

UniPR-3D: Towards Universal Visual Place Recognition with Visual Geometry Grounded Transformer

Tianchen Deng^{1*}, Xun Chen^{3*}, Ziming Li¹, Hongming Shen³, Shuhao Zhai⁴,
Danwei Wang³, Javier Civera², Hesheng Wang¹

¹ Shanghai Jiao Tong University ² I3A, University of Zaragoza, Spain ³ Nanyang Technological University ⁴ University of Macau

Abstract. Visual Place Recognition (VPR) has been traditionally formulated as a single-image retrieval task. Using multiple views offers clear advantages, yet this setting remains relatively underexplored and existing methods often struggle to generalize across diverse environments. In this work, we introduce UniPR-3D, the first VPR architecture that effectively integrates geometry-aware information from multiple views. UniPR-3D builds on a VGGT backbone capable of encoding multi-view 3D representations, which we adapt by designing feature aggregators and fine-tuning them for the place recognition task. To construct our descriptor, we jointly leverage VGGT’s 3D tokens, but also intermediate 2D ones. Based on their distinct characteristics, we design dedicated aggregation modules for 2D and 3D features, allowing our descriptor to capture fine-grained texture patterns while also reasoning across viewpoints. To further enhance generalization, we incorporate both single- and multi-frame aggregation schemes, along with a variable-length sequence retrieval strategy. Our experiments show that UniPR-3D sets a new state of the art, outperforming both single- and multi-view baselines and highlighting the effectiveness of geometry-grounded tokens for VPR. Our code and models will be made publicly available on Github. We will release the code and datasets publicly on <https://github.com/dtc111111/UniPR-3D>.

1 Introduction

Visual Place Recognition (VPR) is a core problem in robotics and computer vision, targeting the recognition of previously visited locations from visual observations [31, 49, 59, 61, 69, 71, 93]. VPR plays a key role in fundamental tasks such as SLAM [13, 22, 23, 28, 82] and absolute pose estimation [44, 67, 79, 95], which underpin a wide range of applications, including autonomous driving, robotic navigation, virtual and augmented reality.

The first two authors contribute equally to this paper. Corresponding author: Hesheng Wang

VPR is typically formulated as a single-view retrieval task, in which a query image is used to retrieve an ordered list of the top- k matching candidates from a database. Existing approaches generally rely on neural backbones to extract features, which are then aggregated using various strategies to form global image descriptors. Images are then represented by their aggregated appearance patterns, which are compared at test time via nearest-neighbor search. Backbone architectures have evolved from ResNet-based [4, 36, 65] to Vision Transformers (ViTs) [41, 85, 96]. Task-specific fine-tuned foundation models, in particular DINOv2 [64], achieve nowadays state-of-the-art performance [7, 38, 39]. Feature aggregation strategies have also advanced considerably, from the seminal NetVLAD [4], inspired by the handcrafted VLAD aggregation [40], to pooling-based methods such as GeM [65], MLP-based aggregations like MixVPR [2], and approaches leveraging optimal transport [39].




Fig. 1: Illustration of a typical single-view VPR model (top) and our multi-view model (bottom). Single-view VPR extracts image features with a deep backbone and aggregates them into an image descriptor. In contrast, our multi-view VPR model employs a VGGT backbone to jointly extract both 2D and 3D tokens from multiple views, followed by tailored aggregation strategies for each token type. Our framework supports both single-frame and variable-length sequence matching, and achieves state-of-the-art performance across standard VPR benchmarks.

distinct aggregation strategies. For 2D tokens, we select the 2D CLS token, 2D register tokens, and 2D patch tokens for aggregation. For the 3D tokens, we discard the camera token (as VPR should be robust to viewpoint changes) and only use the 3D register tokens and 3D patch tokens to ensure stability under viewpoint changes. For patch tokens, we adopt an optimal transport approach [39]

However, most VPR methods share the fundamental limitation of extracting features from a single view. While multiple views provide broader scene coverage and a richer variety of viewpoints, multi-view feature aggregation has always been a challenging and relatively unexplored topic [26]. The recent development of the Visual Geometry Grounded Transformer (VGGT) [84], a large ViT backbone encoding 3D-aware multiview representations, has opened new possibilities for leveraging such input modality in various downstream tasks [45, 75, 76]. Building on this backbone for feature extraction, we introduce in this paper UniPR-3D, the first 3D token-based VPR framework supporting both single-frame and sequence-level place matching.

Our developments work as follows. For single-frame retrieval, we jointly utilize VGGT’s 3D tokens and intermediate 2D tokens to construct the place descriptor. Based on the differences between token types, we design

to derive local patch descriptors, while for register tokens, we employ GeM pooling [65] to generate their local descriptors.

For sequence-level retrieval, we define an anchor frame and multiple support frames, and construct a multi-frame feature aggregation method. We use SALAD in order to aggregate both 2D and 3D patch tokens across frames. For 3D register tokens as well as 2D CLS and register tokens, we design an aggregation that combines GeM pooling with an MLP-based projector. The token projector aligns the feature dimensions across different modalities and frames, enabling our network to process input sequences of arbitrary length. This design allows the model to generalize effectively to sequences with varying lengths during inference.

By incorporating 2D and 3D tokens, UniPR-3D achieves universal VPR capabilities across complex environments. Overall, our contributions are as follows:

- We develop the first 3D token-based VPR method with tailored 2D and 3D feature aggregation methods, enabling universal generalization across diverse environments and supporting both single frame and sequence matching.
- We propose distinct aggregation strategies for different types of tokens. Specifically, register tokens are GeM-pooled, as their number is relatively small, while patch tokens are aggregated using optimal transport to capture fine-grained spatial correspondences across frames.
- We introduce a sequence-level retrieval strategy, with an anchor frame and multiple support frames, and construct a multi-frame feature aggregation method, enabling our network to process input sequences of arbitrary length.

We evaluate UniPR-3D on a comprehensive and diverse set of public datasets under challenging conditions, including temporal variations, spatial changes, and viewpoint shifts. Our results show that our method effectively leverages multi-view information and sets a new state of the art across these benchmarks.

2 Related work

VPR has made remarkable progress in recent years and has become a key component in tasks such as localization [21, 24, 35, 48, 72], planning and control [46, 47, 53, 54, 87], autonomous driving [19, 88, 91, 92] and virtual reality [16, 55–57]. Below, we review the most relevant advances in feature extraction, feature aggregation, and sequence-based retrieval.

Feature Extraction. Early VPR methods relied on handcrafted visual descriptors, using either aggregated local features [40] or holistic global representations [63, 78]. Geometric and temporal consistency checks were often employed to improve robustness in practical systems [29, 62]. With the advent of deep learning, ResNet-based architectures became the standard for feature extraction [4]. More recently, Vision Transformer (ViT)-based foundation models [20] have shown superior representational power and generalization. For example, AnyLoc [41], SALAD [39], UniLGL [73] and so on [50, 52] leverage DINOv2 [64]: AnyLoc uses it directly, while the rest of them fine-tune it for the VPR task.




Fig. 2: Overview. We propose the first VPR method that supports both single-frame and sequence-level place retrieval. Specifically, we use DINOv2 as our visual feature extractor and then utilize the alternating attention blocks of VGGT to derive 3D tokens. The resulting 3D tokens are divided into different groups, each processed with a dedicated aggregation strategy to form the final descriptor. In particular, cls tokens and register tokens are aggregated using GeM pooling, while patch tokens are processed through an optimal transport module, where the Sinkhorn algorithm is applied to compute the assignment matrix. The outputs from these modules are finally concatenated to produce the global descriptor.

These models have demonstrated enhanced robustness to variations in illumination, weather, and viewpoint. Nevertheless, the vast majority of existing methods extract features from a single view, thus failing to exploit the spatial information available across multiple observations. The recent development of the 3D foundation model has opened new opportunities for geometry-aware visual representation learning and has shown strong potential in various downstream tasks [14, 27, 30, 83, 90].

Feature Aggregation. Seminal VPR aggregation schemes include Bag of Visual Words (BoW) [74] and Vector of Locally Aggregated Descriptors (VLAD) [40], both based on clustering handcrafted local descriptors. NetVLAD [4] extended VLAD by embedding it within a deep architecture that jointly learns the convolutional features and the aggregation process. The Generalized Mean Pooling (GeM) [65] of features, proposed later, was a widely used aggregation method combining simplicity and strong retrieval performance. MixVPR [2] introduced a MLP-based aggregation module to capture more expressive feature interactions. Recently, SALAD [39] proposed an aggregation based on optimal transport [12], achieving the current state of the art in feature aggregation for VPR.

Sequence Retrieval. Most existing approaches [25, 62, 70] perform sequence-level retrieval as a post-processing step on top of frame-level similarity scores. In contrast, several learning-based methods explicitly model temporal dependencies by learning sequential descriptors [34, 60] or by aggregating consecutive

reference and query features into unified sequence-level representations [5, 33]. Other works, such as [37], introduce supervised strategies to estimate per-frame sequence-matching confidence, enabling selective fusion of temporal cues. CaseVPR [43] further proposes a hierarchical sequence-to-frame retrieval pipeline that combines coarse sequence retrieval with fine-grained frame alignment for accurate correspondence estimation. Recent methods like CricaVPR [51] and SeqVLAD [60] show promising results. Nevertheless, all of these methods learn to aggregate single-view features without geometric awareness. Despite advances, sequence retrieval methods still exhibit two main limitations. First, they are typically designed with strong inductive biases, whereas our UniPR-3D, built upon VGGT and our learned aggregation modules, adopts a more data-driven formulation, resulting in greater flexibility. Second, their aggregation strategies operate purely along the temporal dimension, which makes them sensitive to variations in speed, frame rate, or sequence sparsity. In contrast, VGGT’s features are geometrically grounded in a 3D-aware multi-view representation, enabling UniPR-3D to achieve greater robustness under such variations. Differently from the existing VPR literature, our architecture is novel in that it demonstrates for the first time an effective and principled fusion of both 2D single-view and 3D geometry-aware multi-view tokens. This allows our framework to process a variable number of images without strong architectural bias.

3 Method

We propose UniPR-3D, a universal visual place recognition framework that supports both frame-to-frame and sequence-to-sequence matching. The input to our method consists of an image sequence $\{\mathbf{I}_i\}$, and for sequence-level matching, we define an anchor frame and multiple support frames. Our model is capable of predicting visual descriptors for both individual frames and entire sequences. The framework is composed of three main modules: (i) Local 3D feature extraction (Sec. 3.1); (ii) Feature assignment and aggregation (Sec. 3.2). (iii) Sequence matching (Sec. 3.3). We elaborate on the entire pipeline of our system in the following subsections.

3.1 Local 3D Feature Extraction

Effective local feature extraction should remain robust under seasonal or day-to-night illumination changes, as well as viewpoint variations. This requires capturing sufficiently stable structural information. However, the most reliable cues for stability lie in the **3D structural information** of the scene, whereas existing methods primarily focus on texture-level feature extraction from a single image.

Existing models such as AnyLoc [41] and SALAD [39] adopt DINOv2 [64] leverage the exceptional representational capabilities of foundation models to extract local features. However, these approaches can only capture 2D image patterns. In UniPR-3D, we propose the first VPR framework based on the 3D

tokens of a VGGT [84] backbone. Our framework follows a ViT-based architecture, in which every input image $\{\mathbf{I}_i\}$ is first divided into patches $m \times m \times c$, which are sequentially processed through transformer blocks to generate different tokens.

Each image is first processed by the DINOv2 encoder to extract the 2D *cls token*, *register token*, and *patch tokens*, which primarily capture the texture-level characteristics of this image. Then we retain only the patch tokens as input to the subsequent alternating attention blocks and initialize two additional tokens: a *camera token*, which encodes the camera extrinsics and intrinsics (including $\{R, T\}$, and the field-of-view parameters), and a *register token*. Through the alternating blocks composed of frame attention and global attention, the model produces three types of 3D tokens: a 3D camera token ($n \times 1$), 3D register tokens ($n \times 4$), and 3D patch tokens ($m \times m$). To ensure viewpoint invariance of the final descriptor, we only preserve the 3D register and 3D patch tokens to capture both geometric and structural information of the environment, forming our 3D geometric descriptor.

It is important to note that UniPR-3D does not require any camera intrinsics or extrinsics as input. All descriptors, including the geometry-aware 3D tokens, are computed exclusively from pure RGB images. By leveraging the intermediate VGGT tokens through spatial alternating attention blocks, our representation naturally captures viewpoint-consistent and temporally stable scene structures without relying on external geometric priors.

3.2 Feature Assignment and Aggregation

Due to the heterogeneous nature of 2D and 3D tokens, we adopt tailored aggregation strategies for each type. In particular, the 2D *cls tokens* ($\mathbf{f}_{cam}$), 2D *register tokens* ($\mathbf{f}_{reg2d}$), and 3D *register tokens* ($\mathbf{f}_{reg3d}$) are limited in number, rendering complex aggregation schemes unnecessary. Instead, we employ Generalized Mean (GeM) pooling [65], preceded and followed by lightweight MLPs, to generate compact and robust global descriptors:

$$\mathbf{d}_{cls2d} = \text{MLP} \left(\left(\frac{1}{N} \sum_{i=1}^N \text{MLP}(\mathbf{f}_{cam}^p) \right)^{\frac{1}{p}} \right), \quad (1)$$

$$\mathbf{d}_{reg2d} = \text{MLP} \left(\left(\frac{1}{N} \sum_{i=1}^N \text{MLP}(\mathbf{f}_{reg2d}^p) \right)^{\frac{1}{p}} \right), \quad (2)$$

$$\mathbf{d}_{reg3d} = \text{MLP} \left(\left(\frac{1}{N} \sum_{i=1}^N \text{MLP}(\mathbf{f}_{reg3d}^p) \right)^{\frac{1}{p}} \right), \quad (3)$$

where p is a learnable pooling parameter controlling the degree of selectivity. This simple yet effective aggregation captures the dominant semantic cues of camera and register tokens with high stability.

In contrast, patch tokens are aggregated using an Optimal Transportation (OT) scheme, which computes a soft matching plan to preserve structural correspondences, following the design of SALAD [39]. We use two randomly initialized fully connected layers to learn each row of the score matrix $\mathbf{S}$ from scratch for 2D and 3D patch token $\mathbf{t}_{2d}, \mathbf{t}_{3d}$:

$$\mathbf{s}_{2d} = \mathbf{W}_{2d_2} (\sigma (\mathcal{W}_{2d_1} (\mathbf{t}_{2d}) + \mathbf{b}_{2d_1})) + \mathbf{b}_{2d_2} \quad (4)$$

$$\mathbf{s}_{3d} = \mathbf{W}_{3d_2} (\sigma (\mathcal{W}_{3d_1} (\mathbf{t}_{3d}) + \mathbf{b}_{3d_1})) + \mathbf{b}_{3d_2} \quad (5)$$

where $\mathbf{W}_{2d_1}, \mathbf{W}_{2d_2}, \mathbf{W}_{3d_1}, \mathbf{W}_{3d_2}$ are the weights and biases of the layers for 2D tokens, and σ is a non-linear activation function. The parameters for the 3D token layers are defined in a similar manner.

Inspired by SALAD [39], we also introduce a dustbin entry for 2D and 3D patch tokens, to which non-informative features can be assigned. We augment the score matrix from $\mathbf{S}$ to $\bar{\mathbf{S}} = [\mathbf{S}, \bar{\mathbf{s}}_{i,m+1}] \in \mathbb{R}_{>0}^{n \times m+1}$, by appending the column $\bar{\mathbf{s}}_{i,m+1}$ representing the feature-to-dustbin relation. Similar to SuperGlue [68], this score is modeled with a single learnable parameter $z \in \mathbb{R}$.

$$\bar{\mathbf{s}}_{i,m+1} = z \mathbf{1}_n \quad (6)$$

where $\mathbf{1}_n = [1, \dots, 1]^\top \in \mathbb{R}^n$ a n -dimensional vector of ones. We follow SuperGlue [68] and SALAD [39], and use the Sinkhorn Algorithm [17] to obtain the optimal assignment $\bar{\mathbf{P}} \in \mathbb{R}^{n \times (m+1)}$:

$$\bar{\mathbf{P}} \mathbf{1}_{m+1} = \boldsymbol{\mu} \quad \text{and} \quad \bar{\mathbf{P}}^\top \mathbf{1}_n = \boldsymbol{\kappa}. \quad (7)$$

The Sinkhorn Algorithm finds the optimal transport assignment between the distributions $\boldsymbol{\mu}$ and $\boldsymbol{\kappa}$ by iteratively normalizing the rows and columns of the score matrix. Finally, we drop the dustbin column to obtain the assignment $\mathbf{P} = [\mathbf{p}_{*,1}, \dots, \mathbf{p}_{*,m}]$. We use a MLP to reduce the dimensionality of the patch tokens for efficiency.

Then the patch descriptor can be computed as: $d_{j,k} = \sum_{i=1}^n P_{i,k} \cdot f_{i,k}$, where $f_{i,k}$ corresponds to the k^{th} dimension of $\mathbf{f}_i$, with $k \in \{1, \dots, l\}$. The final descriptor is composed of the 2D CLS descriptor, 2D register descriptor, 2D patch descriptor, 3D register descriptor and 3D patch descriptor. These five components are concatenated to form the final unified descriptor.

$$\mathbf{d} = [\mathbf{d}_{cls2d}^\top \quad \mathbf{d}_{reg2d}^\top \quad \mathbf{d}_{patch2d}^\top \quad \mathbf{d}_{reg3d}^\top \quad \mathbf{d}_{patch3d}^\top]^\top \quad (8)$$

3.3 Sequence Matching

For frame-to-frame matching, we directly process a single image to extract its corresponding 3D tokens. For sequence-to-sequence matching, we instead define an anchor frame together with multiple support frames. In the VGGT architecture, the first frame defines the world coordinate system, and all tokens are

registered with respect to this anchor. This design, coupled with the training pipeline, makes the first frame essential for preserving spatial consistency across the sequence.

Our multi-frame aggregation strategy is as follows. Existing multi-frame fusion approaches typically operate on fixed-length sequences, requiring identical sequence lengths during training and inference, which limits their generalization. To overcome this constraint, we propose a variable-length sequence fusion framework. Specifically, we design a multi-frame projector based on Generalized Mean (GeM) pooling to aggregate the 2D CLS tokens, 2D register tokens, and 3D register tokens across frames, enabling flexible sequence-level feature fusion.

$$\mathbf{d}_{\text{reg3d}}^M = \text{MLP} \left(\left(\frac{1}{N} \sum_{i=1}^N \text{MLP} \left(\{\mathbf{f}_{\text{reg3d}}^p, \mathbf{f}_{\text{reg3d}}^p, \dots, \mathbf{f}_{\text{reg3d}}^p\} \right) \right)^{\frac{1}{p}} \right), \quad (9)$$

For the multi-frame patch tokens, we adopt an optimal transport formulation for feature aggregation. Unlike the single-frame case, we first cluster the patch tokens belonging to different frames and then compute the assignment matrix using the Sinkhorn algorithm. The final patch descriptor is computed as:

$$\mathbf{d}_{\text{patch2d}}^M = \sum_{i=1}^n P_{i,k}^M \cdot \{f_{i,k}^1, f_{i,k}^2, \dots, f_{i,k}^M\} \quad (10)$$

3.4 Training

We follow the experimental setups of two recent works: SALAD [39] for single-frame retrieval and SeqMatchNet [34] for sequence-level retrieval.

For single-frame training, we use GSV-Cities [1] as it is standard practice in the field. GSV-Cities is a large-scale collection of urban locations from Google Street View. For optimization, we adopt the multi-similarity loss [86] together with the AdamW optimizer [42]. For sequence-level training, since GSV-Cities does not support multi-frame training, we follow the setting of SeqMatchNet [34] and train our model on the Mapillary Street-Level Sequences (MSLS) dataset [89]. MSLS contains over 1.6 million images collected from the Mapillary collaborative mapping platform, covering diverse cities and exhibiting significant variations in viewpoint, weather, and illumination conditions. The training process consists of two stages: in the first stage, we train only the descriptor head; in the second stage, we jointly train the alternating attention blocks of VGGT and the DINOv2 encoder. To ensure a fair and accurate comparison, especially with strong foundation model-based baselines like SALAD [39], we train our models using full-precision (float32). Furthermore, we align the input image resolution across models, which not only standardizes the evaluation but also optimizes the inference latency of our framework.

4 Experiments

To rigorously evaluate the effectiveness of our proposed method, we conduct extensive experiments following the standard evaluation protocols established in SALAD [39] for single-frame retrieval and follow SeqMatchNet [34] for sequence retrieval.

4.1 Implementation Details

Our model architecture follows VGGT [84] with $L = 24$ alternating frame and global attention layers. To accelerate inference, we incorporate FlashAttention-2 [18]. The model is initialized using pre-trained VGGT weights, and is trained for VPR using the AdamW optimizer with a hybrid learning rate schedule: a linear warm-up over the first 0.5 epochs, followed by cosine decay, with a peak learning rate of 1×10^{-6} . All training are conducted on a NVIDIA A100 GPU. The inference speed is tested on a single NVIDIA RTX 4090 GPU. For the 3D backbone, we adopt a LoRA-based fine-tuning strategy, where both the frame attention and global attention blocks are refined. For the cls descriptor and register descriptor, we employ dimensionality reduction, compressing the feature token dimensions to 256. In contrast, the patch descriptor is represented with a dimension of $128 \times 64 = 8192$, with 64 clusters.

Dataset and Metrics For single-frame evaluation, we conduct experiments on a wide range of benchmarks to verify the generalization of our method. We use the validation and challenge partitions of MSLS [89], which consists of dashcam images, Pittsburgh250k-test [81], which images urban scenarios, NordLand [77] for its seasonal variations in train-front imagery in Norway, SPED [15] which records with surveillance cameras, Tokyo 24/7 [80] including differences in light between day and night, and SF-XL [6, 8], containing heavy illumination changes and substantial occlusions. We use Recall@k (R@k) as the evaluation metric for all experiments, following standard practice in the literature. We follow the MixVPR evaluation protocol [2], where retrieval is considered correct if at least one image within 25 meters of the query location (or within two frames for the NordLand dataset) appears among the top-k candidates.

For multi-frame evaluation, we conduct experiments on the Nordland [77], MSLS [89], Oxford1 [58], and Oxford2 [58] datasets. Specifically, for Oxford1 we use 2014-12-16-18-44-24 (winter night) to 2014-11-18-13-20-12 (fall day), sampled every 2 meters. And for Oxford2, 2014-11-14-16-34-33 (fall night) to 2015-11-13-10-28-08 (fall day), sampled every 2 meters. We adopt Recall@k (R@k) as the evaluation metric for all experiments. Differently from the single-frame setting, for the Oxford datasets, directly using a retrieval threshold of 25 meters often leads to saturated results (close to 100%) for existing methods. Therefore, we perform a more detailed evaluation using thresholds of 2 m and 25 m. For the Nordland dataset, retrieval is considered correct if the ground-truth frame within ten frames from the query appears among the top- k retrieved candidates. Following CaseVPR [43], the sequence length S is set to be 5.

Method	MSLS Chall.		MSLS val		Pitts250k		Nordland		SPED		SF-XL v1		SF-XL v2		SF-XL ni.		SF-XL occ.		Tokyo 24/7	
	R@1	R@10	R@1	R@10	R@1	R@10	R@1	R@10	R@1	R@10	R@1	R@10	R@1	R@10	R@1	R@10	R@1	R@10	R@1	R@10
NetVLAD [4]	35.1	51.7	54.5	70.4	85.9	95.0	32.6	53.3	78.7	91.4	40.1	57.7	76.9	91.1	6.7	14.2	9.2	22.4	69.8	82.9
AP-GeM [66]	35.9	53.6	56.0	72.9	80.0	93.5	22.9	42.1	64.6	83.5	37.9	54.1	66.4	84.6	7.5	16.7	5.3	14.5	57.5	77.5
CosPlace [9]	67.2	76.4	85.0	92.6	92.3	98.4	44.2	54.4	75.9	89.1	76.6	85.5	88.8	96.8	23.6	32.8	30.3	44.7	87.3	95.6
MixVPR [2]	64.0	80.7	83.2	91.9	94.3	98.9	58.4	80.0	85.2	94.6	72.5	80.9	88.6	95.0	19.5	30.5	30.3	38.2	87.0	94.0
EigenPlaces [10]	67.4	81.7	85.9	93.1	94.1	98.7	54.2	73.9	82.4	94.7	84.0	90.7	90.8	96.7	23.6	34.5	32.9	52.6	93.0	97.5
AnyLoc [41]	40.2	51.3	58.7	74.5	89.4	98.0	69.5	84.9	80.5	92.5	82.7	88.5	88.9	94.8	20.5	31.9	29.3	39.9	87.6	97.5
SALAD [39]	75.0	91.3	92.2	97.2	95.2	99.2	76.0	90.5	92.1	97.2	90.9	95.9	95.1	98.7	50.3	68.4	51.8	74.1	95.1	98.3
CricaVPR [51]	69.3	85.9	76.7	87.2	92.6	98.3	70.2	85.5	83.1	93.5	62.6	78.9	86.3	96.0	25.8	40.6	27.6	47.4	82.9	93.7
BoQ [3]	71.2	86.9	91.2	96.1	95.0	99.1	70.7	87.5	86.5	95.7	87.9	93.7	88.5	94.9	27.4	37.2	34.7	54.1	97.5	98.7
MegaLoc [7]	73.4	88.7	91.0	95.8	96.4	99.3	76.7	90.9	92.0	95.9	<u>95.3</u>	<u>98.0</u>	94.8	98.5	<u>52.8</u>	<u>73.8</u>	51.3	75.0	96.5	99.4
UniPR 3D	75.5	<u>91.4</u>	<u>92.9</u>	<u>97.6</u>	<u>96.5</u>	<u>99.5</u>	<u>78.4</u>	<u>91.4</u>	<u>92.6</u>	<u>97.7</u>	93.8	96.9	<u>95.7</u>	<u>99.1</u>	52.1	72.8	<u>52.1</u>	<u>75.8</u>	<u>97.6</u>	<u>99.5</u>
UniPR 3D*	75.9	91.8	93.2	97.9	96.6	99.7	78.9	92.4	92.8	97.9	96.3	98.5	95.9	99.3	54.8	74.5	53.9	77.4	97.9	99.7

Table 1: Single-frame matching results. We compare our UniPR-3D against existing single-frame VPR baselines. Our method achieves a noticeable improvement in the recall, highlighting the benefit of combining 2D and 3D tokens for the place retrieval. * indicates same training setting as MegaLoc.

Method	Backbone	Desc. size	MSLS Val (pos=25m)			NordLand (pos=10f)			Oxford1 (pos=2m)			Oxford2 (pos=2m)		
			R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10
SeqSLAM [62]	VGG-16	4096	45.9	58.2	70.4	53.1	67.4	71.9	34.7	51.1	70.2	26.5	36.9	44.0
SeqMatchNet [34]	VGG-16	4096	65.5	77.5	80.3	56.1	71.4	76.9	36.8	43.3	48.3	27.9	38.5	45.3
Delta Descriptors [32]	VGG-16	4096	43.0	58.4	68.6	53.1	67.4	70.2	36.9	55.2	74.1	26.5	46.3	62.8
NetVLAD + FC [4]	ResNet-18	4096	68.5	79.1	84.6	59.3	75.5	80.0	50.1	60.4	71.3	18.1	28.5	35.2
NetVLAD + FC [4]	ResNet-50	4096	71.0	80.5	84.8	64.6	67.2	78.4	58.1	72.4	78.3	20.1	31.5	39.2
SeqPool + CAT	CCT224	384-SL	74.2	83.5	87.9	62.6	65.2	76.4	52.5	65.4	73.3	13.9	26.4	33.2
SeqPool + CAT	CCT384	384-SL	77.8	85.0	88.1	63.5	65.9	76.6	53.5	66.6	74.3	14.5	27.4	34.2
GeM + CAT [65]	ResNet-18	1280	76.8	84.4	89.8	61.2	63.1	74.7	57.5	71.6	79.3	19.5	30.4	38.2
GeM + CAT	ResNet-50	5120	68.6	78.9	84.7	62.2	64.1	75.9	59.6	73.1	81.4	20.4	30.4	38.1
SeqNet [33]	ResNet-18	4096	68.0	77.6	82.5	60.4	62.5	73.1	52.4	65.8	74.8	14.1	25.9	33.3
SeqNet [33]	ResNet-50	4096	71.1	80.2	85.0	61.9	64.4	75.8	57.4	69.8	76.8	16.5	27.9	36.3
JIST [11]	ResNet-18	512	86.6	89.4	91.5	62.3	72.4	88.1	57.2	69.8	77.8	17.1	27.9	35.3
sVPR [94]	ResNet-50	24576	87.6	90.4	92.1	63.5	73.9	88.9	58.3	70.4	78.4	17.8	28.4	36.5
SeqVLAD [60]	CCT384	24576	89.9	92.4	94.1	65.5	75.2	80.0	58.4	72.8	80.8	19.1	29.9	37.3
CaseVPR [43]	DinoV2 Vit-B	10752	<u>91.2</u>	<u>94.1</u>	<u>95.0</u>	<u>84.1</u>	<u>89.9</u>	<u>92.2</u>	<u>90.5</u>	<u>95.2</u>	<u>96.5</u>	<u>72.8</u>	<u>85.8</u>	<u>89.9</u>
UniPR-3D (ours)	VG GT	17152	93.7	95.7	96.9	86.8	91.7	93.8	95.4	98.1	98.7	80.6	90.3	93.9

Table 2: Sequence matching results. We compare our UniPR-3D against existing sequence-level VPR baselines. Our method achieves significantly higher recall than competing approaches, highlighting the advantages of our 3D token-based framework in aggregating image patterns across multiple views. Note that for the Oxford datasets, the retrieval distance is set to 2m.

4.2 Experimental Results

Single-Frame Matching. We compare against a diverse set of baselines, including NetVLAD [4], AP-GeM [66], CosPlace [9], MixVPR [2], EigenPlaces [10], AnyLoc [41], SALAD [39], CricaVPR [51], BoQ [3], and MegaLoc [7]. As shown in Tab. 1, UniPR-3D outperforms all baselines on most datasets and across most metrics. All compared methods are trained under the same setting using the GSV-Cities dataset, and the results of the baseline models are obtained by re-training the official implementations released by the authors. We further evaluate our approach on challenging datasets, including MSLS Challenge, characterized by its large scale, diversity, and closed-set labeling, and Nordland, known for its extreme visual similarity and severe seasonal variations. The experimental results demonstrate that our 3D token-based descriptor is capable of capturing stable 3D structural information within the scene, rather than relying solely on texture cues, leading to significantly improved robustness and generalization. As

Methods	Latency	Oxford1 (pos=2m)			Oxford1 (pos=25m)		
		R@1	R@5	R@10	R@1	R@5	R@10
SeqMatchNet [34]	115	36.8	43.3	48.3	47.2	54.7	63.4
SeqVLAD [60]	129	58.4	72.8	80.8	72.2	80.3	86.1
CaseVPR [43]	75	90.5	95.2	96.5	97.8	98.8	99.3
UniPR-3D (ours)	140	95.4	98.1	98.7	99.3	99.5	99.6

Table 3: Comparison against baselines under varying distance thresholds (2 m and 25 m) in the Oxford dataset [58]. Our method consistently outperforms all baselines across all settings.




Fig. 3: Qualitative sequence matching results on the Oxford dataset [58]. The left column shows two query sequences, while the three right columns present the top-3 candidates retrieved by our UniPR-3D and the baseline CaseVPR [43]. Successful retrievals are framed in green, while erroneous ones are shown in red. UniPR-3D retrieves the correct place even under seasonal, weather, viewpoint, and day-night variations.

expected, our method incurs in higher latency compared to existing approaches due to the incorporation of additional 3D information.

Sequence Matching. We compare against a diverse set of existing sequence-matching VPR methods: SeqSLAM [62], SeqNet [33], SeqMatchNet [34], Delta Descriptors [32], JIST [11], sVPR [94], SeqVLAD [60] and CaseVPR [43]. We retrain all competing models on the MSLS dataset to ensure a fair comparison. As shown in Tab. 2, our method delivers superior performance on multi-frame matching tasks, consistently outperforming prior approaches across all datasets and evaluation metrics. In Tab. 3, we additionally report results on the Oxford dataset under different retrieval distance thresholds (2 m, 25 m). We find that the stricter 2 m threshold more clearly exposes performance differences among methods. Under this setting, UniPR-3D surpasses the previous state of the art



Fig. 4: An illustration explicitly showing the viewpoint-reversed behavior. The two sequences involved are captured under daytime and nighttime conditions, respectively, with the vehicle traversing the same route in opposite directions. As a result, the retrieved images naturally exhibit viewpoint reversal with respect to the queries, leading to front-back flipped perspectives in the visualization.

by more than **10%**. These results clearly show the effectiveness of UniPR-3D in capturing 3D structural cues from multiple views, and maintaining robustness against variations in illumination, weather, and viewpoint. By leveraging spatial information across frames in a geometry-aware manner, our sequence-level framework highlights the significant benefits and future potential of incorporating 3D tokens into VPR systems.

Besides the quantitative results, we also visualize several sequence matching examples in Fig. 3, where we show the query frames together with their top-3 sequences retrieved by UniPR-3D. We further compare our method with the current SOTA sequence matching method CaseVPR [43]. UniPR-3D is able to retrieve correct matches even under challenging conditions, such as severe illumination changes and large viewpoint variations, further confirming the robustness of our 3D token-based descriptor. As visualized in Fig. 3, some retrieved images naturally exhibit a flipped front-back perspective with respect to queries, as shown in Fig. 4. This occurs because the corresponding sequences in the Oxford dataset are captured under daytime and nighttime conditions with the vehicle traversing the same route in opposite directions. Importantly, this viewpoint-reversed behavior does not indicate incorrect retrieval; on the contrary, it strongly reflects the robust place recognition capabilities of our 3D geometry-grounded descriptor under extreme viewpoint and illumination changes. As shown in Fig. 6, we provide t-SNE visualization. Our method produces well-separated and compact clusters, indicating strong discriminative capability and consistent feature structuring. This qualitative comparison further demonstrates the superiority of our 3D token-enhanced representation in capturing robust geometric and semantic information for visual place recognition.

2D cls	2D regis.	2D patch	3D regis.	3D patch OT	3D patch GEM	3D Pose	MSLS val		Oxford 1		Oxford 2		Nordland	
							R@1	R@5	R@1	R@5	R@1	R@5	R@1	R@5
		✓					84.9	87.7	86.8	90.1	71.6	80.2	75.6	79.5
✓	✓	✓					90.4	91.6	91.5	92.9	75.9	88.1	80.4	84.5
			✓				91.9	93.1	92.1	93.8	76.4	88.7	82.5	87.7
			✓	✓			93.7	95.5	95.5	97.6	80.3	89.8	86.4	91.2
✓	✓	✓	✓			✓	92.2	94.5	94.3	95.4	78.8	89.4	85.4	89.8
✓	✓	✓	✓				93.9	95.9	96.1	97.4	78.3	89.3	85.8	90.3
✓	✓	✓	✓	✓			94.1	96.2	96.3	98.9	81.4	91.5	87.4	92.5

Table 4: Ablation studies on 2D–3D token selection, aggregation strategies, and 3D pose injection, and analyze their impact on the final results on additional datasets.




Fig. 5: Activation heatmaps of 3D and 2D features in the Oxford dataset [58]. The left column shows reference images, and the two right columns their corresponding 3D and 2D feature heatmaps. Note how the different feature types fire at different locations in the image, illustrating their complementary nature.

#frames at test time	Oxford2		
	R@1	R@5	R@10
3	70.3	87.2	91.7
5 (training setup)	80.6	90.3	93.9
10	89.1	93.1	95.3
15	92.4	94.9	96.6

Table 5: Results with varying sequence length at test time. UniVPR-3D generalizes for sequence lengths differing from the training setting, maintaining strong performance.

4.3 Ablation Studies

Effect of the different descriptors. Our final descriptor consists of both a texture descriptor and a 3D geometric descriptor. Each descriptor further integrates five distinct types of tokens. To validate the contribution of each token, we first conduct ablation experiments on the 2D and 3D patch descriptors, shown in Tab. 4. The results indicate that both 2D and 3D patch tokens have a significant contribution to the final performance. Specifically, the 2D patch tokens capture rich texture details of the environment, which are particularly beneficial in scenes with sparse structural cues, while the 3D patch tokens encode geometric structures, offering stronger robustness under illumination and weather variations. In addition, the 2D cls and register tokens provide semantic and category-level information that produces an additional slight improvement of the recall.

Effect of 3D Pose Injection. Beyond token selection and aggregation strategies, we also investigated the explicit injection of 3D pose information. Specifically, we experimented with a pose projector using an MLP with Fourier position embeddings to inject relative pose information for sequence matching.




Fig. 6: t-SNE visualization comparing our UniPR-3D descriptor against CaseVPR across different datasets.

However, as shown in Tab. 4, we observed that this design does not improve retrieval performance and even slightly degrades the results. This suggests that our geometry-aware 3D tokens already implicitly capture sufficient spatial relationships, making explicit pose injection redundant.

To further analyze the complementary roles of 2D and 3D tokens, we visualize their relative importance in Fig. 5. The figure illustrates the weights of tokens not assigned to the “dustbin,” revealing the regions that the network deems informative. We can see that the network effectively discards uninformative regions such as skies, roads, and dynamic objects. Moreover, when comparing the two heatmaps, we observe that 2D features primarily focus on texture-rich areas, such as posters, kiosks, and bicycles, whereas 3D features attend more to geometric and structural elements like walls and buildings, reflecting a stronger spatial understanding.

Effect of sequence length. We evaluate the impact of the sequence length to assess the adaptability of our multi-frame feature aggregation module to varying input sizes. The results are presented in Tab. 5. Although UniPR-3D is trained only with sequences of length 5, it generalizes effectively to sequences of different lengths at test time, consistently achieving strong performance across all settings. Moreover, recall consistently improves as more images are aggregated, even beyond the original training length, showing that our aggregation design scales robustly to arbitrary numbers of frames.

4.4 Conclusion

In this paper, we propose UniPR-3D, the first 3D token-based VPR method that supports universal scene recognition and outperforms previous baselines by a substantial margin. Moreover, our framework is capable of handling both frame-to-frame and sequence-level retrievals. We design 2D and 3D feature aggregation strategies tailored to the characteristics of different types of 2D and

3D tokens. For sequence-level retrieval, we further propose a multi-frame feature aggregation method for variable-length sequence matching. The final descriptor is constructed by combining the 2D and 3D descriptors. Our extensive experiments demonstrate the effectiveness of the proposed method, achieving strong performance in both single frame and sequence matching. This work provides new insights for the VPR community, highlighting the potential of transitioning from conventional 2D token-based descriptors to 3D token-based representations.

Acknowledgement This work was supported by National Key R&D Program of China (Grant No.2024YFB4708900). It was also supported in part by the Natural Science Foundation of China under Grant 62225309, U24A20278, 62361166632.

References

1. Ali-bey, A., Chaib-draa, B., Giguere, P.: Gsv-cities: Toward appropriate supervised visual place recognition. *Neurocomputing* **513**, 194–203 (2022) [8](#)
2. Ali-Bey, A., Chaib-Draa, B., Giguere, P.: Mixvpr: Feature mixing for visual place recognition. In: *Proceedings of the IEEE/CVF winter conference on applications of computer vision*. pp. 2998–3007 (2023) [2](#), [4](#), [9](#), [10](#)
3. Ali-Bey, A., Chaib-Draa, B., Giguere, P.: Boq: A place is worth a bag of learnable queries. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 17794–17803 (2024) [10](#)
4. Arandjelovic, R., Gronat, P., Torii, A., Pajdla, T., Sivic, J.: Netvlad: Cnn architecture for weakly supervised place recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 5297–5307 (2016) [2](#), [3](#), [4](#), [10](#)
5. Arroyo, R., Alcantarilla, P.F., Bergasa, L.M., Romera, E.: Towards life-long visual localization using an efficient matching of binary sequences from images. In: *2015 IEEE international conference on robotics and automation (ICRA)*. pp. 6328–6335. IEEE (2015) [5](#)
6. Barbarani, G., Mostafa, M., Bayramov, H., Trivigno, G., Berton, G., Masone, C., Caputo, B.: Are local features all you need for cross-domain visual place recognition? In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6155–6165 (2023) [9](#)
7. Berton, G., Masone, C.: Megaloc: One retrieval to place them all. In: *Proceedings of the Computer Vision and Pattern Recognition Conference Workshops*. pp. 2861–2867 (2025) [2](#), [10](#)
8. Berton, G., Masone, C., Caputo, B.: Rethinking visual geo-localization for large-scale applications. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4878–4888 (2022) [9](#)
9. Berton, G., Masone, C., Caputo, B.: Rethinking visual geo-localization for large-scale applications. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4878–4888 (2022) [10](#)
10. Berton, G., Trivigno, G., Caputo, B., Masone, C.: Eigenplaces: Training viewpoint robust models for visual place recognition. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 11080–11090 (2023) [10](#)

11. Berton, G., Trivigno, G., Caputo, B., Masone, C.: Jist: Joint image and sequence training for sequential visual place recognition. *IEEE Robotics and Automation Letters* **9**(2), 1310–1317 (2023) [10](#), [11](#)
12. Bonneel, N., Digne, J.: A survey of optimal transport for computer graphics and computer vision. In: *Computer Graphics Forum*. vol. 42, pp. 439–460. Wiley Online Library (2023) [4](#)
13. Campos, C., Elvira, R., Rodríguez, J.J.G., Montiel, J.M., Tardós, J.D.: Orb-slam3: An accurate open-source library for visual, visual-inertial, and multimap slam. *IEEE Transactions on Robotics* **37**(6), 1874–1890 (2021) [1](#)
14. Chen, K., Liang, D., Zhou, X., Ding, Y., Liu, X., Wan, P., Bai, X.: Out of sight but not out of mind: Hybrid memory for dynamic video world models. *arXiv preprint arXiv:2603.25716* (2026) [4](#)
15. Chen, Z., Liu, L., Sa, I., Ge, Z., Chli, M.: Learning context flexible attention model for long-term visual place recognition. *IEEE Robotics and Automation Letters* **3**(4), 4015–4022 (2018) [9](#)
16. Cheng, Z., Zhang, X., Di, D., Wei, C., Li, H., Yang, X.: Moca: Modeling object consistency for 3d camera control in video generation. In: *The Fourteenth International Conference on Learning Representations* [3](#)
17. Cuturi, M.: Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems* **26** (2013) [7](#)
18. Dao, T.: Flashattention-2: Faster attention with better parallelism and work partitioning. *arXiv preprint arXiv:2307.08691* (2023) [9](#)
19. Deng, T., Liu, S., Wang, X., Liu, Y., Wang, D., Chen, W.: Prosgnerf: Progressive dynamic neural scene graph with frequency modulated auto-encoder in urban scenes. *arXiv preprint arXiv:2312.09076* (2023) [3](#)
20. Deng, T., Pan, Y., Yuan, S., Li, D., Wang, C., Li, M., Chen, L., Xie, L., Wang, D., Wang, J., Civera, J., Wang, H., Chen, W.: What is the best 3d scene representation for robotics? from geometric to foundation models. *arXiv preprint arXiv:2512.03422* (2025) [3](#)
21. Deng, T., Shen, G., Chen, X., Yuan, S., Shen, H., Peng, G., Wu, Z., Wang, J., Xie, L., Wang, D., Wang, H., Chen, W.: Mcn-slam: Multi-agent collaborative neural slam with hybrid implicit neural scene representation. *arXiv preprint arXiv:2506.18678* (2025) [3](#)
22. Deng, T., Shen, G., Qin, T., Wang, J., Zhao, W., Wang, J., Wang, D., Chen, W.: Plgslam: Progressive neural scene representation with local to global bundle adjustment. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 19657–19666 (2024) [1](#)
23. Deng, T., Shen, G., Xun, C., Yuan, S., Jin, T., Shen, H., Wang, Y., Wang, J., Wang, H., Wang, D., et al.: Mne-slam: Multi-agent neural slam for mobile robots. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 1485–1494 (2025) [1](#)
24. Deng, T., Wang, Y., Xie, H., Wang, H., Guo, R., Wang, J., Wang, D., Chen, W.: Neslam: Neural implicit mapping and self-supervised feature tracking with depth completion and denoising. *IEEE Transactions on Automation Science and Engineering* pp. 1–1 (2025). <https://doi.org/10.1109/TASE.2025.3541064> [3](#)
25. Deng, T., Xie, H., Wang, J., Chen, W.: Long-term visual simultaneous localization and mapping: Using a bayesian persistence filter-based global map prediction. *IEEE Robotics & Automation Magazine* **30**(1), 36–49 (2023) [4](#)
26. Facil, J.M., Olid, D., Montesano, L., Civera, J.: Condition-invariant multi-view place recognition. *arXiv preprint arXiv:1902.09516* (2019) [2](#)

27. Fang, X., Gao, J., Wang, Z., Chen, Z., Ren, X., Lyu, J., Ren, Q., Yang, Z., Yang, X., Yan, Y., Lyu, C.: Dens3r: A foundation model for 3d geometry prediction. arXiv preprint arXiv:2507.16290 (2025) [4](#)
28. Gálvez-López, D., Tardos, J.D.: Bags of binary words for fast place recognition in image sequences. *IEEE Transactions on robotics* **28**(5), 1188–1197 (2012) [1](#)
29. Gálvez-López, D., Tardos, J.D.: Bags of binary words for fast place recognition in image sequences. *IEEE Transactions on robotics* **28**(5), 1188–1197 (2012) [3](#)
30. Gao, J., Wang, Z., Fang, X., Ren, X., Chen, Z., Liu, S., Cheng, Y., Lyu, J., Yang, X., Yan, Y.: More: 3d visual geometry reconstruction meets mixture-of-experts. arXiv preprint arXiv:2510.27234 (2025) [4](#)
31. Garg, S., Fischer, T., Milford, M.: Where is your place, visual place recognition? arXiv preprint arXiv:2103.06443 (2021) [1](#)
32. Garg, S., Harwood, B., Anand, G., Milford, M.: Delta descriptors: Change-based place representation for robust visual localization. *IEEE Robotics and Automation Letters* **5**(4), 5120–5127 (2020) [10](#), [11](#)
33. Garg, S., Milford, M.: Seqnet: Learning descriptors for sequence-based hierarchical place recognition. *IEEE Robotics and Automation Letters* **6**(3), 4305–4312 (2021) [5](#), [10](#), [11](#)
34. Garg, S., Vankadari, M., Milford, M.: Seqmatchnet: Contrastive learning with sequence matching for place recognition & relocalization. In: *Conference on Robot Learning*. pp. 429–443. PMLR (2022) [4](#), [8](#), [9](#), [10](#), [11](#)
35. Gong, Z., Li, X., Tosi, F., Zhang, Y., Mattoccia, S., Wu, J., Poggi, M.: Dino-slam: Dino-informed rgb-d slam for neural implicit and explicit representations. arXiv preprint arXiv:2507.19474 (2025) [3](#)
36. Hausler, S., Garg, S., Xu, M., Milford, M., Fischer, T.: Patch-netvlad: Multi-scale fusion of locally-global descriptors for place recognition. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 14141–14152 (2021) [2](#)
37. Hussaini, S., Fischer, T., Milford, M.: Improving visual place recognition with sequence-matching receptiveness prediction. arXiv preprint arXiv:2503.06840 (2025) [5](#)
38. Izquierdo, S., Civera, J.: Close, but not there: Boosting geographic distance sensitivity in visual place recognition. In: *European Conference on Computer Vision*. pp. 240–257. Springer (2024) [2](#)
39. Izquierdo, S., Civera, J.: Optimal transport aggregation for visual place recognition. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 17658–17668 (2024) [2](#), [3](#), [4](#), [5](#), [7](#), [8](#), [9](#), [10](#)
40. Jégou, H., Douze, M., Schmid, C., Pérez, P.: Aggregating local descriptors into a compact image representation. In: *2010 IEEE computer society conference on computer vision and pattern recognition*. pp. 3304–3311. IEEE (2010) [2](#), [3](#), [4](#)
41. Keetha, N., Mishra, A., Karhade, J., Jatavallabhula, K.M., Scherer, S., Krishna, M., Garg, S.: Anyloc: Towards universal visual place recognition. *IEEE Robotics and Automation Letters* **9**(2), 1286–1293 (2023) [2](#), [3](#), [5](#), [10](#)
42. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: *ICLR* (2015) [8](#)
43. Li, H., Peng, G., Zhang, J., Wen, M., Ma, Y., Wang, D.: Casevpr: Correlation-aware sequential embedding for sequence-to-frame visual place recognition. *IEEE Robotics and Automation Letters* (2025) [5](#), [9](#), [10](#), [11](#), [12](#)
44. Liu, M., Liu, J., Yang, M.Y., Jiang, C., Li, J., Zhang, Y., Wang, H., Nex, F., Cheng, H.: Streamvlo: Streaming visual-lidar odometry with cumulative drift com-

- pensation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 39086–39097 (2026) [1](#)
45. Liu, M., Zhang, D., Liu, J., Cui, J., Xie, H., Chen, G., Ye, H., Yang, M.Y., Nex, F., Cheng, H.: Driveva: Video action models are zero-shot drivers. arXiv preprint arXiv:2604.04198 (2026) [2](#)
 46. Liu, Q., Chen, N., Liu, Z., Wang, H.: Toward learning-based visuomotor navigation with neural radiance fields. IEEE Transactions on Industrial Informatics **20**(6), 8907–8916 (2024). <https://doi.org/10.1109/TII.2024.3378829> [3](#)
 47. Liu, Q., Xin, H., Liu, Z., Wang, H.: Integrating neural radiance fields end-to-end for cognitive visuomotor navigation. IEEE Transactions on Pattern Analysis and Machine Intelligence **46**(12), 11200–11215 (2024). <https://doi.org/10.1109/TPAMI.2024.3455252> [3](#)
 48. Liu, S., Deng, T., Zhou, H., Li, L., Wang, H., Wang, D., Li, M.: Mg-slam: Structure gaussian splatting slam with manhattan world hypothesis. IEEE Transactions on Automation Science and Engineering (2025) [3](#)
 49. Lowry, S., Sünderhauf, N., Newman, P., Leonard, J.J., Cox, D., Corke, P., Milford, M.J.: Visual place recognition: A survey. IEEE Transactions on Robotics **32**(1), 1–19 (2015) [1](#)
 50. Lu, F., Jin, T., Lan, X., Zhang, L., Liu, Y., Wang, Y., Yuan, C.: Selavpr++: Towards seamless adaptation of foundation models for efficient place recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence **48**(3), 2731–2748 (2026) [3](#)
 51. Lu, F., Lan, X., Zhang, L., Jiang, D., Wang, Y., Yuan, C.: Cricavpr: Cross-image correlation-aware representation learning for visual place recognition. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16772–16782 (2024) [5](#), [10](#)
 52. Lu, F., Zhang, L., Lan, X., Dong, S., Wang, Y., Yuan, C.: Towards seamless adaptation of pre-trained models for visual place recognition. In: The Twelfth International Conference on Learning Representations (2024) [3](#)
 53. Lv, J., Ju, X., Wang, C.: Neural network-based nonconservative predefined-time backstepping control for uncertain strict-feedback nonlinear systems. IEEE Transactions on Neural Networks and Learning Systems (2023) [3](#)
 54. Lv, J., Wang, C., Xie, L.: Adaptive distributed observer design for nonlinear multiagent systems. Automatica **183**, 112625 (2026) [3](#)
 55. Ma, Y., Feng, K., Hu, Z., Wang, X., Wang, Y., Zheng, M., He, X., Zhu, C., Liu, H., He, Y., et al.: Controllable video generation: A survey. arXiv preprint arXiv:2507.16869 (2025) [3](#)
 56. Ma, Y., Feng, K., Zhang, X., Liu, H., Zhang, D.J., Xing, J., Zhang, Y., Yang, A., Wang, Z., Chen, Q.: Follow-your-creation: Empowering 4d creation through video inpainting. arXiv preprint arXiv:2506.04590 (2025) [3](#)
 57. Ma, Y., He, Y., Cun, X., Wang, X., Chen, S., Li, X., Chen, Q.: Follow your pose: Pose-guided text-to-video generation using pose-free videos. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 38, pp. 4117–4125 (2024) [3](#)
 58. Maddern, W., Pascoe, G., Linegar, C., Newman, P.: 1 year, 1000 km: The oxford robotcar dataset. The International Journal of Robotics Research **36**(1), 3–15 (2017) [9](#), [11](#), [13](#)
 59. Masone, C., Caputo, B.: A survey on deep visual place recognition. IEEE Access **9**, 19516–19547 (2021) [1](#)
 60. Mereu, R., Trivigno, G., Berton, G., Masone, C., Caputo, B.: Learning sequential descriptors for sequence-based visual place recognition. IEEE Robotics and Automation Letters **7**(4), 10383–10390 (2022) [4](#), [5](#), [10](#), [11](#)

61. Milford, M., Fischer, T.: Going places: Place recognition in artificial and natural systems. *Annual Review of Control, Robotics, and Autonomous Systems* **9** (2025) [1](#)
62. Milford, M.J., Wyeth, G.F.: Seqslam: Visual route-based navigation for sunny summer days and stormy winter nights. In: 2012 IEEE international conference on robotics and automation. pp. 1643–1649. IEEE (2012) [3](#), [4](#), [10](#), [11](#)
63. Murillo, A.C., Singh, G., Kosecka, J., Guerrero, J.J.: Localization in urban environments using a panoramic gist descriptor. *IEEE Transactions on Robotics* **29**(1), 146–160 (2012) [3](#)
64. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193* (2023) [2](#), [3](#), [5](#)
65. Radenović, F., Tolias, G., Chum, O.: Fine-tuning cnn image retrieval with no human annotation. *IEEE transactions on pattern analysis and machine intelligence* **41**(7), 1655–1668 (2018) [2](#), [3](#), [4](#), [6](#), [10](#)
66. Revaud, J., Almazán, J., Rezende, R.S., Souza, C.R.d.: Learning with average precision: Training image retrieval with a listwise loss. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 5107–5116 (2019) [10](#)
67. Sarlin, P.E., Cadena, C., Siegwart, R., Dymczyk, M.: From coarse to fine: Robust hierarchical localization at large scale. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 12716–12725 (2019) [1](#)
68. Sarlin, P.E., DeTone, D., Malisiewicz, T., Rabinovich, A.: Superglue: Learning feature matching with graph neural networks. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 4938–4947 (2020) [7](#)
69. Schubert, S., Neubert, P., Garg, S., Milford, M., Fischer, T.: Visual Place Recognition: A Tutorial. *IEEE Robotics & Automation Magazine* (2023) [1](#)
70. Schubert, S., Neubert, P., Protzel, P.: Fast and memory efficient graph optimization via icm for visual place recognition. In: *Robotics: Science and Systems*. vol. 2 (2021) [4](#)
71. Sferrazza, D., Berton, G., Trivigno, G., Masone, C.: To match or not to match: Revisiting image matching for reliable visual place recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*. pp. 2874–2885 (June 2025) [1](#)
72. Shen, G., Deng, T., Wang, Y., Chen, Y., Shen, Y., Liu, J., Wang, J.: Grs-slam3r: Real-time dense slam with gated recurrent state. *arXiv preprint arXiv:2509.23737* (2025) [3](#)
73. Shen, H., Chen, X., Hui, Y., Wu, Z., Wang, W., Lyu, Q., Deng, T., Wang, D.: Unilgl: Learning uniform place recognition for fov-limited/panoramic lidar global localization. *arXiv preprint arXiv:2507.12194* (2025) [3](#)
74. Sivic, Z., Zisserman, A.: Video google: A text retrieval approach to object matching in videos. In: *Proceedings ninth IEEE international conference on computer vision*. pp. 1470–1477. IEEE (2003) [4](#)
75. Song, W., Chen, J., Ding, P., Zhao, H., Zhao, W., Zhong, Z., Ge, Z., Ma, J., Li, H.: Accelerating vision-language-action model integrated with action chunking via parallel decoding. *arXiv preprint arXiv:2503.02310* (2025) [2](#)
76. Song, W., Zhou, Z., Zhao, H., Chen, J., Ding, P., Yan, H., Huang, Y., Tang, F., Wang, D., Li, H.: Reconvla: Reconstructive vision-language-action model as effective robot perceiver. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 40, pp. 18549–18557 (2026) [2](#)

77. Sünderhauf, N., Neubert, P., Protzel, P.: Are we there yet? challenging seqslam on a 3000 km journey across all four seasons. In: Proc. of workshop on long-term autonomy, IEEE international conference on robotics and automation (ICRA). p. 2013. Citeseer (2013) [9](#)
78. Sünderhauf, N., Protzel, P.: Brief-gist-closing the loop by simple means. In: 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems. pp. 1234–1241. IEEE (2011) [3](#)
79. Taira, H., Okutomi, M., Sattler, T., Cimpoi, M., Pollefeys, M., Sivic, J., Pajdla, T., Torii, A.: Inloc: Indoor visual localization with dense matching and view synthesis. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7199–7209 (2018) [1](#)
80. Torii, A., Arandjelovic, R., Sivic, J., Okutomi, M., Pajdla, T.: 24/7 place recognition by view synthesis. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1808–1817 (2015) [9](#)
81. Torii, A., Sivic, J., Pajdla, T., Okutomi, M.: Visual place recognition with repetitive structures. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 883–890 (2013) [9](#)
82. Tosi, F., Zhang, Y., Gong, Z., Mattoccia, S., Oswald, M.R., Sandstrom, E., Poggi, M.: How nerfs and 3d gaussian splatting are reshaping slam: a survey. IEEE Transactions on Robotics (2026) [1](#)
83. Tu, S., Zhou, X., Liang, D., Jiang, X., Zhang, Y., Li, X., Bai, X.: The role of world models in shaping autonomous driving: A comprehensive survey. arXiv preprint arXiv:2502.10498 (2025) [4](#)
84. Wang, J., Chen, M., Karaev, N., Vedaldi, A., Rupprecht, C., Novotny, D.: Vggt: Visual geometry grounded transformer. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 5294–5306 (2025) [2](#), [6](#), [9](#)
85. Wang, R., Shen, Y., Zuo, W., Zhou, S., Zheng, N.: Transvpr: Transformer-based place recognition with multi-level attention aggregation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13648–13657 (2022) [2](#)
86. Wang, X., Han, X., Huang, W., Dong, D., Scott, M.R.: Multi-similarity loss with general pair weighting for deep metric learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5022–5030 (2019) [8](#)
87. Wang, Y., Fang, Z., Zhao, L., Chen, W.: Learning to tune like an expert: Interpretable and scene-aware navigation via mllm reasoning and cvae-based adaptation. arXiv preprint arXiv:2507.11001 (2025) [3](#)
88. Wang, Y., Zhao, W., Cao, C., Deng, T., Wang, J., Chen, W.: Sfpnet: Sparse focal point network for semantic segmentation on general lidar point clouds. In: European Conference on Computer Vision. pp. 403–421. Springer (2024) [3](#)
89. Warburg, F., Hauberg, S., Lopez-Antequera, M., Gargallo, P., Kuang, Y., Civera, J.: Mapillary street-level sequences: A dataset for lifelong place recognition. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2626–2635 (2020) [8](#), [9](#)
90. Wu, X., Liang, D., Feng, T., Xia, K., Zhang, Y., Li, X., Tan, X., Bai, X.: Generation models know space: Unleashing implicit 3d priors for scene understanding. arXiv preprint arXiv:2603.19235 (2026) [4](#)
91. Yang, C., Pan, H., Wang, J.: Stgcnformer: Spatio-temporal dual-stream graph convolutional networks and transformers for traffic forecasting. IEEE Transactions on Vehicular Technology **74**(10), 15254–15263 (2025). <https://doi.org/10.1109/TVT.2025.3572622> [3](#)

92. Yang, C., Pan, H., Wang, J., Hong, Y.: Trajdiff: Trajectory prediction with diffusion probabilistic models. *IEEE Transactions on Image Processing* pp. 1–14 (2025). <https://doi.org/10.1109/TIP.2025.3640001> 3
93. Zhang, X., Wang, L., Su, Y.: Visual place recognition: A survey from deep learning perspective. *Pattern Recognition* **113**, 107760 (2021) 1
94. Zhao, J., Zhang, F., Cai, Y., Tian, G., Mu, W., Ye, C., Feng, T.: Learning sequence descriptor based on spatio-temporal attention for visual place recognition. *IEEE Robotics and Automation Letters* **9**(3), 2351–2358 (2024). <https://doi.org/10.1109/LRA.2024.3354627> 10, 11
95. Zhou, Q., Maximov, M., Litany, O., Leal-Taixé, L.: The nerfect match: Exploring nerf features for visual localization. In: *European Conference on Computer Vision*. pp. 108–127. Springer (2024) 1
96. Zhu, S., Yang, L., Chen, C., Shah, M., Shen, X., Wang, H.: R2former: Unified retrieval and reranking transformer for place recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 19370–19380 (2023) 2