

LCPR: A Multi-Scale Attention-Based LiDAR-Camera Fusion Network for Place Recognition

Zijie Zhou¹, Jingyi Xu², Guangming Xiong¹, and Junyi Ma^{*}

Abstract—Place recognition is one of the most crucial modules for autonomous vehicles to identify places that were previously visited in GPS-invalid environments. Sensor fusion is considered an effective method to overcome the weaknesses of individual sensors. In recent years, multimodal place recognition fusing information from multiple sensors has gathered increasing attention. However, most existing multimodal place recognition methods only use limited field-of-view camera images, which leads to an imbalance between features from different modalities and limits the effectiveness of sensor fusion. In this paper, we present a novel neural network named LCPR for robust multimodal place recognition, which fuses LiDAR point clouds with multi-view RGB images to generate discriminative and yaw-rotation invariant representations of the environment. A multi-scale attention-based fusion module is proposed to fully exploit the panoramic views from different modalities of the environment and their correlations. We evaluate our method on the nuScenes dataset, and the experimental results show that our method can effectively utilize multi-view camera and LiDAR data to improve the place recognition performance while maintaining strong robustness to viewpoint changes. Our open-source code and pre-trained models are available at <https://github.com/ZhouZijie77/LCPR>.

I. INTRODUCTION

Given an observation from sensors such as LiDARs and cameras, the place recognition task needs to find where the observation was previously collected, which is usually used as the first step of loop closure detection in SLAM systems or global localization [1], [2]. Cameras are the most commonly used sensors in various autonomous driving systems, due to their low cost and ability to provide rich semantic information. Vision-based place recognition has been extensively studied, including early approaches based on hand-crafted descriptor [3], [4], [5] and more recent deep learning-based methods [6], [7], [8]. However, features extracted from images are prone to be affected by weather, season, and lighting conditions, especially in outdoor environments, significantly diminishing the localization performance. Compared to cameras, LiDARs can capture 3D structural information of the environment and are not sensitive to the factors mentioned above, which has been demonstrated by many previous LiDAR-based place recognition works [9], [10], [11], [2]. However, the performance of these LiDAR-based approaches

This work was supported by the National Natural Science Foundation of China under Grant 52372404.

¹Zijie Zhou and Guangming Xiong are with Beijing Institute of Technology, Beijing, 100081, China 3120220445@bit.edu.cn

²Jingyi Xu is with Shanghai Jiao Tong University, 200240, Shanghai 3220200359@bit.edu.cn

^{*}Junyi Ma (corresponding author) is with Beijing Institute of Technology and HAOMO.AI Technology Co., Ltd junyi.ma@bit.edu.cn




Fig. 1. Multi-view images can provide full perspective information of the environment as LiDAR does. LCPR leverages multi-view RGB images and range image from LiDAR as inputs, and utilizes transformer attention to identify the correspondence between two modalities. Localization can be achieved by searching for the nearest neighbor in the database.

still suffers from viewpoint changes, sparsity of point clouds, and lack of texture information.

The combination of LiDARs and cameras can surpass the individual limitations within each modality, which is widely used in autonomous driving perception tasks, such as object detection, semantic segmentation, and depth prediction. In the field of place recognition, some multimodal fusion-based approaches have been proposed, such as AdaFusion [12], MinkLoc++ [13], PIC-Net [14], proving that multimodal fusion can enhance the performance of recognition. However, the limited field of view (FOV) and viewpoint of the image data used by these methods introduce an inherent flaw when the observation angles are altered by mounting positions or driving directions, resulting in an imbalance between features from different modalities for sensor fusion.

In this paper, inspired by the successful application of multimodal fusion in 3D object detection and map segmentation tasks [15], [16], we propose a multi-scale attention-based network LCPR, as shown in Fig. 1, to fuse multimodal sensor information for place recognition. We focus on extracting discriminative and yaw-rotation invariant global descriptors from both the 3D LiDAR point clouds and RGB images from multi-view cameras for fast query-database matching. The environmental features from the two modalities are fused at multiple scales by the proposed Vertically Compressed Transformer Fusion (VCTF) module to fully exploit the panoramic features from LiDARs and cameras as well as their intra- and inter-correlations.

In summary, the main contributions of our work are as follows:

- We propose LCPR, a multi-scale attention-based net-

work designed to generate discriminative and yaw-rotation invariant multimodal descriptors for place recognition, using 3D LiDAR point clouds and 2D RGB images from multi-view cameras as joint inputs.

- A novel multimodal fusion framework and a panoramic fusion module are proposed, utilizing self-attention mechanism of transformers to identify and exploit the correspondences between panoramic features from different modalities.
- We evaluate our proposed LCPR on the nuScenes dataset [17], and the experimental results validate that our method outperforms the current state-of-the-art methods including both the unimodal and multimodal ones.

II. RELATED WORK

The performance of place recognition is strongly related to the adaptability of various types of sensors. Therefore, we make a brief review of place recognition methods based on different sensor setups.

A. Vision-based Place Recognition

Vision-based place recognition (VPR) is usually regarded as an image retrieval problem, where the database image most similar to the current query image is recalled. Traditional image retrieval methods are proposed to extract hand-crafted local feature descriptors. For example, Andreasson [18] uses the simplified SIFT [19] for topological localization for mobile robots. After that, Zhang *et al.* [20] introduce a novel robust motion estimation technique to deal with a larger percentage of mismatches based on SIFT features. In [21], Jégou *et al.* introduce the VLAD descriptor for compact image representation, which shows effectiveness and efficiency in large scale image retrieval tasks. With the rapid development of deep learning, researchers are now more interested in using neural networks to extract high-level place features instead of the hand-crafted ways. In [22], Chen *et al.* first introduce CNNs into the place recognition system. Arandjelović *et al.* [6] propose NetVLAD, a feature aggregation layer which can be easily plugged into any CNN architecture and trained in an end-to-end manner. After that, Radenović *et al.* [23] propose to use 3D reconstructed models to guide the selection of the training data, and propose a novel GeM pooling layer. Based on NetVLAD, Hausler *et al.* [7] propose Patch-NetVLAD which combines the strengths of local and global features. Arcanjo *et al.* [24] design a lightweight network and a voting mechanism to achieve respectable results with low memory usage.

B. LiDAR-based Place Recognition

The active-measurement nature of LiDAR sensors makes LiDAR-based place recognition (LPR) invariant to conditional changes such as illumination, weather and seasons. M2DP proposed by He *et al.* [25] first projects the point cloud to multiple planes and then calculates signatures by the distributions of the projected points on different planes. Scan Context proposed by Kim *et al.* [9] divides the point

cloud into sectors in polar coordinates and encodes the maximum height of laser points in each sector to generate the global descriptor matrix. The successful application of deep learning in processing 3D point clouds has led to more and more learning-based LPR methods. PointNetVLAD by Uy *et al.* [10] builds upon the basic functions of PointNet [26] and NetVLAD [6] for large-scale place recognition. After that, LPDNet proposed by Liu *et al.* [11] combines learning-based and hand-crafted local features to enhance the vanilla PointNet architecture. PCAN by Zhang [27] utilizes the attention mechanism to enforce the network to pay more attention to the task-relevant regions of the environment. More recently, different input forms of point clouds have also been studied and applied to LPR tasks. Komorowski *et al.* [28] utilize sparse voxelized representation of point clouds as the input to the network. OverlapTransformer proposed by Ma *et al.* [2] projects point clouds into range images to tackle the sparsity and disorder of point clouds and introduces the transformer network [29] to achieve high efficiency and yaw-rotation invariance.

C. Fusion-based Place Recognition

Recently, multi-sensor fusion has aroused great interest in the place recognition community. The combination of different sensors can overcome the disadvantages of the unimodal methods. To collaborate point clouds and images, PIC-Net [14] introduces global channel attention to fuse features from different modalities. The Spatial Attention VLAD layer is proposed to select the discriminative points and pixels. Xie *et al.* [30] use PointNet with a trimmed VLAD layer and ResNet-50 [31] to extract LiDAR and visual descriptors separately, which are further fused by FC layers. CORAL [32] by Pan *et al.* constructs an elevation map from a point cloud and projects the image feature to the elevation map to semantically “colorize” the structural features and generate a bi-modal descriptor in the bird’s eye view (BEV) frame. MinkLoc++ proposed by Komorowski *et al.* [13] uses 3D sparse convolutions to extract spatial features from point clouds, and uses pre-trained ResNet blocks [31] for image features. The two descriptors are then concatenated along the channel dimension to form a multimodal descriptor. MMDF by Yu *et al.* [33] uses image semantic features to enhance point cloud features. Pointwise image features are first extracted using the method described in [34], then point cloud features and pointwise image features are aggregated by a convolution layer. To better exploit the discriminativeness of the descriptors from different modalities, Lai *et al.* [12] (AdaFusion) design an attention branch in their network to weight modalities adaptively.

Based on the aforementioned review, none of these methods fully exploit the information from different sensors to find the explicit spatial correspondence and inner feature correlations between 2D and 3D data simultaneously. In addition, the limited FOV of a single monocular camera can lead to false recognition as the viewpoint changes. Therefore, in this paper, we propose a multi-scale attention-based place recognition method that takes multi-view RGB images and

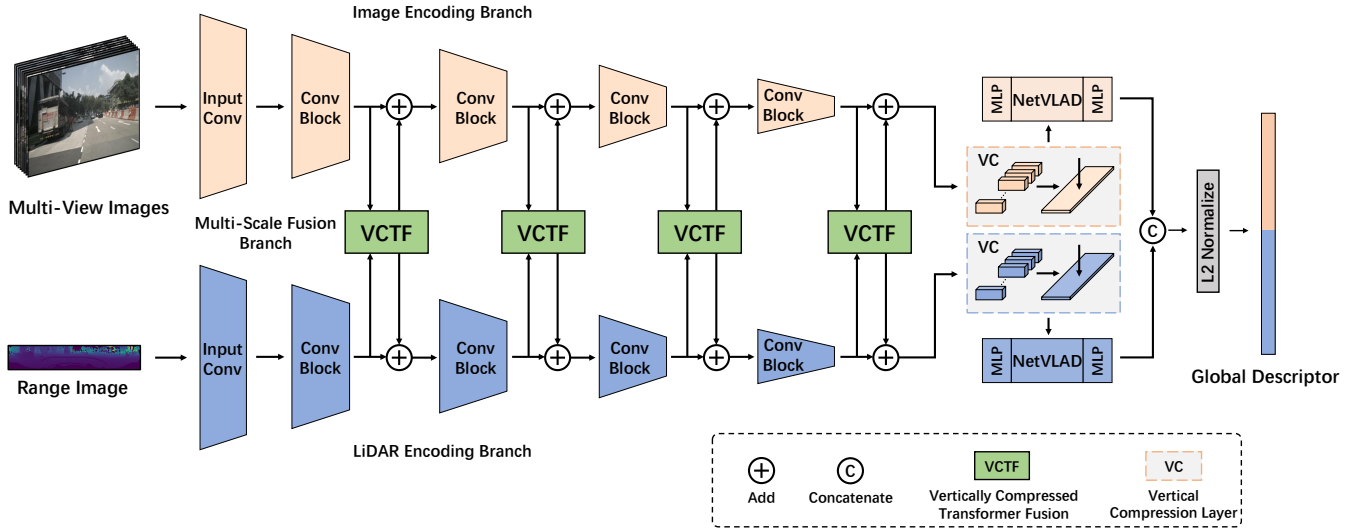


Fig. 2. The overall architecture of LCPR. Multi-view RGB images and one range image are fed into the sibling image encoding (IE) branch and the LiDAR encoding (LE) branch to obtain intermediate features at different resolutions. A set of Vertically Compressed Transformer Fusion (VCTF) modules are employed to fuse these intermediate features at multiple scales. The outputs of the IE and LE branches then pass through the Vertical Compression (VC) layers to obtain the denser panoramic features, which are aggregated using NetVLAD-MLPs combos to generate sub-descriptors. Finally, the global multimodal descriptor is generated by the concatenation of the sub-descriptors, which is further used as a query or reference in the database.

range images from LiDAR scans as input to cover the panoramic view of the environment. We utilize self-attention mechanism to extract the global panoramic context between and within different modalities to generate discriminative and yaw-rotation invariant descriptors.

III. OUR APPROACH

A. Overall Architecture

The overall architecture of our proposed LCPR is depicted in Fig. 2. To address the information imbalance between different modalities of limited FOV cameras and the LiDAR mounted on an autonomous vehicle, our method takes multi-view RGB images and the range image projected from point clouds as input and directly outputs the environmental global descriptor. Inspired by [16], our LCPR mainly consists of an image encoding (IE) branch, a LiDAR encoding (LE) branch, a multi-scale fusion branch, and NetVLAD-MLPs combos [6]. Given the effectiveness of residual structures in processing image features [31], [14], [30], we propose to use residual blocks [31] as the feature extractors in both the IE and LE branches.

To tackle the sparsity and disorder of the point clouds, our method first produces range images using spherical projection on LiDAR point clouds. In particular, the correspondence between a LiDAR point $\mathbf{p} = (x, y, z)^T \in \mathbb{R}^3$ from the point cloud \mathcal{P} and the pixel coordinates $(u, v)^T$ of its corresponding range image \mathcal{R} can be formulated as

$$\begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} \frac{1}{2} [1 - \arctan(y, x)\pi^{-1}] w \\ [1 - (\arcsin(zr^{-1}) + f_{\text{up}}) f^{-1}] h \end{pmatrix}, \quad (1)$$

where $r = \|\mathbf{p}\|_2$ is the range measurement of the LiDAR point, (h, w) represents the height and width of the range image, and $f = f_{\text{up}} + f_{\text{down}}$ is the vertical FOV of the sensor.

The multi-view RGB images and the LiDAR range images are fed into the two encoding branches, the IE branch and the LE branch, respectively. Considering that the low-level features of the network can capture rich texture information, while the high-level features can represent semantic information, we propose and leverage a set of Vertically Compressed Transformer Fusion (VCTF) modules to fuse the intermediate features from both the IE branch and the LE branch at multiple resolutions. The VCTF module employs the self-attention mechanism of transformers [29] to fully incorporate the global panoramic context between and within different modalities. Subsequently, the output of the VCTF module is sent back to the encoding branches, added to the original intermediate feature, and then fed to the next residual block. The outputs of the two encoding branches are each processed by the proposed Vertical Compression (VC) layers respectively, which compress the features in the vertical direction. The compressed features are further aggregated using NetVLAD layers with MLPs to yield two sub-descriptors. Finally, the sub-descriptors from two encoding branches are then concatenated to produce the multimodal global descriptor, which is then used for database creation and fast place retrieval.

B. Vertically Compressed Multimodal Fusion

It is argued that the height dimension contains less information than the width dimension in urban autonomous driving [35]. Therefore, we adopt a vertical compression operation along the height dimension of the intermediate features in the pipeline to reduce the computational cost as well as to achieve yaw-rotation invariance. Here we extend the ideas of the previous work [2] to the multimodal fusion framework.

The architecture of our proposed VCTF module is depicted

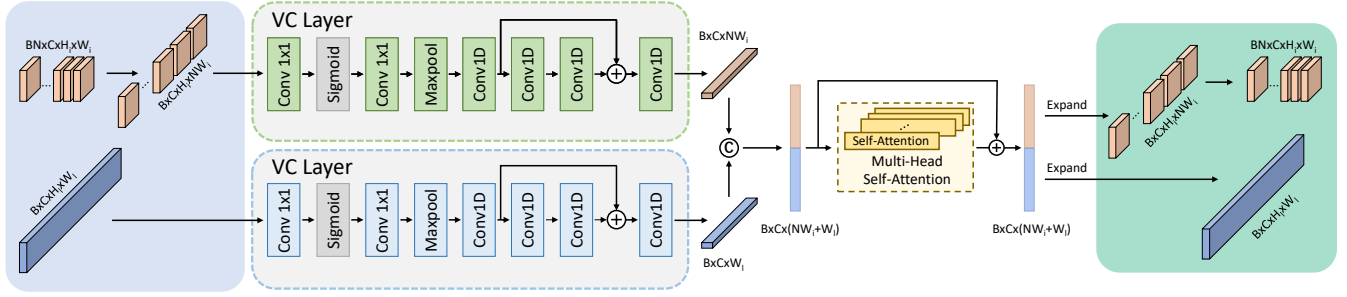


Fig. 3. The architecture of Vertically Compressed Transformer Fusion (VCTF) module. The intermediate feature volumes from the two sibling feature streams are first compressed in the vertical direction. The compressed sentence-like features are concatenated horizontally, and fed into the MHSA module for multimodal fusion. The fused feature is then split and expanded to be sent back to the original feature streams.

in Fig. 3. We denote the input feature volumes from the IE branch and the LE branch as $F_i^{in} \in \mathbb{R}^{B \times N \times C \times H_i \times W_i}$ and $F_l^{in} \in \mathbb{R}^{B \times C \times H_i \times W_i}$ respectively, where N is the number of cameras, and B is the batch size. F_i^{in} is first concatenated panoramically for feature alignment, leading to $F_i^p \in \mathbb{R}^{B \times C \times H_i \times N W_i}$. Then two VC layers are designed to compress F_i^p and F_l^{in} in the vertical direction, respectively. The VC layer consists of two 1×1 convolutions with Sigmoid nonlinearity, a MaxPooling layer for feature compression in the vertical direction, and 1-D convolutions with Batch Normalization and ReLU nonlinearity for feature refinement. The skip connection between the first and the third 1-D convolutions is employed to speed up the convergence.

The output feature volumes of the VC layers with sizes of $B \times C \times N W_i$ and $B \times C \times W_i$ are concatenated along the width dimension to generate a sentence-like feature $F_s \in \mathbb{R}^{B \times C \times (N W_i + W_i)}$. And then we feed F_s into a multi-head self-attention (MHSA) module [29] to extract the correlations between modalities and the global context within each modality simultaneously. The self-attention mechanism can be formulated as

$$\mathcal{A} = \text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V, \quad (2)$$

where Q , K , V represent the queries, keys, and values respectively, d_k is the dimension of keys.

Finally, the output feature volumes of the MHSA module are expanded along the height dimension, so the inputs and outputs of the VC layers have the same size for the fusion with the original intermediate features.

C. Yaw-Rotation Invariance

Our proposed LCPR is designed to generate yaw-rotation invariant global descriptors, ensuring the robustness of our method in viewpoint changes in real-world autonomous driving applications. In this section, we provide a comprehensive mathematical derivation of the yaw-rotation invariance achieved by our network. It has been reported in [10], [2], [36], [37] that the NetVLAD module is permutation-invariant. Hence, here we concentrate on exploring the yaw-rotation equivariance of the modules before NetVLAD in our LCPR.

1) *Residual Unit*: The residual unit in our IE branch and

LE branch can be defined as

$$\mathbf{y} = \sigma(W_s \mathbf{X} + W_2 \sigma(W_1 \mathbf{X})). \quad (3)$$

Here \mathbf{X} and \mathbf{y} are the input and output feature volumes, $\sigma(\cdot)$ denotes the ReLU function, W_1 and W_2 represent two convolution layers, and W_s represents a linear projection for matching dimensions.

Then we right-multiply the input \mathbf{X} with a matrix C to represent a column shift corresponding to a yaw-angle rotation of the vehicle. Since the convolution operation and ReLU function are both translation equivariant, we have

$$\begin{aligned} \mathbf{y}_s &= \sigma(W_s(\mathbf{X}C) + W_2 \sigma(W_1(\mathbf{X}C))) \\ &= \sigma(W_s \mathbf{X}C + W_2 \sigma(W_1 \mathbf{X})C) \\ &= \sigma(W_s \mathbf{X} + W_2 \sigma(W_1 \mathbf{X}))C \\ &= \mathbf{y}C. \end{aligned} \quad (4)$$

As can be seen, the residual unit is translation equivariant, resulting in the two yaw-rotation equivariant encoding branches.

2) *VCTF module*: The VCTF module mainly consists of two parts: the VC layer and the MHSA module. As described in Sec. III-B, the VC layer has a fully convolutional structure, thus possessing yaw-rotation equivariance.

As for the MHSA module, we denote the vertically compressed feature volumes extracted by two VC layers as $\mathbf{X}_i \in \mathbb{R}^{D \times N W_i}$ and $\mathbf{X}_l \in \mathbb{R}^{D \times W_i}$. To avoid ambiguity, here we use D to represent the feature dimension. Then the input of the MHSA module can be denoted as $\mathbf{X} = [\mathbf{X}_i \ \mathbf{X}_l]^T$, which represents the concatenation of the two features along the width dimension. With column shift matrices C_i and C_l , the input becomes

$$\begin{aligned} \mathbf{X}_s &= [\mathbf{X}_i C_i \ \mathbf{X}_l C_l]^T \\ &= \begin{bmatrix} C_i^T & \mathbf{0} \\ \mathbf{0} & C_l^T \end{bmatrix} [\mathbf{X}_i \ \mathbf{X}_l]^T \\ &= C \mathbf{X}. \end{aligned} \quad (5)$$

Here $C \in \mathbb{R}^{2D \times 2D}$ represents the transformation matrix of the input \mathbf{X} resulting from the column shifts of the features

TABLE I
STATISTICS OF OUR DATA ORGANIZATION

Location	Boston-Seaport	SG-OneNorth	SG-HollandVillage
N _{sample}	22103	8104	3427
N _{database}	9686	4796	2433
N _{query_train}	8407	-	-
N _{query_val}	1002	-	-
N _{query_test}	3008	3308	994
Time	Day	Day	Day+Night

N_{sample}, N_{database}, N_{query_train}, N_{query_val}, and N_{query_test} are the numbers of samples, database, training queries, validation queries, and testing queries respectively.

\mathbf{X}_i and \mathbf{X}_l . Then Eq. (2) becomes

$$\begin{aligned}
\mathcal{A}_s &= \text{Attention}(Q_s, K_s, V_s) = \text{softmax} \left(\frac{Q_s K_s^T}{\sqrt{d_k}} \right) V_s \\
&= \text{softmax} \left(\frac{(C\mathbf{X}W_Q)(W_K^T\mathbf{X}^T C^T)}{\sqrt{d_k}} \right) C\mathbf{X}W_V \\
&= C \text{softmax} \left(\frac{(\mathbf{X}W_Q)(W_K^T\mathbf{X}^T)}{\sqrt{d_k}} \right) C^T C\mathbf{X}W_V \\
&= C \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) V \\
&= [\mathbf{Y}_i C_i \quad \mathbf{Y}_l C_l]^T,
\end{aligned} \tag{6}$$

where \mathbf{Y}_i and \mathbf{Y}_l are the output feature volumes of MHSA corresponding to the inputs \mathbf{X}_i and \mathbf{X}_l , respectively.

As can be seen, the column shifts C_i and C_l are preserved from the input to the output, showing that the MHSA module is yaw-rotation equivariant. Consequently, the VCTF module is proven to be yaw-rotation equivariant, leading to the holistic LCPR yaw-rotation equivariant.

D. Network Training

We follow [6] and [38] to use the triplet margin loss to train the network. Specifically, for each training step, we construct a min-batch $(q, p^q, \{n_i^q\})$ containing a query, a positive sample, and several negative samples. Following [38], we define a positive sample as one that is located no more than 9 meters apart from the capture location of the query. Conversely, we define a sample as negative if it is at least 18 meters apart from the capture location of the query.

We denote $f(\cdot)$ as the mapping from the input to its global descriptor. Our goal is to minimize the distance between $f(q)$ and $f(p^q)$ while maximizing the distance between $f(q)$ and $f(n_i^q)$. It has been reported that the network is prone to overfitting in the image domain during training [13]. To mitigate this issue, for each training query, we select a positive sample with the smallest distance between $f(p_i^q)$ and $f(q)$ from the potential positive group to form the mini-batch. The negative samples are randomly selected from the potential negative group. In addition, to increase the training efficiency, we only retain negative samples that satisfy

$$d(f(q), f(n_i^q)) < \text{margin}. \tag{7}$$

Then, the loss function is given by

$$\mathcal{L}_T = \frac{1}{N_{\text{neg}}} \sum_{i=1}^{N_{\text{neg}}} [d(f(q), f(p^q)) - d(f(q), f(n_i^q)) + m]_+, \tag{8}$$

where $[\cdot]_+$ denotes the hinge loss, $d(\cdot)$ is the Euclidean distance, m is a constant giving the margin, and N_{neg} represents the number of the selected negative samples.

IV. EXPERIMENTS

A. Experimental Setup and Implementation Details

Since our method requires both multi-view RGB images with 360° horizontal FOV and LiDAR point clouds as input, we evaluate our method on the nuScenes dataset [17], which comprises 1000 driving scenes, each with a duration of 20 seconds, and carries the full autonomous vehicle sensor suite including six cameras, one LiDAR, and five radars. The dataset is collected from four different locations in Boston and Singapore: Boston Seaport (BS), SG-OneNorth (SON), SG-Queenstown (SQ), and SG-HollandVillage (SHV). Among them, the BS split is the largest, containing 550 scenes, and the others contain 215, 135, and 100 scenes respectively. Due to the limited number of samples and the monotonous environment of the SQ split, we only conduct experiments on the BS, SON, and SHV splits.

We organize the data following the previous work [38]. However, we introduce a modification by creating the *database set* from the entire data split to ensure that the database is up-to-date. The rest of the data is used as the query set and is divided into the *training query set*, the *validation query set*, and the *test query set*. For SON and SHV split, we only divide the data into the *database set* and the *test query set*. Finally, places in the query set that have no ground truth in the *database set* are removed. More details of our data organization can be found in Tab. I

For our proposed LCPR, we use LiDAR range images of size $1 \times 32 \times 1056$, and each input RGB image is resized from $3 \times 1600 \times 900$ to $3 \times 704 \times 256$ resolution for data alignment between the two modalities. We use the residual blocks from ResNet-18 [31] pre-trained on ImageNet [39] in our IE and LE branches. For the MHSA module, we set the number of heads $n_{\text{head}} = 4$ and the embedding dimension $d_k = d_v = d_{\text{model}}$. For the NetVLAD modules, we set the number of clusters $d_{\text{cluster}} = 32$, the output dimension $d_{\text{sub}} = 128$, generating two sub-descriptors with size of 1×128 . The final multimodal descriptor then has a size of 1×256 . For the network training, the ADAM optimizer is used to update the weights of the network. The initial learning rate is set to 1×10^{-5} and decays by a factor of 0.8 every 10 epochs. For each mini-batch, we set the number of negative samples $N_{\text{neg}} = 6$ and $m = 0.5$ for the triplet loss.

In the following experiments, we compare the performance of our LCPR with the state-of-the-art baselines including the vision-based method NetVLAD [6], the radar-based

TABLE II
EVALUATION OF PLACE RECOGNITION PERFORMANCE ON THE BS AND SON DATA SPLIT

Methods	Modality ¹	BS split				SON split			
		Recall@1	Recall@5	Recall@10	max F_1	Recall@1	Recall@5	Recall@10	max F_1
NetVLAD [6]	V	90.26	95.58	96.51	0.9510	96.25	98.07	98.58	0.9808
PointNetVLAD [10]	L	74.30	84.54	87.53	0.8885	98.49	99.43	99.58	0.9931
OverlapTransformer [2]	L	82.05	87.67	88.96	0.9050	98.73	99.70	99.76	0.9937
AutoPlace [38]	R	90.46	97.27	98.40	0.9499	92.99	97.85	98.79	0.9636
AdaFusion [12]	V+L	86.74	95.58	97.17	0.93187	98.70	99.58	99.94	0.9942
PIC-Net [14]	V+L	88.36	95.21	97.17	0.9445	98.13	98.79	99.06	0.9905
MinkLoc++ [13]	V+L	77.76	90.43	93.68	0.8753	91.32	97.28	98.61	0.9556
LCPR (ours)	V+L	97.34	99.24	99.50	0.9873	99.40	99.79	99.88	0.9969

¹ V: Visual, L: LiDAR, V+L: Visual+LiDAR.

TABLE III
OCCLUSION TEST RESULTS (RECALL@1) ON THE BS SPLIT

Occluded Sensor	LCPR	LCPR-I	LCPR-L
–	97.34	93.88	96.21
CAM_FRONT	89.06	81.35	–
CAM_FRONT_RIGHT	88.26	86.84	–
CAM_BACK_RIGHT	87.53	83.71	–
CAM_BACK	91.02	82.85	–
CAM_BACK_LEFT	90.26	84.21	–
CAM_FRONT_LEFT	88.66	87.00	–
LIDAR_TOP	88.66	–	64.69

method AutoPlace [38], the LiDAR-based methods PointNetVLAD [10] and OverlapTransformer [2], and multimodal methods including MinkLoc++ [13], PIC-Net [14], and AdaFusion [12]. We use the released open-source codes of the baseline methods for evaluation except for PIC-Net [14], which is implemented ourselves according to their original paper. All experiments are conducted on a system with an Intel i5-12400F CPU and an Nvidia RTX 3080 GPU.

B. Evaluation for Place Recognition

In order to evaluate the place recognition performance of our proposed method, we first train and test LCPR as well as other state-of-the-art methods on the largest BS split. Furthermore, we also conduct a generalization study with the SON split on all the methods using the weights trained on the BS split. To ensure a fair comparison, we only input single-frame data without RCS histogram Re-Ranking for AutoPlace. In our experiments, AutoPlace and NetVLAD output descriptors of size 1×9216 , while the other baselines produce descriptors of size 1×256 .

Following the previous works [38], [36], we report Recall@N and max F_1 scores as evaluation metrics, which are shown in Tab. II. First, we can see that the recalls of all the methods on the SON split are generally higher than those on the BS split. This can be attributed to the fact that the driving environments in the BS split are much more complex with richer scenes and more observation samples. Our LCPR outperforms all other unimodal and multimodal baselines on both the splits, indicating that our method not only achieves better performance on the previously seen places after training, but also generalizes well to the unseen sce-

narios without fine-tuning. Although AutoPlace can achieve relatively good performance on the BS split, its zero-transfer results on the SON split are not outstanding compared to other methods, indicating its limited generalization ability. Furthermore, compared with AutoPlace and NetVLAD, the descriptors generated by our LCPR are more compact but more discriminative, which means our approach imposes a reduced computational burden on both query retrieval and database storage.

C. Ablation Study

In this section, we investigate the effectiveness of our devised multimodal fusion mechanism. We compare LCPR with its two unimodal baselines, LCPR-L and LCPR-I, which only contain the image encoding branch and LiDAR encoding branch respectively. For a fair comparison, we increase the output dimension of the NetVLAD layer to 256 in each unimodal branch. The evaluation results are shown in the first row of Tab. III. Our proposed multi-scale modality fusion mechanism can effectively exploit the environmental information and further improve the place recognition performance. It can also be observed that the LiDAR range input produces better results than the visual RGB input in our network, suggesting that LiDAR features have a greater contribution to the distinctiveness of our multimodal descriptor. This can be attributed to the RGB input’s deficiency in providing depth information and the impact of raindrops on the camera lens.

D. Robustness Study

We further investigate the robustness of our LCPR in multiple cases with occlusion and changing lighting conditions.

Occlusion: In real-world environments, sensors can be occluded by raindrops and water mist that are attached to them. However, there is a lack of public datasets that contain both extreme weather conditions and visual inputs with a horizontal FOV of 360° . Therefore, we perform manual degradation on the nuScenes dataset to demonstrate the effectiveness of multimodal fusion in dealing with such situations. Specifically, we set all pixels of each camera image to zero in turn to simulate the camera being completely blocked. For LiDAR, we also occlude one-quarter of the range images along the height dimension. As shown in Tab. III, when a certain sensor is disturbed, LCPR can still




Fig. 4. Place recognition results on the SHV split.




Fig. 5. Place recognition results on the BS split when manually adjusting image brightness.

achieve about 90% of Recall@1, while the performance of the two unimodal baselines declines. This demonstrates that LCPR using multimodal inputs has strong robustness against occlusion cases.

Lighting: Variations in lighting conditions significantly affect the visual sensors, which is one of the starting points for introducing multimodal fusion in place recognition. Therefore, to show the robustness of our method in tackling changing lighting conditions, we also evaluate LCPR against baselines that also incorporate visual information of the environment on the SHV and BS split. The SHV split naturally contains many day-to-night repetitive driving trajectories. For the BS split, we manually reduce the image brightness in the query sample with a series of brightness factors to simulate changes in lighting conditions. As shown in Fig. 4 and Fig. 5, multimodal methods generally have better performance than the unimodal method, i.e., NetVLAD with only visual input. The results demonstrate that incorporating the LiDAR point cloud can mitigate the impact of changes in lighting conditions on place recognition tasks. Note that our LCPR still achieves the best performance, indicating its strong robustness against natural lighting condition changes and manual brightness adjustment.

E. Yaw-Rotation Invariance Study

In this experiment, we investigate the yaw-rotation invariance of our proposed descriptor on the BS split. Similar experiment pipelines can be found in the previous works [2], [36], [37], which only focus on the LiDAR-based place recognition. Given that the nuScenes dataset provides multi-view RGB images covering a 360° horizontal perspective, we can extend the pipeline to the multimodal field. For visual input, we first concatenate images of the same frame

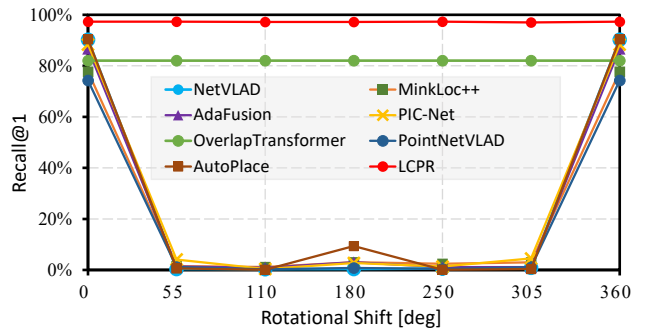


Fig. 6. Study on yaw-rotation invariance.

in the width direction and subsequently apply a circular shift to approximate the yaw-rotation of the vehicle. For point cloud input, we directly rotate the point cloud around the yaw-axis. Considering the shooting angle of the images in the nuScenes dataset, we rotate each query sample by $[55^\circ, 110^\circ, 180^\circ, 250^\circ, 305^\circ, 360^\circ]$ while searching for the places with respect to the same original database. We use Recall@1 as the evaluation metric in this experiment. As illustrated in Fig. 6, both LCPR and OverlapTransformer are not affected by rotation along the yaw-axis, while other baseline methods show a significant drop in performance when a yaw-rotation is applied to the query. Although OverlapTransformer is also yaw-rotation invariant, it still has a lower Recall@1 compared to our LCPR.

F. Runtime and Memory Consumption

In this section, we conduct the evaluation of LCPR's real-time performance and memory usage. The system configuration of this experiment has been clarified in Sec. IV-A. For each query, we accumulate the time consumed in generating the descriptor and retrieve the top-20 candidates from the database to calculate the average runtime. Note that the first 100 queries are used as a warm-up, which are not included in the average runtime. Here we use FAISS library [40] to improve the search efficiency. LCPR takes an average of 16.38 ms to generate a descriptor (12.17 ms) and select the top-20 candidates (4.21 ms) from the database of 9686 samples of the BS split, showing that our method can operate online in real-vehicle applications. Furthermore, our LCPR has a parameter count of 28.97M, which also indicates that our method has a low memory consumption.

V. CONCLUSIONS

In this paper, we present LCPR, a novel and generic multimodal fusion network for place recognition. Our method takes multi-view RGB images from cameras and range images from LiDAR scans as input, and generates global descriptors that are robust to lighting conditions and view-point changes. Benefiting from the self-attention mechanism of the transformer structure, our network can fully exploit the correlations between modalities and the global context within each modality simultaneously, and integrate representations of different modalities. Experimental results show that our descriptor outperforms state-of-the-art unimodal and

multimodal methods on the publicly available dataset and can operate online in real-vehicle applications.

REFERENCES

- [1] P. Yin, S. Zhao, I. Cisneros, A. Abuduweili, G. Huang, M. Milford, C. Liu, H. Choset, and S. Scherer, "General place recognition survey: Towards the real-world autonomy age," *arXiv preprint arXiv:2209.04497*, 2022.
- [2] J. Ma, J. Zhang, J. Xu, R. Ai, W. Gu, and X. Chen, "Overlaptransformer: An efficient and yaw-angle-invariant transformer network for lidar-based place recognition," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 6958–6965, 2022.
- [3] H. Bay, T. Tuytelaars, and L. Van Gool, "Surf: Speeded up robust features," *Lecture notes in computer science*, vol. 3951, pp. 404–417, 2006.
- [4] G. Tang, Z. Liu, and J. Xiong, "Distinctive image features from illumination and scale invariant keypoints," *Multimedia Tools and Applications*, vol. 78, pp. 23 415–23 442, 2019.
- [5] A. Torii, J. Sivic, T. Pajdla, and M. Okutomi, "Visual place recognition with repetitive structures," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2013, pp. 883–890.
- [6] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, "Netvlad: Cnn architecture for weakly supervised place recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 5297–5307.
- [7] S. Hausler, S. Garg, M. Xu, M. Milford, and T. Fischer, "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*, 2021, pp. 14 141–14 152.
- [8] A. Khaliq, S. Ehsan, M. Milford, and K. McDonald-Maier, "Camal: Context-aware multi-layer attention framework for lightweight environment invariant visual place recognition," *arXiv preprint arXiv:1909.08153*, 2020.
- [9] G. Kim and A. Kim, "Scan context: Egocentric spatial descriptor for place recognition within 3d point cloud map," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 4802–4809.
- [10] M. A. Uy and G. H. Lee, "Pointnetvlad: Deep point cloud based retrieval for large-scale place recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 4470–4479.
- [11] Z. Liu, S. Zhou, C. Suo, P. Yin, W. Chen, H. Wang, H. Li, and Y.-H. Liu, "Lpd-net: 3d point cloud learning for large-scale place recognition and environment analysis," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 2831–2840.
- [12] H. Lai, P. Yin, and S. Scherer, "Adafusion: Visual-lidar fusion with adaptive weights for place recognition," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 12 038–12 045, 2022.
- [13] J. Komorowski, M. Wysoczańska, and T. Trzcinski, "Minkloc++: lidar and monocular image fusion for place recognition," in *2021 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2021, pp. 1–8.
- [14] Y. Lu, F. Yang, F. Chen, and D. Xie, "Pic-net: Point cloud and image collaboration network for large-scale place recognition," *arXiv preprint arXiv:2008.00658*, 2020.
- [15] Z. Liu, H. Tang, A. Amini, X. Yang, H. Mao, D. L. Rus, and S. Han, "Befusion: Multi-task multi-sensor fusion with unified bird's-eye view representation," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 2774–2781.
- [16] A. Prakash, K. Chitta, and A. Geiger, "Multi-modal fusion transformer for end-to-end autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 7077–7087.
- [17] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nusenes: A multimodal dataset for autonomous driving," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 621–11 631.
- [18] H. Andreasson and T. Duckett, "Topological localization for mobile robots using omni-directional vision and local features," *IFAC Proceedings Volumes*, vol. 37, no. 8, pp. 36–41, 2004.
- [19] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International journal of computer vision*, vol. 60, pp. 91–110, 2004.
- [20] W. Zhang and J. Kosecka, "Image based localization in urban environments," in *Third international symposium on 3D data processing, visualization, and transmission (3DPVT'06)*. IEEE, 2006, pp. 33–40.
- [21] H. Jégou, M. Douze, C. Schmid, and P. Pérez, "Aggregating local descriptors into a compact image representation," in *2010 IEEE computer society conference on computer vision and pattern recognition*. IEEE, 2010, pp. 3304–3311.
- [22] Z. Chen, O. Lam, A. Jacobson, and M. Milford, "Convolutional neural network-based place recognition," *arXiv preprint arXiv:1411.1509*, 2014.
- [23] F. Radenović, G. Toliás, and O. Chum, "Fine-tuning cnn image retrieval with no human annotation," *IEEE transactions on pattern analysis and machine intelligence*, vol. 41, no. 7, pp. 1655–1668, 2018.
- [24] B. Arcanjo, B. Ferrarini, M. Milford, K. D. McDonald-Maier, and S. Ehsan, "An efficient and scalable collection of fly-inspired voting units for visual place recognition in changing environments," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2527–2534, 2022.
- [25] L. He, X. Wang, and H. Zhang, "M2dp: A novel 3d point cloud descriptor and its application in loop closure detection," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2016, pp. 231–237.
- [26] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 652–660.
- [27] W. Zhang and C. Xiao, "Pcan: 3d attention map learning using contextual information for point cloud based retrieval," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 12 436–12 445.
- [28] J. Komorowski, "Minkloc3d: Point cloud based large-scale place recognition," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2021, pp. 1790–1799.
- [29] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [30] S. Xie, C. Pan, Y. Peng, K. Liu, and S. Ying, "Large-scale place recognition based on camera-lidar fused descriptor," *Sensors*, vol. 20, no. 10, p. 2870, 2020.
- [31] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [32] Y. Pan, X. Xu, W. Li, Y. Cui, Y. Wang, and R. Xiong, "Coral: Colored structural representation for bi-modal place recognition," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 2084–2091.
- [33] X. Yu, B. Zhou, Z. Chang, K. Qian, and F. Fang, "Mmdf: Multi-modal deep feature based place recognition of mobile robots with applications on cross-scene navigation," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 6742–6749, 2022.
- [34] S. Vora, A. H. Lang, B. Helou, and O. Beijbom, "Pointpainting: Sequential fusion for 3d object detection," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 4604–4612.
- [35] H. Zhou, Z. Ge, Z. Li, and X. Zhang, "Matrixvt: Efficient multi-camera to bev transformation for 3d perception," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 8548–8557.
- [36] J. Ma, X. Chen, J. Xu, and G. Xiong, "Seqot: A spatial-temporal transformer network for place recognition using sequential lidar data," *IEEE Transactions on Industrial Electronics*, 2022.
- [37] J. Ma, G. Xiong, J. Xu, and X. Chen, "Cvtnet: A cross-view transformer network for lidar-based place recognition in autonomous driving environments," *IEEE Transactions on Industrial Informatics*, 2023.
- [38] K. Cait, B. Wang, and C. X. Lu, "Autoplacel: Robust place recognition with single-chip automotive radar," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 2222–2228.
- [39] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.
- [40] J. Johnson, M. Douze, and H. Jégou, "Billion-scale similarity search with gpus," *IEEE Transactions on Big Data*, vol. 7, no. 3, pp. 535–547, 2019.