A Fast and Robust Place Recognition Approach for Stereo Visual Odometry Using LiDAR Descriptors

Jiawei Mo¹ and Junaed Sattar²

Abstract—Place recognition is a core component of Simultaneous Localization and Mapping (SLAM) algorithms. Particularly in visual SLAM systems, previously-visited places are recognized by measuring the appearance similarity between images representing these locations. However, such approaches are sensitive to visual appearance change and also can be computationally expensive. In this paper, we propose an alternative approach adapting LiDAR descriptors for 3D points obtained from stereo-visual odometry for place recognition. 3D points are potentially more reliable than 2D visual cues (e.g., 2D features) against environmental changes (e.g., variable illumination) and this may benefit visual SLAM systems in long-term deployment scenarios. Stereo-visual odometry generates 3D points with an absolute scale, which enables us to use LiDAR descriptors for place recognition with high computational efficiency. Through extensive evaluations on standard benchmark datasets, we demonstrate the accuracy, efficiency, and robustness of using 3D points for place recognition over 2D methods.

I. Introduction

Visual SLAM (vSLAM) is an important capability for field robots, especially where GPS signal reception is weak or unavailable (e.g., in urban or underwater settings). In these systems, visual odometry (VO) is used to build a local map and estimate ego-motion to assist in robot navigation. However, significant error can accumulate throughout the process, which causes odometry estimates to diverge from the correct path. Some form of a "loop closure" approach (e.g., Bag-of-Words [1] [2]) is required to recognize previously-visited places and bring non-local constraints into the system to get a globally consistent map and trajectory. Place recognition thus enables loop closures and improves VO accuracy.

Classical place recognition methods for vision-based systems usually rely on 2D images. Each location is represented by an image taken at that place. To determine the possibility of two locations being the same place, the similarity of their corresponding images is evaluated (see Sec. II). However, visual odometry methods provide additional information usable for place recognition purposes. The depth of points (*i.e.*, the distance of these points from the camera) on 2D images can be partially or fully recovered by monocular or multicamera visual odometry, respectively. The 3D structure of the scene can potentially provide important information for place recognition; however, 2D place recognition methods ignore this. The 3D structure is more robust than 2D images in a dynamic environment (*e.g.*, under varying illumination).

The authors are with the Department of Computer Science and Engineering, Minnesota Robotics Institute, University of Minnesota Twin Cities, Minneapolis, MN, USA. $\{^1 \text{moxxx066}, ^2 \text{junaed}\}$ at umn.edu.





(a) Parks Road in spring.

(b) Parks Road in winter.



(c) Holywell Street in spring.

(d) Holywell Street in winter.

Fig. 1: Images from the RobotCar dataset in different seasons. Note the significant changes in appearance.

The motivation is also biological, as humans rely strongly on 3D structures for place recognition [3].

On the other hand, a rich body of literature exists on place recognition methods using 3D points from LiDAR (Light Detection and Ranging) sensors. LiDAR sensors scan the 3D structure of the environment rather than its visual appearance, making LiDAR-based place recognition more robust against environmental changes such as appearance and brightness (see Fig. 1). Another benefit of LiDAR methods is their high computational efficiency, and our evaluations demonstrate this when comparing 2D image-based and 3D LiDAR methods (see Sec. IV).

In this work, we adapt LiDAR place recognition methods, in particular, LiDAR descriptors, into visual odometry systems for place recognition purposes. The goal is to enable accurate and robust place recognition in a computationally efficient way for a vision-based system in a dynamic environment. The proposed approach imitates a LiDAR range scan from 3D points generated by stereo-visual odometry, which enables us to adapt LiDAR descriptors.

Several challenges must be overcome for applying LiDAR-based methods to vision-based systems. First, the 3D points generated by visual odometry are distributed in a frustum due to the much narrower field-of-view of cameras (excluding omnidirectional cameras) compared to most LiDAR sensors. The pose of the frustum changes with the camera, which is not desirable for place recognition. The second challenge is how to (and even if it is necessary to) adapt image intensity information into LiDAR-based methods, as such information is not available to LiDAR sensors. We address these chal-

lenges in this work. To the best of our knowledge, this is the first approach that uses global LiDAR descriptors for place recognition in vSLAM systems. The main contributions of this work, discussed in Sec. III, are as follows:

- Adapting global LiDAR descriptors to a vision-based system for place recognition,
- Achieving high accuracy and robustness against visual appearance changes,
- Achieving lower computational cost over existing approaches.

We evaluate the proposed method on the KITTI dataset [4] and the Oxford RobotCar dataset [5]. We demonstrate the robustness of our method against drastic visual appearance changes across seasons as recorded in the RobotCar dataset, and show that it achieves higher accuracy and computational efficiency over existing methods. Further performance improvement is achieved by augmenting the LiDAR descriptor with image intensity information.

II. RELATED WORK

In the field of vSLAM, ORB-SLAM2 [6] is a recent development that demonstrates high accuracy and computational efficiency. In ORB-SLAM2, loop closure is detected by Bag-of-Words (BoW) using ORB features [7]. A vocabulary tree is used in BoW to speed up feature matching and subsequent place queries. However, if the features are highly repetitive (e.g., plants), BoW may fail; an example is given in Fig. 6. Similarly, LSD-SLAM [8] adopts FAB-MAP [9] for place recognition. Other than BoW, Fisher vectors [10] and VLAD [11] also focus on 2D features. On the other hand, global image descriptors are also used to decide the similarity between images for place recognition. GIST [12] is one example which encodes spatial layout properties (spatial frequencies) of the scene. It exhibits high accuracy if the viewing angle does not significantly change.

Recently, researchers adopted deep learning to place recognition and achieved impressive performance (e.g., NetVLAD [13] and [14]). NetVLAD trained a convolutional neural network to extract learned features and proposed a generalized VLAD layer to describe the image automatically. Their accuracy is promising but their computational cost is usually high so that they are not widely used in real-time vSLAM systems.

Neither BoW nor GIST is robust against visual appearance change, which is not ideal for long term (*e.g.*, from summer to winter) vSLAM applications, in addition to being computationally expensive. In ORB-SLAM2, place recognition runs in a separate execution thread to achieve real-time performance. Direct vSLAM systems (*e.g.*, [15], [16]) have become popular in the past decade, which achieve higher performance in certain scenarios. Adapting BoW into direct vSLAM systems is challenging because features are not selected with the goal of being matched across frames. In LSD-SLAM mentioned above, an additional set of features are detected and matched separately, which are used specifically for place recognition, at a higher computational cost. In LDSO [17], the point selection strategy of its direct

vSLAM system [15] is tuned to flavor features that can be matched across frames to enable BoW. Our proposed approach for place recognition, however, is more elegant for direct vSLAM systems if stereo cameras are available.

A number of 3D place recognition methods have been designed for RGB-D cameras or LiDAR sensors. RGB-D Mapping [18] uses ICP [19] to detect loop closure and RANSAC [20] to get an initial pose for ICP. For LiDAR, place recognition methods can be categorized into local descriptors and global descriptors. Local descriptors use a subset of the points and describe them in a local neighborhood. Examples are Spin Image [21] and SHOT [22]. Spin Image describes a keypoint by a histogram of points lying in each bin of a vertical cylinder centered at that keypoint. SHOT creates a sphere around a keypoint and describes that keypoint by the histogram of normal angles in each bin in the sphere. Global methods describe the **entire set** of points. These methods can be more computationally efficient. Recent development includes NDT [23], M2DP [24], Scan Context [25], and DELIGHT [26]. NDT classifies keypoints into line, plane, and sphere classes according to their neighborhoods. A histogram of these three classes is created to represent the point cloud. M2DP projects points onto multiple planes, and the histogram of point count in each bin on each projection plane is concatenated to get a signature of the point cloud. Scan Context aligns the point cloud to the vertical direction and represents it by the histogram of the maximal height of each bin on the horizontal plane. DELIGHT focuses on LiDAR intensity; the scan sphere is divided into 16 parts and the histogram of LiDAR intensity in each part is concatenated to represent the point cloud.

Cieslewski et. al. [27] looked into the possibility of using the 3D points triangulated from Structure-from-Motion or vSLAM for place recognition. They proposed the NBLD descriptor [27] for the 3D points from a vision-based system. A keypoint is described by its neighborhood points in a vertical cylinder. The point density of each bin in the cylinder is calculated and compared with neighborhoods to create a binary descriptor of that keypoint. Ye et. al. [28] extended NBLD with a neural network. The vertical cylinder of NBLD is created in the same way; however, a neural network is trained to describe the cylinders, instead of calculating the point density. These are novel approaches in adopting point cloud descriptors into vision-based systems for place recognition.

In this work, we adapt global LiDAR descriptors into stereo-visual odometry for robust and efficient place recognition under visual appearance change. Direct vSLAM systems can easily adopt the proposed approach for place recognition without modifying their point selection strategy.

III. METHODOLOGY

Similar to the idea of [27], our method recognizes places based on the 3D points generated by visual odometry. The main difference is that the visual odometry in this work is running on stereo cameras. Specifically, we use SO-DSO [29] as our stereo-visual odometry for its high accuracy and

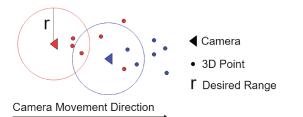


Fig. 2: A demonstration of the points generated by visual odometry, with different colors representing different keyframes. 3D points are coming into the desired range as the camera moves.

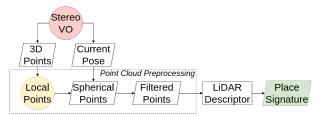


Fig. 3: An overview of the proposed approach. The basis lies in the "Point Cloud Preprocessing" block, where 3D points obtained by stereo VO are used to imitate a LiDAR scan, so that efficient place recognition can be performed.

computational efficiency. To the best of our knowledge, SO-DSO is the only direct stereo-visual odometry that is robust to repetitive textures (Fig. 6). We choose SO-DSO to demonstrate that the proposed method works for direct vS-LAM systems. However, any multi-camera visual odometry is applicable here. Since the 3D points generated by stereo-visual odometry have an absolute scale, we can describe them efficiently using global LiDAR descriptors. The goal is high accuracy and computational efficiency, and robustness to environmental change.

A. Point Cloud Preprocessing

Due to the narrow field-of-view of the cameras, the 3D points generated by stereo-visual odometry are located in a frustum determined by the camera pose. If we apply a global descriptor directly inside the frustum, place recognition will be very sensitive to the viewing angle.

To solve this issue, we propose a simple but effective method that transforms 3D points from stereo-visual odometry in the frustum to an omnidirectional LiDAR-shaped (spherical) 3D point cloud. The proposed method is illustrated in Fig. 3. Stereo-visual odometry generates keyframes with camera poses and associated 3D points. We maintain what we refer to as a local points list. For each incoming keyframe, we store all its 3D points into the list. To imitate a LiDAR scan for the current keyframe, we iterate through the *local points* list: if the distance of the point is within the desired LiDAR range, we transform it to the current keyframe coordinate by current pose, then put it into the spherical points list. Here we assume the camera motion is predominantly in the forward direction so that we continuously have points coming into the desired range to compensate for points leaving the range, as illustrated in





Fig. 4: Snapshots of an imitated LiDAR scan at the end of KITTI seq. 06.

Fig. 2. The *spherical points* may contain duplicate points; for robustness and computational efficiency, we filter them to get the final *filtered points*. An example of an imitated LiDAR scan is given in Fig. 4.

Caching local points enables us to imitate LiDAR scans with denser omnidirectional points. Since visual odometry generates locally accurate camera poses and 3D points, concatenating 3D points transformed from multiple nearby keyframes to imitate a LiDAR scan is feasible.

B. Point Cloud Description

The next step in the proposed method is to describe the filtered points and get a place signature for the keyframe, for which we rely on global LiDAR descriptors. This is preferable for two reasons: the first is for its computational efficiency when describing and matching the point clouds; the second reason is due to the fact that the point clouds we have are generated by visual odometry, they are not as consistent and dense as the ones from a LiDAR. Many local descriptors, such as Spin Image, depend on the surface normal, for which dense point clouds are required, which would be problematic in this case. We choose DELIGHT [26], M2DP [24], and Scan Context [25] as our global descriptors since they are state-of-the-art LiDAR descriptors for place recognition that are robust to sparse and inconsistent point clouds. The high-level ideas of them are illustrated in Fig. 5.

DELIGHT: DELIGHT operates on LiDAR intensities. The LiDAR scan sphere is divided into 16 bins by radius, azimuth, and elevation. Each bin is described by the histogram of LiDAR intensities inside, which are concatenated to form the signature representing the entire LiDAR scan. To make the descriptor less sensitive to rotation and translation, the raw LiDAR scan is aligned to a reference frame obtained by Principal Components Analysis (PCA) [30]. As discussed in [26], there are four versions of the signatures due to the ambiguity of PCA.

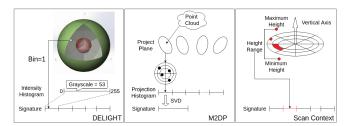


Fig. 5: Demonstration of LiDAR descriptors used in this work: DELIGHT, M2DP, and Scan Context.

Analogous to LiDAR scan intensities, the 3D points from the visual odometry have grayscale intensities from the images. We simply replace LiDAR intensities with grayscale intensities and adapt the DELIGHT descriptor into our system. Each histogram is composed of 256 bins since the grayscale intensity ranges from 0 to 255. Although DELIGHT does not use 3D structural information, we include it in this work to contrast it against M2DP and Scan Context as both these methods use this information. This is done to highlight the value of 3D structures for robust place recognition against visual appearance change.

M2DP: M2DP is a global descriptor that demonstrates high accuracy and efficiency. The point cloud is projected onto multiple planes, and each plane is separated into individual bins by radius and azimuth. The distributions of the projection onto bins are concatenated to form a signature for the point cloud. For computational and memory efficiency, singular-value decomposition (SVD) is used to compress the signature. As in DELIGHT, PCA is used to align the point cloud.

In this work, we augment the M2DP descriptor using grayscale intensity from the visual odometry. Specifically, when projecting the point cloud onto each plane, we not only count the number of points projected onto each bin, but also calculate the average grayscale intensity. Therefore, we have two types of signatures (namely the point count signature and the intensity signature) for each place. To make the intensity signature less sensitive to illumination, we binarize the intensity by comparing it to the global average intensity. The intuition is to highlight the bright bins. To improve accuracy, we adopt the four versions of signature from DELIGHT. We include M2DP in our method as it utilizes 3D structure for place recognition and is not limited to urban scenarios like Scan Context.

Scan Context: Scan Context is a straightforward yet effective descriptor designed for LiDAR scans obtained in urban areas. The LiDAR scan is aligned with respect to the gravitational axis which is measured externally (e.g., with an IMU). Then the horizontal circle plane is separated into multiple bins by radius and azimuth. In each bin, the maximum height is found and concatenated to form a signature for the current place.

To fit Scan Context into our system, we make the following modifications. First, since we want to avoid using additional sensors, we adopt the PCA method used in DELIGHT and M2DP to align the point cloud. Second, due to the PCA

ambiguity, we replace maximum height with height range (maximum height - minimum height). Lastly, we generate the intensity signature as in the modified M2DP. Scan Context is the most efficient and accurate among the three descriptors, as shown in experimental evaluations (Sec. IV).

C. Place Recognition

Based on the place signatures, we are able to determine the similarity between places. We generate a difference matrix by calculating the signature distance from each query place to every place in the reference database. For DELIGHT and M2DP, we take the shortest distance from the query signature to all four possible signatures of the reference place. For Scan Context, even though we can use the same approach as for DELIGHT and M2DP, we compare against all possible yaw angles (as the original Scan Context) for maximal accuracy. The distance of DELIGHT is based on the chi-squared test as described in [26]. For M2DP and Scan Context, the distance is simply the Euclidean distance between normalized (L2norm=1) signatures. As we have two types of signatures (structure signature and intensity signature), we get two individual difference matrices $\mathbf{D_s}$ and $\mathbf{D_i}$. We fuse them in Eq. 1 by normalizing (mean=0, std=1) each row (representing each query) and adding them with a relative weight w_s :

$$\mathbf{D_{fused}} = w_s \cdot N_{row}(\mathbf{D_s}) + N_{row}(\mathbf{D_i}) \tag{1}$$

With the difference matrix, each query place (row) is matched to a reference place with the smallest difference value among all candidates (along the row).

IV. EXPERIMENTAL EVALUATION

To evaluate the proposed method, we compare the results internally among DELIGHT, M2DP, and Scan Context, as well as externally to BoW, GIST, and [28], as real-time performance at a low computational cost is a key aspect of our work.

Implementation

We use the fast (but less accurate) setting (\leq 800 active points, \leq 6 frames in optimization) of SO-DSO to estimate camera poses and generate 3D points. We found that the fast setting is good enough for our purpose and introducing more points decreases the computational efficiency. The outputs of SO-DSO (poses and points) are used by the proposed methods for place recognition.

The authors of M2DP have published their Matlab code, which we re-implement using C++. We similarly re-implement DELIGHT and Scan Context. When preprocessing the point cloud as discussed in Sec. III-A, we set the LiDAR range r=45.0m. For DELIGHT and M2DP, spherical points are filtered in polar coordinate with 1-degree angular resolution. We keep the closest point along each ray originating from the polar center. For Scan Context, however, we filter points in the Cartesian coordinate with $1.5m \times 0.75m \times 1.5m$ resolution. We switch to the Cartesian coordinate and assign a higher resolution along the vertical axis because Scan Context focuses on height. For the KITTI

Method		***			GIST
				0.893 0.788	
				0.011 0.012	
				0.867 0.809	
Seq. 06	0.916 0.531	0.946 0.671	0.897 0.679	0.968 0.963	0.925 0.729
Seq. 07	0.000 0.000	0.000 0.000	0.000 0.000	0.713 0.627	0.350 0.149

TABLE I: AUC (first number) and maximal recall at 100% precision (second number) on KITTI dataset.



Fig. 6: Snapshot of KITTI seq. 02 at a revisited place.

dataset [4], there are 2706.2 (1610.6) points on average in each *filtered points* using polar (Cartesian) filtering. For DELIGHT, the radius of the inner/outer sphere is set to 10/45 meters, respectively. The parameters of M2DP and Scan Context are set to default values. For more details, our implementations are available online¹. When fusing difference matrices in Eq. 1, we set a higher weight $w_s=2$ to structure because structure is more reliable than grayscale intensity.

The implementation of BoW comes from ORB-SLAM2; the implementation of GIST is available online².

Evaluation

For place recognition accuracy, we focus on the area under the precision-recall curve (AUC) and the maximal recall at 100% precision (no false positives, *i.e.*, no errors) as two indices. Larger AUC means more places are recognized with fewer errors. To some extent, AUC reflects the discrimination power of a place recognition algorithm. Larger maximal recall at 100% precision indicates that more places are recognized before making any mistakes, it is important because a single false positive might significantly affect the accuracy of the entire SLAM algorithm. In the ideal case, both AUC and maximal recall should be 1. Other than accuracy, computational efficiency is also evaluated. We take apart each algorithm and compare the components individually. The accuracy and efficiency are evaluated on KITTI dataset [4] and Oxford RobotCar Dataset [5].

KITTI Dataset

The KITTI dataset is one of the most influential datasets for benchmarking autonomous driving research. The odometry dataset comes with 22 stereo sequences. However, only the first 11 sequences have ground-truth publicly available. Among them, sequences {00, 02, 05, 06, 07} have loop closure segments, which are used in this section.

Accuracy: When computing precision and recall, two places are considered to be revisited places if their distance

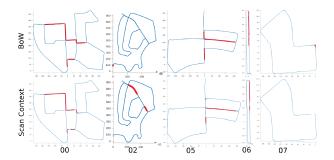


Fig. 7: Places recognized (marked as red) by BoW and by Scan Context on KITTI dataset at 100% precision.

Method	DELI.	M2DP	S.C.	BoW	GIST
Imitate LiDAR Scan (c++)	1.151	1.204	0.692	-	-
Desc. extraction (c++)	0.082	46.10	0.123	37.41	160.0
Query descriptor (Matlab)	l	ı		1	ı
Total	104.4	50.72	8.149	152.4	161.1

TABLE II: Run time analysis in milliseconds.

is smaller than 10 meters, which is relatively small as the distance between trajectories can be small (*e.g.*, sequence 06 in Fig. 7).

Table I reports the accuracy of each algorithm. BoW achieves the best accuracy in all sequences other than sequence 02. Each sequence in the KITTI dataset is recorded continuously in a short period of time, there is not much visual appearance change. Hence, feature matching is robust and BoW works perfectly. For sequence 02, the BoW approach fails to recognize places because the revisited places are occupied with repetitive textures (i.e., plants in Fig. 6), for which feature matching is unreliable. The accuracy of the adapted LiDAR approaches is not as good as BoW or GIST but their AUCs are still fairly high on sequences {00, 02, 05, 06}. Scan Context achieves the best overall accuracy among the adapted LiDAR approaches. Sequence 07 is special because there is only a small segment of loop closure when the vehicle comes back to the starting place. Thus, the accuracy of BoW and GIST is not very reliable. None of the proposed 3D methods detects any loop because there is not enough overlapped trajectory to accumulate 3D points.

Fig. 7 plots the loops detected by BoW and Scan Context. Scan Context fails to recognize places in short segments. This is because the proposed method needs enough overlapped trajectory to accumulate 3D points. This is a limitation of the proposed method. However, it recognized places with repetitive textures in sequence 2. Nevertheless, we showed that our implementation of each algorithm works on the KITTI dataset (at least for BoW and GIST; we will further validate the proposed 3D approaches using the RobotCar dataset).

Efficiency: Table II reports the run-time required to query a place in the database. The test platform is based on an Intel i7-6700 with 16GB of RAM. The run-time is calculated in sequence 06.

Point filtering in Scan Context is slightly faster because of the simple Cartesian coordinate filtering. For descriptor

¹https://github.com/IRVLab/so_dso_place_recognition

²http://lear.inrialpes.fr/software

Tests	Spr.	Spr.	Spr.	Spr.	Sum.	Sum.	Sum.	Fall	Fall	Win.
	Spr.	Sum.								
Dates	05-19	05-19	05-19	05-19	08-13	08-13	08-13	10-30	10-30	02-10
	05-22	08-13	10-30	02-10	07-14	10-30	02-10	11-28	02-10	12-12

TABLE III: Test sequences on RobotCar dataset.

generation, DELIGHT is the fastest due to its straightforward mechanism. Scan Context is the second-fastest because calculating the height range is also efficient. M2DP is the slowest one among all the adapted 3D methods. The entire set of points is projected onto multiple planes, followed by SVD compression, which is computationally expensive. For BoW and GIST, their high accuracy on the KITTI dataset is achieved at a high computational cost. For place query, GIST is the quickest since it simply calculates the Euclidean distance between two descriptors. This is followed by M2DP, for which we have calculated all four descriptors of each place, the distance between two places is simply the smallest Euclidean distance. Scan Contest is slightly slower because the query descriptor is matched against all possible yaws. DELIGHT and BoW are much slower. The chi-squared test in DELIGHT is computationally expensive. For Bow, we extract 4000 features for each image and use a vocabulary tree with about 1 million nodes. Hence, both descriptor extraction and matching of BoW are computationally expensive.

Scan Context achieves the highest overall efficiency that can run in real-time in most vSLAM systems. Although BoW and GIST achieve higher accuracy, they are much slower than Scan Context.

RobotCar Dataset

The RobotCar dataset is challenging for place recognition since visual appearance and brightness changes drastically. Snapshots of RobotCar are shown in Fig. 1. The testing pairs are given in Table III, covering all combinations of seasons.

Accuracy: Since the authors of [28] have not published their code, we evaluate our algorithms using the same settings for fair comparisons. Specifically, we use the same segment as illustrated in Fig. 5(a) of [28] for testing. When computing precision and recall, we also use 25 meters as the GPS distance threshold.

Table IVa shows the AUC of each algorithm running the tests in Table III. Since we cannot re-run the tests in [28], data of [28] and NBLD are taken directly from [28], which are just for reference; whereas the rows marked "DELI.", "M2DP", and "S.C." represent our approaches adapting these three global descriptors. Table IVb illustrates the maximal recall with 100% precision. Scan Context ("S.C.") achieves both the highest AUC and the highest recall in most tests. Scan Context depends on the height range for place recognition. In the RobotCar dataset, the maximum/minimum height is usually from buildings/ground. Therefore, the height range is not very sensitive to seasonal change. GIST behaves best in the rest of the tests. It works well because the viewing angle is mostly unchanged in the RobotCar dataset. M2DP is the next-best performing approach, which has a relatively high accuracy when there is

Tests	Spr.	Spr.	Spr.	Spr.	Sum.	Sum.	Sum.	Fall	Fall	Win.
Tests	Spr.	Sum.	Fall	Win.	Sum.	Fall	Win.	Fall	Win.	Win.
[28]	0.774	0.736	0.589	0.419	0.764	0.557	0.489	0.599	0.443	0.597
NBLD	0.651	0.700	0.611	0.351	0.672	0.496	0.379	0.454	0.351	0.491
DELI.	0.869	0.677	0.445	0.040	0.836	0.612	0.008	0.498	0.003	0.014
M2DP	0.900	0.851	0.498	0.322	0.853	0.519	0.276	0.540	0.349	0.541
S.C.	0.956	0.944	0.782	0.729	0.928	0.779	0.618	0.644	0.491	0.814
BoW	0.558	0.342	0.208	0.300	0.305	0.418	0.371	0.002	0.293	0.001
GIST	0.932	0.918	0.679	0.778	0.914	0.694	0.738	0.003	0.606	0.000

(a) AUC.

Tests	Spr.	Spr.	Spr.	Spr.	Sum.	Sum.	Sum.	Fall	Fall	Win.
	Spr.					Fall				
DELI.						l				
M2DP										
S.C.	0.758	0.558	0.408	0.322	0.685	0.415	0.325	0.346	0.247	0.519
BoW	0.032	0.021	0.023	0.031	0.005	0.034	0.100	0.000	0.043	0.000
GIST	0.794	0.377	0.242	0.176	0.503	0.242	0.156	0.000	0.109	0.000

(b) Maximal recall at 100% precision.

TABLE IV: Place recognition accuracy on RobotCar dataset.

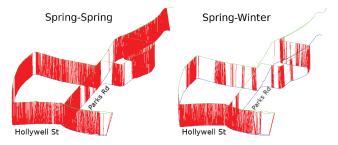


Fig. 8: Places recognized (marked as red) by Scan Context at 100% precision.

less visual appearance change (e.g., Test a: Spring-Spring). However, in different seasons, the trees along the streets have vastly different appearances (e.g., for losing leaves) as illustrated in Fig. 1. M2DP projects all 3D points to get a signature, so its accuracy drops. DELIGHT suffers more from visual appearance change between seasons because it purely depends on grayscale intensity, which indicates the importance of 3D structure for place recognition. BoW has the worst performance; a potential reason is that feature matching is sensitive to changing scene factors such as trees and traffic (vehicular and pedestrian).

Fig. 8 shows the places recognized by Scan Context at 100% precision. For the easy case (Spring-Spring), Scan Context recognizes most of the places. For the challenging case (Spring-Winter), with the reference of Fig. 1, most places along Holywell street are correctly recognized since it is occupied mostly with buildings; but most places along Parks road are not recognized because there are many trees (maximum height) on both sides of the road.

Fig. 9 shows the decrease in accuracy of each algorithm against seasonal visual appearance change. For AUC, Scan Context and GIST outperform the rest in terms of robustness; for maximal recall, BoW has an abnormal curve because all of its maximal recall values are very low (Table IVb). Other than that, Scan Context is the most robust one. Therefore, we conclude that Scan Context is robust against visual appearance change.

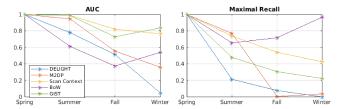


Fig. 9: Robustness against seasonal visual appearance change, using spring as query season. Values are normalized by Spring-Spring.

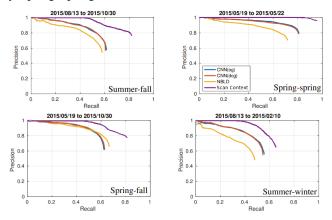


Fig. 10: Precision-recall curves of Scan Context compared with that of [28].

Fig. 10 shows the comparison of the precision-recall curve between the adapted Scan Context and the methods in [28]. It validates that Scan Context adapted in this work outperforms [28] in accuracy.

Intensity Contribution: Table V shows the AUC and maximal recall of Scan Context using structure only, intensity only, and both structure and intensity. Scan Context with intensity only performs poorly on the RobotCar dataset, since grayscale intensity changes drastically throughout different seasons. However, augmenting the structure descriptor with intensity information clearly improves the maximal recall, even though it does not obviously improve the AUC.

We do not include the efficiency comparison on the RobotCar dataset since we use the identical setup with the KITTI dataset and the results are similar.

Use Case Analysis

After the experiments, we claim that the use case of the proposed method is for forward-moving vehicles (for accumulating points) equipped with stereo cameras in visually

Tests	Spr.	Spr.	Spr.	Spr.	Sum.	Sum.	Sum.	Fall	Fall	Win.
	Spr.	Sum.							Win.	
Ct	0.955	0.940	0.762	0.699	0.931	0.753	0.610	0.652	0.500	0.778
Structure	0.270	0.216	0.390	0.154	0.279	0.066	0.105	0.147	0.049	0.134
Intensity	0.834	0.645	0.344	0.112	0.831	0.390	0.086	0.290	0.096	0.478
intensity	0.230	0.050	0.039	0.021	0.151	0.057	0.027	0.056	0.027	0.032
Fused	0.956	0.944	0.782	0.729	0.928	0.779	0.681	0.644	0.491	0.814
	0.758	0.558	0.408	0.322	0.685	0.415	0.325	0.346	0.247	0.519

TABLE V: AUC (top sub-rows) and maximal recall (bottom sub-rows) at 100% precision of Scan Context with structure and/or grayscale intensity.

changing environments (*e.g.*, RobotCar dataset), where the proposed approach recognizes place with high accuracy, efficiency, and robustness. It also works with repetitive texture (*e.g.*, Sequence 02 of KITTI dataset). Additionally, adopting the proposed approach is easier than adopting BoW for direct vSLAM.

The conventional BoW works on individual images, and there is no forward-moving constraint. It achieves higher accuracy, especially for small loop segments (*e.g.*, Fig. 7) when there is not much visual appearance change.

V. CONCLUSIONS

In this paper, we propose a novel place recognition approach for stereo-visual odometry. Instead of 2D image similarity, we depend on the 3D points generated by the visual odometry to determine the correlation between places. The 3D points from stereo systems with an absolute scale are used to imitate LiDAR scans which are fed into three global LiDAR descriptors, which are DELIGHT, M2DP, and Scan Context. We augment the descriptors with grayscale intensity information. Experiments on the KITTI dataset and RobotCar dataset show the accuracy, efficiency, and robustness of the proposed method.

For the next step, we will integrate the proposed method into state-of-the-art stereo-visual odometry algorithms for loop closure detection, and quantify the performance improvement in visually challenging environments. Furthermore, we intend to extend the proposed approach with deep learning and compare to learning-based place recognition approaches.

ACKNOWLEDGMENT

This work was partially supported by the Minnesota Robotics Institute Seed (MnRI) Grant.

REFERENCES

- [1] J. Sivic and A. Zisserman, "Video google: A text retrieval approach to object matching in videos," in *null*, p. 1470, IEEE, 2003.
- [2] D. Gálvez-López 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.
- [3] R. Epstein and N. Kanwisher, "A Cortical Representation of the Local Visual Environment," *Nature*, vol. 392, no. 6676, p. 598, 1998.
- [4] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [5] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, "1 Year, 1000km: The Oxford RobotCar Dataset," *The International Journal of Robotics Research*, vol. 36, no. 1, pp. 3–15, 2017.
- [6] 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.
- [7] E. Rublee, V. Rabaud, K. Konolige, and G. R. Bradski, "ORB: An Efficient Alternative to SIFT or SURF.," in *ICCV*, vol. 11, p. 2, Citeseer, 2011.
- [8] J. Engel, T. Schöps, and D. Cremers, "LSD-SLAM: Large-Scale Direct Monocular SLAM," in *European conference on computer vision*, pp. 834–849, Springer, 2014.
- [9] M. Cummins and P. Newman, "FAB-MAP: Probabilistic Localization and Mapping in the Space of Appearance," *The International Journal* of Robotics Research, vol. 27, no. 6, pp. 647–665, 2008.
- [10] F. Perronnin and C. Dance, "Fisher Kernels on Visual Vocabularies for Image Categorization," in 2007 IEEE Conference on Computer Vision & Pattern Recognition, pp. 1–8, IEEE, 2007.

- [11] H. Jégou, M. Douze, C. Schmid, and P. Pérez, "Aggregating Local Descriptors into a Compact Image Representation," in CVPR 2010-23rd IEEE Conference on Computer Vision & Pattern Recognition, pp. 3304–3311, IEEE Computer Society, 2010.
- [12] A. Oliva and A. Torralba, "Building the Gist of a Scene: The Role of Global Image Features in Recognition," *Progress in brain research*, vol. 155, pp. 23–36, 2006.
- [13] 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*, pp. 5297–5307, 2016.
- [14] Z. Chen, A. Jacobson, N. Sünderhauf, B. Upcroft, L. Liu, C. Shen, I. Reid, and M. Milford, "Deep Learning Features at Scale for Visual Place Recognition," in 2017 IEEE International Conference on Robotics and Automation (ICRA), pp. 3223–3230, IEEE, 2017.
- [15] J. Engel, V. Koltun, and D. Cremers, "Direct Sparse Odometry," *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 3, pp. 611–625, 2017.
- [16] C. Forster, M. Pizzoli, and D. Scaramuzza, "Svo: Fast semi-direct monocular visual odometry," in 2014 IEEE international conference on robotics and automation (ICRA), pp. 15–22, IEEE, 2014.
- [17] X. Gao, R. Wang, N. Demmel, and D. Cremers, "Ldso: Direct sparse odometry with loop closure," in 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 2198–2204, IEEE, 2018.
- [18] P. Henry, M. Krainin, E. Herbst, X. Ren, and D. Fox, "RGB-D mapping: Using Kinect-Style Depth Cameras for Dense 3D Modeling of Indoor Environments," *The International Journal of Robotics Research*, vol. 31, no. 5, pp. 647–663, 2012.
- [19] P. J. Besl and N. D. McKay, "Method for Registration of 3-D Shapes," in *Sensor Fusion IV: Control Paradigms and Data Structures*, vol. 1611, pp. 586–607, International Society for Optics and Photonics, 1992.
- [20] M. A. Fischler and R. C. Bolles, "Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.
- [21] A. E. Johnson and M. Hebert, "Using Spin Images for Efficient Object Recognition in Cluttered 3D Scenes," *IEEE Transactions on pattern* analysis and machine intelligence, vol. 21, no. 5, pp. 433–449, 1999.
- [22] F. Tombari, S. Salti, and L. Di Stefano, "A Combined Texture-Shape Descriptor for Enhanced 3D Feature Matching," in 2011 18th IEEE international conference on image processing, pp. 809–812, IEEE, 2011
- [23] M. Magnusson, The Three-Dimensional Normal-Distributions Transform: An Efficient Representation For Registration, Surface Analysis, and Loop Detection. PhD thesis, Örebro universitet, 2009.
- [24] 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), pp. 231–237, IEEE, 2016.
- [25] 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), pp. 4802–4809, IEEE, 2018.
- [26] K. P. Cop, P. V. Borges, and R. Dubé, "DELIGHT: An Efficient Descriptor for Global Localisation using LiDAR Intensities," in 2018 IEEE International Conference on Robotics and Automation (ICRA), pp. 3653–3660, IEEE, 2018.
- [27] T. Cieslewski, E. Stumm, A. Gawel, M. Bosse, S. Lynen, and R. Siegwart, "Point Cloud Descriptors for Place Recognition using Sparse Visual Information," in 2016 IEEE International Conference on Robotics and Automation (ICRA), pp. 4830–4836, IEEE, 2016.
- [28] Y. Ye, T. Cieslewski, A. Loquercio, and D. Scaramuzza, "Place Recognition in Semi-Dense Maps: Geometric and Learning-Based Approaches," in *Proc. Brit. Mach. Vis. Conf.*, pp. 72–1, 2017.
- [29] J. Mo and J. Sattar, "Extending Monocular Visual Odometry to Stereo Camera System by Scale Optimization," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS). To appear, November 2019. arXiv preprint arXiv:1905.12723.
- [30] F. Tombari, S. Salti, and L. Di Stefano, "Unique Signatures of Histograms for Local Surface Description," in *European Conference* on Computer Vision, pp. 356–369, Springer, 2010.