

Towards Test-time Efficient Visual Place Recognition via Asymmetric Query Processing

Jaeyoon Kim*, Yoonki Cho*, Sung-Eui Yoon

Korea Advanced Institute of Science and Technology (KAIST)

Abstract

Visual Place Recognition (VPR) has advanced significantly with high-capacity foundation models like DINOv2, achieving remarkable performance. Nonetheless, their substantial computational cost makes deployment on resource-constrained devices impractical. In this paper, we introduce an efficient asymmetric VPR framework that incorporates a high-capacity gallery model for offline feature extraction with a lightweight query network for online processing. A key challenge in this setting is ensuring compatibility between these heterogeneous networks, which conventional approaches address through computationally expensive k-NN-based compatible training. To overcome this, we propose a geographical memory bank that structures gallery features using geolocation metadata inherent in VPR databases, eliminating the need for exhaustive k-NN computations. Additionally, we introduce an implicit embedding augmentation technique that enhances the query network to model feature variations despite its limited capacity. Extensive experiments demonstrate that our method not only significantly reduces computational costs but also outperforms existing asymmetric retrieval techniques, establishing a new aspect for VPR in resource-limited environments.

Introduction

Visual Place Recognition (VPR) is a fundamental task in computer vision, playing a crucial role in applications such as navigation (Mirowski et al. 2018; Bresson et al. 2017), augmented reality (Garg, Fischer, and Milford 2021), SLAM (Cadena et al. 2016), and geo-localization systems (Lowry et al. 2015; Berton, Masone, and Caputo 2022). It is typically formulated as an image retrieval problem, where a model determines the location of a query image by identifying the most similar geo-tagged reference images in a database (*i.e.*, gallery). Recent advances in VPR have been driven by foundation models such as DINOv2 (Oquab et al. 2024), which have demonstrated remarkable capabilities for rich visual representations (Lu et al. 2024b,a; Qiu et al. 2024; Lu et al. 2024c). Nonetheless, these impressive performance gains come at a substantial computational

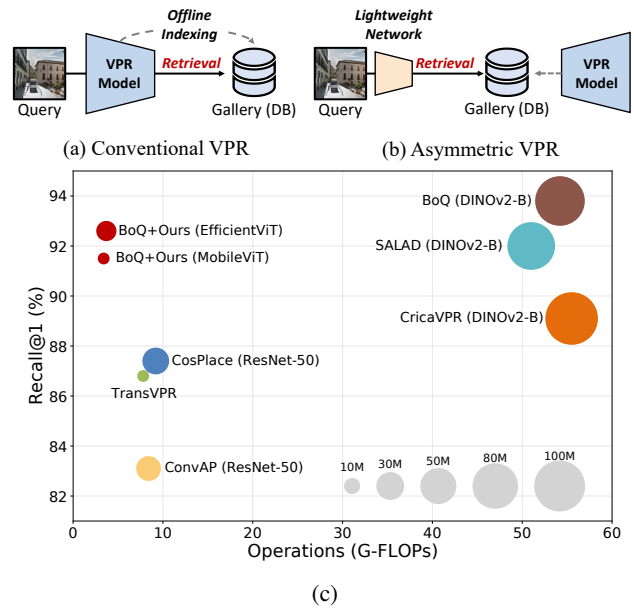


Figure 1: **Comparison of conventional and asymmetric VPR approaches.** (a) Conventional VPR uses the same model for both query and gallery processing. (b) Our asymmetric approach uses a lightweight network for query processing while maintaining a high-capacity model for gallery indexing. (c) Comparison of various VPR methods on Recall@1 vs. computational cost (G-FLOPs) for online query processing on the MSLS validation dataset. Circle sizes indicate the number of model parameters. Our method achieves competitive accuracy with significantly less computation, suitable for resource-constrained devices.

cost, posing a significant challenge for resource-limited platforms – deploying such computationally intensive models on resource-constrained devices remains impractical.

To remedy this issue, we explore *asymmetric* retrieval techniques for VPR as a promising and practical solution. Unlike conventional approaches where the same network processes both query and gallery images, asymmetric retrieval employs separate models with different capacities for each task (Budnik and Avrithis 2021; Duggal et al. 2021). While asymmetric retrieval has been explored in

*These authors contributed equally.

other domains (Budnik and Avrithis 2021; Du et al. 2024; Shoshan et al. 2024; Suma and Tolia 2023), it remains largely unexplored in VPR, where foundation model-based approaches (Izquierdo and Civera 2024; Tzachor et al. 2025) are now mainstream but operations on edge devices (*e.g.*, robots) are also required. Fig. 1 illustrates the difference between conventional VPR and our asymmetric framework, showing competitive accuracy with significantly reduced computational requirements. By leveraging a high-capacity model to precompute and index rich gallery features offline, this approach enables a lightweight query network to efficiently retrieve relevant results during online query processing on resource-constrained environments.

A key challenge in this asymmetric setting is that the query network must produce embeddings that are *compatible* with the gallery’s representation space, despite differences in model capacity and architecture. Due to the limited capacity of lightweight query models, directly regressing gallery embeddings is insufficient for achieving feature compatibility. Recent methods address this limitation by leveraging contextual cues, such as similarity distributions or ranking orders among k -nearest neighbors (k -NN) (Wu et al. 2022; Xie et al. 2024; Wu et al. 2023) relationships. However, this approach introduces substantial overhead – precomputing and storing k -NN information for large-scale gallery embeddings incurs significant computational costs that scale with the size of the database. In VPR scenarios involving millions of images from diverse locations, these k -NN computations become particularly prohibitive.

In this paper, we introduce an efficient asymmetric framework for VPR that addresses these challenges by leveraging the geospatial structure inherent in visual place recognition. Unlike prior approaches that rely on k -NN-based contextual information, we propose a *geographical memory bank* that organizes gallery features based on their geographic locations, constructing location-specific statistical representations by aggregating embeddings from images captured within the same vicinity. For each distinct location, our method computes and stores statistical moments of high-capacity feature distributions, creating a compact yet informative representation of place-specific visual characteristics. During training, these geographic feature statistics serve as supervisory signals in our asymmetric contrastive learning, guiding the lightweight query network to produce compatible embeddings without requiring exhaustive k -NN searches. Additionally, we develop an *implicit embedding augmentation* technique that enhances the query network’s ability to model feature variations by incorporating location-specific covariance information, enabling effective alignment with rich gallery representations despite limited capacity. Extensive experiments on standard VPR benchmarks show that our method outperforms state-of-the-art asymmetric retrieval approaches while significantly reducing computational costs for online query processing.

Related Work

Visual Place Recognition (VPR) has evolved through significant advancements in feature aggregation techniques. Early approaches rely on aggregating hand-crafted features

(SIFT (Lowe 2004), SURF (Bay et al. 2008)) using methods like Bag of Words (Angeli et al. 2008) and VLAD (Jégou et al. 2010). The emergence of deep learning introduces CNN-based approaches, with NetVLAD (Arandjelovic et al. 2016) pioneering differentiable aggregation for end-to-end training, followed by its variants (Zhang, Cao, and Wu 2021; Yu et al. 2019; Jin Kim, Dunn, and Frahm 2017). Further developments include GeM (Radenović, Tolia, and Chum 2018) pooling in CosPlace (Berton, Masone, and Caputo 2022) and EigenPlaces (Berton et al. 2023), and MLP-based aggregation in MixVPR (Ali-Bey, Chaib-Draa, and Giguere 2023), all maintaining reasonable computational efficiency.

Recently, foundation models like DINOv2 (Oquab et al. 2024), pre-trained with self-supervised learning on diverse visual data, has demonstrated remarkable capabilities with rich visual representations for VPR. AnyLoc (Keetha et al. 2023) pioneers the use of DINOv2 features combined with traditional aggregation techniques like VLAD. Following this paradigm shift, several studies (Lu et al. 2024b,a; Tzachor et al. 2025; Lu et al. 2024c; Qiu et al. 2024) have explored different strategies for employing foundation models’ power for VPR. SALAD (Izquierdo and Civera 2024) fine-tunes the pre-trained DINOv2 model specifically for VPR, introducing an optimal transport-based aggregation. BoQ (Ali-bey, Chaib-draa, and Giguère 2024) introduces a novel transformer-based aggregation technique that learns a set of global queries to probe local features via cross-attention. However, the high computational cost of foundation models poses deployment challenges, as their substantial inference overhead limits practical use on resource-constrained devices. Our work relaxes these limitations by developing an efficient asymmetric framework that balances performance with computational practicality.

Compatible Training aims to encode feature embeddings that are interoperable with pre-existing models, facilitating seamless integration across systems for efficient updates (Shen et al. 2020; Wang et al. 2020; Cui et al. 2024; Zhou et al. 2022). This approach significantly reduces the computational burden associated with re-indexing (*i.e.*, backfilling) process when transitioning from older models to newer ones (Su et al. 2022; Zhang et al. 2022a; Pan et al. 2023; Bai et al. 2022; Shen et al. 2020; Ramanujan et al. 2022; Meng et al. 2021). Recent studies (Zhang et al. 2022b; Wan et al. 2022) have explored compatible training in more challenging open-set scenarios, where models adapt to dynamic data distributions while maintaining interoperability.

Building upon this, asymmetric retrieval has emerged as a specialized paradigm within compatible training, deploying a lightweight model on the query side and a more powerful model on the gallery side to balance computational efficiency with retrieval accuracy (Budnik and Avrithis 2021; Duggal et al. 2021; Du et al. 2024). Budnik and Avrithis (2021) first introduced an asymmetric metric learning (AML) framework and demonstrated that a simple regression loss can distill the embedding spaces. Recent methods have advanced compatibility between query and gallery models by leveraging k -nearest neighbor (k -NN) information in various ways. Specifically, CSD (Wu et al. 2022) transfers contextual similarity knowledge through k -NN re-

relationships, while D3still (Xie et al. 2024) employs decoupled differential distillation based on pairwise similarity across k-NN embeddings. MSP preserves similarity ranking consistency via monotonic mapping, and its companion approach ROP utilizes smooth Heaviside approximations to constrain sorting results (Wu et al. 2023). In contrast, our method leverages an efficient geographical memory bank with asymmetric contrastive learning, eliminating the need for exhaustive k-NN computations.

Method

The overview of the proposed *asymmetric* VPR framework is illustrated in Fig. 2. It employs a high-capacity network for gallery feature extraction (performed offline) while utilizing a significantly lighter network for query feature extraction (executed at inference time), enabling efficient retrieval under resource-constrained environments. In the following subsection, we formally define the problem setting and the constraints imposed by this asymmetric framework.

Problem Formulation

Let $f_g : \mathcal{X} \rightarrow \mathbb{R}^d$ be a pre-trained, high-capacity gallery network that has generated a set of features $\mathcal{F}_g = \{\mathbf{g}_i\}_{i=1}^N$ extracted from gallery images $\mathcal{X}_g = \{x_i\}_{i=1}^N$, where $\mathbf{g}_i = f_g(x_i)$, indexed for retrieval. The query network $f_q : \mathcal{X} \rightarrow \mathbb{R}^d$ is a lightweight model designed for efficient inference on resource-constrained devices.

A key challenge in the asymmetric setting—where a lightweight query network must interoperate with a high-capacity gallery network—is that f_q must generate embeddings compatible with \mathcal{F}_g , despite differences in model capacity and architecture. Training to achieve this compatibility is inherently constrained to using *only the gallery dataset \mathcal{X}_g and their precomputed features \mathcal{F}_g* , without modifying the gallery network or altering the stored embeddings—a practical requirement for real-world deployment where gallery features are precomputed at scale.

Given a query image x , we denote its query embedding as $\mathbf{q} = f_q(x)$. Retrieval is performed by ranking the similarity between \mathbf{q} and the gallery embeddings in \mathcal{F}_g . To enable effective retrieval, the following ranking constraint generally needs to be satisfied (Duggal et al. 2021):

$$\text{sim}(\mathbf{q}, \mathbf{g}^+) \gg \text{sim}(\mathbf{q}, \mathbf{g}^-), \quad (1)$$

where \mathbf{g}^+ represents the gallery embedding of a relevant image (corresponding to the same place as the query), \mathbf{g}^- represents the gallery embedding of an irrelevant image (different place), and $\text{sim}(\cdot, \cdot)$ is a similarity function, typically dot product between L2-normalized embeddings. Satisfying this constraint is challenging since f_q and f_g have different architectures and capacities, potentially leading to disparate embedding spaces.

Asymmetric Contrastive Learning

To improve feature compatibility, recent methods have commonly leveraged contextual information from k-nearest neighbor embeddings (Wu et al. 2022, 2023; Xie et al. 2024). However, performing k-NN searches and storing this

information for all images in a large-scale database incurs a non-negligible computational cost. To address this, we propose an efficient asymmetric contrastive learning approach that exploits the geographic information naturally available in VPR tasks. Our approach encourages the query network to produce embeddings that align with those of the gallery network while remaining distinct from features of different places, leveraging global geographic structures rather than local nearest neighbors.

Unlike other retrieval tasks, VPR inherently includes geolocation metadata, which our method leverages without relying on exhaustive nearest neighbor computations. Specifically, we define our geographical memory bank as a set of M representative centroids:

$$\mathcal{M} = \{\mathbf{c}_j\}_{j=1}^M \subset \mathbb{R}^d, \quad (2)$$

where each centroid \mathbf{c}_j represents the average of gallery features from \mathcal{F}_g that belong to the same geographic location. These centroids can be effectively precomputed using GPS coordinates, providing an efficient representation of the feature distribution across different locations.

Given this memory bank, we formulate our asymmetric learning objective as a contrastive learning problem, using geographic centroids as negative samples. For a given gallery image x with its gallery network embedding \mathbf{g} and query network embedding \mathbf{q} , we compute our asymmetric contrastive learning loss by:

$$\mathcal{L}_{\text{asym}} = -\log \frac{e^{\mathbf{q} \cdot \mathbf{g} / \tau}}{e^{\mathbf{q} \cdot \mathbf{g} / \tau} + \sum_{j \in \mathcal{N}(x)} e^{\mathbf{q} \cdot \mathbf{c}_j / \tau}}, \quad (3)$$

where τ is a temperature parameter that controls the sharpness of the distribution, $\mathcal{N}(x)$ denotes the set of indices for memory bank entries that do not correspond to the same location as image x , and \cdot denotes the dot product.

By leveraging geographical information inherent in VPR datasets, our approach eliminates the need for exhaustive k-NN computations, significantly reducing computational overhead while scaling efficiently with database size. Additionally, our geographically informed memory bank captures global feature relationships, leveraging priors tailored for VPR to enhance feature compatibility.

Implicit Embedding Augmentation

While our asymmetric contrastive learning framework considers geographical feature relationships, it can still face challenges caused by the capacity gap between the query network and the gallery network. Specifically, the lightweight query model may struggle to fully capture the complex distribution and variability of features within individual locations that the high-capacity gallery model can encode. To address this, we introduce an implicit embedding augmentation technique that exploits location-specific feature variations, enhancing the compatibility to approximate the rich representations of the gallery network.

Embedding augmentation is typically modeled using multivariate normal distribution (Han et al. 2023; Wang et al. 2019; Li et al. 2021), which we adopt in our formulation:

$$\tilde{\mathbf{g}} \sim \mathcal{N}(\mathbf{g}, \gamma \Sigma), \quad (4)$$

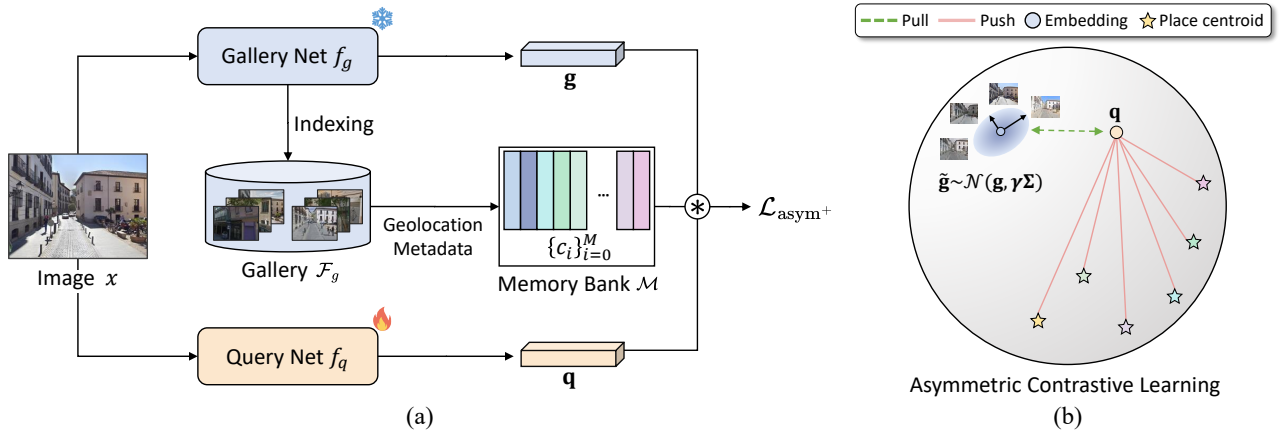


Figure 2: **Overview of the proposed asymmetric VPR framework.** (a) A high-capacity gallery network f_g extracts and stores embeddings offline, while a lightweight query network f_q is trained to generate compatible embeddings for efficient online retrieval. The gallery network remains fixed during training while only the query network is optimized. The geographical memory bank \mathcal{M} structures gallery features \mathcal{F}_g using geolocation metadata, storing place centroids $\{c_i\}_{i=0}^M$ to efficiently represent feature distributions across locations for compatibility training. (b) Illustration of asymmetric contrastive learning with implicit embedding augmentation. The query embedding \mathbf{q} is trained to align with its corresponding gallery embedding \mathbf{g} while being pushed away from negative centroids. To enhance feature compatibility, augmented gallery embeddings $\tilde{\mathbf{g}} \sim \mathcal{N}(\mathbf{g}, \gamma \Sigma)$ incorporate location-specific covariance, improving the query network to model feature variations despite its limited capacity.

where $\Sigma \in \mathbb{R}^{d \times d}$ represents the covariance matrix of features in location of \mathbf{g} , and γ is a scaling hyperparameter.

To incorporate this augmentation into our training framework, we modify our contrastive learning loss (Eq. 3) by replacing the single gallery embedding with multiple augmented samples. For each gallery image, we can generate K augmented embeddings $\{\tilde{\mathbf{g}}^1, \tilde{\mathbf{g}}^2, \dots, \tilde{\mathbf{g}}^K\}$ from the defined multivariate normal distribution. We then compute the explicit augmentation loss for the augmented samples:

$$\mathcal{L}_{\text{asym}^+}^{\text{exp}} = -\frac{1}{K} \sum_{k=1}^K \log \frac{e^{\mathbf{q} \cdot \tilde{\mathbf{g}}^k / \tau}}{e^{\mathbf{q} \cdot \tilde{\mathbf{g}}^k / \tau} + \sum_{j \in \mathcal{N}(x)} e^{\mathbf{q} \cdot \mathbf{c}_j / \tau}}. \quad (5)$$

Proposition 1. *The explicit augmentation loss can be transformed into an implicit form as $K \rightarrow \infty$.*

Proof. When $K \rightarrow \infty$, Eq. 5 becomes an expectation:

$$\bar{\mathcal{L}}_{\text{asym}^+} = \mathbb{E}_{\tilde{\mathbf{g}}} \left[\log \left(1 + \sum_{j \in \mathcal{N}(x)} e^{\mathbf{q} \cdot \mathbf{c}_j / \tau - \mathbf{q} \cdot \tilde{\mathbf{g}} / \tau} \right) \right]. \quad (6)$$

Since the logarithm function is concave, we can apply Jensen's inequality:

$$\begin{aligned} \bar{\mathcal{L}}_{\text{asym}^+} &\leq \log \left(1 + \sum_{j \in \mathcal{N}(x)} \mathbb{E}_{\tilde{\mathbf{g}}} \left[e^{\mathbf{q} \cdot \mathbf{c}_j / \tau - \mathbf{q} \cdot \tilde{\mathbf{g}} / \tau} \right] \right) \\ &= \log \left(1 + \sum_{j \in \mathcal{N}(x)} e^{\mathbf{q} \cdot \mathbf{c}_j / \tau - \mathbf{q} \cdot \mathbf{g} / \tau + \frac{\gamma}{2\tau^2} \mathbf{q}^T \Sigma \mathbf{q}} \right), \end{aligned} \quad (7)$$

where moment-generating function of multivariate Gaussian is used as: $\mathbb{E}[e^{-\mathbf{q} \cdot \tilde{\mathbf{g}} / \tau}] = e^{-\mathbf{q} \cdot \mathbf{g} / \tau + \frac{\gamma}{2\tau^2} \mathbf{q}^T \Sigma \mathbf{q}}$. \square

For simplicity, our final implicit embedding augmentation loss can be defined by using the upper bound in Eq. 7:

$$\mathcal{L}_{\text{asym}^+} = -\log \frac{e^{\mathbf{q} \cdot \mathbf{g} / \tau}}{e^{\mathbf{q} \cdot \mathbf{g} / \tau} + \sum_{j \in \mathcal{N}(x)} e^{\mathbf{q} \cdot \mathbf{c}_j / \tau + (\gamma/2\tau^2) \mathbf{q}^T \Sigma \mathbf{q}}}. \quad (8)$$

The formulation introduces a variance-based regularization term $(\gamma/2\tau^2) \mathbf{q}^T \Sigma \mathbf{q}$ for the negative terms, incorporating covariance information without requiring explicit sampling.

Regularization effect of covariance guidance. Using the eigendecomposition, the quadratic form becomes the following.

$$\begin{aligned} \mathbf{q}^T \Sigma \mathbf{q} &= \mathbf{q}^T \mathbf{V} \Lambda \mathbf{V}^T \mathbf{q} \\ &= (\mathbf{V}^T \mathbf{q})^T \Lambda (\mathbf{V}^T \mathbf{q}) \\ &= \sum_j \lambda_j (v_j^T \mathbf{q})^2, \end{aligned} \quad (9)$$

where λ_j and v_j are the j -th eigenvalue and eigenvector, respectively. To minimize Eq. 8, the regularization term must also be minimized during the optimization process, since it is attached to the negative terms. It results in an optimization behavior where the query model takes stronger penalties along directions with larger eigenvalues (*i.e.*, higher variance) in the gallery's feature space.

These high-variance directions capture less discriminative intra-place variations, such as changing viewpoints or lighting. By reducing sensitivity to these, our approach guides the lightweight query model to prioritize stable, place-distinctive features, enabling efficient asymmetric retrieval despite its limited capacity.

Experiments

Experimental Setup

Training and Evaluation Protocols. For training, we utilize GSV-Cities (Ali-bey, Chaib-draa, and Giguère 2022) dataset, a large-scale collection of 560k images from 67k distinct locations spanning a 14-year period, sourced from Google Street View. To construct the gallery embeddings, we preprocess the entire database by extracting features from all images using a gallery network. These precomputed gallery features are stored offline and remain fixed throughout training, aligning with the practical constraints of real-world deployment. We then utilize the geolocation metadata (*i.e.*, GPS coordinates) to compute place centroids by averaging the gallery embeddings from the same location. Simultaneously, we compute the covariance matrices for each place, which are used for implicit embedding augmentation.

For evaluation, we use five standard VPR benchmarks: Pitts250k (Torii et al. 2013), MSLS Validation (Warburg et al. 2020), Tokyo24/7 (Torii et al. 2015), Nordland (Zafar et al. 2021), and, AmsterTime (Yildiz et al. 2022) which cover diverse challenges such as viewpoint shifts, seasonal variations, and illumination changes. Following prior VPR literature (Lu et al. 2024b,a; Izquierdo and Civera 2024), we evaluate performance using the Recall@k (R@k) as an evaluation metric, which measures the percentage of queries with at least one correct match among the top-K retrieved candidates. Positive matches are determined using a 25-meter threshold based on geographic coordinates, except for Nordland, where a match is considered correct if retrieved within three frames of the ground truth position, and AmsterTime which utilizes its provided ground truth pairs.

Network Architectures. For the gallery models, we use SALAD (Izquierdo and Civera 2024) and BoQ (Ali-bey, Chaib-draa, and Giguère 2024), both trained on GSV-Cities and built on DINOv2-B, providing strong visual representations with embedding dimensions of 8448 and 12288, respectively. These high-capacity models impose substantial computational demands, making them a practical choice for our asymmetric setting, where gallery features are precomputed offline. For the query models, we use two popular lightweight backbones: MobileViTv2 (Mehta and Rastegari 2022) and EfficientViT-B2 (Cai et al. 2023), designed for mobile and resource-constrained devices. To extract features with these lightweight backbones, we maintain architectural consistency by employing the same type of aggregator on both gallery and query sides in a symmetric manner. For instance, when SALAD serves as the gallery model, its aggregation architecture is also used in the query model.

Implementation Details. We train our models using the AdamW (Loshchilov and Hutter 2017) optimizer with an initial learning rate of 5×10^{-4} and cosine learning rate decay, reaching a minimum learning rate of 1×10^{-4} . Training is performed for 15 epochs with a batch size of 64. All images are resized to 322×322 for both training and evaluation. We empirically set the temperature parameter $\tau = 0.05$ for asymmetric contrastive learning, and the scaling parameter $\gamma = 15$ for implicit embedding augmentation. Utilizing the full covariance matrix during training incurs GPU mem-

ory overhead due to the high dimensionality of the feature space. To mitigate this, we approximate it using only the diagonal elements of the covariance matrix, significantly reducing computational costs while preserving essential statistical information. We implement our framework in PyTorch (Paszke et al. 2019) and utilize a single RTX-4090 GPU for all experiments.

Comparison with State-of-the-Arts

We evaluate our method against state-of-the-art asymmetric retrieval approaches, including CSD (Wu et al. 2022), ROP (Wu et al. 2023), MSP (Wu et al. 2023), and D3still (Xie et al. 2024), across standard VPR benchmarks. The results are summarized in Table 1. The table is structured into two primary sections, each dedicated to one of the aggregators: SALAD (Izquierdo and Civera 2024) and BoQ (Ali-bey, Chaib-draa, and Giguère 2024). In each section, the upper part details the symmetric retrieval performance. This includes the result of the high-capacity model itself (e.g., BoQ with DINOv2-B), which can be regarded as the performance upper bound. Additionally, it provides baseline results for lightweight models (MobileViTv2 (Mehta and Rastegari 2022) and EfficientViT-B2 (Cai et al. 2023)) when used symmetrically for both query and gallery processing. The remainder of each section presents the core asymmetric retrieval results. In this configuration, the high-capacity gallery model is fixed, while different lightweight query networks are trained for compatibility.

As shown in the table, our method achieves superior performance (R@1) against existing asymmetric retrieval approaches across all datasets and query-gallery model configurations. Specifically, when using BoQ as the gallery model and EfficientViT-B2 as the query network, it achieves strong retrieval performance, reaching 95.4% R@1 on Pitts250k, 92.3% on MSLS Val, and 92.7% on Tokyo24/7. Additionally, our method demonstrates consistent improvements on Nordland and AmsterTime, achieving 74.6% and 48.6%, respectively, both of which present diverse environmental variations inherently challenging for VPR. These results highlight the effectiveness of leveraging the geographical prior knowledge inherent in VPR, bridging the capacity gap and improving compatibility between lightweight query networks and high-capacity gallery models. By incorporating this structure through the geographical memory bank with implicit embedding augmentation, our method enhances feature compatibility while mitigating the computational overhead of k-NN-based techniques.

Analysis and Ablation Study

We conduct extensive analysis and ablation studies to validate the effectiveness of our method and design choices. Unless otherwise specified, all experiments use BoQ as a gallery model and EfficientViT-B2 as a query model.

Computational Efficiency Analysis. We analyze the computational efficiency for online query processing of different backbone and aggregator configurations in our asymmetric VPR framework, with the results summarized in Table 2. Compared to high-capacity gallery networks, lightweight

Method	Query Network	Pitts250k		MSLS Val		Tokyo24/7		Nordland		AmsterTime	
		R@1	R@5	R@1	R@5	R@1	R@5	R@1	R@5	R@1	R@5
<i>Training with SALAD of DINOv2 as gallery model</i>											
SALAD [†] (Izquierdo and Civera 2024)	DINOv2-B	95.1	98.6	92.0	96.4	94.6	97.5	76.0	89.2	58.5	78.9
MobileViTv2-SALAD [†]	MobileViTv2	87.9	95.1	76.6	85.0	70.5	83.5	18.3	28.8	28.5	45.6
EfficientViT-B2-SALAD [†]	EfficientViT-B2	92.1	97.1	85.3	91.4	83.2	92.1	26.5	39.1	38.5	59.5
CSD (Wu et al. 2022)		88.1	95.5	87.3	94.7	75.2	87.3	38.8	56.2	34.9	54.8
ROP (Wu et al. 2023)		86.5	94.5	87.3	94.3	69.2	81.6	32.9	50.4	34.8	56.0
MSP (Wu et al. 2023)	MobileViTv2	88.4	95.6	88.5	94.5	76.2	87.9	41.0	57.0	34.8	55.5
D3still (Xie et al. 2024)		86.5	94.0	87.8	94.6	71.4	84.1	41.3	57.5	33.3	53.0
Ours		89.0	95.2	89.3	95.0	76.5	87.6	49.3	68.0	37.6	56.1
CSD (Wu et al. 2022)		91.2	97.0	88.8	95.0	82.2	91.7	53.6	71.6	37.4	57.3
ROP (Wu et al. 2023)		89.1	96.1	87.8	94.6	72.1	86.0	44.5	61.7	35.1	55.6
MSP (Wu et al. 2023)	EfficientViT-B2	91.6	96.7	88.2	95.4	85.7	94.2	55.5	73.2	41.8	62.2
D3still (Xie et al. 2024)		91.9	97.0	89.7	95.7	85.4	94.3	58.2	75.8	40.9	61.3
Ours		92.4	97.5	91.2	95.8	86.7	94.0	61.7	77.6	43.0	64.5
<i>Training with BoQ of DINOv2 as gallery model</i>											
BoQ [†] (Ali-bey, Chaib-draa, and Giguère 2024)	DINOv2-B	96.6	99.1	93.8	96.8	96.5	97.8	81.3	92.5	63.0	81.6
MobileViTv2-BoQ [†]	MobileViTv2	92.8	98.0	85.0	92.3	85.1	93.3	49.4	65.5	38.7	58.0
EfficientViT-B2-BoQ [†]	EfficientViT-B2	94.3	98.3	87.7	93.2	85.7	92.7	51.6	67.3	41.6	63.8
CSD (Wu et al. 2022)		92.5	97.5	90.9	95.7	85.7	93.7	62.6	79.5	42.6	64.2
ROP (Wu et al. 2023)		90.7	96.8	89.6	95.4	81.3	89.5	56.9	75.8	39.6	62.9
MSP (Wu et al. 2023)	MobileViTv2	93.0	97.6	90.0	95.4	87.0	93.3	64.9	81.1	41.7	63.0
D3still (Xie et al. 2024)		93.2	97.5	90.4	95.5	87.0	93.7	66.5	82.2	41.7	64.3
Ours		93.4	97.8	91.1	96.1	89.2	94.3	67.8	82.8	43.8	67.0
CSD (Wu et al. 2022)		94.6	98.4	91.4	96.5	90.5	94.9	68.7	82.4	48.1	69.6
ROP (Wu et al. 2023)		91.6	97.1	90.5	96.0	84.1	91.4	60.3	77.1	37.4	58.6
MSP (Wu et al. 2023)	EfficientViT-B2	94.3	98.1	91.4	95.8	90.5	94.3	69.0	85.0	46.5	67.9
D3still (Xie et al. 2024)		95.0	98.4	91.9	95.9	91.7	95.9	72.8	86.9	46.7	69.9
Ours		95.4	98.5	92.3	95.9	92.7	96.5	74.6	88.3	48.6	71.3

Table 1: Comparison with state-of-the-art asymmetric retrieval methods on standard VPR benchmarks. † denotes performance using the same model for both query and gallery embeddings (*i.e.*, symmetric retrieval).

	Backbone	FLOPs		Params		Latency	
		(G)	rel.	(M)	rel.	(ms)	rel.
SALAD	DINOv2-B	51.0	100%	88.0	100%	61.5	1.0×
	EfficientViT-B2	3.7	7.3%	15.8	18.0%	17.8	3.5×
	MobileViTv2	3.4	6.7%	5.4	6.1%	15.2	4.0×
BoQ	DINOv2-B	54.2	100%	95.2	100%	62.4	1.0×
	EfficientViT-B2	4.4	8.1%	22.3	23.4%	21.1	3.0×
	MobileViTv2	4.2	7.7%	12.1	12.7%	17.3	3.6×

Table 2: Computational efficiency comparison of high-capacity and lightweight backbones.

query networks offer substantial efficiency improvements, highlighting the effectiveness of asymmetric VPR framework. For instance, EfficientViT-B2 requires only 7.3% of the FLOPs and 18.0% of the parameters compared to DINOv2-B when used with SALAD, while enabling 3.5× faster embedding extraction time. MobileViTv2 achieves even greater efficiency, requiring only 6.7% of FLOPs and 6.1% of parameters, leading to 4.0× faster inference. With similar efficiency gains with BoQ, these significant reductions in computational demands make real-time VPR feasible on mobile and edge devices with limited capabilities.

Additionally, we analyze the asymmetric training efficiency of our method compared to CSD (Wu et al. 2022), a representative k-NN-based method, in Table 3. Our approach reduces the gallery embedding precomputation time

Method	Precomputation (min)	Training Iteration (sec)	GPU Memory (GB)
CSD	1392.76	1.34	23.9
Ours	0.26	0.19	17.5

Table 3: Computational efficiency comparison of asymmetric training between CSD (Wu et al. 2022) and our method.

from 1392.76 minutes to just 0.26 minutes. This significant reduction is achieved by replacing exhaustive k-NN computations with a geographical memory bank leveraging the geolocation metadata inherent in VPR datasets. Training iterations are also faster, as our method removes k-NN processing overhead and lowers GPU memory consumption. These improvements are particularly significant for large-scale VPR applications, where the computational cost of k-NN based methods grows with database size, while our geolocation-based approach maintains consistent efficiency.

Ablation Study. To evaluate the effectiveness of each component in our framework, we conduct an ablation study, with the results summarized in Table 4. Specifically, removing implicit embedding augmentation (*i.e.*, applying $\mathcal{L}_{\text{asym}}$ in Eq. 3) leads to a consistent performance drop across all benchmarks, with R@1 decreasing by up to 4.3% on the challenging Nordland dataset, demonstrating its role in bridging the capacity gap between query and gallery networks. We further compare our implicit augmentation with explicit embedding augmentation in Eq. 5, where K is set

Method	Pitts250k		MSLS Val		Tokyo24/7		Nordland		AmsterTime	
	R@1	R@5	R@1	R@5	R@1	R@5	R@1	R@5	R@1	R@5
Ours	95.4	98.5	92.3	95.9	92.7	96.5	74.6	88.3	48.6	71.3
<i>without</i> implicit embedding augmentation	94.3	98.1	91.2	95.7	90.2	95.6	70.3	85.3	46.9	68.9
<i>with</i> explicit embedding augmentation	94.7	98.4	91.5	96.2	91.4	95.2	73.3	86.8	47.7	70.8
<i>with</i> queue-based memory bank	94.2	97.9	91.4	96.1	90.2	95.9	72.0	85.8	47.1	69.7

Table 4: Ablation study with BoQ as a gallery model and EfficientViT-B2 as a query model, comparing our method to variants without implicit embedding augmentation, with explicit embedding augmentation, and with a queue-based memory bank.

Aggregator	Method	Tokyo24/7	Nordland	AmsterTime
MLP (33.4M)	CSD	81.6	39.3	36.1
	MSP	81.0	40.1	34.8
	D3still	86.9	37.2	35.5
	Ours	89.1	43.6	38.7
Symmetric (Ours) (22.3M)	CSD	90.5	68.7	48.1
	MSP	90.5	69.0	46.5
	D3still	91.7	72.8	46.7
	Ours	92.7	74.6	48.6

Table 5: Ablation study (R@1%) of aggregation strategies between MLP-based and Symmetric (Ours) architectures.

to 10. While explicit augmentation provides some improvement over the baseline, it remains less effective than the implicit formulation, highlighting the advantage of leveraging feature statistics. Additionally, replacing our geographical memory bank with a queue-based memory bank, which maintains a fixed-size queue of past embeddings as a dynamic dictionary (He et al. 2020), results in overall performance degradation, with the drop of 2.6% R@1 on Nordland. This highlights the advantage of structured geospatial priors, which are particularly crucial for VPR, over dynamically updated feature queues. Overall, these results validate the effectiveness of each component in our proposed framework for asymmetric VPR performance.

Impact of Aggregation Architecture. We apply the same type of aggregator as the gallery network to extract features from the lightweight query network, ensuring better compatibility with the gallery model. To analyze the impact of maintaining architectural consistency of aggregators, we compare our symmetric aggregation approach with an alternative setup that employs an MLP-based transformation to align with the gallery embedding space. Despite having significantly more parameters (33.4M vs. 22.3M), the MLP-based approach leads to lower asymmetric retrieval performance, as shown in Table 5. This finding underscores the importance of structural alignment in feature aggregation for asymmetric VPR, suggesting that architectural consistency between query and gallery models plays a crucial role in improving feature compatibility.

Qualitative Analysis. To provide deeper insights into the effectiveness of our method, we conduct qualitative analysis using Grad-CAM (Selvaraju et al. 2017) with the geographical memory bank as a linear classifier. As shown in Fig. 3, different methods exhibit distinct attention patterns, revealing key insights about feature alignment. The high-capacity

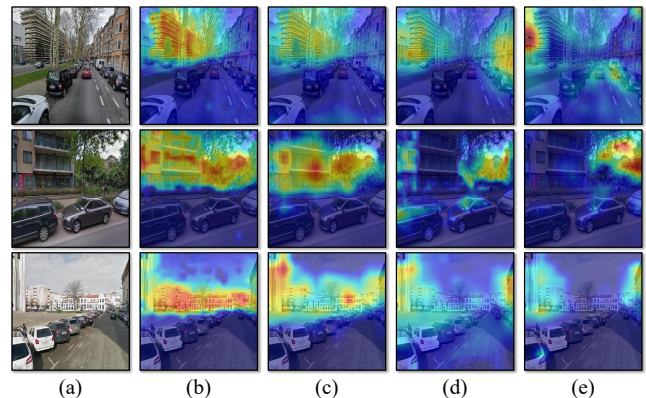


Figure 3: Grad-CAM (Selvaraju et al. 2017) visualization using geographical memory bank as a classifier. (a) Original image; (b) Gallery model (*i.e.*, BoQ); (c) Query model by ours; (d) Query model by ours without implicit embedding augmentation; (e) Query model by CSD (Wu et al. 2022).

gallery model (Fig.3(b)) demonstrates rich attention patterns capturing fine-grained scene details. Our method (Fig.3(c)) successfully emulates these patterns, achieving strong alignment with the gallery model’s focus regions while maintaining efficiency. In contrast, our variant without implicit embedding augmentation (Fig.3(d)) exhibits scattered attention patterns, highlighting the critical role of our augmentation technique in bridging the capacity gap. The k-NN-based CSD (Fig.3(e)) shows markedly different attention patterns with less structured emphasis regions, suggesting insufficient feature compatibility. This analysis supports our quantitative findings, showing that our method achieves better feature compatibility than existing asymmetric approaches.

Conclusion

In this paper, we introduced an efficient asymmetric framework for visual place recognition (VPR), where a lightweight query network is trained to be compatible with a high-capacity gallery model, enabling efficient retrieval in resource-constrained environments. We proposed a geographical memory bank that eliminates the need for exhaustive k-NN computations, and an implicit embedding augmentation technique that bridges the capacity gap between query and gallery models. Extensive experimental results demonstrated that our method not only significantly reduces computational costs compared to conventional VPR methods but also consistently outperforms existing approaches.

Acknowledgements

This work was supported by the Institute of Information & communications Technology Planning & Evaluation (IITP) grant (No. RS-2025-25443318, Physically-grounded Intelligence: A Dual Competency Approach to Embodied AGI through Constructing and Reasoning in the Real World; No. RS-2023-00237965, Recognition, Action and Interaction Algorithms for Open-world Robot Service) and the National Research Foundation of Korea (NRF) grant (No. RS-2023-00208506), both funded by the Korea government (MSIT). Prof. Sung-Eui Yoon is a corresponding author (e-mail: sungeui@kaist.edu).

References

- Ali-bey, A.; Chaib-draa, B.; and Giguère, P. 2022. Gsv-cities: Toward appropriate supervised visual place recognition. *Neurocomputing*.
- Ali-Bey, A.; Chaib-Draa, B.; and Giguere, P. 2023. Mixvpr: Feature mixing for visual place recognition. In *WACV*.
- Ali-bey, A.; Chaib-draa, B.; and Giguère, P. 2024. BoQ: A Place is Worth a Bag of Learnable Queries. In *CVPR*.
- Angeli, A.; Filliat, D.; Doncieux, S.; and Meyer, J.-A. 2008. Fast and incremental method for loop-closure detection using bags of visual words. *IEEE Transactions on Robotics*.
- Arandjelovic, R.; Gronat, P.; Torii, A.; Pajdla, T.; and Sivic, J. 2016. NetVLAD: CNN architecture for weakly supervised place recognition. In *CVPR*.
- Bai, Y.; Jiao, J.; Lou, Y.; Wu, S.; Liu, J.; Feng, X.; and Duan, L.-Y. 2022. Dual-tuning: Joint prototype transfer and structure regularization for compatible feature learning. *IEEE Transactions on Multimedia*.
- Bay, H.; Ess, A.; Tuytelaars, T.; and Van Gool, L. 2008. Speeded-up robust features (SURF). *Computer Vision and Image Understanding*.
- Berton, G.; Masone, C.; and Caputo, B. 2022. Rethinking visual geo-localization for large-scale applications. In *CVPR*.
- Berton, G.; Trivigno, G.; Caputo, B.; and Masone, C. 2023. EigenPlaces: Training Viewpoint Robust Models for Visual Place Recognition. In *ICCV*.
- Bresson, G.; Alsayed, Z.; Yu, L.; and Glaser, S. 2017. Simultaneous localization and mapping: A survey of current trends in autonomous driving. *IEEE Transactions on Intelligent Vehicles*.
- Budnik, M.; and Avrithis, Y. 2021. Asymmetric metric learning for knowledge transfer. In *CVPR*.
- Cadena, C.; Carlone, L.; Carrillo, H.; Latif, Y.; Scaramuzza, D.; Neira, J.; Reid, I.; and Leonard, J. J. 2016. Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on Robotics*.
- Cai, H.; Li, J.; Hu, M.; Gan, C.; and Han, S. 2023. Efficientvit: Lightweight multi-scale attention for high-resolution dense prediction. In *ICCV*.
- Cui, Z.; Zhou, J.; Wang, X.; Zhu, M.; and Peng, Y. 2024. Learning Continual Compatible Representation for Re-indexing Free Lifelong Person Re-identification. In *CVPR*.
- Du, Y.; Wang, M.; Zhou, W.; Hui, S.; and Li, H. 2024. Image2Sentence based Asymmetrical Zero-shot Composed Image Retrieval. In *ICLR*.
- Duggal, R.; Zhou, H.; Yang, S.; Xiong, Y.; Xia, W.; Tu, Z.; and Soatto, S. 2021. Compatibility-aware heterogeneous visual search. In *CVPR*.
- Garg, S.; Fischer, T.; and Milford, M. 2021. Where is your place, visual place recognition? In *IJCAI*.
- Han, K.; Gong, S.; Huang, Y.; Wang, L.; and Tan, T. 2023. Clothing-change feature augmentation for person re-identification. In *CVPR*.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *CVPR*.
- Izquierdo, S.; and Civera, J. 2024. Optimal transport aggregation for visual place recognition. In *CVPR*.
- Jégou, H.; Douze, M.; Schmid, C.; and Pérez, P. 2010. Aggregating local descriptors into a compact image representation. In *CVPR*.
- Jin Kim, H.; Dunn, E.; and Frahm, J.-M. 2017. Learned contextual feature reweighting for image geo-localization. In *CVPR*.
- Keetha, N.; Mishra, A.; Karhade, J.; Jatavallabhula, K. M.; Scherer, S.; Krishna, M.; and Garg, S. 2023. Anyloc: Towards universal visual place recognition. *IEEE Robotics and Automation Letters*.
- Li, P.; Li, D.; Li, W.; Gong, S.; Fu, Y.; and Hospedales, T. M. 2021. A simple feature augmentation for domain generalization. In *ICCV*.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Lowe, D. G. 2004. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*.
- Lowry, S.; Sünderhauf, N.; Newman, P.; Leonard, J. J.; Cox, D.; Corke, P.; and Milford, M. J. 2015. Visual place recognition: A survey. *IEEE Transactions on Robotics*.
- Lu, F.; Lan, X.; Zhang, L.; Jiang, D.; Wang, Y.; and Yuan, C. 2024a. CricaVPR: Cross-image Correlation-aware Representation Learning for Visual Place Recognition. In *CVPR*.
- Lu, F.; Zhang, L.; Lan, X.; Dong, S.; Wang, Y.; and Yuan, C. 2024b. Towards Seamless Adaptation of Pre-trained Models for Visual Place Recognition. In *ICLR*.
- Lu, F.; Zhang, X.; Ye, C.; Dong, S.; Zhang, L.; Lan, X.; and Yuan, C. 2024c. SuperVLAD: Compact and Robust Image Descriptors for Visual Place Recognition. In *NeurIPS*.
- Mehta, S.; and Rastegari, M. 2022. Separable Self-attention for Mobile Vision Transformers. *Transactions on Machine Learning Research*.
- Meng, Q.; Zhang, C.; Xu, X.; and Zhou, F. 2021. Learning compatible embeddings. In *ICCV*.
- Mirowski, P.; Grimes, M.; Malinowski, M.; Hermann, K. M.; Anderson, K.; Teplyashin, D.; Simonyan, K.; Zisserman, A.; Hadsell, R.; et al. 2018. Learning to navigate in cities without a map. *NIPS*.

- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H. V.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; HAZIZA, D.; Massa, F.; El-Nouby, A.; Assran, M.; Ballas, N.; Galuba, W.; Howes, R.; Huang, P.-Y.; Li, S.-W.; Misra, I.; Rabbat, M.; Sharma, V.; Synnaeve, G.; Xu, H.; Jegou, H.; Mairal, J.; Labatut, P.; Joulin, A.; and Bojanowski, P. 2024. DI-NOv2: Learning Robust Visual Features without Supervision. *Transactions on Machine Learning Research*.
- Pan, T.; Xu, F.; Yang, X.; He, S.; Jiang, C.; Guo, Q.; Qian, F.; Zhang, X.; Cheng, Y.; Yang, L.; et al. 2023. Boundary-aware backward-compatible representation via adversarial learning in image retrieval. In *CVPR*.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*.
- Qiu, Q.; Zhang, S.; Gao, H.; Yang, H.; Ying, H.; Wang, W.; and He, X. 2024. EMVP: Embracing Visual Foundation Model for Visual Place Recognition with Centroid-Free Probing. In *NeurIPS*.
- Radenović, F.; Tolias, G.; and Chum, O. 2018. Fine-tuning CNN image retrieval with no human annotation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Ramanujan, V.; Vasu, P. K. A.; Farhadi, A.; Tuzel, O.; and Pouransari, H. 2022. Forward compatible training for large-scale embedding retrieval systems. In *CVPR*.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*.
- Shen, Y.; Xiong, Y.; Xia, W.; and Soatto, S. 2020. Towards backward-compatible representation learning. In *CVPR*.
- Shoshan, A.; Linial, O.; Bhonker, N.; Hirsch, E.; Zamir, L.; Kviatkovsky, I.; and Medioni, G. 2024. Asymmetric image retrieval with cross model compatible ensembles. In *WACV*.
- Su, S.; Zhang, B.; Ge, Y.; Xu, X.; Wang, Y.; Yuan, C.; and Shan, Y. 2022. Privacy-preserving model upgrades with bidirectional compatible training in image retrieval. *arXiv preprint arXiv:2204.13919*.
- Suma, P.; and Tolias, G. 2023. Large-to-small image resolution asymmetry in deep metric learning. In *WACV*.
- Torii, A.; Arandjelovic, R.; Sivic, J.; Okutomi, M.; and Pajdla, T. 2015. 24/7 place recognition by view synthesis. In *CVPR*.
- Torii, A.; Sivic, J.; Pajdla, T.; and Okutomi, M. 2013. Visual place recognition with repetitive structures. In *CVPR*.
- Tzachor, I.; Lerner, B.; Levy, M.; Green, M.; Shalev, T. B.; Habib, G.; Samuel, D.; Zailer, N. K.; Shimshi, O.; Darshan, N.; et al. 2025. Effovpr: Effective foundation model utilization for visual place recognition. In *ICLR*.
- Wan, T. S.; Chen, J.-C.; Wu, T.-Y.; and Chen, C.-S. 2022. Continual learning for visual search with backward consistent feature embedding. In *CVPR*.
- Wang, C.-Y.; Chang, Y.-L.; Yang, S.-T.; Chen, D.; and Lai, S.-H. 2020. Unified representation learning for cross model compatibility. In *BMVC*.
- Wang, Y.; Pan, X.; Song, S.; Zhang, H.; Huang, G.; and Wu, C. 2019. Implicit semantic data augmentation for deep networks. *NeurIPS*.
- Warburg, F.; Hauberg, S.; Lopez-Antequera, M.; Gargallo, P.; Kuang, Y.; and Civera, J. 2020. Mapillary street-level sequences: A dataset for lifelong place recognition. In *CVPR*.
- Wu, H.; Wang, M.; Zhou, W.; and Li, H. 2023. A General Rank Preserving Framework for Asymmetric Image Retrieval. In *ICLR*.
- Wu, H.; Wang, M.; Zhou, W.; Li, H.; and Tian, Q. 2022. Contextual similarity distillation for asymmetric image retrieval. In *CVPR*.
- Xie, Y.; Lin, Y.; Cai, W.; Xu, X.; Zhang, H.; Du, Y.; and He, S. 2024. D3still: Decoupled Differential Distillation for Asymmetric Image Retrieval. In *CVPR*.
- Yildiz, B.; Khademi, S.; Siebes, R. M.; and Van Gemert, J. 2022. Amstertime: A visual place recognition benchmark dataset for severe domain shift. In *ICPR*.
- Yu, J.; Zhu, C.; Zhang, J.; Huang, Q.; and Tao, D. 2019. Spatial pyramid-enhanced NetVLAD with weighted triplet loss for place recognition. *IEEE transactions on Neural Networks and Learning Systems*.
- Zaffar, M.; Garg, S.; Milford, M.; Kooij, J.; Flynn, D.; McDonald-Maier, K.; and Ehsan, S. 2021. Vpr-bench: An open-source visual place recognition evaluation framework with quantifiable viewpoint and appearance change. *International Journal of Computer Vision*.
- Zhang, B.; Ge, Y.; Shen, Y.; Li, Y.; Yuan, C.; XU, X.; Wang, Y.; and Shan, Y. 2022a. Hot-Refresh Model Upgrades with Regression-Free Compatible Training in Image Retrieval. In *ICLR*.
- Zhang, B.; Ge, Y.; Shen, Y.; Su, S.; Wu, F.; Yuan, C.; Xu, X.; Wang, Y.; and Shan, Y. 2022b. Towards universal backward-compatible representation learning. In *IJCAI*.
- Zhang, J.; Cao, Y.; and Wu, Q. 2021. Vector of locally and adaptively aggregated descriptors for image feature representation. *Pattern Recognition*.
- Zhou, D.-W.; Wang, F.-Y.; Ye, H.-J.; Ma, L.; Pu, S.; and Zhan, D.-C. 2022. Forward compatible few-shot class-incremental learning. In *CVPR*.