

DOLG: Single-Stage Image Retrieval with Deep Orthogonal Fusion of Local and Global Features

Min Yang*, Dongliang He*,[†], Miao Fan, Baorong Shi,
Xuetong Xue, Fu Li, Errui Ding, Jizhou Huang[†]
Baidu Inc., China

{yangmin09, hedongliang01, fanmiao, shibaorong}@baidu.com,
{xuexuetong, lifu, dingerrui, huangjizhou01}@baidu.com

Abstract

Image Retrieval is a fundamental task of obtaining images similar to the query one from a database. A common image retrieval practice is to firstly retrieve candidate images via similarity search using global image features and then re-rank the candidates by leveraging their local features. Previous learning-based studies mainly focus on either global or local image representation learning to tackle the retrieval task. In this paper, we abandon the two-stage paradigm and seek to design an effective single-stage solution by integrating local and global information inside images into compact image representations. Specifically, we propose a Deep Orthogonal Local and Global (DOLG) information fusion framework for end-to-end image retrieval. It attentively extracts representative local information with multi-atrious convolutions and self-attention at first. Components orthogonal to the global image representation are then extracted from the local information. At last, the orthogonal components are concatenated with the global representation as a complementary, and then aggregation is performed to generate the final representation. The whole framework is end-to-end differentiable and can be trained with image-level labels. Extensive experimental results validate the effectiveness of our solution and show that our model achieves state-of-the-art image retrieval performances on Revisited Oxford and Paris datasets.¹

1. Introduction

Image retrieval is an important task in computer vision, and its main purpose is to find out the images from a large-scale database that are similar to a query one. It is extensively studied by designing various handcrafted features [25, 6, 50]. Owing to the development of deep learn-

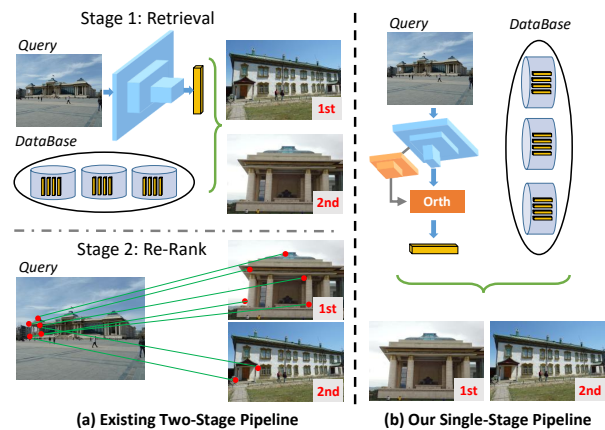


Figure 1: Illustration of current two-stage and our single-stage image retrieval. Previous methods (a) firstly obtain candidates similar to the query from the database via global deep representation, and then local descriptors are extracted for leveraged re-ranking. Our method (b) aggregates global and local features via an orthogonal fusion to generate the final compact descriptor, and then single-shot similarity search is performed.

ing technologies, great progress has been achieved recently [1, 29, 37, 9]. Representations (also named as descriptors) of images, which are used to encode image contents and measure their similarities, play a central role in this task. In the literature of learning-based solutions, two types of image representations are widely explored. One is global feature [4, 3, 44, 1] which serves as high-level semantic image signature and the other one is local feature [5, 36, 29, 18] which can comprise discriminative geometry information about specific image regions. Generally, the global feature can be learned to be invariant to viewpoint and illumination, while local features are more sensitive to local geometry and textures. Therefore, previous state-of-the-art solutions

* Equal contribution. [†] Corresponding authors.

¹Codes: [PaddlePaddle Implementation](#).

[38, 29, 9] always work in a two-stage paradigm. As shown in Figure 1(a), candidates are retrieved via global feature with high recall, and then re-ranking is performed with local features to further improve precision.

In this paper, we also concentrate on the field of image retrieval with deep networks. Though state-of-the-art performance has been achieved by previous two-stage solutions, they need to rank images twice, and the second re-ranking stage is conducted using the expensive RANSAC [13] or AMSK [42] for spatial verification with local features. More importantly, errors exist inevitably in both stages. Two-stage solutions would suffer from error accumulation which can be a bottleneck for further performance improvement. To alleviate these problems, we abandon the two-stage framework and attempt to find an effective unified single-stage image retrieval solution, which is shown in Figure 1(b). Previous wisdom has implied that global features and local features are two complementary and essential elements for image retrieval. Intuitively, integrating local features and global features into a compact descriptor can achieve our goal. A satisfying local and global fusion scheme can take advantage of both types of features to mutually boost each other for single-stage retrieval. Besides, error accumulation can be avoided. Therefore, we technically answer how to design an effective global and local fusion mechanism for end-to-end single-stage image retrieval.

Specifically, we proposed a **Deep Orthogonal Local and Global feature fusion model (DOLG)**. It consists of a local and a global branch for learning two types of features jointly and an orthogonal fusion module to combine them. In detail, the local components orthogonal to the global feature are decomposed from the local features. Subsequently, the orthogonal components are concatenated with the global feature as a complementary part. Finally, it is aggregated into a compact descriptor. With our orthogonal fusion, the most critical local information can be extracted and redundant components to the global information are eliminated, such that local and global components can be mutually reinforced to produce final representative descriptor with objective-oriented training. To enhance local feature learning, inspired by lessons from prior research, the local branch is equipped with multi-atrous convolutions [10] and self-attention [29] mechanisms to attentively extract representative local features. We think alike FP-Net [31] in terms of orthogonal feature space learning, but DOLG aims at complementary fusion of features in orthogonal spaces. Extensive experiments on Revisited Oxford and Pairs [32] show the effectiveness of our framework. DOLG also achieves state-of-the-art performance on both datasets. To summarize, our main contributions are as follows:

- We propose to retrieve images in a single-stage paradigm with a novel orthogonal global and local feature fusion framework, which can generate a com-

pact representative image descriptor and is end-to-end learnable.

- In order to attentively extract discriminative local features, a module with multi-atrous convolution layers followed by a self-attention module is designed for improving our local branch.
- Extensive experiments are conducted and comprehensive analysis is provided to validate the effectiveness of our solution. Our single-stage method significantly outperforms previous two-stage state-of-the-art ones.

2. Related Work

2.1. Local feature

Prior to deep learning, SIFT [25] and SURF [6] are two well-known hand-engineered local features. Usually such local features are combined with KD trees [7], vocabulary trees [28] or encoded by aggregation methods such as [49, 22] for (approximate) nearest neighbor search. Spatial verification via matching local features with RANSAC [13] to re-rank candidate retrieval results [2, 30] are also shown to significantly improve precision. Recently, driven by the development of deep learning, remarkable progresses have been made in learning local features from images such as [48, 16, 15, 5, 36, 29, 18]. Comprehensive reviews of deep local feature learning can be found in [51, 12]. Among these methods, the state-of-the-art local feature learning framework DELF [29], which proposes an attentive local feature descriptor for large-scale image retrieval, is closely related to our work. One of the design choices of our local branch, namely attentive feature extraction, is inspired by its merit. However, DELF uses only a single-scale feature map and ignores various object scales inside natural images. Our local branch is designed to simulate the image pyramid trick used in SIFT [25] by multi-atrous convolution layers [10].

2.2. Global feature

Conventional solutions obtain global feature by aggregating local features by BoW [39, 33], Fisher vectors [24] or VLAD [23]. Later, aggregated selective match kernels (ASMK) [42] attempts to unify aggregation-based techniques with matching-based approaches such as Hamming Embedding [21]. In deep learning era, global feature is obtained by such differentiable aggregation operations as sum-pooling [43] and GeM pooling [34]. To train deep CNN models, ranking based triplet [8], quadruplet [11], angular [46] and listwise [35] losses or classification based losses [45, 14] are proposed. With these innovations, nowadays, most high performing global features are obtained with deep CNNs for image retrieval [4, 3, 44, 1, 17, 34, 35, 29, 27, 9]. In our work, we leverage lessons from previous studies to use ArcFace loss [14] in the training phase and

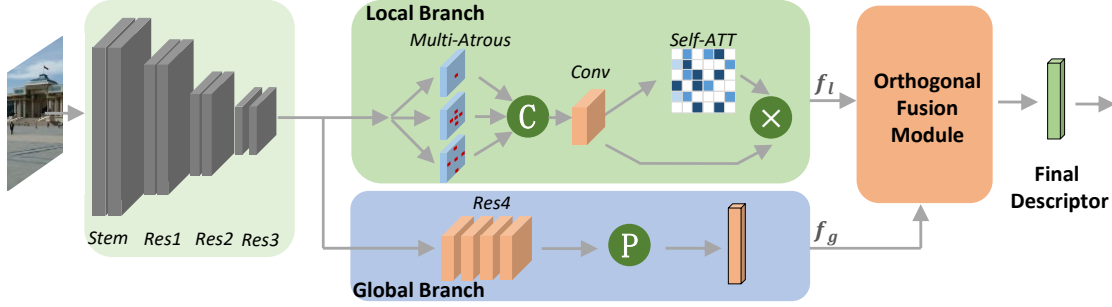


Figure 2: Block diagram of our deep orthogonal local and global (DOLG) information fusion framework. Taking ResNet [19] for illustration, we build a local branch and a global branch after Res3. The local branch uses multi-atrous layers to simulate spatial pyramid to take into consideration of scale variations among images. Self-attention is leveraged for importance modeling following lessons of existing works [29, 9]. The global branch generates a descriptor, which is fed into an orthogonal fusion module together with the local features for integrating both types of features into a final compact descriptor. “P”, “C” and “X” denote pooling, concatenation and element-wise multiplication, respectively.

to explore different pooling schemes for performance improvement. Our model also generates a compact descriptor, meanwhile, it explicitly considers fusing local and global features in an orthogonal way.

2.3. Joint local and global CNN features

It is natural to consider local and global features jointly, because feature maps from an image representation model can be interpreted as local visual words [38, 40]. Joint learning local matching and global representation may be beneficial for both sides. Therefore, distilling pre-trained local feature [15] and global feature [1] into a compact descriptor is proposed in [37]. DELG [9] takes a step further and proposes to jointly train local and global features in an end-to-end manner. However, DELG still works in a two-stage fashion. Our work is essentially different from [29, 9] and we propose orthogonal global and local fusion in order to perform accurate single-stage image retrieval.

3. Methodology

3.1. Overview

Our DOLG framework is depicted in Figure 2. Following [29, 9], it is built upon state-of-the-art image recognition model ResNet [19]. The global branch is kept the same as the original ResNet except that 1) the global averaging pooling is replaced by the GeM pooling [34]; 2) a FC layer is used to reduce feature dimension when generating the global representation $f_g \in R^{C \times 1}$. Specifically, let us denote the output feature map of Res4 as $f_4 \in R^{C_4 \times h \times w}$, then the GeM pooling can be formalized as

$$f_{g,c} = \left(\frac{1}{hw} \sum_{(i,j)} f_{4,(c,i,j)}^p \right)^{1/p}, \quad (1)$$

where $p > 0$ is a hyper-parameter and $p > 1$ pushes the output to focus more on salient feature points. In this paper, we follow the setting of DELG [9] and empirically set it to be 3.0. To jointly extract local descriptors, a local branch is appended after the Res3 block of ResNet. Our local branch consists of multiple atrous convolution layers [10] and a self-attention module. Then, a novel orthogonal fusion module is designed for aggregating f_g and the local feature tensor $f_l \in R^{C \times H \times W}$ obtained by the local branch. After orthogonal fusion, a final compact descriptor, where local and global information is well integrated, is generated.

3.2. Local Branch

The two major building blocks of our local branch are the multi-atrous convolution layers and the self-attention module. The former building block is to simulate feature pyramid which can handle scale variations among different image instances, and the latter building block is leveraged to performance importance modeling. The detailed network configurations of this branch is shown in Figure 3. The multi-atrous module contains three dilated convolution layers to obtain feature maps with different spatial receptive field and a global average pooling branch. These features are concatenated and then processed by a 1×1 convolution layer. The output feature map is then delivered to the self-attention module for further modeling the importance of each local feature point. Specifically, its input is firstly processed using a 1×1 conv-bn module, then the subsequent feature is normalized and modulated by an attention map generated via a 1×1 convolution layer followed by the SoftPlus operation.

3.3. Orthogonal Fusion Module

The working flow of our orthogonal fusion module is shown in Figure 4a. It takes f_l and f_g as inputs and then

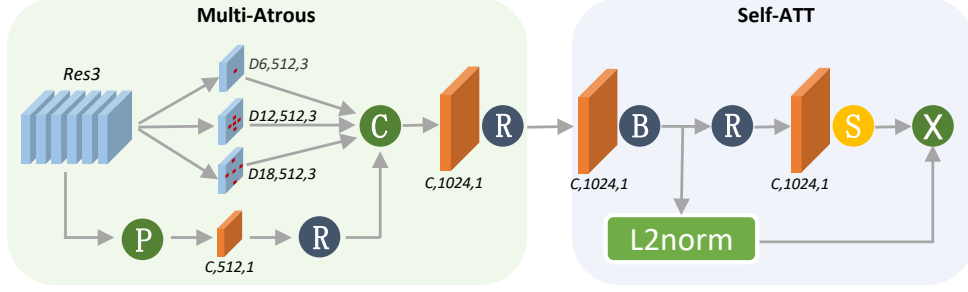
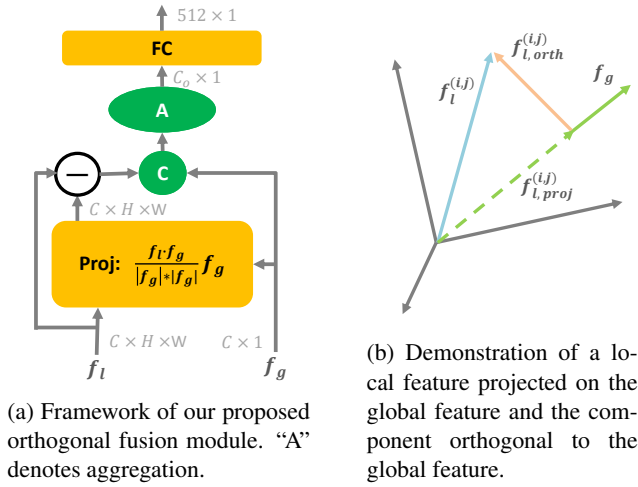


Figure 3: Configurations of our local branch. “Ds, c, k” denotes dilated convolution with rate s , output channel number c and kernel size k . “C,c, k” means vanilla convolution. “R”, “B” and “S” denote ReLU, BN and Softplus, respectively.



calculates the projection $f_{l,proj}^{(i,j)}$ of each local feature point $f_l^{(i,j)}$ onto the global feature f_g . Mathematically, the projection can be formulated as:

$$f_{l,proj}^{(i,j)} = \frac{f_l^{(i,j)} \cdot f_g}{|f_g|^2} f_g, \quad (2)$$

where $f_l^{(i,j)} \cdot f_g$ is dot product operation and $|f_g|^2$ is the L_2 norm of f_g :

$$f_l^{(i,j)} \cdot f_g = \sum_{c=1}^C f_{l,c}^{(i,j)} f_{g,c} \quad (3)$$

$$|f_g|^2 = \sum_{c=1}^C (f_{g,c})^2. \quad (4)$$

As demonstrated in Figure 4b, the orthogonal component is the difference between the local feature and its projection vector, therefore, we can obtain the component orthogonal to f_g by:

$$f_{l,orth}^{i,j} = f_l^{(i,j)} - f_{l,proj}^{(i,j)}. \quad (5)$$

In this way, a $C \times H \times W$ tensor where each point is orthogonal to f_g can be extracted. Afterwards, we append

to each point of this tensor with the $C \times 1$ vector f_g and then the new tensor is aggregated to be a $C_o \times 1$ vector. Finally, a fully connected layer is used to produce a 512×1 descriptor. Typically, C equals 1024 in ResNet [19]. Here, we simply leverage the pooling functionality to aggregate the concatenated tensor, that is to say, “A” in Figure 4a is pooling in our current implementation. Actually, it can be designed to be other learnable modules to aggregate the tensor. We will further analysis on this in Section 4 and 5.

3.4. Training Objective

Following DELG [9], the training of our method involves only one L_2 -normalized N class prediction head $\hat{\mathcal{W}} \in R^{512 \times N}$ and just needs image-level labels. ArcFace margin loss [14] is used to train the whole network:

$$L = -\log \left(\frac{\exp(\gamma \times AF(\hat{\omega}_t^T \hat{f}_g, 1))}{\sum_n \exp(\gamma \times AF(\hat{\omega}_n^T \hat{g}, y_n))} \right) \quad (6)$$

where $\hat{\omega}_i$ refers to the i_{th} row of $\hat{\mathcal{W}}$ and \hat{f}_g is the L_2 -normalized version of f_g . y is the one-hot label vector and t is the groundtruth class index ($y_t = 1$). γ is a scale factor. AF denotes the ArcFace-adjusted cosine similarity and it can be calculated as $AF(s, c)$:

$$AF(s, c) = \begin{cases} \cos(a \cos(s) + m), & \text{if } c = 1 \\ s, & \text{if } c = 0 \end{cases} \quad (7)$$

where s is the cosine similarity, m is the ArcFace margin and $c = 1$ means this is the groundtruth truth class.

4. Experiments

4.1. Implementation Details

Datasets and Evaluation metric Google landmarks dataset V2 (GLDv2) [47] is developed for large-scale and fine-grained landmark instance recognition and image retrieval. It contains a total of 5M images of 200K different instance tags. It is collected by Google to raise the challenges faced by the landmark identification system under

real industrial scenarios as much as possible. Researchers from the Google Landmark Retrieval Competition 2019 further cleaned and revised the GLDv2 to be GLDv2-clean. It contains a total of 1,580,470 images and 81,313 classes. This dataset is used to train our models. To evaluate our model, we mainly use Oxford and Paris datasets with revisited annotations [32], referred to be Roxf and Rpar in the following, respectively. There are 4,993 (6,322) images in the Roxf (Rpar) dataset and a different query set for each, both with 70 images. In order for a fair comparison with state-of-the-art methods [29, 9, 27], mean average precision (mAP) is used as our evaluation metric on the Medium and Hard splits of both datasets. mAP provides a robust measurement of retrieval quality across recall levels and has shown to have good discrimination and stability.

Implementation details All the experiments in this paper are trained based on GLDv2-clean dataset. We randomly divide 80% of the dataset for training and the rest 20% for validation. ResNet50 and ResNet101 are mainly used for experiments. Models are initialized from ImageNet pre-trained weights. The images first undergo augmentations by randomly cropping / distorting the aspect ratio; then, they are resized to 512×512 resolution. We use batch size of 128 to train our models on 8 V100 GPUs with 16G memory per card asynchronously for 100 epochs. One complete training phase takes about 3.8 days for ResNet50 and 6.3 days for ResNet101. SGD optimizer with momentum of 0.9 is used. Weight decay factor is set to 0.0001 and cosine learning rate decay strategy is adopted. Note that we train our models with 5 warming-up epochs and the initial learning rate is 0.05. For the ArcFace margin loss, we empirically set the margin m as 0.15 and the ArcFace scale γ as 30. For GeM pooling, we fix the parameter p as 3.0.

As for feature extraction, following previous works [29, 9], we use an image pyramid at inference time to produce multi-scale representations. Specifically, we use 5 scales, *i.e.*, 0.3535, 0.5, 0.7071, 1.0, 1.4142, to extract final compact feature vectors. To fuse these multi-scale features, we firstly normalize them such that their L_2 norm equals 1, then the normalized features are averaged and finally a L_2 normalization is applied to produce the final descriptor.

4.2. Results

4.2.1 Comparison with State-of-the-art Methods

We divide the previous state-of-the-art methods into three groups: (1) local feature aggregation and re-ranking; (2) global feature similarity search; (3) global feature search followed by re-ranking with local feature matching and spatial verification (SP). From some point of view, our method belongs to the global feature similarity search group. The results are summarized in Table 1 and we can see that our solution consistently outperforms existing solutions.

Comparison with local feature based solutions. In the

local feature aggregation group, besides DELF [29], it is worth mentioning that current work R50-How [43] provides a manner for learning local descriptors with ASMK [42] and outperforms DELF. It achieves a boost up to 3.4% on Roxf-Medium and 1.4% on Rpar-Medium. However, the complexity of this work is considerable, where $n=2000$ shows it finally uses 2000 strongest local keypoints. Our method outperforms it by up 1.1% on Roxf-Medium and 8.21% on Rpar-Medium with the same ResNet50 backbone. For the hard samples, our R50-DOLG achieves 58.82% and 77.7% in mAP on the Roxf and Rpar respectively, which is significantly better than 56.9% and 62.4% achieved by R50-How. The results show that our single-stage model is better than existing local feature aggregation methods which are enhanced by a second re-ranking stage.

Comparison with global feature based solutions. Our method completes image retrieval with single-stage and the global feature based solutions do the same. It can be found the global feature learned by DELG [9] performs the best. Especially when the models are trained using the GLDv2-clean dataset. Our models are also trained on this dataset and they are validated to be better than DELG. The performance is significantly improved by our solution. For example, with Res50 backbone, the mAP is 80.5% *v.s.* 77.51% on Roxf-Medium and 58.82% *v.s.* 54.76% on Roxf-Hard. Please note that, our R50-DOLG performs better than R101-DELG. These results well demonstrate the superiority of our framework.

Comparison with global+local feature based solutions. In the solutions where global feature is followed by a local feature re-ranking, R50/101-DELG is still the existing state-of-the-art method. Compared with the best result of DELG, our method R50-DOLG outperforms the R50-DELG with a boost of up to 1.42% on Roxf-Medium, 1.03% on Rpar-Medium, 0.42% on Roxf-Hard and 1.5% on Rpar-Hard. Our R101-DOLG outperforms R101-DELG with a boost of up to 0.3% on Roxf-Medium, 3.82% on Rpar-Medium and 7.5% on Rpar-Hard. From these results, we can see, although 2-stage solutions can well promote their single stage counterparts, our solution combining both local and global information is a better choice.

Comparison in mP@10. We compare mP@10 in Table 2. It shows the mP@10 performances of DOLG are better than 2-stage DELG^r on both RPar and Roxf. Such results validate our single-stage solution is more precise than state-of-the-art 2-stage DELG, owing to the advantages of end-to-end training and free of error accumulation.

“+1M” distractors. From Table 1, DOLG and 2-stage DELG^r outperform the official 2-stage DELG by a large margin. This is reasonable. Firstly, the DELG^r and our DOLG are both trained for 100 epochs while the official DELG is only trained for 25 epochs, so the original DELG features are not so robust, (w/o 1M distractors, DELG-

Method	Medium				Hard			
	Roxf	+1M	Rpar	+1M	Roxf	+1M	Rpar	+1M
<i>(A) Local feature aggregation + re-ranking</i>								
HesAff-rSIFT-ASMK* +SP[42]	60.6	46.80	61.40	42.30	36.70	26.90	35.00	16.80
HesAff-HardNet-ASMK* +SP[26]	65.60	-	65.20	-	41.10	-	38.50	-
HesAff-rSIFT-ASMK*+SP→R[34]-GeM+DFS[20]	79.10	74.30	91.00	85.90	52.70	48.70	81.00	73.20
DELf-ASMK* +SP[29, 32]	67.80	53.80	76.90	57.30	43.10	31.20	55.40	26.40
DELf-R-ASMK* +SP[41]	76.00	64.00	80.20	59.70	52.40	38.10	58.60	58.60
R50-How-ASMK,n=2000[43]	79.40	65.80	81.60	61.80	56.90	38.90	62.40	33.70
<i>(B) Global features</i>								
R101-R-MAC[17]	60.90	39.30	78.90	54.80	32.40	12.50	59.40	28.00
R101-GeM ↑[38]	65.30	46.10	77.30	52.60	39.60	22.20	56.60	24.80
R101-GeM-AP[35]	67.50	47.50	80.10	52.50	42.80	23.20	60.50	25.10
R101-GeM-AP (GLDv1) [35]	66.30	-	80.20	-	42.50	-	60.80	-
R152-GeM[34]	68.70	-	79.70	-	44.20	-	60.30	-
ResNet101-GeM+SOLAR† [27]	69.90	53.50	81.60	59.20	47.90	29.90	64.50	33.40
R50-DELG[9]	69.70	55.00	81.60	59.70	45.10	27.80	63.40	34.10
R50-DELG (GLDv2-clean)[9]	73.60	60.60	85.70	68.60	51.00	32.70	71.50	44.40
R50-DELG (GLDv2-clean)^r[9]	77.51	74.80	87.90	77.3	54.76	50.40	73.82	61.01
R101-DELG[9]	73.20	54.80	82.40	61.80	51.20	30.30	64.70	35.50
R101-DELG (GLDv2-clean)[9]	76.30	63.70	86.60	70.60	55.60	37.50	72.40	46.90
<i>(C) Global features + Local feature re-ranking</i>								
R101-GeM↑+DSM[38]	65.30	47.60	77.40	52.80	39.20	23.20	56.20	25.00
R50-DELG[9]	75.10	61.10	82.30	60.50	54.20	36.80	64.90	34.80
R50-DELG (GLDv2-clean)[9]	78.30	67.20	85.70	69.60	57.90	43.60	71.00	45.70
R50-DELG (GLDv2-clean)^r[9]	79.08	75.90	88.78	77.69	58.40	52.40	76.20	61.60
R101-DELG[9]	78.50	62.70	82.90	62.60	59.30	39.30	65.50	37.00
R101-DELG (GLDv2-clean)[9]	81.20	69.10	87.20	71.50	64.00	47.50	72.80	48.70
R50-DOLG (GLDv2-clean)	80.50	76.58	89.81	80.79	58.82	52.21	77.7	62.83
R101-DOLG (GLDv2-clean)	81.50	77.43	91.02	83.29	61.10	54.81	80.30	66.69

Table 1: Results (% mAP) of different solutions are obtained following the Medium and Hard evaluation protocols of Roxf and Rpar. “*” means feature quantization is used and “†” means second-order loss is added into SOLAR. “GLDv1”, “GLDv2” and “GLDv2-clean” mark the difference in training dataset. ^r denotes our re-implementation. State-of-the-art performances are marked bold and ours are summarized in the bottom. The underlined numbers are the best performances.

Model	Roxf-M	Roxf-H	Rpar-M	Rpar-H
R50-DELG ^r	90.79	69.00	95.57	92.00
R50-DOLG	92.52	71.14	98.43	93.71

Table 2: Results of mP@10 of different methods.

Global^r outperforms DELG-Global by 3.9 points in mAP on Roxf-M and Re-ranking on Rpar is even slightly worse than DELG-Global). When a huge amount of distractors exist, less robust global and local feature will result in severer error accumulation (DELG-Global^r > 2-stage DELG with “+1M”). As a consequence, significant performance gap appears between our re-implemented DELG and its official version. From the last two rows, we see DOLG still

outperforms 2-stage DELG^r when +1M distractors exist.

Qualitative Analysis. We showcase top-10 retrieval results of a query image in Figure 5. We can see that state-of-the-art methods with global feature will result in many false positives which are semantically similar to the query. With re-ranking, some false positives can be eliminated but those with similar local patterns still exist. Our solution combines global and local information and is end-to-end optimized, so it is better at figuring out true positives.

4.2.2 Ablation Studies

To empirically verify some of our design choices, ablation experiments are conducted using the Res50 backbone.

Where to Fuse. To check which block is better for the

Location	Roxf			Rpar		
	E	M	H	E	M	H
Global only	90.65	78.21	56.31	95.65	89.00	76.17
Fuse f4-only	92.08	79.39	58.13	95.93	89.92	77.92
Fuse f3-only	93.17	80.50	58.82	95.95	89.81	77.70
both f3&f4	92.34	79.41	57.08	96.01	89.78	77.69

Table 3: Experimental results of DOLG variants where the orthogonal fusion is performed at different locations.

Pooling		Roxf			Rpar		
Global	Ortho	E	M	H	E	M	H
GeM	GeM	92.62	78.28	55.30	96.20	89.50	76.99
AVG	AVG	92.20	78.14	56.14	95.86	89.25	76.32
GeM	AVG	93.17	80.50	58.82	95.95	89.81	77.70
AVG	GeM	89.63	73.48	44.88	94.67	86.76	72.98

Table 4: Differences when different pooling functions are used. ‘‘AVG’’ means ordinary global average pooling.

global and local orthogonal integration, we provide empirical results to verify our choice. Specifically, shallow layers are known to be not appropriate for local feature representations [29, 9], thus we mainly check the res3 and res4 block. We have implemented DOLG variants where the local branch(es) is (are) originated from f_4 only (both f_3 and f_4). Hence, fusing f_3 , f_4 and f_g means there are two orthogonal fusion branches based on Res3 and Res4, and the two orthogonal tensors generated from the two fusion branches are concatenated with f_g and pooled. The results are summarized in Table 3. We can see that 1) without local branch, the global only setting performs worse. 2) Fusing f_3 or f_4 or *both* f_3 & f_4 can improve the perform of ‘‘Global only’’. Fusing f_3 obviously outperforms fusing f_4 on Roxf although it is slightly worse on Rpar. Fusing both f_3 and f_4 does not provide improvement over f_3 -only but it is better than f_4 -only. The above phenomena is reasonable. f_3 is of sufficient spatial resolution and its network depth is also sufficient, so it is better than f_4 to serve as local features. *both* f_3 & f_4 will make the model more complicated. Besides, f_g is derived from f_4 as well, then *both* f_3 & f_4 setting may put more emphasis on f_4 , therefore degrading the overall performance. Overall speaking, f_3 -only is the best.

Impact of Poolings. In this experiment, we study how GeM pooling [34] and average pooling will make a difference to our overall framework. We report results of DOLG when the pooling function of the global branch and the orthogonal fusion module alters. With other settings kept the same, the performances of R50-DOLG are presented in Table 4. It is interesting to see that using GeM pooling for the global branch while using average pooling for the orthogonal fusion module results in the best combination.

Impact of Each Component in the Local Branch. A multi-atrous block and self-attention block are designed in our local branch to simulate the spatial feature pyramid by dilated convolution layers [10] and to model the local feature importance with attention mechanism [29], respectively. We provide experimental results to validate the contribution of each of these components by removing individual component from the whole framework. The performance is shown in Table 5. It is clear that fusing the local features helps to improve the overall performance significantly. The mAP is improved from 78.2% to 80.5% and 89.0% to 89.8% on Roxf-Medium and Rpar-Medium, respectively. When Multi-Atrous module is removed, the performance will slightly drop on the Medium and Hard splits, especially for the hard split. For example, mAP is decreased from 58.82% to 58.36% and 77.7% to 76.52% on Roxf-Hard and Rpar-Hard, respectively. However, for easy cases, Multi-Atrous will make the performance slightly worse, but this make little difference because the mAP is already very high and the retrieval performance drop is very limited for easy case. Such results validate the effectiveness of Mutli-Atrous module. When the self-attention module is removed the performance also notably drops, which is consistent with results obtained by [9].

Verification of the Orthogonal Fusion. In the orthogonal fusion module, we propose to decompose the local features into two components, one is parallel to the global feature f_g and the other is orthogonal to f_g . Then we fuse the complementary orthogonal component and f_g . To show such orthogonal fusion is a better choice, we conduct experiments by removing the orthogonal decomposition procedure shown in Figure 4a and concatenate the f_l and f_g directly. We also try fusing f_l and f_g by Hadamard product (also known as element-wise product), which is usually used to fuse two vectors. We can find from the empirical results (see Table 6) that among the three fusion schemes, our proposed orthogonal fusion performs the best. Such experimental results are also within our expectation. With orthogonal fusion, the information relevant to the global feature f_g is excluded from each local feature point $f_l^{(i,j)}$. In this way, the output local feature points are the most informative one and are orthogonal to f_g . Not only will they provide

Config	Roxf			Rpar		
	E	M	H	E	M	H
w/o Local	90.65	78.21	56.31	95.65	89.00	76.17
w/o MultiAtrous	93.48	80.48	58.36	96.66	89.27	76.52
w/o Self-ATT	90.64	78.15	55.34	95.73	89.48	77.16
Full Model	93.17	80.50	58.82	95.95	89.81	77.70

Table 5: Ablation experiments on components of the local branch in our framework.

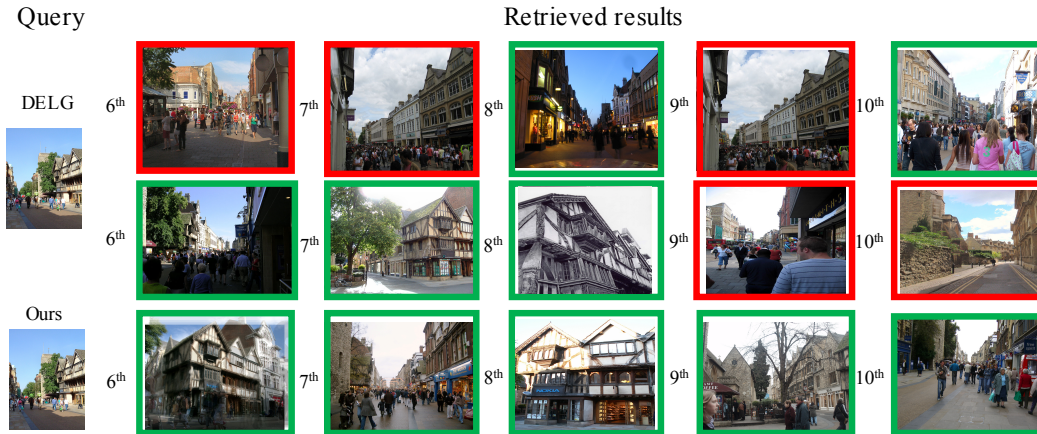


Figure 5: Demonstration of top-10 retrieved results. The top-5 retrieved images are all correct and are excluded in this figure. Results of DELG global, DELG global+local and our DOLG are shown from top to bottom. Green and red boxes denote positive and negative images, respectively.

Method	Roxf			Rpar		
	E	M	H	E	M	H
Concatenation [†]	91.29	78.40	56.55	95.88	89.37	76.80
Hadamard	92.21	79.20	56.76	95.94	89.91	77.40
orthogonal	93.17	80.50	58.82	95.95	89.81	77.70

Table 6: Comparison of orthogonal fusion with other fusion strategies. Concatenation and Hadamard product are explored. [†] with $m = 2.0, \gamma = 30$ for the ArcFace margin loss, otherwise the training does not converge.

complementary information to better describe a image, but also they will not put extra emphasis on global feature f_g because of their irrelevance.

5. Discussions

Here, we would like to have some discussion on our **current implementations** and **model complexity**. First of all, we have not extensively studied and tuned on many hyper-parameters, such as p for GeM, γ and m for ArcFace margin loss, and the dilation rate s settings of dilated convolution layers. Instead, we directly follow the practices in DELG [9] and ASPP [10]. We do so in order to show the effectiveness of our proposed building blocks instead of tuning for better models, although we believe tuning on these parameters may obtain better performance. The other thing worthy of mention is the orthogonal fusion module. We pay our attention to developing single stage solution by aggregating orthogonal local and global information. The design choice of the aggregation operation denoted as “A” in Figure 4a is simply chosen from GeM and average pooling for proof-of-concept purpose. Note that average pooling is a

linear operation, in this case, the orthogonal fusion module is equivalent to pool the local feature at first and then perform the projection and subtraction, so its computation can be further simplified. In short, our current orthogonal fusion module is sufficiently simple yet effective. We believe exploring more complicated learning based aggregation “A” in Figure 4a is promising and it is left as our future work. As for the complexity, compared to DELG [9] and DELF [29], extra computational cost comes from the Multi-Atrous module and the orthogonal fusion module. The former one is composed of a few dilated convolution layers meanwhile the latter one can currently be reduced to $Pool(f_i) - (Pool(f_i) \cdot f_g) / |f_g|^2$. Therefore, the overhead of our solution is quite limited. Besides, our retrieval process can be finished in a single-stage.

6. Conclusion

In this paper, we make the first attempt to fuse local and global features in an orthogonal manner for effective single-stage image retrieval. We have designed a novel local feature learning branch, where multi-atrous module is leveraged to simulate spatial feature pyramid to handle scale variation among images and self-attention module is adopted to perform significance modeling for each local descriptor. We also design a novel orthogonal fusion module in order to combine complementary local and global information, to mutually reinforce each other and produce a representative final descriptor via objective-oriented training. Extensive experimental results have been shown for proof-of-concept purpose, and we also significantly improve state-of-the-art performance on Roxf and Rpar.

References

- [1] Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Netvlad: Cnn architecture for weakly supervised place recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5297–5307, 2016. 1, 2, 3
- [2] Yannis Avrithis and Giorgos Tolias. Hough pyramid matching: Speeded-up geometry re-ranking for large scale image retrieval. *International journal of computer vision*, 107(1):1–19, 2014. 2
- [3] Artem Babenko and Victor Lempitsky. Aggregating local deep features for image retrieval. In *Proceedings of the IEEE international conference on computer vision*, pages 1269–1277, 2015. 1, 2
- [4] Artem Babenko, Anton Slesarev, Alexandr Chigorin, and Victor Lempitsky. Neural codes for image retrieval. In *European conference on computer vision*, pages 584–599. Springer, 2014. 1, 2
- [5] Vassileios Balntas, Edgar Riba, Daniel Ponsa, and Krystian Mikolajczyk. Learning local feature descriptors with triplets and shallow convolutional neural networks. In *Bmvc*, volume 1, page 3, 2016. 1, 2
- [6] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool. Speeded-up robust features (surf). *Computer vision and image understanding*, 110(3):346–359, 2008. 1, 2
- [7] Jeffrey S Beis and David G Lowe. Shape indexing using approximate nearest-neighbour search in high-dimensional spaces. In *Proceedings of IEEE computer society conference on computer vision and pattern recognition*, pages 1000–1006. IEEE, 1997. 2
- [8] Michael M Bronstein, Alexander M Bronstein, Fabrice Michel, and Nikos Paragios. Data fusion through cross-modality metric learning using similarity-sensitive hashing. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 3594–3601. IEEE, 2010. 2
- [9] Bingyi Cao, André Araujo, and Jack Sim. Unifying deep local and global features for image search. In *European Conference on Computer Vision*, pages 726–743. Springer, 2020. 1, 2, 3, 4, 5, 6, 7, 8
- [10] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. 2, 3, 7, 8
- [11] Weihua Chen, Xiaotang Chen, Jianguo Zhang, and Kaiqi Huang. Beyond triplet loss: a deep quadruplet network for person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 403–412, 2017. 2
- [12] Wei Chen, Yu Liu, Weiping Wang, Erwin Bakker, Theodoros Georgiou, Paul Fieguth, Li Liu, and Michael S Lew. Deep image retrieval: A survey. *arXiv preprint arXiv:2101.11282*, 2021. 2
- [13] Random Sample Consensus. A paradigm for model fitting with applications to image analysis and automated cartography. *MA Fischler, RC Bolles*, 6:381–395, 1981. 2
- [14] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2019. 2, 4
- [15] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 224–236, 2018. 2, 3
- [16] Mihai Dusmanu, Ignacio Rocco, Tomas Pajdla, Marc Pollefeys, Josef Sivic, Akihiko Torii, and Torsten Sattler. D2-net: A trainable cnn for joint detection and description of local features. *arXiv preprint arXiv:1905.03561*, 2019. 2
- [17] Albert Gordo, Jon Almazan, Jerome Revaud, and Diane Larlus. End-to-end learning of deep visual representations for image retrieval. *International Journal of Computer Vision*, 124(2):237–254, 2017. 2, 6
- [18] Kun He, Yan Lu, and Stan Sclaroff. Local descriptors optimized for average precision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 596–605, 2018. 1, 2
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016. 3, 4
- [20] Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, Teddy Furon, and Ondrej Chum. Efficient diffusion on region manifolds: Recovering small objects with compact cnn representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2077–2086, 2017. 6
- [21] Herve Jegou, Matthijs Douze, and Cordelia Schmid. Hamming embedding and weak geometric consistency for large scale image search. In *European conference on computer vision*, pages 304–317. Springer, 2008. 2
- [22] Hervé Jégou, Matthijs Douze, Cordelia Schmid, and Patrick Pérez. Aggregating local descriptors into a compact representation. 2
- [23] Hervé Jégou, Matthijs Douze, Cordelia Schmid, and Patrick Pérez. Aggregating local descriptors into a compact image representation. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 3304–3311. IEEE, 2010. 2
- [24] Hervé Jégou, Florent Perronnin, Matthijs Douze, Jorge Sánchez, Patrick Pérez, and Cordelia Schmid. Aggregating local image descriptors into compact codes. *IEEE transactions on pattern analysis and machine intelligence*, 34(9):1704–1716, 2011. 2
- [25] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004. 1, 2
- [26] Dmytro Mishkin, Filip Radenovic, and Jiri Matas. Repeatability is not enough: Learning affine regions via discriminability. *Springer, Cham*, 2018. 6
- [27] Tony Ng, Vassileios Balntas, Yurun Tian, and Krystian Mikolajczyk. Solar: Second-order loss and attention for image retrieval. *Arxiv*, 2020. 2, 5, 6

- [28] David Nister and Henrik Stewenius. Scalable recognition with a vocabulary tree. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 2161–2168. Ieee, 2006. [2](#)
- [29] Hyeonwoo Noh, Andre Araujo, Jack Sim, Tobias Weyand, and Bohyung Han. Large-scale image retrieval with attentive deep local features. In *Proceedings of the IEEE international conference on computer vision*, pages 3456–3465, 2017. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [30] James Philbin, Ondrej Chum, Michael Isard, Josef Sivic, and Andrew Zisserman. Object retrieval with large vocabularies and fast spatial matching. In *2007 IEEE conference on computer vision and pattern recognition*, pages 1–8. IEEE, 2007. [2](#)
- [31] Qi Qin, Wenpeng Hu, and Bing Liu. Feature projection for improved text classification. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8161–8171, 2020. [2](#)
- [32] Filip Radenović, Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondřej Chum. Revisiting oxford and paris: Large-scale image retrieval benchmarking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5706–5715, 2018. [2](#), [5](#), [6](#)
- [33] Filip Radenović, Giorgos Tolias, and Ondřej Chum. Cnn image retrieval learns from bow: Unsupervised fine-tuning with hard examples. In *European conference on computer vision*, pages 3–20. Springer, 2016. [2](#)
- [34] Filip Radenović, Giorgos Tolias, and Ondřej Chum. Fine-tuning cnn image retrieval with no human annotation. *IEEE transactions on pattern analysis and machine intelligence*, 41(7):1655–1668, 2018. [2](#), [3](#), [6](#), [7](#)
- [35] Jerome Revaud, Jon Almazán, Rafael S Rezende, and Cesar Roberto de Souza. Learning with average precision: Training image retrieval with a listwise loss. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5107–5116, 2019. [2](#), [6](#)
- [36] Jerome Revaud, Philippe Weinzaepfel, César De Souza, Noe Pion, Gabriela Csurka, Johann Cabon, and Martin Humenberger. R2d2: repeatable and reliable detector and descriptor. *arXiv preprint arXiv:1906.06195*, 2019. [1](#), [2](#)
- [37] Paul-Edouard Sarlin, Cesar Cadena, Roland Siegwart, and Marcin Dymczyk. From coarse to fine: Robust hierarchical localization at large scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12716–12725, 2019. [1](#), [3](#)
- [38] Oriane Siméoni, Yannis Avrithis, and Ondrej Chum. Local features and visual words emerge in activations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11651–11660, 2019. [2](#), [3](#), [6](#)
- [39] Josef Sivic and Andrew Zisserman. Video google: A text retrieval approach to object matching in videos. In *Computer Vision, IEEE International Conference on*, volume 3, pages 1470–1470. IEEE Computer Society, 2003. [2](#)
- [40] Hajime Taira, Masatoshi Okutomi, Torsten Sattler, Mircea Cimpoi, Marc Pollefeys, Josef Sivic, Tomas Pajdla, and Akihiko Torii. Inloc: Indoor visual localization with dense matching and view synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7199–7209, 2018. [3](#)
- [41] Marvin Teichmann, Andre Araujo, Menglong Zhu, and Jack Sim. Detect-to-retrieve: Efficient regional aggregation for image search. *arXiv*, 2019. [6](#)
- [42] Giorgos Tolias, Yannis Avrithis, and Hervé Jégou. Image search with selective match kernels: aggregation across single and multiple images. *International Journal of Computer Vision*, 116(3):247–261, 2016. [2](#), [5](#), [6](#)
- [43] Giorgos Tolias, Tomas Jenicek, and Ondřej Chum. Learning and aggregating deep local descriptors for instance-level recognition. In *European Conference on Computer Vision*, pages 460–477. Springer, 2020. [2](#), [5](#), [6](#)
- [44] Giorgos Tolias, Ronan Sicre, and Hervé Jégou. Particular object retrieval with integral max-pooling of cnn activations. *arXiv preprint arXiv:1511.05879*, 2015. [1](#), [2](#)
- [45] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5265–5274, 2018. [2](#)
- [46] Jian Wang, Feng Zhou, Shilei Wen, Xiao Liu, and Yuanqing Lin. Deep metric learning with angular loss. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2593–2601, 2017. [2](#)
- [47] Tobias Weyand, Andre Araujo, Bingyi Cao, and Jack Sim. Google landmarks dataset v2-a large-scale benchmark for instance-level recognition and retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2575–2584, 2020. [4](#)
- [48] Kwang Moo Yi, Eduard Trulls, Vincent Lepetit, and Pascal Fua. Lift: Learned invariant feature transform. In *European conference on computer vision*, pages 467–483. Springer, 2016. [2](#)
- [49] Yin Zhang, Rong Jin, and Zhi-Hua Zhou. Understanding bag-of-words model: a statistical framework. *International Journal of Machine Learning and Cybernetics*, 1(1-4):43–52, 2010. [2](#)
- [50] Yan-Tao Zheng, Ming Zhao, Yang Song, Hartwig Adam, Ulrich Buddemeier, Alessandro Bissacco, Fernando Brucher, Tat-Seng Chua, and Hartmut Neven. Tour the world: building a web-scale landmark recognition engine. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1085–1092. IEEE, 2009. [1](#)
- [51] Wengang Zhou, Houqiang Li, and Qi Tian. Recent advance in content-based image retrieval: A literature survey. *arXiv preprint arXiv:1706.06064*, 2017. [2](#)