(L)}] \in \mathcal{R}^{L \times C}$. Specifically, we first map visual tokens \mathbf{T} to Queries ($\mathbf{Q}_s \in \mathcal{R}^{L \times C}$), Keys ($\mathbf{K}_s \in \mathcal{R}^{L \times C}$) and Values ($\mathbf{V}_s \in \mathcal{R}^{L \times C}$) with three $C \times C$ -dimensional learnable matrices. After that, the similarity $\mathbf{S} \in \mathcal{R}^{L \times L}$ between visual tokens is computed through

$$\mathbf{S}(\mathbf{Q}_s, \mathbf{K}_s) = \text{SOFTMAX}\left(\frac{\mathbf{Q}_s \mathbf{K}_s^T}{\sqrt{C}}\right) \in \mathcal{R}^{L \times L}, \quad (6)$$

where the normalized similarity $\mathbf{S}_{i,j}$ models the correlation between different visual tokens \mathbf{t}^i and \mathbf{t}^j . To focus on multiple semantically related visual tokens simultaneously, we calculate similarity \mathbf{S} with Multi-Head Attention (MHA).

In MHA, different projection matrices for Queries, Keys, and Values are used for different heads, and these matrices map visual tokens to different subspaces. After that, $\mathbf{S}^{(i)}$ of each head, calculated by Eq. (6), is used to aggregate semantically related visual tokens. MHA then concatenates and fuses the outputs of different heads using the learnable projection $\mathbf{W}_M \in \mathcal{R}^{C \times C}$. Formally,

$$\begin{aligned} \mathbf{T}_s^{(i)} &= \text{DROPOUT}(\mathbf{S}^{(i)} \mathbf{V}_s^{(i)}), \text{ for } i = 1, 2, \dots, N, \\ \mathbf{T}_s &= \text{CONCAT}(\mathbf{T}_s^{(1)}, \mathbf{T}_s^{(2)}, \dots, \mathbf{T}_s^{(N)}) \mathbf{W}_M, \end{aligned} \quad (7)$$

where N is head number and $\mathbf{T}_s^{(i)}$ is output of the i -th head. Finally, \mathbf{T}_s is normalized via Layer Normalization and added to original \mathbf{T} to produce the relation aware visual tokens:

$$\mathbf{T}_r = \mathbf{T} + \text{LAYERNORM}(\mathbf{T}_s). \quad (8)$$

Visual token enhancement. To further enhance the existing visual tokens, we next extract features from \mathbf{F}^c with the

cross-attention mechanism. As shown in Fig. 2, we first flatten \mathbf{F}^c into a sequence $[\mathbf{f}_c^1, \mathbf{f}_c^2, \dots, \mathbf{f}_c^{HW}] \in \mathcal{R}^{HW \times C}$. Then, with different fully-connected (FC) layers, \mathbf{T}_r is mapped to Queries ($\mathbf{Q}_c \in \mathcal{R}^{L \times C}$) and \mathbf{F}^c is mapped to Keys ($\mathbf{K}_c \in \mathcal{R}^{HW \times C}$) and Values ($\mathbf{V}_c \in \mathcal{R}^{HW \times C}$), respectively. The similarity between the visual tokens \mathbf{T}_r and the original local feature \mathbf{F}^c is calculated as

$$\mathbf{S}(\mathbf{Q}_c, \mathbf{K}_c) = \text{SOFTMAX}\left(\frac{\mathbf{Q}_c \mathbf{K}_c^T}{\sqrt{C}}\right) \in \mathcal{R}^{L \times HW}. \quad (9)$$

Here, the similarity $\mathbf{S}_{i,j}$ indicates the probability that the j -th local feature \mathbf{f}_c^j in \mathbf{F}^c should be assigned to the i -th visual token, which is different from the meaning \mathbf{S} in Eq. (6). Then, the weighted sum of \mathbf{F}^c and \mathbf{S} is added to \mathbf{T}_r to produce the updated visual tokens:

$$\begin{aligned} \mathbf{T}_c &= \text{DROPOUT}(\mathbf{S} \mathbf{V}_c), \\ \mathbf{T}_{update} &= \mathbf{T}_r + \text{LAYERNORM}(\mathbf{T}_c). \end{aligned} \quad (10)$$

As in Eq. (7), MHA is also used to calculate the similarity.

We stack N refinement blocks to obtain more discriminative visual tokens. The refined visual tokens $\mathbf{T}_{update} \in \mathcal{R}^{L \times C}$ come from the output of the last block of our model. We concatenate the different visual tokens \mathbf{T}_{update} into a global descriptor and a fully-connected layer is adopted to reduce its dimension to d :

$$\mathbf{f}_g = \text{CONCAT}(\mathbf{t}_{update}^{(1)}, \mathbf{t}_{update}^{(2)}, \dots, \mathbf{t}_{update}^{(L)}) \mathbf{W}_g, \quad (11)$$

where $\mathbf{W}_g \in \mathcal{R}^{LC \times d}$ is the weight of the FC layer.

Training Objectives

Following DELG (Cao, Araujo, and Sim 2020), ArcFace margin loss (Deng et al. 2019) is adopted to train the whole model. The ArcFace improves the normalization of the classifier weight vector $\hat{\mathbf{W}}$ and the interval of the additive angles m so as to enhance the separability between classes and meantime enhance the compactness within class. It has shown excellent results for global descriptor learning by inducing smaller intra-class variance. Formally,

$$L = -\log \left(\frac{\exp(\gamma \times AF(\hat{\mathbf{w}}_k^T \hat{\mathbf{f}}_g, 1))}{\sum_n \exp(\gamma \times AF(\hat{\mathbf{w}}_n^T \hat{\mathbf{f}}_g, y_n))} \right), \quad (12)$$

where $\hat{\mathbf{w}}_i$ refers to the i -th row of $\hat{\mathbf{W}}$ and $\hat{\mathbf{f}}_g$ is the L_2 -normalized \mathbf{f}_g . y is the one-hot label vector and k is the ground-truth class index ($y_k = 1$). γ is a scale factor. AF denotes the adjusted cosine similarity and it is calculated as:

$$AF(s, c) = (1 - c) \times s + c \times \cos(\text{acos}(s) + m), \quad (13)$$

where s is the cosine similarity, m is the ArcFace margin, and c is a binarized value that denotes whether it is the ground-truth category.

Experiments

Experimental Setup

Training dataset. The clean version of Google landmarks dataset V2 (GLDv2-clean) (Weyand et al. 2020) is used for

METHOD	MEDIUM				HARD			
	$\mathcal{R}Oxf$	$\mathcal{R}Oxf+\mathcal{R}1M$	$\mathcal{R}Par$	$\mathcal{R}Par+\mathcal{R}1M$	$\mathcal{R}Oxf$	$\mathcal{R}Oxf+\mathcal{R}1M$	$\mathcal{R}Par$	$\mathcal{R}Par+\mathcal{R}1M$
<i>(A) Local feature aggregation</i>								
HesAff-rSIFT-ASMK*+SP	60.60	46.80	61.40	42.30	36.70	26.90	35.00	16.80
DELf-ASMK*+SP(GLDv1-noisy)	67.80	53.80	76.90	57.30	43.10	31.20	55.40	26.40
DELf-R-ASMK*+SP(GLDv1-noisy)	76.00	64.00	80.20	59.70	52.40	38.10	58.60	58.60
R50-HOW-ASMK*(SfM-120k)	79.40	65.80	81.60	61.80	56.90	38.90	62.40	33.70
R101-HOW-VLAD†(GLDv2-clean)	73.54	60.38	82.33	62.56	51.93	33.17	66.95	41.82
R101-HOW-ASMK*†(GLDv2-clean)	80.42	70.17	85.43	68.80	62.51	45.36	70.76	45.39
<i>(B) Global features + Local feature re-ranking</i>								
R101-GeM+DSM	65.30	47.60	77.40	52.80	39.20	23.20	56.20	25.00
R50-DELG+SP(GLDv2-clean)	78.30	67.20	85.70	69.60	57.90	43.60	71.00	45.70
R101-DELG+SP(GLDv2-clean)	81.20	69.10	87.20	71.50	64.00	47.50	72.80	48.70
R101-DELG+SP†(GLDv2-clean)	81.78	70.12	88.46	76.04	64.77	49.36	76.80	53.69
<i>(C) Global features</i>								
R101-R-MAC(NC-clean)	60.90	39.30	78.90	54.80	32.40	12.50	59.40	28.00
R101-R-MAC†(GLDv2-clean)	75.14	61.88	85.28	67.37	53.77	36.45	71.28	44.01
R101-GeM-AP(GLDv1-noisy)	67.50	47.50	80.10	52.50	42.80	23.20	60.50	25.10
R101-NetVLAD†(GLDv2-clean)	73.91	60.51	86.81	71.31	56.45	37.92	73.61	48.98
R50-DELG(GLDv2-clean)	73.60	60.60	85.70	68.60	51.00	32.70	71.50	44.40
R50-DELG†(GLDv2-clean)	76.40	64.52	86.74	70.71	55.92	38.60	72.60	47.39
R101-DELG(GLDv2-clean)	76.30	63.70	86.60	70.60	55.60	37.50	72.40	46.90
R101-DELG†(GLDv2-clean)	78.55	66.02	88.58	73.65	60.89	41.75	76.05	51.46
R101-SOLAR(GLDv1-noisy)	69.90	53.50	81.60	59.20	47.90	29.90	64.50	33.40
R101-SOLAR†(GLDv2-clean)	79.65	67.61	88.63	73.21	59.99	41.14	76.15	50.98
R50-Ours(GLDv2-clean)	80.53	68.29	87.55	73.90	62.14	43.36	73.80	53.32
R101-Ours_{PQ8}(GLDv2-clean)	82.02	70.06	89.16	75.58	65.90	46.52	78.07	54.46
R101-Ours_{PQ1}(GLDv2-clean)	82.30	70.51	89.33	76.65	66.62	47.43	78.55	55.90
R101-Ours(GLDv2-clean)	82.28	70.52	89.34	76.66	66.57	47.27	78.56	55.90

Table 1: mAP comparison against existing methods on the full benchmark. R101: ResNet101; R50: ResNet50; +SP: spatial verification; “*”: binarized local features; “†”: our re-implementation. Training datasets are shown in brackets. PQ8 and PQ1 denote PQ quantization using 8 and 1-dimensional subspaces, respectively. Black bold: best results.

training. It is first collected by Google and further cleaned by researchers from the Google Landmark Retrieval Competition 2019. It contains a total of 1,580,470 images and 81,313 classes. We randomly divide it into two subsets ‘train’/‘val’ with 80%/20% split. The ‘train’ split is used for training model, and the ‘val’ split is used for validation.


Evaluation datasets and metrics. Revisited versions of the original Oxford5k (Philbin et al. 2007) and Paris6k (Philbin et al. 2008) datasets are used to evaluate our method, which are denoted as $\mathcal{R}Oxf$ and $\mathcal{R}Par$ (Radenović et al. 2018) in the following. Both datasets contain 70 query images and additionally include 4,993 and 6,322 database images, respectively. Mean Average Precision (mAP) is used as our evaluation metric on both datasets with Medium and Hard protocols. Large-scale results are further reported with the $\mathcal{R}1M$ dataset, which contains one million distractor images.

Training details. All models are pre-trained on ImageNet. For image augmentation, a 512×512 -pixel crop is taken from a randomly resized image and then undergoes random color jittering. We use a batch size of 128 to train our model on 4 NVIDIA RTX 3090 GPUs for 30 epochs, which


takes about 3 days. SGD is used to optimize the model, with an initial learning rate of 0.01, a weight decay of 0.0001, and a momentum of 0.9. A linearly decaying scheduler is adopted to gradually decay the learning rate to 0 when the desired number of steps is reached. The dimension d of the global feature is set as 1024. For the ArcFace margin loss, we empirically set the margin m as 0.2 and the scale γ as 32.0. Refinement block number N is set to 2. Test images are resized with the larger dimension equal to 1024 pixels, preserving the aspect ratio. Multiple scales are adopted, *i.e.*, $[1/\sqrt{2}, 1, \sqrt{2}]$. L_2 normalization is applied for each scale independently, then three global features are average-pooled, followed by another L_2 normalization. We train each model 5 times and evaluate the one with median performance on the validation set.

Results on Image Retrieval

Setting for fair comparison. Commonly, existing methods are compared under different settings, *e.g.*, training set, backbone network, feature dimension, loss function, *etc.* This may affect our judgment on the effectiveness of the pro-



(a) Visualization of attention maps



(b) Case study

Figure 3: Qualitative examples. (a) Visualization of the attention maps associated with different visual tokens for eight images. $\#i$ denotes the i -th visual token. (b) Detailed analysis of the top-2 retrieval results of the “hertford” query in the $\mathcal{R}Oxf$ dataset. The 2nd visual token focus on the content of the query image in the target image, which is boxed in red.

posed method. In Tab. 1, we re-train several methods under the same settings (using GLDv2-clean dataset and ArcFace loss, 2048 global feature dimension, ResNet101 as backbone), marked with †. Based on this benchmark, we fairly compare the mAP performance of various methods and ours.

Comparison with the state of the art. Tab. 1 compares our approach extensively with the state-of-the-art retrieval methods. We divide the previous methods into three groups:

(1) *Local feature aggregation.* The current state-of-the-art local aggregation method is R101-HOW. We outperform it in mAP by 1.86% and 4.06% on the $\mathcal{R}Oxf$ dataset and by 3.91%, and 7.80% on the $\mathcal{R}Par$ dataset with Medium and Hard protocols, respectively. The results show that our aggregation method is better than existing local feature aggregation methods based on large visual codebook.

(2) *Global single-pass.* When trained with GLDv2-clean, R101-SOLAR achieves the best performance mostly. When using ResNet101 as the backbone, the comparison between our method and it in mAP is 82.28% vs. 79.65%, 66.57% vs. 59.99% on the $\mathcal{R}Oxf$ dataset and 89.34% vs. 88.63%, 78.56% vs. 76.15% on the $\mathcal{R}Par$ dataset with Medium and Hard protocols, respectively. These results well demonstrate the superiority of our framework.

(3) *Global feature followed by local feature re-ranking.* We outperform the best two-stage method (R101-DELG+SP) in mAP by 0.50%, 1.80% on the $\mathcal{R}Oxf$ dataset and 0.88%, 1.76% on the $\mathcal{R}Par$ datasets with Medium and Hard protocols, respectively. Although 2-stage solutions well promote their single-stage counterparts, our method that aggregates local features into a compact descriptor is a better option.

METHOD	RET. (s)	EXT. (ms)	MEM. (GB)	
			$\mathcal{R}Oxf + \mathcal{R}1M$	$\mathcal{R}Par + \mathcal{R}1M$
DELG-RASMK*	1.5341	1410	27.6	27.8
DELG-ASMK*	0.5732	176	10.3	10.4
HOW-VLAD	0.4047	263	7.6	7.6
HOW-ASMK*	0.7123	257	14.3	14.4
DELG	0.4189	109	7.6	7.6
DELG+SP*	49.3821	259	22.6	22.7
Ours	0.2871	125	3.9	3.9
Ours_{PQ1}	0.2217	128	1.0	1.0
Ours_{PQ8}	0.1042	126	0.1	0.1

Table 2: Extraction (EXT.), retrieval (RET.) latency and memory footprint (MEM.) on a single thread GPU (RTX 3090) / CPU (Intel Xeon CPU E5-2640 v4 @ 2.40GHz).

Qualitative results. We visualize the spatial attention generated by the cross-attention layer of the last refinement block in the Fig. 3 (a). Although there is no direct supervision, different visual tokens are associated with different visual patterns. Most of these patterns focus on the foreground building and remain consistent across images, which implicitly enable pattern alignment. *e.g.*, the 3rd visual token reflects the semantics of “the upper edge of the window”.

We further select the top-2 results of the “hertford” query from the $\mathcal{R}Oxf$ dataset for the case study. As shown in Fig. 1, when the query object only occupies a small part of the target image, the state-of-the-art methods with global features return false positives which are semantically similar to the query. Our approach uses visual tokens to distinguish different visual patterns, which has the capability of regional matching. In Fig. 3 (b), the 2nd visual token corresponds to the visual pattern described by the query image.

Speed and memory costs. In Tab. 2, we report retrieval latency, feature extraction latency and memory footprint on $\mathcal{R}1M$ for different methods. Compared to the local feature aggregation approaches, global features have a smaller memory footprint. To perform spatial verification, “R101-DELG+SP” needs to store a large number of local features, which requires about 485 GB of memory. Our method uses 4 tokens to represent the image, generating a 1024-dimensional global feature, which requires 3.9 GB memory. This is further compressed with PQ quantization (Jégou, Douze, and Schmid 2011). As shown in Tab. 1 and Tab. 2, the compressed features greatly reduce the memory footprint with only a small performance loss. Our method appears to be a good solution in the performance-memory trade-off.

The extraction of global features is faster, since the extraction of local features usually requires scaling the image to more scales. Our aggregation method requires tokenization and iterative enhancement, which is slightly slower than direct spatial pooling, *e.g.*, 125 ms for ours vs. 109 ms for “R101-DELG”. The average retrieval latency of our method on $\mathcal{R}1M$ is 0.2871 seconds, which demonstrates the potential of our method for real-time image retrieval.

Ablation Study

Verification of different components. In Tab. 3, we provide experimental results to validate the contribution of the three components in our framework, by adding individual components to the baseline framework. When the tokenizer is adopted, there is a significant improvement in overall performance. mAP increases from 77.0% to 79.8% on \mathcal{ROxf} -Medium and 56.0% to 62.5% on \mathcal{ROxf} -Hard. This indicates that dividing local features into groups according to visual patterns is more effective than direct global spatial pooling. From the 3rd and last row, the performance is further enhanced when the refinement block is introduced, which shows that enhancing the visual tokens with the original features further makes them more discriminative.

LFSA	TOKENIZER	REFINEMENT	MEDIUM		HARD	
			\mathcal{ROxf}	\mathcal{RPar}	\mathcal{ROxf}	\mathcal{RPar}
			77.0	86.6	56.0	73.0
	✓		79.8	88.2	62.5	76.0
✓	✓		80.4	88.4	63.0	76.3
	✓	✓	81.3	89.2	65.0	78.5
✓	✓	✓	82.3	89.3	66.6	78.6

Table 3: Ablation studies of different components. We use R101-SPoC as the baseline and incrementally add tokenizer, Local Feature Self-Attention (LFSA) and refinement block.

Impact of each component in the refinement block. The role of the different components in the refinement block is shown in Tab. 4. By removing the individual components, we find that modeling the relationship between different visual words before and further enhancing the visual tokens using the original local features demonstrate the effectiveness in enhancing the aggregated features.

SELF-ATT	CROSS-ATT	MEDIUM		HARD	
		\mathcal{ROxf}	\mathcal{RPar}	\mathcal{ROxf}	\mathcal{RPar}
		80.4	88.4	63.0	76.3
	✓	81.3	89.3	63.5	78.2
✓		80.9	88.5	62.8	77.5
✓	✓	82.3	89.3	66.6	78.5

Table 4: Analysis of components in the refinement block.

Impact of tokenizer type. In Tab. 5, we compare our *Atten-based* tokenizer with the other two tokenizers: (1) *Vq-Based*. We directly define visual tokens as a matrix $T \in \mathcal{R}^{L \times C}$. It is randomly initialized and further updated by a moving average operation in one mini-batch. See the appendix for details. (2) *Learned*. It is similar to the *Vq-Based* method, except that T is set as the network parameters, learned during training. Our method achieves the best performance. We use the attention mechanism to generate visual tokens directly from the original local features. Compared with the other two, our approach obtains more discriminative visual tokens with a better capability to match different images.

TOKENIZER TYPE	MEDIUM		HARD	
	\mathcal{ROxf}	\mathcal{RPar}	\mathcal{ROxf}	\mathcal{RPar}
VQ-BASED	79.4	87.7	62.2	75.9
LEARNED	81.1	87.8	63.7	76.2
ATTEN-BASED	82.3	89.3	66.6	78.5

Table 5: mAP comparison of different variants of tokenizers.

Impact of token number. The granularity of visual tokens is influenced by their number. As shown in Tab. 6, as L increases, mAP performance first increases and then decreases, achieving the best at $L = 4$. This is due to the lack of capability to distinguish local features when the number of visual tokens is small; conversely, when the number is large, they are more fine-grained and noise may be introduced when grouping local features.

TOKEN NUMBER	MEDIUM		HARD	
	\mathcal{ROxf}	\mathcal{RPar}	\mathcal{ROxf}	\mathcal{RPar}
$L=1$	79.8	87.9	60.4	75.7
$L=2$	80.3	88.7	62.3	76.3
$L=3$	81.6	89.4	64.9	78.5
$L=4$	82.3	89.3	66.6	78.5
$L=6$	81.0	88.2	62.5	78.9
$L=8$	79.3	87.1	61.8	76.6

Table 6: mAP comparison of visual tokens number L .

Conclusion

In this paper, we propose a joint local feature learning and aggregation framework, which generates compact global representations for images while preserving the capability of regional matching. It consists of a tokenizer and a refinement block. The former represents the image with a few visual tokens, which is further enhanced by the latter based on the original local features. Extensive experiments demonstrate that the proposed method achieves superior performance on image retrieval benchmark datasets. In the future, we will extend the proposed aggregation method to a variety of existing local features, which means that instead of directly performing local feature learning and aggregation end-to-end, local features of images are first extracted using existing methods and further aggregated with our method.

Acknowledgements

This work was supported in part by the National Key R&D Program of China under contract 2018YFB1402605, in part by the National Natural Science Foundation of China under Contract 62102128, 61822208 and 62172381, and in part by the Youth Innovation Promotion Association CAS under Grant 2018497. It was also supported by the GPU cluster built by MCC Lab of Information Science and Technology Institution, USTC.

References

- Arandjelovic, R.; Gronat, P.; Torii, A.; Pajdla, T.; and Sivic, J. 2016. NetVLAD: CNN architecture for weakly supervised place recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 5297–5307.
- Babenko, A.; and Lempitsky, V. 2015. Aggregating Deep Convolutional Features for Image Retrieval. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 1269–1277.
- Babenko, A.; Slesarev, A.; Chigorin, A.; and Lempitsky, V. 2014. Neural codes for image retrieval. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 584–599.
- Bay, H.; Tuytelaars, T.; and Van Gool, L. 2006. SURF: Speeded up robust features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 404–417.
- Cao, B.; Araujo, A.; and Sim, J. 2020. Unifying deep local and global features for image search. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 726–743.
- Chum, O.; Philbin, J.; Sivic, J.; Isard, M.; and Zisserman, A. 2007. Total recall: Automatic query expansion with a generative feature model for object retrieval. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1–8.
- Deng, J.; Guo, J.; Xue, N.; and Zafeiriou, S. 2019. ArcFace: Additive Angular Margin Loss for Deep Face Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Dmytro Mishkin, J. M., Filip Radenovic. 2018. Repeatability Is Not Enough: Learning Discriminative Affine Regions via Discriminability. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Fan, Y.; Ryota, H.; Yusuke, M.; Steven, L.; and Shin'ichi, S. 2019. Efficient Image Retrieval via Decoupling Diffusion into Online and Offline Processing. *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Gordo, A.; Almazan, J.; Revaud, J.; and Larlus, D. 2017. End-to-end learning of deep visual representations for image retrieval. *International Journal of Computer Vision (IJCV)*, 237–254.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Jégou, H.; Douze, M.; Sánchez, J.; Pérez, P.; and Schmid, C. 2011. Aggregating local image descriptors into compact codes. *IEEE transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 1704–1716.
- Jegou, H.; Douze, M.; and Schmid, C. 2008. Hamming Embedding and Weak Geometric Consistency for Large Scale Image Search. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 304–317.
- Jégou, H.; Douze, M.; and Schmid, C. 2011. Product quantization for nearest neighbor search. *IEEE transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 117–128.
- Kalantidis, Y.; Mellina, C.; and Osindero, S. 2016. Cross-dimensional weighting for aggregated deep convolutional features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 685–701.
- Lowe, D. G. 2004. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision (IJCV)*, 91–110.
- Mishchuk, A.; Mishkin, D.; Radenovic, F.; and Matas, J. 2017. Working hard to know your neighbor's margins: Local descriptor learning loss. In *Conference and Workshop on Neural Information Processing Systems (NeurIPS)*.
- Ng, T.; Baltas, V.; Tian, Y.; and Mikołajczyk, K. 2020. SOLAR: second-order loss and attention for image retrieval. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 253–270.
- Noh, H.; Araujo, A.; Sim, J.; Weyand, T.; and Han, B. 2017. Large-scale image retrieval with attentive deep local features. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 3456–3465.
- Perronnin, F.; Liu, Y.; Sánchez, J.; and Poirier, H. 2010. Large-scale image retrieval with compressed Fisher vectors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 3384–3391.
- Philbin, J.; Chum, O.; Isard, M.; Sivic, J.; and Zisserman, A. 2007. Object retrieval with large vocabularies and fast spatial matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1–8.
- Philbin, J.; Chum, O.; Isard, M.; Sivic, J.; and Zisserman, A. 2008. Lost in quantization: Improving particular object retrieval in large scale image databases. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1–8.
- Radenović, F.; Iscen, A.; Toliás, G.; Avrithis, Y.; and Chum, O. 2018. Revisiting oxford and paris: Large-scale image retrieval benchmarking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 5706–5715.
- Radenović, F.; Toliás, G.; and Chum, O. 2018. Fine-tuning CNN image retrieval with no human annotation. *IEEE transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 1655–1668.
- Revaud, J.; Almazán, J.; Rezende, R. S.; and Souza, C. R. d. 2019. Learning with average precision: Training image retrieval with a listwise loss. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 5107–5116.
- Sivic, J.; and Zisserman, A. 2003. Video Google: A text retrieval approach to object matching in videos. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 1470–1470.
- Teichmann, M.; Araujo, A.; Zhu, M.; and Sim, J. 2019. Detect-to-retrieve: Efficient regional aggregation for image search. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 5109–5118.
- Tian, Y.; Yu, X.; Fan, B.; Wu, F.; Heijnen, H.; and Baltas, V. 2019. SOSNet: Second Order Similarity Regularization for Local Descriptor Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Toliás, G.; Avrithis, Y.; and Jégou, H. 2013. To aggregate or not to aggregate: Selective match kernels for image search. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1401–1408.
- Toliás, G.; Jeníček, T.; and Chum, O. 2020. Learning and aggregating deep local descriptors for instance-level recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 460–477.
- Toliás, G.; Sicre, R.; and Jégou, H. 2016. Particular Object Retrieval With Integral Max-Pooling of CNN Activations. In *International Conference on Learning Representations (ICLR)*, 1–12.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 5998–6008.
- Weyand, T.; Araujo, A.; Cao, B.; and Sim, J. 2020. Google landmarks dataset v2-a large-scale benchmark for instance-level recognition and retrieval. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2575–2584.

Xu, J.; Cunzhao, S.; Chengzuo, Q.; Chunheng, W.; and Baihua, X. 2018. Unsupervised Part-Based Weighting Aggregation of Deep Convolutional Features for Image Retrieval. *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.

Zhou, W.; Lu, Y.; Li, H.; Song, Y.; and Tian, Q. 2010. Spatial coding for large scale partial-duplicate web image search. In *Proceedings of the ACM international conference on Multimedia (MM)*, 511–520.