

ModaLink: Unifying Modalities for Efficient Image-to-PointCloud Place Recognition

Weidong Xie, Lun Luo, Nanfei Ye, Yi Ren, Shaoyi Du*, Minhang Wang,
Jintao Xu, Rui Ai, Weihao Gu, Xieyuanli Chen*

Abstract—Place recognition is an important task for robots and autonomous cars to localize themselves and close loops in pre-built maps. While single-modal sensor-based methods have shown satisfactory performance, cross-modal place recognition that retrieving images from a point-cloud database remains a challenging problem. Current cross-modal methods transform images into 3D points using depth estimation for modality conversion, which are usually computationally intensive and need expensive labeled data for depth supervision. In this work, we introduce a fast and lightweight framework to encode images and point clouds into place-distinctive descriptors. We propose an effective Field of View (FoV) transformation module to convert point clouds into an analogous modality as images. This module eliminates the necessity for depth estimation and helps subsequent modules achieve real-time performance. We further design a non-negative factorization-based encoder to extract mutually consistent semantic features between point clouds and images. This encoder yields more distinctive global descriptors for retrieval. Experimental results on the KITTI dataset show that our proposed methods achieve state-of-the-art performance while running in real time. Additional evaluation on the HAOMO dataset covering a 17 km trajectory further shows the practical generalization capabilities. We have released the implementation of our methods as open source at: <https://github.com/haomo-ai/ModaLink.git>.

I. INTRODUCTION

Place recognition is a key module of autonomous vehicle navigation systems. It locates sensor data, such as images and point clouds, on a pre-built map or database for global localization and loop closure detection. LiDAR-based place recognition utilizes precise structural information of point clouds to generate discriminative descriptors [1]–[3]. Since LiDAR sensors are expensive and bulky, deploying such methods is not cost-effective. As an alternative solution, visual place recognition is competitive due to its lower price and smaller size [4], [5] of cameras. However, it is not robust to illumination changes and view variations [6]. Since different sensors have varying advantages, it can occur in real robotic applications that the data queried by one robot

* corresponding authors.

W. Xie and S. Du are with National Key Laboratory of Human-Machine Hybrid Augmented Intelligence, National Engineering Research Center for Visual Information and Applications, and Institute of Artificial Intelligence and Robotics, Xi'an Jiaotong University, Xi'an 710049, China, xieweidongxjtu@stu.xjtu.edu.cn and dushaoyi@xjtu.edu.cn. W. Xie is also with Haomo.AI Technology Co., Ltd. L. Luo, M. Wang, J. Xu, R. Ai, W. Gu, N. Ye are with Haomo.AI Technology Co., Ltd. X. Chen is with National University of Defense Technology. Y. Ren is with Carnegie Mellon University.

This work was supported by the National Science and Technology Major Project of China under Grant No. 2020AAA0108102, the National Natural Science Foundation of China under Grant Nos. 62327808 and 62088102.

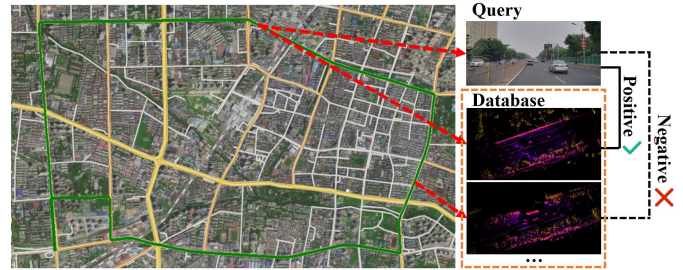


Fig. 1: The purpose of ModaLink is to retrieve the most positive corresponding point cloud of the query image in a pre-built large-scale point cloud database.

has different modalities compared to the database recorded by another, making cross-modal place recognition highly necessary. For instance, autonomous vehicles equipped with cheap cameras can use such methods to locate input images on large-scale point cloud databases and obtain accurate localization results, as shown in Fig. 1.

Generating accurate global descriptors that represent scene information is crucial for successful place recognition. However, designing these descriptors for cross-modal place recognition is very challenging, since data from different sensors such as LiDAR and camera are vastly different. Specifically, the LiDAR captures depth information by measuring the laser time of flight (TOF), while the camera collects photometric data by responding to the ambient light through CMOS. As both sensors capture different types of data, it is difficult to develop a unified feature extractor. To handle this issue, some methods [7], [8] transform different types of data into intermediate formats, such as bird's eye views (BEVs) and range images. They use depth estimation algorithms to transform images into 3D points and align them with the LiDAR point cloud. Although these methods perform well on public datasets, they cannot run in real time due to the high computational requirements of the depth estimation module.

In this work, we present a new approach that uses an FoV transformation module to bring together the data of images and point clouds. Instead of relying on depth estimation, our module converts 3D point clouds into 2D depth images. This makes it possible for subsequent modules to run in real-time, without performing time-consuming depth estimation. To further improve the distinctiveness of descriptors, we propose a non-negative factorization-based encoder to mine the implicit common features in an unsupervised way and use a shared-weight network to learn global descriptors based on the unified data formats. We thoroughly evaluated our

ModaLink on both the KITTI and our self-collected dataset covering a 17 km trajectory. The experimental results demonstrate that our method achieves state-of-the-art performance in cross-modal place recognition.

In summary, our main contributions are as follows:

- We propose a lightweight cross-modal place recognition method called ModaLink based on FoV transformation while using only monocular images.
- We introduce a Non-Negative Matrix Factorization-based module to extract extra potential semantic features to improve the distinctiveness of descriptors.
- Extensive experimental results on the KITTI and a self-collected dataset show that our proposed method can achieve state-of-the-art performance while running in real-time at about 30Hz.

II. RELATED WORK

In this section, we first review the place recognition methods based on single-modal data. We then discuss the recent development of the approaches to achieve image-to-PointCloud place recognition.

A. Visual Place Recognition (VPR)

Early visual place recognition methods [9], [10] concentrate on developing hand-engineered features [11]–[13]. These methods assume that similar environmental structures will result in comparable feature distributions. However, this assumption may not always apply for VPR, particularly when the appearances of environments vary across seasons. With the development of CNN architectures, CNN-based VPR methods have been developed. These methods transform the image features into localization descriptors, which improves the VPR performance by designing better local feature extractors and feature aggregators. For example, NetVLAD [4] extracts local features by CNN and designs a network to learn the cluster centers. Patch-NetVLAD [14] follows NetVLAD and leverages the strengths of both the global and local features to generate more distinct patch-level features. By learning from large datasets, these methods have exhibited better robustness and generalization ability in long-term [14] and large-scale [15] tasks. More recently, transformers have been broadly used in the VPR problem. Yang et al. [16] combine CNN and transformers and build a CNN+transformer structure. TransGeo [5] introduces the vision transformer [17] structure and explicitly removes the invalid parts of the input images. With the help of transformers, these methods extract more distinct descriptors. However, those methods focus only on VPR and cannot localize the visual descriptors on point cloud maps.

B. LiDAR Place Recognition (LPR)

In recent years, LiDAR-based place recognition has rapidly developed due to the ability of LiDAR to obtain accurate depth information and its robustness to changes in illumination and view. Following the setup of the visual counterparts NetVLAD [4], PointNetVLAD [18] utilizes the point-based encoder, i.e. PointNet [19] for local feature

extraction, and then generates global features with NetVLAD and achieve promising LPR results. However, it cannot capture local geometric structures due to its approach of treating each point independently. Thus, subsequent methods primarily focus on exploiting neighborhood information [20], [21]. LiDAR point clouds are disordered and massive, thus hard to process using advanced vision learning techniques. Therefore, some methods project point clouds into images and encode the images into descriptors. Kim et al. [22] propose the scan context descriptor by projecting point clouds into an ego-centric coordinate system. This system partitions the ground space into bins according to both azimuthal and radial directions. OverlapNet [1] and its later develops [2], [3], [23] project point clouds into range images and adopt a siamese network to learn the overlap between a pair of images. BVMatch [24] projects point clouds into BEV images and extracts the global BVFT feature for pose estimation. It shows that the projected image is a quite effective representation of point cloud and these methods [24], [25] achieve state-of-the-art performance in terms of retrieval recall, robustness, and generalization ability.

C. Image to Point Cloud Place Recognition

Although place recognition based on single-modal data has been extensively studied, identifying images on a point-cloud map remains a difficult task. As a result, there are fewer related works in comparison to the single-modal-based counterparts. One straightforward solution is to jointly train a 2D neural network for images and a 3D network for point clouds to create shared embeddings [26]. However, this approach does not generalize well to unseen environments. Bernreiter et al. [27] proposed to project images and point clouds into unit spheres and then extract features through such unit representations. This method requires multiple images as input, which is not always available. There are also approaches [28], [29] detecting local features from images and point clouds to find cross-modal correspondences for place retrieval. 2D3D-MatchNet [28] designed a CNN network to cross-match the features in both the images and ground-removed point cloud submaps, while P2-Net [30] selected a batch-hard detector to produce shared embedding between different sensor modalities. However, they perform only on a small scale and can hardly be extended to a large-scale environment well due to the heavy point cloud representation. Another intuitive solution is to convert two types of data into the same modality. To this end, LC² [8] exploited depth estimation to transform both images and point clouds into depth images. I2P-Rec [7] adopted monocular and stereo depth estimation to convert 2D images into 3D point clouds and then utilized BEV-based methods to achieve LPR. However, the photometric rendering process is also computationally expensive. The stereo matching and monocular depth estimation modules require labeled data for training and are usually of low running speed. Unlike the previous approaches, we design a system to achieve place recognition without depth estimation modules.

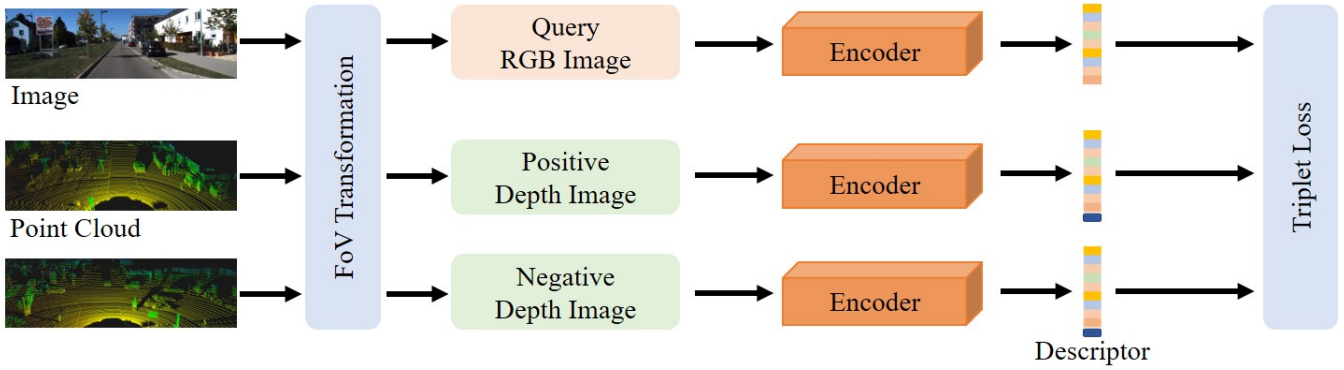


Fig. 2: The training framework of our proposed ModaLink framework. Point clouds are converted to depth images via projection. Then, the query image and depth images are cropped into the same overlap of FoV. Based on the depth completion module, sparse depth images are upsampled to dense depth images. Then, global place descriptors are generated by a shared-weight encoder. Finally, we adopt triplet loss for supervision.

III. IMAGE-TO-POINTCLOUD PLACE RECOGNITION BASED ON FOV TRANSFORMATION

The goal of this paper is to develop a method for fast and accurate image-to-pointcloud place recognition. Specifically, it can be defined as: given an image, the task is to retrieve the corresponding point cloud from the database that recorded at the same place. To this end, we propose ModaLink based on a novel FoV transformation module as shown in Fig. 2. For a query image, we find its positive and negative point clouds from the training set. Then, we use an FoV transformation module to unify the data format. We use a shared-weight encoder based on NMF to generate global descriptors. Finally, triplet loss is adopted for supervision. We elaborate on our designs next.

A. FoV Transformation

The objective of the FoV transformation module is to unify images and point clouds into a similar data format. As illustrated in Fig. 3, the FoV transformation module first aligns images and point clouds by projecting the point clouds into depth images with the extrinsic and intrinsic parameters. It then crops both depth and RGB images into the same FoV. After obtaining the sparse depth images, it completes the sparse depth image into a denser one for more distinctive feature extraction. More details for each step are as follows:

Point Cloud to Image Projection. We first project the point cloud P into depth image D to reduce the gap between different modalities. LiDAR point clouds are provided with Cartesian coordinates (x, y, z) . According to the principles of pin-hole camera geometry, we project each point $p(x, y, z)$ in the point cloud into the image plane with its coordinates (u, v) as

$$[u, v]^T = \mathbf{K}(\mathbf{R}[x, y, z]^T + \mathbf{t}), \quad (1)$$

where \mathbf{K} is the camera's intrinsic matrix, and \mathbf{R} and \mathbf{t} correspond to the extrinsic rotation matrix and translation vector, respectively. The corresponding pixel intensity $D(u, v)$ of the depth image is calculated as the depth of point and can be formulated as $D(u, v) = \sqrt{x^2 + y^2 + z^2}$. If two points are projected onto the same pixel of the image, we choose the one with the smaller depth and discard the farther one.

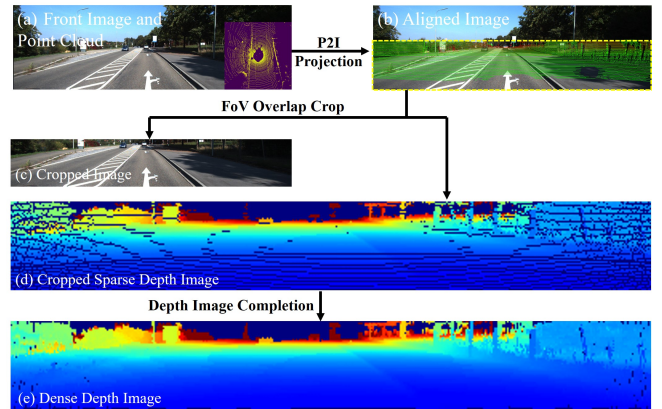


Fig. 3: Visualization of image and point cloud alignment based on intrinsic and extrinsic matrix. After alignment, we crop the images based on FoV overlap and complete sparse depth images to generate dense depth images.

FoV Overlap Crop. The sizes of the point cloud depth image and RGB image are usually different since LiDAR scanners and cameras have different FoV. Typically, multi-beam LiDAR sensors have a 360-degree field of view, whereas the cameras are limited to perceiving a cone region in their frontal view. To deal with this, we crop the RGB images and depth images into a fixed-size window, the max elevation angle of which is limited to 5 degrees according to the FoV overlap of the two specific sensors.

Depth Image Completion. In urban driving scenarios, most objects stand on the ground plane. In this situation, vertically adjacent points in point clouds are more likely to come from the same object [31]. We take advantage of this feature to generate interpolating points according to vertically adjacent pixels to achieve depth image completion. Fig. 4 provides a description of this completion procedure. As shown in Fig. 4(a), the two blue lines are adjacent LiDAR scanning lines and the blue points are two points in the point cloud. We project these two points onto the depth image as two pixels. Then, there is a fissure pixel between the above two pixels obtained by projection. We interpolate the fissure pixel and back-project it into the 3D space to obtain the red




Fig. 4: Three cases of depth image completion. The red point represents a pixel on the image plane where no LiDAR point is projected.




Fig. 5: A demonstration of depth image interpolation.

points. Ideally, the red dot should be on the surface of the object. Let D denote the depth image of the point cloud P , and $D(u', v')$ is a fissure pixel to which no LiDAR point corresponds. We interpolate the missing pixel $D(u', v')$ with $D(u' + i, v')$ and $D(u' - j, v')$ through linear interpolating,

$$D(u', v) = \frac{jD(u' + i, v) + iD(u' - j, v)}{i + j}, \quad (2)$$

where $i, j > 0$.

Noted that, due to the occlusion between the background and the foreground objects, two adjacent pixels in the depth image may be nonadjacent in the source point cloud. This phenomenon will cause the interpolation point (noted as the red point) to appear between the foreground and the background instead of actually being on the surface of the foreground object, leading to incorrect interpolation results, as shown in Fig. 4 (b)(c). To solve this problem, we set a threshold σ . If $|D(u + i, v) - D(u - j, v)| > \sigma$, we regard the larger one as the background point, and to keep the foreground object's outline, we make $D(u, v) = \min(D(u + i, v), D(u - j, v))$. Fig. 5 demonstrates an example of the above-mentioned depth image completion method.

B. Novel NMF-Based Encoder Network

After obtaining similar LiDAR and camera data representations, we design a shared weight descriptor encoder to further encode depth images and RGB images into distinctive global descriptors for cross-modal place recognition. As illustrated in Fig. 6, we first employ a CNN to extract local features. We then use the Non-negative Matrix Factorization module to extract additional semantic features, mining potential semantic consistency between the image and its corresponding point cloud. Two NetVLAD networks aggregate local and semantic features and generate global descriptors. Each module is detailed as follows:

Local Feature Extraction. For an RGB or depth image input x , we first extract local feature map $f(x) \in \mathbb{R}^{H \times W \times C}$ using a CNN. In this representation, H and W denote the

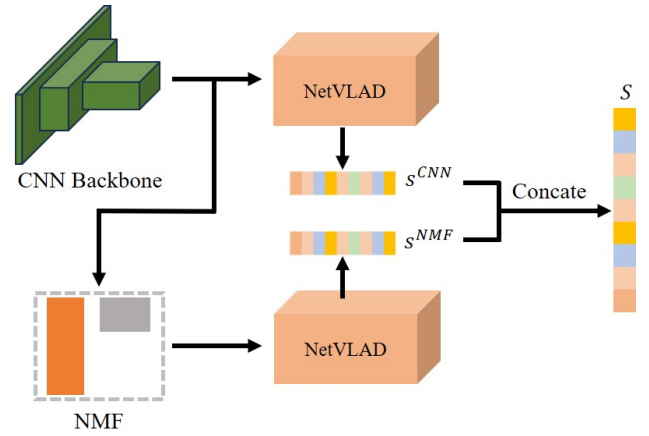


Fig. 6: The structure of the shared-weight global descriptor encoder. It contains a CNN backbone, an NMF module and two NetVLAD networks.

size of the output from the last CNN layer, and C represents the dimension of the feature space. We use ResNet34 [32] as our local feature extractor.

NMF Semantic Feature Extraction. Semantic information is theoretically helpful to determine whether two observations are apparently similar through high-level scene understanding [33]. However, the use of semantic segmentation networks can be computationally inefficient and require large amounts of manually annotated data for training. Employing such techniques may lead to slow inference speeds and introduce errors in semantic estimation, which can negatively affect place recognition. In ModaLink, we introduce a lightweight Non-negative Matrix Factorization (NMF) based encoder. This module has inherent clustering property [34] and can mine latent semantic information in an unsupervised way. By exploiting such implicit semantic information, our NMF-based encoder can produce more distinctive descriptors while requiring no semantic labels.

NMF is a matrix decomposition method. It factorizes a non-negative feature matrix A into a non-negative cluster matrix P and a non-negative orthogonal matrix Q . Here non-negative means that the matrices have no negative elements. NMF achieves an approximation of $A \approx PQ$ by minimizing the error matrix, i.e.

$$\min \|A - PQ\|_F, P \geq 0, Q \geq 0, \quad (3)$$

where $\|\cdot\|_F$ means the Frobenius Norm. By imposing an orthogonality constraint on Q , where $QQ^T = I$, the aforementioned minimization is equivalent to K-means mathematically. Each row of Q can be interpreted as a cluster centroid in C dimensions due to the devised orthonormal constraints of NMF. Each row of the matrix P corresponds to the spatial position of a pixel from the matrix A . It represents the pixel feature as a particular cluster. NMF has a built-in clustering property that automatically groups the columns of matrix A into clusters [34]. Each row of the matrix Q can be viewed as a cluster centroid of C dimensions, which




Fig. 7: The process of the NMF module.




Fig. 8: Brief illustration of the clustering effect of NMF.

corresponds to a coherent object among views. As a result, the cluster matrix P can serve as a feature vector.

Since the feature map activated by ReLU inherently consists of non-negative elements, we use NMF to factorize the feature map A to generate explicit semantic features. Fig. 7 shows the factorization procedure. The extracted feature map from the local feature extraction module has dimensions (H, W, C) , and these feature maps are concatenated and reshaped into an (NHW, C) matrix A . Each row of A represents a feature vector of an observation. By solving NMF [34], the A matrix is factorized into a (NHW, K) matrix P and a (K, C) matrix Q , where K represents the number of semantic clusters, and $A \approx PQ$. P captures the similarity between each pixel and the clustered semantic objects, and are then employed as our semantic features.

Fig. 8 illustrates examples of extracted semantic feature maps with our proposed encoder network, where the NMF exhibits a semantic clustering property. For similar input RGB images, NMF can generate similar semantic feature maps. For dissimilar RGB images, NMF can extract similar features from the parts of the image with the same semantics.

Global descriptor generation. To generate global place descriptors, we finally leverage NetVLAD [4] to cluster local features. The local feature map $f(x)$ serves as the input of NetVLAD and is then transformed into a partial global localization descriptor $s^{CNN} \in \mathbb{R}^{D_{CNN}}$, where D_{CNN} is the dimension of the descriptor. Following the NMF module, $f(x)$ is clustered into $g(x) \in \mathbb{R}^{H \times W \times K}$. Another NetVLAD transformation is applied to $g(x)$, resulting in $s^{NMF} \in \mathbb{R}^{D_{NMF}}$. Finally, the concatenation of s^{CNN} and s^{NMF} forms the ultimate global place descriptor s .

C. Training And Inference

Before training, we construct a keyframe database to represent a global point cloud map. We find positive and

negative matches for the query image according to ground truth distance. Because in the subsequent experiments, when the positioning result is within 10 meters of the ground truth, we regard it as a correct positioning. So, in the training process, we regard the point clouds within the stricter 5-meter range as the point cloud located in the same place, and accordingly, for each query image, we regard the point clouds within 5 meters away from the query as positive samples, while other point clouds as negative samples.

For supervision, we adopt lazy-triplet loss [35] to guide the model encoding the camera and depth images into descriptors that map the same location close in the feature space, while different locations farther in the feature space:

$$\mathcal{L}_{triplet} = \sum_i^n \max(m + d(s_q, s_p) - d(s_q, s_n^i), 0), \quad (4)$$

where m is the constant margin set as 0.3. s_q , s_p , and s_n^i denote the learned descriptors from the query image, positive depth image, and the i -th negative depth image. The distance between descriptors is defined by $d(x, y) = \|x - y\|_2$.

The global pose information of the prior database can be obtained using a SLAM system or RTK information. We sample LiDAR keyframes and generate their global descriptors using our Modalink at intervals of 5 meters. During the inference phase, we use NMF-Encoder to extract the global descriptor for the incoming query image and employ the Faiss approach [36] for fast retrieval of the best-matched LiDAR frame. The global pose of the query image is then determined according to the corresponding position of the matched LiDAR frame.

IV. EXPERIMENTS

In this section, we evaluate the performance of our ModaLink for cross-modal on point cloud maps place recognition. Based on the KITTI [37] and our self-collected dataset, we evaluate the performance of ModaLink and draw comparisons with baseline methods [7], [26]. Finally, we conduct ablation studies to analyze crucial components and assess the runtime requirements of our method.

A. Datasets

KITTI. The KITTI contains a large number of synchronized and calibrated point clouds and images. It is widely used to verify autonomous driving-related algorithms. We use the images and point clouds of 0-3000 frames in sequence 00 for training and the rest of the frames in sequence 00 for validation. We use the 02, 05, 06, and 08 sequences of its odometry subset for testing.

HAOMO Dataset. To assess the generalization capability of ModaLink, we further use the HAOMO Dataset [2]. This dataset was collected by the HAOMO.AI company in China, with a vehicle equipped with a LiDAR sensor (128 beams), a wide-angle camera, and an RTK GNSS. The HAOMO dataset encompasses a trajectory length of over 17.4 km. The scenes include an overpass and a town. It comprises 25,278 pairs of point clouds and images. Fig. 9 displays some sample images from HAOMO dataset.

I: THE RECALL@1 AND RECALL @1% OF PLACE RECOGNITION METHODS ON THE KITTI DATASET.

Sequence Recall	00		02		05		06		08	
	@1	@1%	@1	@1%	@1	@1%	@1	@1%	@1	@1%
Baseline	58.0	-	4.0	-	10.6	-	22.3	-	5.3	-
MIM-Points	53.3	71.6	6.9	57.1	14.8	68.3	20.0	44.4	12.5	68.6
PSM-Points*	58.0	72.3	8.2	53.6	16.0	71.1	32.1	50.0	14.7	75.0
LEA-Points*	58.8	72.6	9.1	61.3	21.8	80.0	31.6	71.6	18.0	76.3
MIM-I2P-Rec	74.3	88.6	46.3	89.7	49.9	89.1	41.9	81.5	42.1	88.4
PSM-I2P-Rec*	81.7	97.4	47.6	92.9	63.8	94.1	52.1	88.5	53.5	92.1
LEA-I2P-Rec*	92.6	99.7	77.0	98.4	83.4	98.8	55.5	91.5	69.4	96.4
Ours	98.0	100	70.5	97.7	91.3	99.3	87.4	100	84.4	99.9

* indicates that these methods require stereo images as input.



Fig. 9: Some sample images from HAOMO Dataset.

B. Evaluation Metrics

For evaluation, we adopt the recall at Top-N and the recall at Top-1% as metrics. When performing the matching, we extract global descriptors of query images and retrieve Top-N nearest matches from the point cloud database. Following the setup of I2P-Rec [7], We regard a match between a query image and the point cloud as true positive when their geometric distance is less than a threshold of $t = 10$ meters. Recall at Top-N is computed as the proportion of correct predictions among all positive instances in a dataset considering the Top-N prediction for each sample.

C. Implementation Setup

Our ModaLink employs a pretrained ResNet34 as the backbone, with the number of NMF clusters $K = 16$. NetVLAD modules have a cluster number of $d_k = 64$. The intermediate feature channels of the ResNet and NMF branches are set to 256 and K , respectively. The output descriptor of ModaLink is a float vector of length $256 \times 64 + K \times 64$. We use data augmentation to enhance the model’s robustness to variations. Specifically, we incorporate random rotation within $\pm 5^\circ$ and position shifts within $\pm 0.1m$ to the point cloud before projection. More implementation details can be found at our open-source repository¹.

D. Performance on the KITTI

We assess the place recognition performance on the KITTI dataset, and compare our method with baseline [26] and I2P-Rec [7]. The baseline method uses 2D CNN and 3D DNN to process 2D images and 3D point clouds separately, and then generate descriptors with NetVLAD [4]. I2P-Rec uses

¹<https://github.com/haomo-ai/ModaLink.git>

II: GENERALIZATION ON THE HAOMO DATASET.

	Recall@1	Recall@1%
Trained on KITTI	35.5	75.5
Refined on HAOMO	70.9	98.2

III: PERFORMANCE OF MODALINK AGAINST DIFFERENT THRESHOLDS ON HAOMO DATASET.

Threshold(m)	5	10	15	20	25	30	35	40	45	50
Recall@1	69.1	70.9	71.8	72.6	73.3	73.6	73.9	74.2	74.6	74.9

monocular depth estimation, i.e. MIM [38] or stereo matching, i.e. PSM [39] and LEA [40] to recover point clouds from RGB images, and generates descriptors based on the BEV images of clouds. For a more comprehensive evaluation, in I2P-Rec, Zheng et al. [7] also extract global descriptors from the recovered clouds with PointNetVLAD [18]. They denote this method using different depth estimation algorithms as MIM-Points, PSM-Points, and LEA-Points.

As shown in Table I, our proposed ModaLink exhibits a significant improvement, outperforming stereo-image-based and depth-estimation-based I2P-Rec in most scenarios. Fig. 10 visually depicts the recall rates at Top-N, where ModaLink consistently outperforms other methods. Note that our method only uses monocular images as queries while outperforming methods exploiting stereo images.

E. Performance on the HAOMO dataset

HAOMO dataset was collected along busy streets with significant interference and utilized this dataset for evaluation. To validate the generalization ability of our method, we designate 30% of the LiDAR frames and corresponding images for training, reserving the remainder for testing. As shown in Table II, ModaLink achieves a Top-1 recall rate of 35.5% while only trained on the KITTI and a Top-1 recall rate of 70.9% while refined on the HAOMO dataset, validating its generalization ability. As shown in Table III, after further training on 30% of the HAOMO dataset, ModaLink performs satisfactorily under different positive thresholds.

Additionally, we explore the sensitivity of different methods against the ground truth threshold t on both datasets. As depicted in Fig. 11(a), recall@1 increases as t varies from 5m to 50m, with our ModaLink maintaining a consistently high recall@1. As shown in Fig. 11(b), ModaLink also has a satisfactory performance on other sequences which further approves its generalization ability.

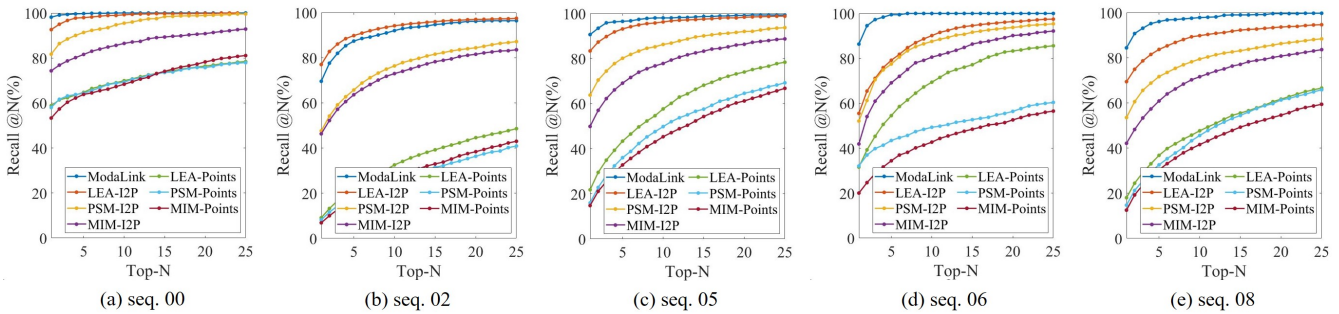


Fig. 10: The Top-N recall rate on the KITTI dataset.

IV: ABLATION STUDY OF DIFFERENT COMPONENTS IN OUR MODALINK.

NMF	DC	00	02	05	06	08
		26.1	12.9	32.8	64.5	22.1
✓		83.4	69.5	83.4	55.5	69.4
	✓	93.1	62.8	88.3	85.8	77.1
✓	✓	98.0	70.5	91.3	87.4	84.4

V: COMPARISON OF DIFFERENT NUMBERS OF SEMANTIC CLUSTERS K .

Clusters	00	02	05	06	08
3	89.6	38.1	75.7	76.5	60.8
5	93.5	60.8	87.4	85.3	73.9
16	98.0	70.5	91.3	87.4	84.4
64	94.1	63.7	86.7	84.7	77.1

F. Ablation Study

We conduct detailed ablation studies for the depth completion (DC) in FoV Transformation Module and the NMF on the KITTI dataset.

Depth Completion in FoV Transformation Module. We compare our ModaLink with the variants of not using depth completion in the FoV transformation module. Comparing four different setups, as shown in Table IV, ModaLink with DC outperforms the setups without DC for place recognition. The results show that dense depth images can help generate more distinctive place descriptors.

NMF Segmentation Module. As shown in Table IV, NMF+DC outperforms the other setups for place recognition. The results show that our NMF module increased the performance significantly, and potential semantic information lead to higher efficiency.

NMF Cluster Number. The number of clusters, denoted as K , is a crucial hyperparameter for NMF. To explore its effect, we conducted experiments with different cluster numbers. The results, as presented in Table V, indicate that a higher number of clusters does not consistently result in superior performance. An ideal K is observed to be around 16 for ModaLink.

G. Runtime

This experiment evaluates the runtime requirements of our method. All experiments are conducted on a computer

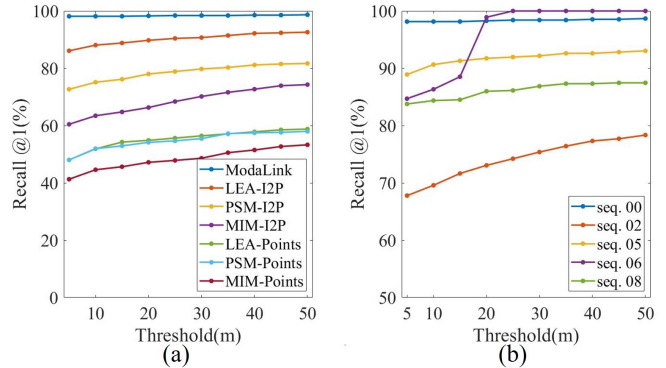


Fig. 11: The Top-1 recall with respect to different threshold on seq 00 and the Top-1 recall with respect to different sequence and threshold for ModaLink.

equipped with an Intel Xeon Platinum 8362 CPU and an Nvidia A100 SXM4 GPU.

The NMF-Encoder takes an average of 22.51ms to encode an image into a descriptor. The descriptor retrieval takes averages 10.94ms. Overall, our ModaLink takes 33.45ms to localize an image and achieve an FPS of about 30Hz. In contrast, the depth estimation part of I2P-Rec, namely LEA-Stereo, consumes at least 0.3s to estimate the depth for a pair of stereo images. This showcases the evident advantage in the speed of ModaLink.

V. CONCLUSION

We presented ModaLink, a novel framework for cross-modal place recognition, specifically addressing Image-to-PointCloud tasks in autonomous vehicles. Our approach introduces a Field of View (FoV) transformation module for real-time performance and eliminates the need for computationally intensive depth estimation modules. We also design a Non-negative Matrix Factorization (NMF) based module to improve the distinctiveness of descriptors. Experimental results on the KITTI dataset showcase the state-of-the-art performance of ModaLink. Extensive evaluations on the HAOMO dataset further validate the system's practical generalization capabilities in diverse environments. The open-source implementation of our methods is available, contributing a fast solution to cross-modal place recognition for autonomous vehicles.

REFERENCES

- [1] X. Chen, T. Läbe, A. Milioto, T. Röhling, O. Vysotska, A. Haag, J. Behley, C. Stachniss, and F. Fraunhofer, "OverlapNet: Loop closing for LiDAR-based SLAM," in *Proc. of Robotics: Science and Systems*, 2020.
- [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] J. Ma, G. Xiong, J. Xu, and X. Chen, "CVTNet: A cross-view transformer network for place recognition using LiDAR data," *arXiv preprint arXiv: 2302.01665*, 2023.
- [4] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, "NetVLAD: CNN architecture for weakly supervised place recognition," in *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5297–5307, 2016.
- [5] S. Zhu, M. Shah, and C. Chen, "TransGeo: Transformer is all you need for cross-view image geo-localization," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1162–1171, 2022.
- [6] K. Zywanowski, A. Banaszczyk, and M. R. Nowicki, "Comparison of camera-based and 3D LiDAR-based place recognition across weather conditions," in *2020 16th International Conference on Control, Automation, Robotics and Vision (ICARCV)*, pp. 886–891, IEEE, 2020.
- [7] S. Zheng, Y. Li, Z. Yu, B. Yu, S.-Y. Cao, M. Wang, J. Xu, R. Ai, W. Gu, L. Luo, *et al.*, "I2P-Rec: Recognizing images on large-scale point cloud maps through bird's eye view projections," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1395–1400, IEEE, 2023.
- [8] A. J. Lee, S. Song, H. Lim, W. Lee, and H. Myung, "(LC)²: LiDAR-camera loop constraints for cross-modal place recognition," *IEEE Robotics and Automation Letters*, vol. 8, no. 6, pp. 3589–3596, 2023.
- [9] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardos, "ORB-SLAM: A versatile and accurate monocular SLAM system," *IEEE Transactions on Robotics*, vol. 31, no. 5, pp. 1147–1163, 2017.
- [10] R. Mur-Artal and J. D. Tardós, "ORB-SLAM2: An open-source SLAM system for monocular, stereo, and RGB-D cameras," *IEEE transactions on robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.
- [11] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "ORB: An efficient alternative to SIFT or SURF," in *International Conference on Computer Vision*, pp. 2564–2571, IEEE, 2011.
- [12] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International journal of computer vision*, vol. 60, pp. 91–110, 2004.
- [13] D. Galvez-Lpez and J. D. Tardos, "Bags of binary words for fast place recognition in image sequences," *IEEE Transactions on Robotics*, vol. 28, no. 5, pp. 1188–1197, 2012.
- [14] S. Hausler, S. Garg, M. Xu, M. Milford, and T. Fischer, "PatchNetVLAD: Multi-scale fusion of locally-global descriptors for place recognition," 2021.
- [15] J. Zhu, Y. Ai, B. Tian, D. Cao, and S. Scherer, "Visual place recognition in long-term and large-scale environment based on CNN feature," in *2018 IEEE Intelligent Vehicles Symposium (IV)*, pp. 1679–1685, IEEE, 2018.
- [16] H. Yang, X. Lu, and Y. Zhu, "Cross-view geo-localization with evolving transformer," *arXiv preprint arXiv:2107.00842*, 2021.
- [17] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," 2021.
- [18] M. Angelina Uy and G. Hee Lee, "PointNetVLAD: Deep point cloud based retrieval for large-scale place recognition," in *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4470–4479, IEEE, 2018.
- [19] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "PointNet: Deep learning on point sets for 3D classification and segmentation," in *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 652–660, 2017.
- [20] Z. Liu, S. Zhou, C. Suo, P. Yin, W. Chen, H. Wang, H. Li, and Y. Liu, "LPD-Net: 3D point cloud learning for large-scale place recognition and environment analysis," in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 2831–2840, 2019.
- [21] Y. Xia, Y. Xu, S. Li, R. Wang, J. Du, D. Cremers, and U. Stilla, "SOE-Net: A self-attention and orientation encoding network for point cloud based place recognition," in *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 11343–11352, 2021.
- [22] G. Kim and A. Kim, "Scan Context: Egocentric spatial descriptor for place recognition within 3D point cloud map," in *IEEE International Conference on Intelligent Robots and Systems*, pp. 4802–4809, IEEE, 2018.
- [23] X. Chen, T. Läbe, A. Milioto, T. Röhling, J. Behley, and C. Stachniss, "Overlapnet: A siamese network for computing lidar scan similarity with applications to loop closing and localization," *Autonomous Robots*, pp. 1–21, 2022.
- [24] L. Luo, S.-Y. Cao, B. Han, H.-L. Shen, and J. Li, "BVMatch: Lidar-based place recognition using bird's-eye view images," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 6076–6083, 2021.
- [25] L. Luo, S. Zheng, Y. Li, Y. Fan, B. Yu, S.-Y. Cao, J. Li, and H.-L. Shen, "BEVPlace: Learning LiDAR-based place recognition using bird's eye view images," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 8700–8709, 2023.
- [26] D. Cattaneo, M. Vaghi, S. Fontana, A. L. Ballardini, and D. G. Sorrenti, "Global visual localization in LiDAR-maps through shared 2D-3D embedding space," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4365–4371, IEEE, 2020.
- [27] L. Bernreiter, L. Ott, J. Nieto, R. Siegwart, and C. Cadena, "Spherical multi-modal place recognition for heterogeneous sensor systems," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1743–1750, IEEE, 2021.
- [28] M. Feng, S. Hu, M. H. Ang, and G. H. Lee, "2D3D-Matchnet: Learning to match keypoints across 2D image and 3D point cloud," in *2019 International Conference on Robotics and Automation (ICRA)*, pp. 4790–4796, IEEE, 2019.
- [29] K. Chen, H. Yu, W. Yang, L. Yu, S. Scherer, and G.-S. Xia, "I2D-Loc: Camera localization via image to lidar depth flow," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 194, pp. 209–221, 2022.
- [30] B. Wang, C. Chen, Z. Cui, J. Qin, C. X. Lu, Z. Yu, P. Zhao, Z. Dong, F. Zhu, N. Trigoni, *et al.*, "P2-Net: Joint description and detection of local features for pixel and point matching," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 16004–16013, 2021.
- [31] Y. Liu, G. Chen, and A. Knoll, "Globally optimal vertical direction estimation in Atlanta world," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 4, pp. 1949–1962, 2020.
- [32] 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*, pp. 770–778, 2016.
- [33] R. Dubé, D. Dugas, E. Stumm, J. Nieto, R. Siegwart, and C. Cadena, "SegMatch: Segment based place recognition in 3D point clouds," in *IEEE International Conference on Robotics and Automation*, pp. 5266–5272, IEEE, 2017.
- [34] C. Ding, X. He, and H. D. Simon, "On the equivalence of nonnegative matrix factorization and spectral clustering," in *Proceedings of the 2005 SIAM international conference on data mining*, pp. 606–610, SIAM, 2005.
- [35] F. Schroff, D. Kalenichenko, and J. Philbin, "Facenet: A unified embedding for face recognition and clustering," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 815–823, 2015.
- [36] 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.
- [37] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the KITTI vision benchmark suite," in *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3354–3361, 2012.
- [38] J. H. Z. Z. H. H. Y. C. Zhenda Xie, Zigang Geng, "Revealing the dark secrets of masked image modeling," *arXiv preprint arXiv:2205.13543*, 2022.
- [39] J.-R. Chang and Y.-S. Chen, "Pyramid stereo matching network," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5410–5418, 2018.
- [40] X. Cheng, Y. Zhong, M. Harandi, Y. Dai, X. Chang, H. Li, T. Drummond, and Z. Ge, "Hierarchical neural architecture search for deep stereo matching," *Advances in Neural Information Processing Systems*, vol. 33, 2020.