

# BoQ: A Place is Worth a Bag of Learnable Queries

Amar Ali-bey\*      Brahim Chaib-draa      Philippe Giguère  
 Department of Computer Science and Software Engineering  
 Université Laval, Québec, Canada

amar.ali-bey.1@ulaval.ca, {brahim.chaib-draa, philippe.giguere}@ift.ulaval.ca

## Abstract

In visual place recognition, accurately identifying and matching images of locations under varying environmental conditions and viewpoints remains a significant challenge. In this paper, we introduce a new technique, called *Bag-of-Queries (BoQ)*, which learns a set of global queries, designed to capture universal place-specific attributes. Unlike existing techniques that employ self-attention and generate the queries directly from the input, BoQ employ distinct learnable global queries, which probe the input features via cross-attention, ensuring consistent information aggregation. In addition, this technique provides an interpretable attention mechanism and integrates with both CNN and Vision Transformer backbones. The performance of BoQ is demonstrated through extensive experiments on 14 large-scale benchmarks. It consistently outperforms current state-of-the-art techniques including NetVLAD, MixVPR and EigenPlaces. Moreover, despite being a global retrieval technique (one-stage), BoQ surpasses two-stage retrieval methods, such as Patch-NetVLAD, TransVPR and R2Former, all while being orders of magnitude faster and more efficient. The code and model weights are publicly available at <https://github.com/amaralibey/Bag-of-Queries>.

## 1. Introduction

Visual Place Recognition (VPR) consists of determining the geographical location of a place depicted in a given image, by comparing its visual features to a database of previously visited places. The dynamic and ever-changing nature of real-world environments pose significant challenges for VPR [33, 57]. Factors such as varying lighting conditions, seasonal changes and the presence of dynamic elements such as vehicles and pedestrians introduce considerable variability into the appearance of a place. Additionally, changes in viewpoint and image scale can expose previ-




Figure 1. Recall@1 performance comparison between our proposed technique, Bag-of-Queries (BoQ), and current state of the art methods, Conv-AP [3], CosPlace [11], MixVPR [4] and EigenPlaces [12]. ResNet-50 is used as backbone for all techniques. BoQ consistently achieves better performance in various environment conditions such as viewpoint changes (Pitts-250k [44], MapillarySLS [50]), seasonal changes (Nordland [53]), historical locations (AmsterTime [51]) and extreme lightning and weather conditions (SVOX [10]).

ously obscured areas, further complicating the recognition process. These challenges are exacerbated by the operational constraints of VPR systems, which often need to operate in real-time and under limited memory. Consequently, there is a compelling need for efficient algorithms capable of generating compact yet robust representations.

With the rise of deep learning, numerous VPR-specific neural networks have been proposed [2–4, 6, 11, 20, 28, 30, 41], leveraging Convolutional Neural Networks (CNNs) to extract high-level features, followed by aggregation layers that consolidate these features into a single global descriptor. Such end-to-end trainable architectures have been in-

\*Corresponding author

strumental in enhancing the efficiency and performance of VPR systems.

Recently, Vision Transformers (ViT) [18] have demonstrated remarkable performance in various computer vision tasks, including image classification [15], object detection [14, 31] and semantic segmentation [58]. Their success can be attributed to their self-attention mechanism, which captures global dependencies between distant parts of the image [21]. However, current Transformer-based VPR techniques [46, 49, 54, 59], often rely on *reranking* [8], a post-processing step aimed at refining the initial set of candidate locations identified through global descriptor search. The reranking process is usually done with geometric verification (*e.g.* RANSAC) on stored local patch tokens, which is computationally and memory intensive. Moreover, the global retrieval process in Transformer-based methods, whether through specific aggregation of local patches [46, 54] or using the class token of the ViT [59], has yet to reach the performance levels of non-Transformer-based approaches like MixVPR [4] and CosPlace [58].

In this paper, we bridge this performance gap, by introducing a novel Transformer-based aggregation technique, called Bag-of-Queries (BoQ), that learns a set of embeddings (global queries) and employs a cross-attention mechanism to probe local features coming from the backbone network. This approach enables each global query to consistently seek relevant information uniformly across input images. This is in contrast with self-attention-based techniques [46, 54], where the queries are dynamically derived from the input itself. Furthermore, BoQ is designed for end-to-end training, thus seamlessly integrating with both conventional CNN and ViT backbones. Its effectiveness is validated through extensive experiments on multiple large-scale benchmarks, consistently outperforming state-of-the-art techniques, including MixVPR [4], EigenPlaces [12], and NetVLAD [6]. More importantly, BoQ, a single-stage (global) retrieval method that does not employ reranking, outperforms two-stage retrieval methods like TransVPR [46] and R2Former [59], while being orders of magnitude more efficient in terms of computational and memory resources.

## 2. Related Works

Early visual place recognition (VPR) techniques relied on hand-crafted local features such as SIFT [32], SURF [9] and ORB [40], which were aggregated into global descriptors using techniques like Bag-of-Words (BoW) [19, 38, 44] or Vector of Locally Aggregated Descriptors (VLAD) [5, 26]. BoW involves learning a *visual vocabulary* (or set of clusters), where each *visual word* (or cluster) represents a specific visual characteristic. Images are then represented as histograms, with bins indicating the frequency of each visual word. VLAD extends on this by capturing first-order

statistics by accumulating the differences between local descriptors and their respective cluster centers.

**CNN-based VPR.** The advent of deep learning marked a significant shift in VPR techniques, with various aggregation architectures proposed. Arandjelovic *et al.* [6] introduced NetVLAD, a trainable version of VLAD that integrates with CNN backbones, achieving superior performance over traditional methods. Following its success, many researchers proposed variants such as CRN [28], SPE-VLAD [52], Gated NetVLAD [55] and SSL-VLAD [35]. Other approaches focus on regions of interests within feature maps [7, 43, 56]. Another key aggregation method in image retrieval is the Generalized Mean (GeM) [39], which is a learnable form of global pooling. Building on GeM, Berton *et al.* [58] have recently introduced CosPlace, enhancing GeM by integrating it with a linear projection layer, achieving remarkable performance for VPR. Ali-bey *et al.* [4] introduced MixVPR, an all-MLP aggregation technique that achieved state-of-the-art performance on multiple benchmarks. Another facet of VPR is the training procedure, where the focus is on the loss function [2, 3, 12, 29].

**Transformer-based VPR.** The Transformer architecture was initially introduced for natural language processing [45], and later adapted to Vision Transformers (ViT) for computer vision applications [18]. In VPR, they have recently been used as backbone combined with MixVPR [24, 25] or NetVLAD [48] resulting in a performance boost compared to CNN backbones. Furthermore, AnyLoc [27] used features extracted from a foundation model, called DINOv2 [36], combined with unsupervised aggregation methods such as VLAD, resulting in notable performance on a wide range of benchmarks.

Recently, two-stage retrieval strategy has gained prominence in Transformers-based VPR. It starts with global retrieval using global descriptors to identify top candidates, followed by a computationally-intensive reranking phase that refines the results based on local features. In this context, Wang *et al.* [46] introduced TransVPR, which uses CNNs for feature extraction and Transformers for attention fusion to create global image descriptors, with additional patch-level descriptors for geometric verification. On the other hand, Zhang *et al.* [54] proposed ETR, a Transformer-based reranking technique that uses a CNN backbone for local and global descriptors. Their method leverages cross-attention between the local features of two images to cast the reranking as classification. Recently, Zhu *et al.* [59] proposed a unified framework integrating global retrieval and reranking, solely using transformers. For global retrieval, it employs the class token and utilizes other image tokens as local features for the reranking module. These two-stage techniques showed great performance, but at the ex-

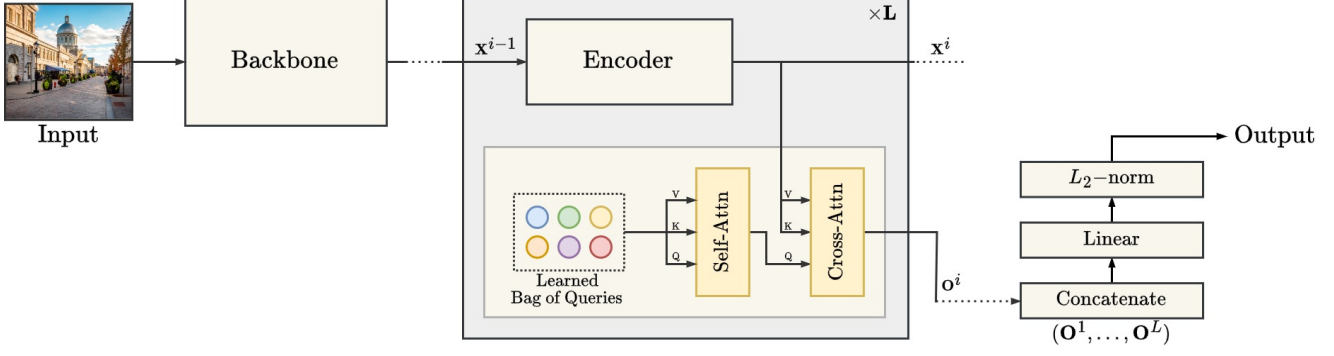


Figure 2. Overall architecture of the Bag-of-Queries (BoQ) model. The input image is first processed by a backbone network to extract its local features, which are then sequentially refined in a cascade of Encoder units. Each BoQ block contains a set of learnable queries  $\mathbf{Q}$  (Learned Bag of Queries), which undergo self-attention to integrate their shared information. The refined features  $\mathbf{X}^i$  are then processed through cross-attention with  $\mathbf{Q}$  for selective aggregation. Outputs from all BoQ blocks ( $\mathbf{O}^1, \mathbf{O}^2, \dots, \mathbf{O}^L$ ) are concatenated and linearly projected. The final global descriptor is L2-normalized to optimize it for subsequent similarity search.

pense of more computation and memory overheads. However, their global retrieval performance is still very limited compared to one-stage methods such as MixVPR [4] and CosPlace [11]. In this work, we propose, Bag-of-Queries (BoQ), a Transformer-based aggregation technique for global retrieval, that outperforms existing state-of-the-art without relying on reranking. This makes BoQ particularly suitable for applications where computational resources are limited, yet high accuracy and efficiency are essential.

### 3. Methodology

In visual place recognition, effective aggregation of local features is crucial for accurate and fast global retrieval. To address this, we propose the *Bag-of-Queries* (BoQ) technique, a novel aggregation architecture that is end-to-end trainable and surprisingly simple, as depicted in Fig. 2.

Our technique rely on Multi-Head Attention (MHA) mechanism [45], which takes three inputs, queries ( $\mathbf{q}$ ), keys ( $\mathbf{k}$ ) and values ( $\mathbf{v}$ ), and linearly project them into multiple parallel heads. The output, for each head, is computed as follows:

$$\text{MHA}(\mathbf{q}, \mathbf{k}, \mathbf{v}) = \text{softmax}\left(\frac{\mathbf{q} \cdot \mathbf{k}^T}{\sqrt{d}}\right)\mathbf{v}. \quad (1)$$

In this mechanism, the queries  $\mathbf{q}$  play a crucial role. They act as a set of filters, determining which parts of the input (represented by keys  $\mathbf{k}$  and values  $\mathbf{v}$ ) are most relevant. The attention scores, derived from the dot product between queries and keys, essentially indicate to the model the degree of *attention* to give to each part of the input. We refer to self-attention when  $\mathbf{q}$ ,  $\mathbf{k}$  and  $\mathbf{v}$  are derived from the same input, *e.g.*  $\text{MHA}(\mathbf{x}, \mathbf{x}, \mathbf{x})$ . In contrast, cross-attention is the scenario where the query comes from a different source than the key and value.

Given an input image  $I \in \mathbb{R}^{3 \times H \times W}$ , we first process it through a backbone network, typically a pre-trained Convolutional Neural Network (CNN) or Vision Transformer (ViT), to extract its high-level features. For CNN backbones, we apply a  $3 \times 3$  convolution to reduce their dimensionality, whereas, in the case of a ViT backbone, we apply a linear projection for the same purpose. We regard the result as a sequence of  $N$  local features of dimension  $d$ , such as:  $\mathbf{X}^0 = [\mathbf{x}_1^0, \mathbf{x}_2^0, \dots, \mathbf{x}_N^0]$ . We then process  $\mathbf{X}^0$  through a sequence of Transformer-Encoder units and BoQ blocks. Each encoder transforms its input features, and passes the result to its subsequent BoQ block as follows:

$$\mathbf{X}^i = \text{Encoder}^i(\mathbf{X}^{i-1}). \quad (2)$$

Here,  $\mathbf{X}^{i-1}$  denotes the feature input to the  $i$ -th Encoder unit, and  $\mathbf{X}^i$  represents the transformed output, which becomes the input for the next block in the pipeline.

Each BoQ block contains a fixed set (*a bag*) of  $M$  learnable queries, denoted as  $\mathbf{Q}^i = [\mathbf{q}_1^i, \mathbf{q}_2^i, \dots, \mathbf{q}_M^i]$ . These queries are trainable parameters of the model, independent of the input features (*e.g.*, `nn.Parameter` in PyTorch code), not to be confused with the term “*query images*”, which refers to the test images used for benchmarking.

Before using  $\mathbf{Q}^i$  to aggregate information in the  $i^{\text{th}}$  BoQ block, we first apply self-attention between them:

$$\mathbf{Q}^i = \text{MHA}(\mathbf{Q}^i, \mathbf{Q}^i, \mathbf{Q}^i) + \mathbf{Q}^i. \quad (3)$$

The self-attention operation allows the learnable queries to integrate their shared information during the training phase. Next, we apply a cross-attention between  $\mathbf{Q}^i$  and the input features  $\mathbf{X}^i$  of the current BoQ block:

$$\mathbf{O}^i = \text{MHA}(\mathbf{Q}^i, \mathbf{X}^i, \mathbf{X}^i). \quad (4)$$

This allows the learnable queries to dynamically assess the importance of each input feature by computing the rele-

vance scores (attention weights) between  $\mathbf{Q}^i$  and  $\mathbf{X}^i$  and aggregate them into the output  $\mathbf{O}^i$ .

Following this, the outputs from each BoQ block are concatenated to form a single output.

$$\mathbf{O} = \text{Concat}(\mathbf{O}^1, \mathbf{O}^2, \dots, \mathbf{O}^L). \quad (5)$$

This ensures that the final descriptor  $\mathbf{O}$  combines information aggregated at different levels of the network, forming a rich and hierarchical representation. The final step involves reducing the dimensionality of the output  $\mathbf{O}$  using one or two successive linear projections with weights  $\mathbf{W}_1$  and  $\mathbf{W}_2$ , similar to the approach in [3, 4].

$$\text{Output} = \mathbf{W}_2(\mathbf{W}_1\mathbf{O})^T. \quad (6)$$

This is followed by  $L_2$ -normalization, to bring the global descriptor to the unit hypersphere, which optimizes the similarity search. The training is performed using pair-based (or triplet-based) loss functions that are widely used in the literature [3].

**Relations to other methods.** The Bag-of-Queries (BoQ) approach bears conceptual resemblance to Detection Transformer (DETR) model [14] in its use of object queries. However, BoQ is fundamentally different, as its global queries are exclusively used to extract and aggregate local feature information from the input; these queries do not contribute directly to the final representation, as demonstrated by the absence of any residual connection between the global queries and the output of the cross-attention. This is in contrast to DETR, where the object queries are integral part to the output, upon which object detection and classification is done.

In comparison to NetVLAD [6], which employs a set of learned cluster centers, and aggregates local features by calculating the weighted sum of their distances to each center. BoQ, leverages cross-attention between its learned queries and the input local feature to dynamically assess their relevance.

## 4. Experiments

In this section, we conduct comprehensive experiments to demonstrate the effectiveness of our technique, Bag-of-Queries (BoQ), compared to current state-of-the-art methods. First, we describe the datasets used, and then the implementation details. This is followed by the evaluation metrics. We then provide a detailed comparative analysis of performance, ablation studies and qualitative results.

### 4.1. Datasets

Our experiments are performed on existing benchmarks that present a wide range of real-world challenges for VPR systems. Tab. 1 provides a brief summary of these benchmarks. MapillarySLS (MSLS) [50] was collected using

Dataset name	# quer.	# ref.	Variations		
			Viewpoint	Season	Day/Night
MSLS [50]	740	18.9k	×	×	×
Pitts250k [44]	8.2k	84k	×	×	×
Pitts30k [44]	6.8k	10k	×	×	×
AmsterTime [51]	1231	1231	×	×	×
Eynsham [16]	24k	24k	×		
Nordland* [53]	2760	27.6k		×	×
Nordland** [42]	27.6k	27.6k		×	×
St-Lucia [34]	1464	1549			
SVOX [10]	14.3k	17.2k	×	×	×
SPED [53]	607	607		×	×

Table 1. Overview of VPR benchmarks used in our experiments, highlighting their number of query images (# quer.) and reference images (# ref.), as well as their variations in terms of viewpoint, season, and day/night. The number of × marks indicates the degree of variations present (× means low, ××× means high).

dashcams and features a broad range of viewpoints and lighting changes, testing the system’s adaptability to typical variations in ground-level VPR. Pitts250k [44] and Pitts30k datasets, extracted from Google Street View, exhibit significant changes in viewpoint, which tests the system’s ability to maintain recognition across varying angles and perspectives. AmsterTime [51] presents a unique challenge with historical grayscale images as queries and contemporary color images as references, covering temporal variations over several decades. Eynsham [16] consists entirely of grayscale images, adding complexity due to the absence of color information. Nordland [42] dataset was collected over four seasons, using a camera mounted on the front of a train. It encompasses scenes ranging from snowy winters to sunny summers, presenting the challenge of extreme appearance changes due to seasonal variations. Note that two variants of Nordland are used in VPR benchmarking: one uses a subset of 2760 summer images as queries against the entire winter sequence as reference images (marked with \*), while the other uses the entire winter sequence as query images against the entire summer sequence as reference images (marked with \*\*). SVOX [10] stands out with its comprehensive coverage of weather conditions, including overcast, rainy, snowy, and sunny images, to test adaptability to meteorological changes. Additionally, SVOX Night subset focuses on nocturnal scenes, a challenging scenario for most VPR systems.

### 4.2. Implementation details

**Architecture.** For our experiments, we primarily employ pre-trained ResNet-50 backbones for feature extraction [23]. Importantly, our proposed Bag-of-Queries (BoQ) technique is architecturally agnostic and seamlessly integrates with various CNNs or Vision Transformer (ViT) backbones. When employing a ResNet, we crop it at the

second last residual to preserve a higher resolution of local features. Additionally, we freeze all but the last residual block and add a  $3 \times 3$  convolution to reduce the number of channels, which decreases the computational and memory footprint of the self-attention mechanism. Implementing the BoQ model is straightforward with popular deep learning frameworks such as PyTorch [37] and TensorFlow [1], both providing highly optimized implementations of self-attention and cross-attention, which are the basic blocks of our technique.

**Training.** We train BoQ following the standard framework of GSV-Cities [3], which provides a highly accurate dataset of 560k images depicting 67k places. For the loss function, we use Multi-Similarity loss [47], as it has been shown to perform best for visual place recognition [3]. We use batches containing 120 – 200 places, each depicted by 4 images resulting in mini-batches of 480 – 800 images. We use the AdamW optimizer with a learning rate of 0.0002 – 0.0004 depending the batch-size, and a weight decay of 0.001. The learning rate is multiplied by 0.3 after each 5 – 10 epochs. Finally, we train for a maximum of 40 epochs (including 3 epochs of linear warmup) using images resized to  $320 \times 320$ . Note that while current techniques such as NetVLAD [6], CosPlace [58] and EigenPlaces [12] train on images of size  $480 \times 640$ —effectively triple the pixel count of our chosen resolution—our BoQ model still achieves better performance, underscoring its effectiveness.

**Evaluation metrics.** We follow the same evaluation metric of existing literature [2–4, 6, 11, 12, 28, 46, 50, 53], where the recall@ $k$  is measured. Recall@ $k$  is defined as the percentage of query images for which at least one of the top- $k$  predicted reference images falls within a predefined threshold distance, which is typically 25 meters. However, for the Nordland benchmark, which features aligned sequences, the threshold is 10 frame counts for Nordland\*\* and only 1 frame for Nordland\*, which makes the latter more challenging.

### 4.3. Comparison with previous works

In this section, we present an extensive comparison of our proposed method, Bag-of-Queries (BoQ), with a wide range of existing state-of-the-art VPR aggregation techniques. This includes global average pooling (AVG) [6], Generalized Mean (GeM) [39], NetVLAD [6] alongside its recent variants, SPE-NetVLAD [52] and Gated NetVLAD [55]. Importantly, we compare against recent cutting-edge techniques, such as ConvAP [3], CosPlace [58], EigenPlaces [12] and MixVPR [4], which currently hold the best scores in most existing benchmarks, setting a high standard for BoQ to measure up against in the field of VPR.

Moreover, despite BoQ being a global retrieval technique which does not employ reranking, we compare it with

current state-of-the-art two-stage retrieval techniques [13, 17, 22, 46, 59]. These methods leverage geometric verification to enhance recall@ $k$  performance at the expense of increased computation and memory overhead.

**Results analysis.** In Tab. 2, we present a comparative analysis across several challenging datasets, focusing on Recall@1 (R@1), Recall@5 (R@5), and Recall@10 (R@10). We also offer insight into the performance of each method with regard to the top-ranked retrieved images.

- On Pitts250k-test, BoQ outperforms all other methods with a marginal yet notable improvement as it is the first global technique that reaches the 95% R@1 threshold on this benchmark.
- On MSLS-val and SPED, BoQ outperforms the second best methods (EigenPlaces and MixVPR) by 2.2 and 1 percentage points respectively.
- On the challenging Nordland\* benchmark, known for its extreme seasonal variations, BoQ significantly outperforms all compared methods, indicating its robustness to severe environmental changes. Surpassing CosPlace and Eigenplaces by 20, and MixVPR by 16 percentage points, which translates to over 450 additional accurate retrievals within the top-1 results.

In Tab. 3, we follow the same evaluation style proposed by Berton *et al.* [12], where benchmarks are categorized into multi-view datasets, where images are oriented in various angles, and frontal-view dataset where images are mostly forward facing. Top-1 retrieval performance (R@1) is reported for each technique.

- For the multi-view benchmarks, BoQ achieves highest R@1 on AmsterTime and Eynsham, and second best compared to EigenPlaces on Pitts30k-test. Both techniques show robustness to viewpoint changes in urban environments, while facing challenges with the AmsterTime dataset, which includes decades-old historical imagery.
- On frontal-view benchmarks, particularly on Nordland\*\* (which comprises 27.6k queries), BoQ achieves R@1 of 85.5%, which is 9.3 and 13 percentage points more than MixVPR, and CosPlace respectively. This translates to BoQ correctly retrieving an additional 2480 and 3800 images within the top-1 results.
- On SVOX, BoQ achieves best results under all conditions, especially on SVOX night, where it achieves 85.2% R@1, outperforming the second best method (MixVPR) by 20.8 percentage points. This highlights BoQ’s potential under low-light conditions.

Furthermore, we compare BoQ to two-stage retrieval techniques in Tab. 4, by reporting performance on MSLS-val and Pitts30k-test. Although our technique does not perform reranking, its global retrieval performance surpasses that of existing two-stage techniques, including Patch-NetVLAD (2021), TransVPR (2022) and R2Former (2023). Importantly, this makes our approach more efficient in terms of

Method	Dim.	Pitts250k-test			MSLS-val			SPED			Nordland*		
		R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10
AVG [6]	2048	78.3	89.8	92.6	73.5	83.9	85.8	58.8	77.3	82.7	15.3	27.4	33.9
GeM [39]	2048	82.9	92.1	94.3	76.5	85.7	88.2	64.6	79.4	83.5	20.8	33.3	40.0
NetVLAD [6]	32768	90.5	96.2	97.4	82.6	89.6	92.0	78.7	88.3	91.4	32.6	47.1	53.3
SPE-NetVLAD [52]	163840	89.2	95.3	97.0	78.2	86.8	88.8	73.1	85.5	88.7	25.5	40.1	46.1
Gated NetVLAD [55]	32768	89.7	95.9	97.1	82.0	88.9	91.4	75.6	87.1	90.8	34.4	50.4	57.7
Conv-AP [3]	4096	92.4	97.4	98.4	83.4	90.5	92.3	80.1	90.3	93.6	38.2	54.8	61.2
CosPlace [11]	2048	92.3	97.4	98.4	87.4	<u>93.8</u>	94.9	75.3	85.9	88.6	54.4	69.8	75.9
MixVPR [4]	4096	<u>94.2</u>	<u>98.2</u>	<u>98.9</u>	88.0	92.7	94.6	<u>85.2</u>	<u>92.1</u>	94.6	<u>58.4</u>	<u>74.6</u>	<u>80.0</u>
EigenPlaces [12]	2048	94.1	97.9	98.7	<u>89.2</u>	93.6	<u>95.0</u>	82.4	91.4	<u>94.7</u>	54.2	68.0	73.9
<b>BoQ (Ours)</b>	4096	<b>95.0</b>	<b>98.4</b>	<b>99.1</b>	91.1	<b>94.8</b>	95.7	85.4	93.1	95.4	69.5	83.4	87.0
<b>BoQ (Ours)</b>	16384	<b>95.0</b>	98.3	99.0	<b>91.4</b>	94.5	<b>96.1</b>	<b>86.2</b>	<b>94.4</b>	<b>96.1</b>	<b>74.4</b>	<b>86.1</b>	<b>89.8</b>

Table 2. Comparison of our technique, Bag-of-Queries (*BoQ*), to existing state-of-the-art methods. Best results are in shown in **bold** and second best are underlined. All methods use a pre-trained ResNet-50 backbone and follow identical training procedures on the GSV-Cities dataset, except for CosPlace and EigenPlaces where authors’ pre-trained weights demonstrated superior performance and are thus used here.

Method	Backbone	Multi-view datasets				Frontal-view datasets					
		AmsterTime	Eynsham	Pitts30k	Nordland**	St Lucia	SVOX Night	SVOX Overcast	SVOX Rain	SVOX Snow	SVOX Sun
NetVLAD [6]	VGG-16	16.3	77.7	85.0	13.1	64.6	8.0	66.4	51.5	54.4	35.4
SFRS [20]	VGG-16	29.7	72.3	89.1	16.0	75.9	28.6	81.1	69.7	76.0	54.8
CosPlace [11]	VGG-16	38.7	88.3	88.4	58.5	95.3	44.8	88.5	85.2	89.0	67.3
EigenPlaces [12]	VGG-16	38.0	89.4	89.7	54.5	95.4	42.3	89.4	83.5	89.2	69.7
Conv-AP [3]	ResNet-50	33.9	87.5	90.5	62.9	99.7	43.4	91.9	82.8	91.0	80.4
CosPlace [11]	ResNet-50	47.7	90.0	90.9	71.9	99.6	50.7	92.2	87.0	92.0	78.5
MixVPR [4]	ResNet-50	40.2	89.4	91.5	<u>76.2</u>	99.6	<u>64.4</u>	<u>96.2</u>	<u>91.5</u>	<u>96.8</u>	84.8
EigenPlaces [12]	ResNet-50	<u>48.9</u>	<u>90.7</u>	<b>92.5</b>	71.2	99.6	58.9	93.1	90.0	93.1	<u>86.4</u>
<b>BoQ (Ours)</b>	ResNet-50	<b>53.0</b>	<b>91.5</b>	<u>92.4</u>	<b>85.5</b>	<b>99.9</b>	<b>85.2</b>	<b>98.3</b>	<b>96.4</b>	<b>98.4</b>	<b>96.5</b>

Table 3. Recall@1 comparison across multi-view and frontal-view datasets for various techniques, including our Bag-of-Queries (BoQ). Our method achieves state-of-the-art performance on most benchmarks, demonstrating robustness to extreme weather and illumination changes, particularly on the challenging SVOX dataset.

memory and computation: over  $30\times$  faster retrieval time compared to R2Former, which is currently the fastest two-stage VPR technique.

#### 4.4. Ablation studies

**Number of learnable queries.** We present an ablation study in Tab. 5 to provide insights into the impact of varying the number of learnable queries,  $Q$ , on the performance of BoQ. Using ResNet-50 as a backbone and incorporating two BoQ blocks, we observe that performance improves as the number of queries increases. We observe that in less diverse environments, such as Pitts30k, incorporating additional queries yields marginal performance improvements (going from 8 to 64 queries brings as little as 0.2% R@1 performance gain). However, as demonstrated by the results on AmsterTime, which contains highly diverse images

(spanning decades), the model benefits from adding more queries. This is in accordance with the intuition behind BoQ where each query implicitly learns a distinct universal pattern.

**Number of BoQ blocks.** To assess the influence of the number of BoQ blocks within our architecture, we conduct experiments by varying  $L$ , the number BoQ blocks utilized. We use a ResNet-18 backbone followed by  $L$  BoQ blocks, each comprising a bag of 32 learnable queries. We report R@1 performance for each setup. The results, presented in Tab. 6, demonstrate that even when paired with a lightweight backbone like ResNet-18, BoQ remains competitive against state-of-the-art methods such as CosPlace [11] and MixVPR [4], which use ResNet-50 backbone. Best overall performance is observed when employing 4 BoQ blocks.

	Extrac. (ms)	Rerank. (ms)	MSLS-val		Pitts30k	
			R@1	R@5	R@1	R@5
SP-SuperGlue [17]	160	7500	78.1	81.9	87.2	94.8
DELG [13]	190	35200	83.2	90.0	89.9	95.4
Patch-NetVLAD [22]	1300	7400	79.5	86.2	88.7	94.5
TransVPR [46]	45	3200	86.8	91.2	89.0	94.9
R2Former [59]	31	400	89.7	<b>95.0</b>	91.1	95.2
<b>BoQ (Ours)</b>	7	–	<b>91.4</b>	94.5	<b>92.4</b>	<b>96.2</b>

Table 4. **Comparing against two-stage retrieval techniques** in terms of performance and latency of feature extraction and reranking (when applicable). The first 5 techniques use a second refinement pass (matching) to re-rank the top 100 candidates in order to improve retrieval performance. BoQ (ours) does not employ re-ranking, which makes it orders of magnitude faster than the fastest two-stage technique.

Size of $Q$	MSLS-val		Pitts30k-val		AmsterTime	
	R@1	R@5	R@1	R@5	R@1	R@5
4	86.9	92.7	93.1	97.9	42.7	61.4
8	88.1	93.4	93.9	98.4	44.3	63.8
16	88.7	93.4	94.0	98.5	46.2	68.4
32	90.6	94.0	94.1	98.6	48.9	69.0
64	91.3	94.3	94.5	98.9	52.0	70.1

Table 5. **Ablation on the number of learnable queries.** We vary the size of  $Q$  in each BoQ block. ResNet-50 is used as backbone, with two BoQ blocks. Performance increases with the number of queries until stagnation at 64 queries.

# blocks ( $\times L$ )	# params (M)	MSLS	Nordland	AmsterTime	SVOX Night
1	2.5	87.1	56.7	40.5	62.8
2	4.3	88.0	61.1	42.2	68.6
4	8.0	88.1	62.7	44.0	71.4
8	15.4	87.2	59.4	43.0	70.1

Table 6. **Ablation on the number of BoQ blocks used.** Recall@1 performance is reported for each configuration. ResNet-18 is used as backbone. The total number of parameters (in millions) is reported.

**Backbone architecture.** In Tab. 7 we present a comparative performance of BoQ coupled with different ResNet backbone architectures. For each backbone, we crop it at the second last residual block and follow it with two BoQ blocks, each comprising 64 learnable queries. The total number of parameters (trainable + non trainable) is provided. The empirical results showcase that when BoQ is coupled with ResNet-34, a relatively lightweight backbone, it achieves competitive performance compared to NetVLAD, CosPlace and EigenPlaces coupled with ResNet-50 backbone. Interestingly, using ResNet-101, a relatively deeper backbone, we did not achieve better performance than ResNet-50,

Backbone	# params (M)	MSLS-val		SPED		Nordland*	
		R@1	R@5	R@1	R@5	R@1	R@5
ResNet-18	7.1	88.4	92.6	83.5	92.7	65.9	79.3
ResNet-34	12.5	89.5	93.1	85.5	92.6	64.1	80.1
ResNet-50	14.6	91.4	94.5	86.2	94.4	74.4	86.1
ResNet-101	42.9	90.5	95.4	85.6	93.8	71.9	84.9

Table 7. **Comparing performance with different backbones.** Each backbone is cropped at the second last residual block. Two BoQ blocks are used, with 64 learnable queries in each. We show the total number of parameters. BoQ achieves state-of-the-art performance with only a ResNet-34 backbone, which highlights its potential use for real-time scenarios.

which could be attributed to memory constraints necessitating a smaller training batch size. Further experiments with different backbones are provided in the Supplementary.

**Ablation on the self-attention.** Applying self-attention on ( $Q^i$ ) brings more stability/performance by adding context between the queries. As shown in Tab. 8, we train a BoQ model with ResNet-18 backbone for 35 epochs. Adding self-attention between the learnable queries brings consistent performance improvement across various benchmarks. Note that the self-attention’s output can be cached during eval/test to avoid recomputation at every test iteration.

	MSLS-val	Pitts30k	Pitts250k	Nordland
BoQ w/out self-att.	87.1	90.7	92.3	56.4
<b>BoQ with self-att.</b>	<b>88.4</b>	<b>91.5</b>	<b>93.3</b>	<b>65.9</b>

Table 8. Ablation on the use of self-attention between the global learned queries ( $Q^i$ ). Recall@1 is reported.

## 4.5. Visualization of the learned queries

Fig. 3 illustrates the cross-attention weights between the input image and a subset of learned queries within the BoQ blocks. We highlight four distinct queries (among 64 in total) from the second BoQ block, in order to understand their individual aggregation characteristics. Observing vertically, we can see how the input image is seen through the “lenses” of each learned query. The aggregation is realized through the product of the cross-attention weights with the input feature maps, resulting in a single aggregated descriptor per query. Notice that the role of each query is to “scan”—via cross-attention—the input features and generate the aggregation weights.

Observing horizontally, the diversity in attention patterns across the learned queries becomes apparent. For instance, the first query (top row) appears to concentrate on fine-grained details within the feature maps, as indicated by the intense localized areas. In contrast, the second query (sec-




Figure 3. Visualization of the cross attention weights between the input images and the learned queries. The three examples are from Nordland, Pitts30k and MSLS datasets, respectively. We selected four queries (among 64) from the second BoQ block of a trained network. Vertically, we can see how the input image is aggregated by each query. The aggregation is done through the product of the weight with the input feature maps, resulting in one aggregated descriptor per query. Horizontally, we can see in each line how each query spans the input image. For example, the first query looks more for fine grained details, while the second looks more for large areas in the input images.

ond row) shows a preference for larger regions in the input features, suggesting a more holistic capture of scene attributes.

## 5. Conclusion

In this work, we introduced *Bag-of-Queries (BoQ)*, a novel aggregation technique for visual place recognition, which is based on the use of learnable global queries to probe local features via a cross-attention mechanism, allowing for robust and discriminative feature aggregation. Our extensive experiments on 14 different large-scale benchmarks demonstrate that BoQ consistently outperforms current state-of-the-art techniques, particularly in handling complex variations in viewpoint, lighting, and seasonal changes. Further-

more, BoQ being a global (one-stage) retrieval technique, it outperforms existing two-stage retrieval methods that employ geometric verification for re-ranking, all while being orders of magnitude faster, setting a new standard for VPR research.

In future work, building upon the concepts discussed in Sec. 3, the spatial information carried in the output of the last encoder, denoted as  $\mathbf{X}^L$ , presents an opportunity for further enhancement through the application of special reranking strategies.

**Acknowledgement:** This work has been supported by The Fonds de Recherche du Québec Nature et technologies (FRQNT 254912). We gratefully acknowledge the support of NVIDIA Corporation with the donation of a Quadro RTX 8000 GPU used for our experiments.

## References

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *USENIX Symposium on Operating Systems Design and Implementation*, pages 265–283, 2016. 5
- [2] Amar Ali-bey, Brahim Chaib-draa, and Philippe Giguère. Global proxy-based hard mining for visual place recognition. In *British Machine Vision Conference (BMVC)*, 2022. 1, 2, 5
- [3] Amar Ali-bey, Brahim Chaib-draa, and Philippe Giguère. GSV-CITIES: Toward Appropriate Supervised Visual Place Recognition. *Neurocomputing*, 513:194–203, 2022. 1, 2, 4, 5, 6
- [4] Amar Ali-bey, Brahim Chaib-draa, and Philippe Giguère. MixVPR: Feature mixing for visual place recognition. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 2998–3007, 2023. 1, 2, 3, 4, 5, 6
- [5] Relja Arandjelovic and Andrew Zisserman. All about VLAD. In *IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1578–1585, 2013. 2
- [6] Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Pa-jdla, and Josef Sivic. NetVLAD: CNN architecture for weakly supervised place recognition. In *IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5297–5307, 2016. 1, 2, 4, 5, 6
- [7] Artem Babenko and Victor Lempitsky. Aggregating local deep features for image retrieval. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1269–1277, 2015. 2
- [8] Giovanni Barbarani, Mohamad Mostafa, Hajali Bayramov, Gabriele Trivigno, Gabriele Berton, Carlo Masone, and Barbara Caputo. Are local features all you need for cross-domain visual place recognition? In *IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6154–6164, 2023. 2
- [9] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. SURF: Speeded up robust features. In *European Conference on Computer Vision (ECCV)*, pages 404–417, 2006. 2



- [10] Gabriele Berton, Valerio Paolicelli, Carlo Masone, and Barbara Caputo. Adaptive-attentive geolocalization from few queries: A hybrid approach. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 2918–2927, 2021. 1, 4
- [11] Gabriele Berton, Carlo Masone, and Barbara Caputo. Rethinking visual geo-localization for large-scale applications. In *IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4878–4888, 2022. 1, 3, 5, 6
- [12] Gabriele Berton, Gabriele Trivigno, Barbara Caputo, and Carlo Masone. Eigenplaces: Training viewpoint robust models for visual place recognition. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11080–11090, 2023. 1, 2, 5, 6
- [13] Bingyi Cao, Andre Araujo, and Jack Sim. Unifying deep local and global features for image search. In *European Conference on Computer Vision (ECCV)*, pages 726–743, 2020. 5, 7
- [14] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European Conference on Computer Vision (ECCV)*, pages 213–229, 2020. 2, 4
- [15] Chun-Fu Richard Chen, Quanfu Fan, and Rameswar Panda. CrossViT: Cross-attention multi-scale vision transformer for image classification. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 357–366, 2021. 2
- [16] Mark Cummins. Highly scalable appearance-only SLAM-FAB-MAP 2.0. *Robotics: Science and Systems (RSS)*, 2009. 4
- [17] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In *IEEE International Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 224–236, 2018. 5, 7
- [18] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *International Conference on Learning Representations (ICLR)*, 2020. 2
- [19] Dorian Gálvez-López and Juan D Tardos. Bags of binary words for fast place recognition in image sequences. *IEEE Transactions on Robotics*, 28(5):1188–1197, 2012. 2
- [20] Yixiao Ge, Haibo Wang, Feng Zhu, Rui Zhao, and Hongsheng Li. Self-supervising fine-grained region similarities for large-scale image localization. In *European Conference on Computer Vision (ECCV)*, pages 369–386, 2020. 1, 6
- [21] Amin Ghiasi, Hamid Kazemi, Eitan Borgnia, Steven Reich, Manli Shu, Micah Goldblum, Andrew Gordon Wilson, and Tom Goldstein. What do vision transformers learn? a visual exploration. *arXiv preprint arXiv:2212.06727*, 2022. 2
- [22] Stephen Hausler, Sourav Garg, Ming Xu, Michael Milford, and Tobias Fischer. Patch-NetVLAD: Multi-scale fusion of locally-global descriptors for place recognition. In *IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14141–14152, 2021. 5, 7
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4
- [24] Qingfeng Hou, Jun Lu, Haitao Guo, Xiangyun Liu, Zhihui Gong, Kun Zhu, and Yifan Ping. Feature relation guided cross-view image based geo-localization. *Remote Sensing*, 15(20):5029, 2023. 2
- [25] Gaoshuang Huang, Yang Zhou, Xiaofei Hu, Chenglong Zhang, Luying Zhao, Wenjian Gan, and Mingbo Hou. Dino-mix: Enhancing visual place recognition with foundational vision model and feature mixing. *arXiv preprint arXiv:2311.00230*, 2023. 2
- [26] 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
- [27] Nikhil Keetha, Avneesh Mishra, Jay Karhade, Krishna Murthy Jatavallabhula, Sebastian Scherer, Madhava Krishna, and Sourav Garg. Anyloc: Towards universal visual place recognition. *arXiv preprint arXiv:2308.00688*, 2023. 2
- [28] Hyo Jin Kim, Enrique Dunn, and Jan-Michael Frahm. Learned contextual feature reweighting for image geo-localization. In *IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3251–3260, 2017. 1, 2, 5
- [29] María Leyva-Vallina, Nicola Strisciuglio, and Nicolai Petkov. Data-efficient large scale place recognition with graded similarity supervision. In *IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 23487–23496, 2023. 2
- [30] Liu Liu, Hongdong Li, and Yuchao Dai. Stochastic attraction-repulsion embedding for large scale image localization. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2570–2579, 2019. 1
- [31] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10012–10022, 2021. 2
- [32] David G Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision (IJCV)*, 60(2):91–110, 2004. 2
- [33] Carlo Masone and Barbara Caputo. A survey on deep visual place recognition. *IEEE Access*, 9:19516–19547, 2021. 1
- [34] Michael Milford and G. Wyeth. Mapping a suburb with a single camera using a biologically inspired slam system. *IEEE Transactions on Robotics*, 24:1038–1053, 2008. 4
- [35] Jiwei Nie, Joe-Mei Feng, Dingyu Xue, Feng Pan, Wei Liu, Jun Hu, and Shuai Cheng. A training-free, lightweight global image descriptor for long-term visual place recognition toward autonomous vehicles. *IEEE Transactions on Intelligent Transportation Systems*, 2023. 2
- [36] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafranec, Vasil Khalidov, Pierre Fernandez,

- Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. [2](#)
- [37] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 8026–8037, 2019. [5](#)
- [38] James Philbin, Ondrej Chum, Michael Isard, Josef Sivic, and Andrew Zisserman. Object retrieval with large vocabularies and fast spatial matching. In *IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, 2007. [2](#)
- [39] 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](#), [5](#), [6](#)
- [40] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. ORB: An efficient alternative to sift or surf. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2564–2571, 2011. [2](#)
- [41] Zachary Seymour, Karan Sikka, Han-Pang Chiu, Supun Samarasekera, and Rakesh Kumar. Semantically-aware attentive neural embeddings for long-term 2D visual localization. In *British Machine Vision Conference (BMVC)*, 2019. [1](#)
- [42] Niko Sünderhauf, Peer Neubert, and Peter Protzel. Are we there yet? challenging SeqSLAM on a 3000 km journey across all four seasons. In *Workshop on Long-Term Autonomy, IEEE International Conference on Robotics and Automation (ICRA)*, 2013. [4](#)
- [43] Giorgos Tolias, Ronan Sicre, and Hervé Jégou. Particular object retrieval with integral max-pooling of CNN activations. *arXiv preprint arXiv:1511.05879*, 2015. [2](#)
- [44] Akihiko Torii, Josef Sivic, Tomas Pajdla, and Masatoshi Okutomi. Visual place recognition with repetitive structures. In *IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 883–890, 2013. [1](#), [2](#), [4](#)
- [45] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems (NeurIPS)*, 30, 2017. [2](#), [3](#)
- [46] Ruotong Wang, Yanqing Shen, Weiliang Zuo, Sanping Zhou, and Nanning Zheng. TransVPR: Transformer-based place recognition with multi-level attention aggregation. In *IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13648–13657, 2022. [2](#), [5](#), [7](#)
- [47] Xun Wang, Xintong Han, Weilin Huang, Dengke Dong, and Matthew R Scott. Multi-similarity loss with general pair weighting for deep metric learning. In *IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5022–5030, 2019. [5](#)
- [48] Yuwei Wang, Yuanying Qiu, Peitao Cheng, and Junyu Zhang. Hybrid CNN-transformer features for visual place recognition. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(3):1109–1122, 2023. [2](#)
- [49] Yuwei Wang, Yuanying Qiu, Peitao Cheng, and Junyu Zhang. Transformer-based descriptors with fine-grained region supervisions for visual place recognition. *Knowledge-Based Systems*, 280:110993, 2023. [2](#)
- [50] Frederik Warburg, Soren Hauberg, Manuel López-Antequera, Pau Gargallo, Yubin Kuang, and Javier Civera. Mapillary street-level sequences: A dataset for lifelong place recognition. In *IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2626–2635, 2020. [1](#), [4](#), [5](#)
- [51] Burak Yildiz, Seyran Khademi, Ronald Maria Siebes, and Jan van Gemert. Amstertime: A visual place recognition benchmark dataset for severe domain shift. *arXiv preprint arXiv:2203.16291*, 2022. [1](#), [4](#)
- [52] Jun Yu, Chaoyang Zhu, Jian Zhang, Qingming Huang, and Dacheng Tao. Spatial pyramid-enhanced NetVLAD with weighted triplet loss for place recognition. *IEEE Transactions on Neural Networks and Learning Systems*, 31(2):661–674, 2019. [2](#), [5](#), [6](#)
- [53] Mubariz Zaffar, Sourav Garg, Michael Milford, Julian Kooij, David Flynn, Klaus McDonald-Maier, and Shoab Ehsan. VPR-Bench: An open-source visual place recognition evaluation framework with quantifiable viewpoint and appearance change. *International Journal of Computer Vision (IJCV)*, pages 1–39, 2021. [1](#), [4](#), [5](#)
- [54] Hao Zhang, Xin Chen, Heming Jing, Yingbin Zheng, Yuan Wu, and Cheng Jin. ETR: An efficient transformer for re-ranking in visual place recognition. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 5665–5674, 2023. [2](#)
- [55] Jian Zhang, Yunyin Cao, and Qun Wu. Vector of locally and adaptively aggregated descriptors for image feature representation. *Pattern Recognition*, 116:107952, 2021. [2](#), [5](#), [6](#)
- [56] Qieshi Zhang, Zhenyu Xu, Yuhang Kang, Fusheng Hao, Ziliang Ren, and Jun Cheng. Distilled representation using patch-based local-to-global similarity strategy for visual place recognition. *Knowledge-Based Systems*, 280:111015, 2023. [2](#)
- [57] Xiwu Zhang, Lei Wang, and Yan Su. Visual place recognition: A survey from deep learning perspective. *Pattern Recognition*, 113:107760, 2021. [1](#)
- [58] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6881–6890, 2021. [2](#), [5](#)
- [59] Sijie Zhu, Linjie Yang, Chen Chen, Mubarak Shah, Xiaohui Shen, and Heng Wang. R2Former: Unified retrieval and reranking transformer for place recognition. In *IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19370–19380, 2023. [2](#), [5](#), [7](#)