# BVMatch: Lidar-Based Place Recognition Using Bird's-Eye View Images

Lun Luo<sup>®</sup>, Si-Yuan Cao, Bin Han, Hui-Liang Shen<sup>®</sup>, and Junwei Li

Abstract—Recognizing places using Lidar in large-scale environments is challenging due to the sparse nature of point cloud data. In this letter we present BVMatch, a Lidar-based frame-to-frame place recognition framework, that is capable of estimating 2D relative poses. Based on the assumption that the ground area can be approximated as a plane, we uniformly discretize the ground area into grids and project 3D Lidar scans to bird's-eye view (BV) images. We further use a bank of Log-Gabor filters to build a maximum index map (MIM) that encodes the orientation information of the structures in the images. We analyze the orientation characteristics of MIM theoretically and introduce a novel descriptor called bird's-eye view feature transform (BVFT). The proposed BVFT is insensitive to rotation and intensity variations of BV images. Leveraging the BVFT descriptors, we unify the Lidar place recognition and pose estimation tasks into the BVMatch framework. The experiments conducted on three large-scale datasets show that BVMatch outperforms the state-of-the-art methods in terms of both recall rate of place recognition and pose estimation accuracy.

Index Terms-Localization, range sensing, SLAM.

## I. INTRODUCTION

**P**LACE recognition is a crucial ability for the autonavigating robots to perform long-term simultaneous localization and mapping (SLAM) tasks [1]. Benefiting from the local image features [2], [3] and the bag-of-features techniques [4], [5], a number of image-based place recognition methods [6], [7] have been introduced and perform satisfactorily in some environments. However, their performance may degrade in case of illumination change and viewpoint variation, due to the nature of imaging mechanism. Compared with images, Lidar scans have larger perception fields and are insensitive to illumination changes. In road scenes where the ground area can be approximately regarded as a plane, Lidar scans are also insensitive to orientations. These advantages mean that the Lidar-based place recognition methods may have better performance than the image-based ones. Actually, current Lidar-based methods [8]-[11] have shown the superiority of Lidar scans in terms of place retrieval ability in large-scale environments.

Manuscript received March 17, 2021; accepted June 16, 2021. Date of publication June 22, 2021; date of current version July 2, 2021. This letter was recommended for publication by Associate Editor A. Nuechter and Editor S. Behnke upon evaluation of the reviewers' comments. This work was supported by the Aviation Science Foundation under Grant #20181976. (*Corresponding author: Hui-Liang Shen.*)

Lun Luo, Si-Yuan Cao, and Bin Han are with the College of Information Science and Electronic Engineering, Zhejiang University, Hangzhou 310027, China (e-mail: luolun@zju.edu.cn; karlcao@hotmail.com; 21931020@zju.edu.cn).

Hui-Liang Shen and Junwei Li are with the College of Information Science and Electronic Engineering, Zhejiang University, Hangzhou 310027, China, and also with the Ningbo Research Institute, Zhejiang University, Ningbo 315100, China (e-mail: shenhl@zju.edu.cn; lijunwei7788@zju.edu.cn).

Digital Object Identifier 10.1109/LRA.2021.3091386

In this letter we present BVMatch, a novel Lidar-based place recognition method that is capable of estimating 2D relative poses. Based on the assumption that the ground area can be regarded as a plane in road scenes, BVMatch divides the ground plane into uniform grids and accumulates the scanned points in each grid to form the BV image. This orthogonal projection procedure makes the transform of the BV image pair a rigid transform, which is similar to the 2D transform of the Lidar scan pair. However, we note that BV image inherits the sparsity of Lidar scans and suffers severe intensity distortion that cannot be handled by classic image features such as SIFT [2]. To tackle this problem, we introduce the bird's-eye view feature transform (BVFT) that is invariant to intensity and rotation changes of BV images. We build BVFT based upon the maximum index map (MIM) of Log-Gabor filter responses [12], and theoretically analyze its characteristics on orientation shift. Based on BVFT, BVMatch uses the bag-of-words (BoW) approach for place recognition and uses RANSAC [13] for relative pose estimation. In summary, the main contributions of this work are:

- We propose a novel local descriptor called BVFT that is insensitive to intensity and rotation variation of BV images, with which the BVMatch framework can unify the Lidar place recognition and pose estimation tasks.
- We theoretically prove that BVFT can achieve rotation invariance by shifting the orientations in the local MIM.
- We experimentally validate on three large-scale datasets that BVMatch outperforms the state-of-the-arts on both place recognition and pose estimation.

## II. RELATED WORK

The Lidar-based place recognition methods can be approximately classified into two categories, i.e. the methods that directly utilize raw point clouds and the methods that use images as intermediate representations.

The first category focuses on extracting local features or global descriptors from raw point clouds. Rusu *et al.* [14] present fast point feature histograms (FPFH) to align point cloud pairs. Tombari *et al.* [15] introduce signatures of histograms for local surface description (SHOT) for surface matching. Bosse and Zlot [16] extract Gestalt keypoints and descriptors from point clouds and use keypoints voting to recognize places. These local features exploit local characteristics of point cloud using geometric measures such as normals and curvatures, while BVFT encodes the structure information in BV images. Instead of using local descriptors, SegMatch [17] extracts higher-level segments

2377-3766 © 2021 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information. belonging to partial or full objects, and then matches these segments for place recognition using the learned descriptors. OneShot [18] follows SegMatch but extracts efficient segments from a single Lidar frame. These two segment-based methods adopt the frame-to-map matching framework while our BV-Match is a frame-to-frame method. M2DP [19] projects a point cloud to multiple 2D planes and generates a density signature for points in each of the planes. The singular value decomposition (SVD) components of the signature are then used to compute a global descriptor. PointNetVLAD [8] leverages PointNet [20] to extract features of point clouds and uses NetVLAD [5] to generate global descriptors. Zhang and Xiao [10] introduce the point contextual attention network (PCAN) that enables the network to pay more attention to the task-relevent features when generating global descriptors. LPD-Net [9] adopts an adaptive local feature extraction module to extract the local structures and uses a graph-based neighborhood aggregation module to discover the spatial distribution of local features. These learning based methods do not use local keypoints and thus they cannot estimate relative poses. DH3D [11] uses the 3D local feature encoder and detector to extract local descriptors. It embeds the descriptors to a global feature for place recognition and align the matched Lidar pairs using RANSAC. In comparison, our BVFT is handcrafted and does not need training.

The second category projects Lidar scans to images for place recognition. Steder et al. [21] extract local features from range images of Lidar scans. Unlike BV image, the range image is not Euclidean in nature since it is generated with the polar projection. Kim et al. [22] partition the ground space into bins according to both azimuthal and radial directions. Then they propose the scan context descriptor in a 2D matrix form. They further extend scan context to three channels, and introduce the concept of scan context image (SCI) [23]. Place recognition is achieved by classifying the SCIs using a convolutional network. However, the pose estimation problem remains unsolved. Cao et al. [24] project point clouds to cylindrical images and leverage Gabor filters to describe the contours of scenes. A histogram based descriptor is generated for every Lidar scan. Similar to the image formations used in the aforementioned methods, the content of cylindrical image is complex due to the special projection procedure. Thus it cannot estimate relative poses as well. To solve the problem, OverlapNet [25] adopts a siamese network to estimate an overlap of range images and provides a relative yaw angle estimate of Lidar pair. However, it is incapable of estimating translation. On the contrary, BVMatch uses a simple orthogonal projection procedure that makes the transforms of BV image pair and Lidar scan pair similar, with which BVMatch can localize query Lidar scans and produce 2D poses.

We note that BV image is similar to the occupancy grid map [26] since they both partition space into grids. There are several methods [27], [28] designed for the occupancy grid map matching in the literature. Our BVMatch differs from these methods in two aspects. First, they target on two different types of representation. More specifically, the BV image is built with point cloud density while the occupancy grid map is formed by occupancy probabilities. Second, BVMatch is a frame-to-frame method leveraging single 3D Lidar scans while [27], [28] are map-to-map based methods.

## III. BV IMAGE BASED PLACE RECOGNITION

The BVMatch method consists of the place recognition step and the pose estimation step as illustrated in Fig. 1. In the place recognition step, BVMatch generates the BV image and performs frame retrieval with the BoW approach. In the pose estimation step, BVMatch matches the BV image pair using RANSAC and reconstructs the coarse 2D pose of the Lidar pair with a similar transform. At last, BVMatch uses iterative closest point (ICP) refinement with the 2D pose as an initial guess to accurately align the Lidar pair.

In the following, we first introduce the BV image generation mechanism and then present the coarse 2D pose reconstruction of Lidar scan pairs. Finally, we demonstrate the offline dictionary and the keyframe database creation.

## A. BV Image Generation

There are several types of BV images, such as the maximum height map used in place recognition [22] and the density map used in object detection [29]. Both maps summarize the vertical shape of surrounding structures. The maximum height map uses the coordinate values of points with the maximal height, while the density map leverages the point cloud density. The maximum height map is sensitive to poses because the coordinate values of points can vary severely when robot moves. In contrast, the density map is more robust because the density of point cloud does not depend on specific points. For this reason, we use the latter as our BV image representation in this work.

Let  $\mathcal{P} = \{P_i | i = 1, \dots, N_p\}$  be a point cloud formed by points  $P_i = (x_i, y_i, z_i)$  and  $N_p$  the number of points in the cloud. Suppose that the point cloud is collected in road scenes. The x-axis is pointing to the right, the y-axis is pointing forward, and the z-axis is pointing upward. In this coordinate system, the x-y plane is the ground plane. Given a point cloud  $\mathcal{P}$ , we first use a voxel grid filter with the leaf size of g meters to evenly distribute the points. Then we discretize the ground space into grids with resolution of g meters. The point cloud density is the number of points located in each grid. We consider a [-C m, C m] cubic window centered at the coordinate origin. Then BV image B(u, v) is a matrix of size  $\lceil \frac{2C}{g} \rceil \times \lceil \frac{2C}{g} \rceil$ . The BV intensity B(u, v) is defined as

$$B(u,v) = \frac{\min(N_g, N_m)}{N_m},\tag{1}$$

where  $N_g$  denotes the number of points in the grid at position (u, v) and  $N_m$  the normalization factor.  $N_m$  is set to be the 99th percentile of the point cloud density.

BV image is a compressed representation of point cloud and describes the 2.5D structural information of an egocentric environment. It ignores the point distribution along the z-axis while keeping the rigid structures on the x-y plane. We find that the poles, facades, and guideposts in road scenes usually form edges in the BV image. These features have good repeatability and remain stable when the robot moves. We adopt FAST [30]



Fig. 1. BVMatch framework. In the place recognition step, the dictionary and keyframe database are built offline. Given a query frame, BVMatch generate the corresponding BV image and then extract keypoints and BVFT descriptors. Place recognition is achieved by the bag-of-words approach. In the pose estimation step, BVMatch uses RANSAC to match the BV image pair and reconstructs the 2D relative pose of the Lidar pair. Finally, BVMatch uses the 2D pose as initial guess of iterative closest point (ICP) to align the Lidar pair accurately.

for feature detection and use the keypoints for registration in the BVMatch pipeline.

#### B. Pose Reconstruction

As BV image discretizes the ground space uniformly, the transform of the Lidar scan pair  $(\mathcal{P}_i, \mathcal{P}_j)$  is similar to that of the BV image pair  $(B_i(u, v), B_j(u, v))$ . After obtaining the transform of  $(B_i(u, v), B_j(u, v))$  we have  $B_i(u, v) = B_j(u, 'v')$ , with the coordinate

$$u' = \cos(\theta)u + \sin(\theta)v + t_u$$
$$v' = -\sin(\theta)u + \cos(\theta)v + t_v, \tag{2}$$

where  $(t_u, t_v, \theta)$  are transform parameters. The transform matrix  $T_{ij}$  of the pair  $(\mathcal{P}_i, \mathcal{P}_j)$  is

$$T_{ij} = \begin{pmatrix} \cos(\theta) & \sin(\theta) & gt_u \\ -\sin(\theta) & \cos(\theta) & gt_v \\ 0 & 0 & 1 \end{pmatrix}$$
(3)

where g denotes the leaf size of the voxel grid filter.

#### C. Dictionary and Keyframe Database

BVMatch leverages the bag-of-words approach to extract global descriptors and uses a keyframe database to detect the best match frame of a query Lidar scan.

Bag-of-words approaches assume that similar structures in the environment produce similar distributions of features. To learn this distribution in our scene, we extract a rich set of BVFT descriptors from the training BV images. We then apply K-means clustering to obtain a total of b clusters. Each cluster is a word, and the centroids of these words form a bag-of-words dictionary. We use the dictionary to encode a point cloud in terms of words by classifying the BVFT descriptors of the point cloud which words they are. To weight those words that are frequent and less discriminative, we use frequency-inverse document frequency (TF-IDF) [4]. Finally, we get a global descriptor with size b for every point cloud.

Keyfame database stores BV images with their global poses and descriptors. A robot traverses a specific place and collects Lidar scans along the way. By building a map of the place using SLAM or GPS information, every Lidar scan collected in this traversal is tagged with a global pose. We extract keyframe Lidar scans every S meters the robot moves and generate a global descriptor for every keyframe. The keyframe database is built using all these global descriptors, poses, and BV images.

#### IV. PROPOSED BVFT DESCRIPTOR

Although BV image preserves the vertical structures that are stable in a scene, it suffers severe intensity distortion due to the sparsity nature of Lidar scans. To extract distinct local descriptors, we first leverage Log-Gabor filters to compute the local responses of BV images. We then construct a maximum index map (MIM) [12] originally used for multi-modal image matching. Finally, we build bird's-eye view feature transform (BVFT) that is insensitive to intensity and rotation variations of BV images.

## A. Maximum Index Map (MIM)

For simplicity, we use the polar coordinate to represent images in the following. The polar coordinate  $(\rho, \theta)$  of a Euclidean coordinate (u, v) is defined as

$$\rho = \sqrt{(u - \bar{u})^2 + (v - \bar{v})^2}$$
  

$$\theta = \arctan(v - \bar{v}, u - \bar{u}), \qquad (4)$$

where  $(\bar{u}, \bar{v})$  denotes the center position of an image.



Fig. 2. A simple example illustrating the MIM rotation invariance.  $\mathcal{P}_{\alpha}$  is the rotated point cloud of  $\mathcal{P}$ . *B* and *MIM* are the BV image and MIM of  $\mathcal{P}$ , respectively. *B*<sub> $\alpha$ </sub> and *MIM*<sub> $\alpha$ </sub> are the BV image and MIM of  $\mathcal{P}_{\alpha}$ , respectively. It can be observed that the content of *MIM* and rotated *MIM*<sub> $\alpha$ </sub> are quite different. By shifting (i), the difference map (e) tends to be zero which confirms the rotation invariance design.

We build MIM based on Log-Gabor filters [31]. The 2D Log-Gabor filter in frequency domain is defined as

$$\mathcal{L}(f,\omega,s,o) = \exp\left(\frac{-(\log(f/f_s))^2}{2(\log(\sigma_f/f_s))^2}\right) \exp\left(\frac{-(\omega-\omega_o)^2}{2\sigma_\omega^2}\right),\tag{5}$$

where  $f_s$  and  $\omega_o$  are the center frequencies at scale s and orientation o,  $\sigma_f$  and  $\sigma_\omega$  are width parameters. For a Log-Gabor filter set of  $N_s$  scales and  $N_o$  orientations, the filters are designed to evenly cover the spectrum. In our implementation, we follow the scales setting in [12]. We select the orientations from the set  $\mathcal{O} = \{0, \pi/N_o, 2\pi/N_o, \ldots, (N_o - 1)\pi/N_o\}$  and hence  $\omega_o = o\pi/N_o$ . Note that the frequency part of the 2D log-Gabor filter is isotropic, which is essential for BVFT design.

We use a bank of Log-Gabor filters of  $N_s$  scales and  $N_o$  orientations to build the MIM. For a Log-Gabor filter at scale s and orientation o, Let  $L(\rho, \theta, s, o)$  be the corresponding filter of (5) in the spatial domain,  $A(\rho, \theta, s, o)$  be the Log-Gabor amplitude response of  $B(\rho, \theta)$  at orientation o and scale s, we have

$$A(\rho, \theta, s, o) = ||B(\rho, \theta) * L(\rho, \theta, s, o)||_2, \tag{6}$$

where \* represents the convolution operation. Then the Log-Gabor amplitude response at orientation o is

$$A(\rho, \theta, o) = \sum_{s} A(\rho, \theta, s, o).$$
(7)

The maximum index map (MIM) is a map of the orientations with the maximal Log-Gabor responses, that is

$$MIM(\rho, \theta) = \arg\max_{o} A(\rho, \theta, o).$$
(8)

## B. BVFT Descriptor

The Log-Gabor filters, widely employed in image processing, are jointly localized in position, orientation and spatial frequency [31]. In BVFT, we employ Log-Gabor filters to capture the orientation information of the rigid vertical structures in scenes using (8). This makes MIM insensitive to intensity distortion of BV images. However, MIM is still sensitive to rotation. In the following we show that, by shifting the orientations in the MIM with respect to a specified dominant orientation, rotation invariance can be achieved.

Fig. 2(a) is a randomly selected Lidar scan  $\mathcal{P}$  from the Oxford RobotCar Radar dataset [32]. Fig. 2(f) is the rotated point cloud  $\mathcal{P}_{\alpha}$  of  $\mathcal{P}$  with a rotation angle  $\alpha$  about the z-axis. Fig. 2(g) and (b) are the corresponding BV images  $B_{\alpha}$  and B related as

$$B_{\alpha}(\rho,\theta) = B(\rho,\theta+\alpha). \tag{9}$$

Fig. 2(h) and (c) are  $MIM_{\alpha}$  and MIM respectively. It can be observed that the contents of B and  $B_{\alpha}$  are the same, while the contents of MIM and  $MIM_{\alpha}$  are quite different. To eliminate the rotation difference, we rotate  $MIM_{\alpha}$  by  $-\alpha$  and obtain Fig. 2(i). It is observed the difference of two MIMs shown in Fig. 2(d) is significant. To obtain rotation invariance, this difference should be eliminated.

We compute the Log-Gabor amplitude response of  $B_{\alpha}$  with orientation *o*, and have

$$A_{\alpha}(\rho, \theta, o) = \sum_{s} ||B_{\alpha}(\rho, \theta) * L(\rho, \theta, s, o)||_{2}$$
$$= \sum_{s} ||B(\rho, \theta + \alpha) * L(\rho, \theta, s, o)||_{2}.$$
(10)

Since the frequency part of the 2D log-Gabor filter is isotropic in the frequency domain (see (5)), every filter at scales s can be obtained by rotating other filter at the same scale by some angle. When  $\alpha$  is an angle within the orientation set  $\mathcal{O}$ ,

$$\mathcal{L}(f,\omega,s,o) = \mathcal{L}(f,\omega+\alpha,s,o_{\alpha}), \tag{11}$$

where

$$o_{\alpha} = \operatorname{mod}\left(o - \alpha \frac{N_o}{\pi}, N_o\right). \tag{12}$$

6080

Here we use a mod operation on the orientation because the  $N_o$  orientations form a ring structure [12]. Since 2D Fourier transform is rotation invariant, the 2D spatial Log-Gabor filter has the same property as the frequency one,

$$L(\rho, \theta, s, o) = L(\rho, \theta + \alpha, s, o_{\alpha}).$$
(13)

Substituting (13) to (10) yields

$$A_{\alpha}(\rho, \theta, o) = \sum_{s} ||B(\rho, \theta + \alpha) * L(\rho, \theta + \alpha, s, o_{\alpha})||_{2}$$
$$= A(\rho, \theta + \alpha, o_{\alpha}).$$
(14)

This means that the Log-Gabor response image  $A_{\alpha}(\rho, \theta, o)$  can be obtained by rotating  $A(\rho, \theta + \alpha, o_{\alpha})$  by  $\alpha$ . Substituting (14) to (8) yields

$$MIM_{\alpha}(\rho, \theta) = \arg\max_{o} A_{\alpha}(\rho, \theta, o)$$
  
=  $\arg\max_{o} A(\rho, \theta + \alpha, o_{\alpha})$   
=  $\operatorname{mod}\left(MIM(\rho, \theta + \alpha) - \alpha \frac{N_{o}}{\pi}, N_{o}\right), (15)$ 

which reveals that when the point cloud is rotated by any angle within the orientation set  $\mathcal{O}$ , its MIM can be obtained by a simple circle shift operation on the MIM of the unrotated one. Fig. 2 validates the above analysis. Fig. 2(j) is the rotated and shifted  $MIM_{\alpha}$  using (15). It can be observed from Fig. 2(e) that the difference map is very small, in despite of the quantization error of orientation set  $\mathcal{O}$ .

We further extend the above analysis onto image patches and design rotation invariant local descriptors. For every detected keypoint, we find its dominant orientation. We build a local histogram h(o) of pixel values over a square MIM patch  $patch(\rho, \theta)$ with  $J \times J$  pixels centered at the keypoint. When computing the histogram, the increments are weighted by a Gaussian window function centered at the keypoint with mean (0,0) and standard deviation (J/2, J/2). Suppose that the peak of the histogram is at orientation  $o_m$ , i.e.,  $o_m = \arg \max_o h(o)$ , the dominant orientation is

$$\beta = \pi \frac{o_m}{N_o}.$$
 (16)

Then we rotate the patch by  $\beta$  and shift the patch,

$$\operatorname{patch}_{\beta}(\rho, \theta) = \operatorname{mod}(\operatorname{patch}(\rho, \theta + \beta) - o_m, N_o).$$
(17)

With this treatment, the local patch is rotation invariant. We divide the patch into  $l \times l$  sub-grids and build a distribution histogram for each sub-grid. These histograms are concatenated to form a BVFT feature vector of size  $l \times l \times N_o$ .

The above analysis is conducted under the assumption that the rotation angles are within the orientation set  $\mathcal{O}$ . However, (11) and (15) still hold when the rotation angles are within the set  $\hat{\mathcal{O}} = \{\pi, (N_o + 1)\pi/N_o, (N_o + 2)\pi/N_o, \dots, (2N_o - 1)\pi/N_o\}$ . Thus, we cannot determine whether the dominant orientation is  $\pi \frac{o_m}{N_o}$  or  $(2\pi - \pi \frac{o_m}{N_o})$  given  $o_m$ . To avoid this ambiguity, we rotate every shifted MIM patch by  $\pi$  and assign every keypoint with an additional descriptor generated using this patch. When the rotation angle is not within  $\mathcal{O}$  or  $\hat{\mathcal{O}}$ , it is obvious that (15) does not hold. In fact, BVFT can not cover continuous orientation because the Log-Gabors are band-pass filters. In practice, we have found that BVFT has good matching ability in continuous orientation, and we will show that our BVFT has excellent performance for the place recognition problem in the experiment section.

## V. EXPERIMENTS

We compare the place recognition capability of our BVMatch with M2DP [19], PointNetVLAD [8], PCAN [10], LPD-Net [9] and DH3D [11], among which the last four methods are deep learning ones. We compare the pose estimation performance with SIFT [2] (e.g., extracting SIFT descriptors from BV images and using RANSAC for pose estimation), OverlapNet [25] and DH3D. Note that DH3D can perform both place recognition and pose estimation like BVMatch. The source codes of all the competing methods are publicly available on the websites.<sup>1</sup> For the sake of fairness, we fine-tuned all the deep learning methods on the training sets of the individual datasets. For M2DP and SIFT, we use the parameters provided by the authors. For BVMatch, we empirically set the parameters as follows: voxel grid size q = 0.4 meters, Log-Gabor filter frequency scales  $N_s = 4$  and orientations  $N_o = 6$ , local patch size J = 96, number of subgrids  $l \times l = 6 \times 6$ , and size of bag-of-words dictionary b = 10000. Accordingly, the size of BVFT descriptor is  $l \times l \times N_o = 216$ .

#### A. Datasets

We conduct the experiments on three long-term and largescale datasets: Oxford RobotCar dataset [32], Oxford RobotCar Radar dataset [33] and NCLT dataset [34].

1) Oxford RobotCar dataset: The Oxford RobotCar dataset was created using a 2D LMS LiDAR sensor mounted on a car that repeatedly drives through Oxford at different times traversing a 10 km route. It captures many different combinations of weather, traffic, and pedestrians over one year. However, it only provides 2D Lidar scans. To make a 3D point cloud with enough points, we use the relative pose between a previous scan and a recent scan to accumulate sequential 2D Lidar scans with a trajectory length of 80 m. The relative pose is obtained from GPS/IMU readings. We select 45 sequences from the dataset for evaluating (the same sequences used in PointNetVLAD [8]).

2) Oxford RobotCar Radar dataset: The Oxford RobotCar Radar dataset was created at the same place as the Oxford Robot-Car dataset in 7 days. It provides sparse 3D Lidar point clouds generated by two Velodyne32-VLP LiDAR sensors mounted on the left and right sides of a car. In this work, we only use the data of the left Lidar. Since the sequences collected on the same day are quite similar, we randomly select a sequence from each day and get 7 sequences.

<sup>1</sup>https://github.com/LiHeUA/M2DP https://github.com/mikacuy/pointnetvlad https://github.com/XLechter/PCAN https://github.com/Suoivy/LPD-net https://github.com/JuanDuGit/DH3D https://github.com/Suoivy/OverlapNet

TABLE I TRAINING AND TEST SETS OF THE DATASETS

Dataset: NCLT											
	Train Set		Test Set								
Seq.	2013-02-23,2013-04-05		2012-01-15,	2012-02-04,	2012-03-17,	2012-05-26,	2012-06-15,	2012-08-20,	2012-09-28,	2012-10-28,	2012-11-16
Frames	2050	2036	1402	1050	1062	1155	711	1103	1012	1040	873
Dataset: Oxford Radar RobotCar											
	Train	Train Set Test Set									
Seq.	2019-01-10-11-46,2	2019-01-11-12-26	2019-01-14	-12-05-52, 2	019-01-15-13	3-53-14, 2019	9-01-16-14-1:	5-33, 2019-0	1-17-14-03-0	0, 2019-01-1	8-12-42-34
Frames	2334	2133	118	1	1164		1132		1222	1	164



Average recalls with respect to the N-number of top candidates on different datasets. Fig. 3.

TABLE II AVERAGE RECALLS AT TOP-1% AND TOP-1

	Oxf	ord	Oxford	Radar	NCLT		
	@1%	@1	@1%	@1	@1%	@1	
M2DP [19]	85.8	76.5	76.4	50.9	57.5	36.2	
PN-VLAD [8]	96.8	86.7	94.7	86.6	78.3	62.8	
PCAN [10]	97.0	87.7	93.8	80.3	82.2	59.0	
LPD-Net [9]	97.4	91.4	95.5	90.0	90.7	72.5	
DH3D [11]	97.8	88.7	94.6	78.2	83.0	59.4	
BVMatch (Ours)	98.8	95.7	99.2	93.9	95.2	83.6	

3) NCLT dataset: The NCLT dataset was created at the University of Michigan North Campus using a Velodyne32-HDL LiDAR sensor with varying routes. It provides sparse 3D Lidar point clouds. We use 11 sequences from the dataset to evaluate the methods.

Note that the three datasets are of different characteristics. Specifically, the Lidar scans in the Oxford Radar RobotCar dataset and the NCLT dataset are sparser than those in the Oxford RobotCar dataset due to the usage of different types of Lidar sensor. We split the sequences in every dataset into training and testing sequences. For the Oxford RobotCar dataset, we use 15 sequences collected in 2014 for training and 30 sequences collected in 2015 for test. For the Oxford RobotCar Radar dataset and the NCLT dataset, the sequence partition is summarized in Table I. All the training sequences are sampled every 2 meters, while the test sequences are sampled every 10 meters. Point cloud data in a [-50 m, 50 m] cubic window of a frame are used for all methods.

## B. Place Recognition

We first train the bag-of-words dictionary in each dataset with the training sequences. We then build keyframe databases with

100 100 \*\*\*\*\*\*\*\* 90 90 Average Recall (%) Average Recall (%) 80 80 70 70 M2PD M2PD 60 PointNetVLAD 60 PointNetVLAD PCAN PCAN 50 50 LPD-Net LPD-Net DH3D DH3D 40 40 BVMatch (Ours BVMatch (Ours) 30 30 10 15 20 25 10 15 20 25 N-number of Top Candidate N-number of Top Candidate (a) Oxford Radar RobotCar (b) NCLT

M2PD

PCAN

DH3D

15

(c) NCLT

10

LPD-Net

PointNetVLAD

BVMatch (Ours)

25

Fig. 4. Evaluation of generalization ability in terms of average recalls with respect to the N-number of top candidates. The methods are trained on the Oxford RobotCar dataset and tested on the Oxford Radar Robotcar and NCLT datasets.

TABLE III AVERAGE RECALLS AT TOP-1% AND TOP-1 ON THE OXFORD RADAR ROBOTCAR AND NCLT DATASETS FOR GENERALIZATION ABILITY EVALUATION

	Oxford	Radar	NCLT		
	@1%	@1	@1%	@1	
M2DP [19]	76.4	50.9	57.5	36.2	
PN-VLAD [8]	71.8	53.6	64.2	44.8	
PCAN [10]	69.0	41.3	81.2	55.9	
LPD-Net [9]	93.2	76.9	78.4	57.1	
DH3D [11]	85	47.7	74.3	40.3	
BVMatch (Ours)	95.3	70.3	67.8	54.9	

global descriptors for the test sequences. For every keyframe in a sequence, we use the Euclidean distance of the global descriptors to retrieve the best match from other sequences. The match is positive if the ground truth distance of the matched frame is less than t meters. We use Top-N recall to evaluate the place recognition ability. Fig. 3 shows the average recall results at Top-N with the ground truth threshold distance t = 25 meters. Table II illustrates the average recalls at Top-1 and 1%. It can

		Ox	ford Rob	otCar		NCLT				
	RTE (m)		RRE (deg)		SR (%)	RTE (m)		RRE (deg)		SR (%)
	Mean	Std.	I. Mean Std.			Mean	Std.	Mean	Std.	
SIFT [2]	0.77	0.47	0.53	0.47	82.1	0.69	0.46	1.49	1.23	53.3
OverlapNet [25]	-					-		2.43	1.42	15.4
DH3D [11]	0.43	0.31	1.23	0.89	99.9	1.16	0.54	3.43	1.06	14.8
BVMatch (Ours)	0.51	0.33	0.41	0.42	99.1	0.57	0.38	1.08	1.00	94.5

 TABLE IV

 Pose Estimation Results on the Oxford Robotcar and the NCLT Dataset



Fig. 5. Average recalls with respect to varying orientations on the two Oxford datasets.



Fig. 6. Parameter study. (a) Average recalls at Top-1 with respect to varying distance thresholds on the Oxford RobotCar dataset. (b) Average recalls at Top-1 with respect to orientation numbers  $N_o$  and dictionary size b on the the Oxford Radar RobotCar dataset.

be observed that our BVMatch outperforms other methods in every dataset.

To validate the generalization ability of the methods, we train the deep learning based methods and the bag-of-words dictionary on the Oxford Robotcar dataset and evaluate the methods on the other two datasets. Fig. 4 illustrates that our BVMatch outperforms the competitors on the Oxford Radar dataset while is of moderate performance on the NCLT dataset. Table III further shows that BVMatch ranks first at Top-1% and ranks second at Top-1 on the Oxford Radar dataset, but does not perform very well on the NCLT dataset.

Robustness to orientation variance is important since the robot may have arbitrary headings in real-world applications. Since BVFT is rotation invariant, BVMatch is inherently orientation invariant. We evaluate place retrieval ability of the methods with varying orientations on the sequences 2015-02-03 and 2015-02-10 in the Oxford RobotCar dataset and on the sequences 2019-01-14-12-05-52 and 2019-01-15-13-53-14 in the Oxford Radar RobotCar dataset. The average recall at Top-1 at every rotation angle is shown in Fig. 5. It is observed that the recall rates of DH3D are relatively stable while other deep learning methods drop dramatically when the angle deviates from 0°. Our BVMatch performs closely to DH3D on the Oxford RobotCar dataset, but much better than DH3D on the Oxford Radar RobotCar dataset in which the Lidar scans are more sparse. Note that the recall curve of BVMatch is fluctuant due to the angle quantization of the orientation set  $\mathcal{O}$ .

An extension evaluation on retrieval ability with respect to the distance threshold t is shown in Fig. 6(a). It is clear that BVMatch performs better than other methods under various distance threshold settings.

We conducted an experiment on the Oxford Radar RobotCar dataset to investigate the influence of Log-Gabor filter orientations  $N_o$  and the dictionary size b. Fig. 6(b) shows that when  $N_o$  and b increase, the average recall at Top-1 grows and saturate at about  $N_o = 6$  and b = 10000.

#### C. Pose Estimation

We evaluate the performance of pose estimation on the sequences 2015-02-03 and 2015-02-10 of the Oxford RobotCar dataset and on the sequences 2012-02-04 and 2012-01-15 of the NCLT dataset. In the evaluation, we retrieve the best match of each query frame in a sequence (e.g., 2015-02-03) from the other sequence in the same dataset (e.g., 2015-02-03) from the other sequence in the same dataset (e.g., 2015-02-10). We obtain 752 matched pairs from the Oxford dataset and 969 pairs from the NCLT dataset. Since the ground truth poses in the Oxford dataset are biased, we align the Lidar pairs by ICP and regard the aligned results as the ground truth.

Our BVMatch uses the BVFT descriptor with RANSAC to find the relative poses. We compared BVMatch with SIFT [2], OverlapNet [25] and DH3D [11]. Since OverlapNet is only suitable for single 3D scans and cannot use the data accumulated from 2D Lidar scans in the Oxford dataset, we only evaluate OverlapNet on the NCLT dataset. For fair comparison, we fine-tuned OverlapNet and DH3D on the datasets. We compute the relative translation error (RTE) and the relative rotation error (RRE). We regard the pose as a successful estimation when the RTE and RRE are below 2 meters and 5°, respectively. Table IV shows that, on the Oxford dataset, our BVMatch achieves the best RRE while performs slightly worse than DH3D on the RTE and SR metrics. On the NCLT dataset, BVMatch outperforms the competitive methods by a large margin. This further validates the superiority of BVMatch in handling sparse Lidar scans.

#### D. Runtime Evaluation

We implemented the BVFT descriptor generation and RANSAC using C++ and implemented BVMatch framework using Python. Our platform running the experiments is a desktop computer equipped with an Intel Quad-Core 3.40 GHz i5-7500 CPU and 16 GB RAM. The average time cost for each frame is 0.29 seconds to extract BVFT descriptors, 0.02 seconds to generate a global descriptor, 0.05 seconds to perform retrieval, and 0.23 seconds to register BV image pair using RANSAC. The total time cost for each frame is 0.59 seconds.

## VI. CONCLUSION

This letter has presented BVMatch, a novel Lidar-based frame-to-frame place recognition method. The method employs bird's-eye view (BV) image as the intermediate representation of point cloud, and introduces the BVFT descriptor to perform match. Compared with the state-of-the-arts, BVMatch is more efficient in place recognition and is able to estimate 2D relative poses. In our future work, we will focus on further reducing its time complexity and improving its generalization ability.

#### REFERENCES

- C. Cadena *et al.*, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Trans. Robot.*, vol. 32, no. 6, pp. 1309–1332, Dec. 2016.
- [2] D. Lowe, "Distinctive image features from scale-invariant keypoints," Int. J. Comput. Vis., vol. 20, pp. 91–110, Nov. 2004.
- [3] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "ORB: An efficient alternative to SIFT or SURF," in *Proc. Int. Conf. Comput. Vis.*, 2011, pp. 2564–2571.
- [4] D. Nister and H. Stewenius, "Scalable Recognition With a Vocabulary Tree," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, vol. 2, 2006, pp. 2161–2168.
- [5] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, "NetVLAD: CNN architecture for weakly supervised place recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 5297–5307.
- [6] D. Galvez-Lpez and J. D. Tardos, "Bags of binary words for fast place recognition in image sequences," *IEEE Trans. Robot.*, vol. 28, no. 5, pp. 1188–1197, Oct. 2012.
- [7] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardos, "ORB-SLAM: A versatile and accurate monocular SLAM system," *IEEE Trans. Robot.*, vol. 31, no. 5, pp. 1147–1163, Oct. 2015.
- [8] M. A. Uy and G. Hee Lee, "PointNetVLAD: Deep point cloud based retrieval for large-scale place recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 4470–4479.
- [9] Z. Liu *et al.*, "LPD-Net: 3D point cloud learning for large-scale place recognition and environment analysis," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019, pp. 2831–2840.
- [10] W. Zhang and C. Xiao, "PCAN: 3D attention map learning using contextual information for point cloud based retrieval," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 12428–12437.
- [11] J. Du, R. Wang, and D. Cremers, "DH3D: Deep hierarchical 3D descriptors for robust large-scale 6DoF relocalization," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 744–762.
- [12] J. Li, Q. Hu, and M. Ai, "RIFT: Multi-modal image matching based on radiation-variation insensitive feature transform," *IEEE Trans. Image Process.*, vol. 29, pp. 3296–3310, Dec. 2019.

- [13] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, pp. 381–395, 1981.
- [14] R. B. Rusu, N. Blodow, and M. Beetz, "Fast point feature histograms (FPFH) for 3 d registration," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2009, pp. 3212–3217.
- [15] F. Tombari, S. Salti, and L. Di Stefano, "Unique signatures of histograms for local surface description," in *Proc. Eur. Conf. Comput. Vis.*, Springer, 2010, pp. 356–369.
- [16] M. Bosse and R. Zlot, "Place recognition using keypoint voting in large 3D lidar datasets," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2013, pp. 2677– 2684.
- [17] R. Dubé, A. Cramariuc, D. Dugas, J. I. Nieto, R. Siegwart, and C. Cadena, "Segmap: 3D segment mapping using data-driven descriptors," in *Proc. Robot.: Sci. Syst.*, Jun. 2018.
- [18] S. Ratz, M. Dymczyk, R. Siegwart, and R. Dubé, "OneShot Global Localization: Instant LiDAR-Visual Pose Estimation," 2020, arXiv:2003.13279.
- [19] L. He, X. Wang, and H. Zhang, "M2DP: A novel 3D point cloud descriptor and its application in loop closure detection," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2016, pp. 231–237.
- [20] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "PointNet: Deep learning on point sets for 3D classification and segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 652–660.
- [21] B. Steder, M. Ruhnke, S. Grzonka, and W. Burgard, "Place recognition in 3 d scans using a combination of bag of words and point feature based relative pose estimation," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2011, pp. 1249–1255.
- [22] G. Kim and A. Kim, "Scan context: Egocentric spatial descriptor for place recognition within 3D point cloud map," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2018, pp. 4802–4809.
- [23] G. Kim, B. Park, and A. Kim, "1-day learning, 1-year localization: Longterm LIDAR localization using scan context image," *IEEE Robot. Automat. Lett.*, vol. 4, no. 2, pp. 1948–1955, Apr. 2019.
- [24] F. Cao, F. Yan, S. Wang, Y. Zhuang, and W. Wang, "Season-invariant and viewpoint-tolerant LiDAR place recognition in GPS denied environments," *IEEE Trans. Ind. Electron.*, vol. 68, no. 1, pp. 563–574, Jan. 2021.
- [25] X. Chen et al., "OverlapNet: Loop closing for LiDAR-based SLAM," in Proc. Robot.: Sci. Syst., Jun. 2020.
- [26] A. Elfes, "Using occupancy grids for mobile robot perception and navigation," *Computer*, vol. 22, no. 6, pp. 46–57, 1989.
- [27] J.-L. Blanco, J. Gonzalez, and J.-A. Fernandez-Madrigal, "A new method for robust and efficient occupancy grid-map matching," in *Proc. Iberian Conf. Pattern Recognit. Image Anal.*, 2007, pp. 194–201.
- [28] M. Rapp, T. Giese, M. Hahn, J. Dickmann, and K. Dietmeyer, "A featurebased approach for group-wise grid map registration," in *Proc. IEEE Int. Conf. Intell. Transp. Syst.*, 2015, pp. 511–516.
- [29] X. Chen, H. Ma, J. Wan, B. Li, and T. Xia, "Multi-view 3D object detection network for autonomous driving," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 1907–1915.
- [30] E. Rosten and T. Drummond, "Machine learning for high-speed corner detection," in *Eur. Conf. Comput. Vis.*, Springer, 2006, pp. 430–443.
- [31] S. Fischer, F. Šroubek, L. Perrinet, R. Redondo, and G. Cristóbal, "Selfinvertible 2D log-gabor wavelets," *Int. J. Comput. Vis.*, vol. 75, no. 2, pp. 231–246, 2007.
- [32] D. Barnes, M. Gadd, P. Murcutt, P. Newman, and I. Posner, "The oxford RobotCar dataset: A radar extension to the oxford RobotCar dataset," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2020, pp. 6433–6438.
- [33] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, "1 year, 1000 km: The oxford RobotCar dataset," *Int. J. Robot. Res.*, vol. 36, no. 1, pp. 3–15, 2017.
- [34] N. Carlevaris-Bianco, A. K. Ushani, and R. M. Eustice, "University of michigan north campus long-term vision and lidar dataset," *Int. J. Robot. Res.*, vol. 20, pp. 91–110, Nov. 2016.