

# Spectral Geometric Verification: Re-Ranking Point Cloud Retrieval for Metric Localization

Kavisha Vidanapathirana<sup>1,2</sup>, Peyman Moghadam<sup>1,2</sup>, Sridha Sridharan<sup>2</sup>, Clinton Fookes<sup>2</sup>

**Abstract**—In large-scale metric localization, an incorrect result during retrieval will lead to an incorrect pose estimate or loop closure. Re-ranking methods propose to take into account all the top retrieval candidates and re-order them to increase the likelihood of the top candidate being correct. However, state-of-the-art re-ranking methods are inefficient when re-ranking many potential candidates due to their need for resource intensive point cloud registration between the query and each candidate. In this work, we propose an efficient spectral method for geometric verification (named SpectralGV) that does not require registration. We demonstrate how the optimal inter-cluster score of the correspondence compatibility graph of two point clouds represents a robust fitness score measuring their spatial consistency. This score takes into account the subtle geometric differences between structurally similar point clouds and therefore can be used to identify the correct candidate among potential matches retrieved by global similarity search. SpectralGV is deterministic, robust to outlier correspondences, and can be computed in parallel for all potential candidates. We conduct extensive experiments on 5 large-scale datasets to demonstrate that SpectralGV outperforms other state-of-the-art re-ranking methods and show that it consistently improves the recall and pose estimation of 3 state-of-the-art metric localization architectures while having a negligible effect on their runtime. The open-source implementation and trained models are available at: <https://github.com/csiro-robotics/SpectralGV>.

**Index Terms**—Localization, Recognition, SLAM

## I. INTRODUCTION

ACCURATE 6 Degrees of Freedom (DoF) metric localization is an essential component in many applications in embodied intelligence. For an embodied agent (mobile robot, autonomous car etc.) to autonomously operate in any environment, it must first be able to estimate its pose. The task of robust global localization in large-scale environments remains an open problem as current methods struggle to differentiate between structurally similar places with subtle differences. When scaling this problem to city-scale environments (or larger), this task is made efficient through a hierarchical two-step formulation. This includes the retrieval-based place recognition step, which estimates a set of coarse place candidates, and a pose estimation step which estimates

the alignment between the query and the top place candidate to obtain the 6DoF pose.

In the point cloud domain, the tasks of place recognition and 6DoF alignment (*i.e.*, registration) have been explored largely in isolation [1]–[3]. Recently, several methods have combined the two tasks to offer a complete solution for metric localization [4]–[6]. They follow a correspondence-based approach, where local-features are used to estimate correspondences between the query point cloud and the top-candidate returned by global-descriptor matching, and the pose is obtained using the estimated correspondences. The top place candidate retrieved via global descriptor matching may not always be correct, and an incorrect place candidate will result in an infeasible pose estimation task.

To address such limitations in retrieval, re-ranking methods take into consideration many top retrieval candidates and re-order them such that correct candidates will be ranked higher. Re-ranking is performed using pre-defined criteria specific to each method and often involves utilizing additional information complementary to the global descriptors. Fig. 1 shows a scenario where an incorrect top-1 candidate can be replaced with the correct candidate via re-ranking. Re-ranking has become popular in other domains of information retrieval [7], [8] but has limited utility in point cloud retrieval due to current re-ranking methods either lacking robustness or being inefficient in handling large-scale outdoor point clouds. Current methods are not suited to efficiently utilize the precise 3D geometry information present in point clouds.

In this paper, we first demonstrate that geometric-verification is necessary to ensure robust point cloud re-ranking, but these methods are inefficient due their reliance on resource intensive point cloud registration. To address this, we introduce an efficient spectral method for geometric verification based re-ranking, named SpectralGV. Using the correspondence compatibility graph introduced in Spectral Matching [9], we demonstrate how the optimal inter-cluster score of this graph represents a robust confidence score on the spatial consistency of two point clouds, and utilize it to formulate registration-free geometric verification.

SpectralGV allows integration with any architecture that extracts both local and global features for a given point cloud, without modifying the architectures and without further training of learning-based methods. It is an architecture agnostic method and shows no bias towards datasets allowing it to generalize well. We integrate our method with 3 state-of-the-art architectures and demonstrate improvements in recall and pose estimation across 5 large-scale datasets in all evaluation settings. SpectralGV outperforms other re-ranking methods

Manuscript received: September 22, 2022; Revised December 27, 2022; Accepted February, 21, 2023.

This paper was recommended for publication by Editor Cesar Cadena upon evaluation of the Associate Editor and Reviewers' comments. This work was supported by the CSIRO's MLAI FSP.

<sup>1</sup> Robotics and Autonomous Systems Group, DATA61, CSIRO, Australia. [firstname.lastname@data61.csiro.au](mailto:firstname.lastname@data61.csiro.au)

<sup>2</sup> Signal Processing, AI and Vision Technologies (SAIVT), Queensland University of Technology (QUT), Brisbane, Australia. E-mails: {s.sridharan, c.fookes}@qut.edu.au

Digital Object Identifier (DOI): see top of this page.

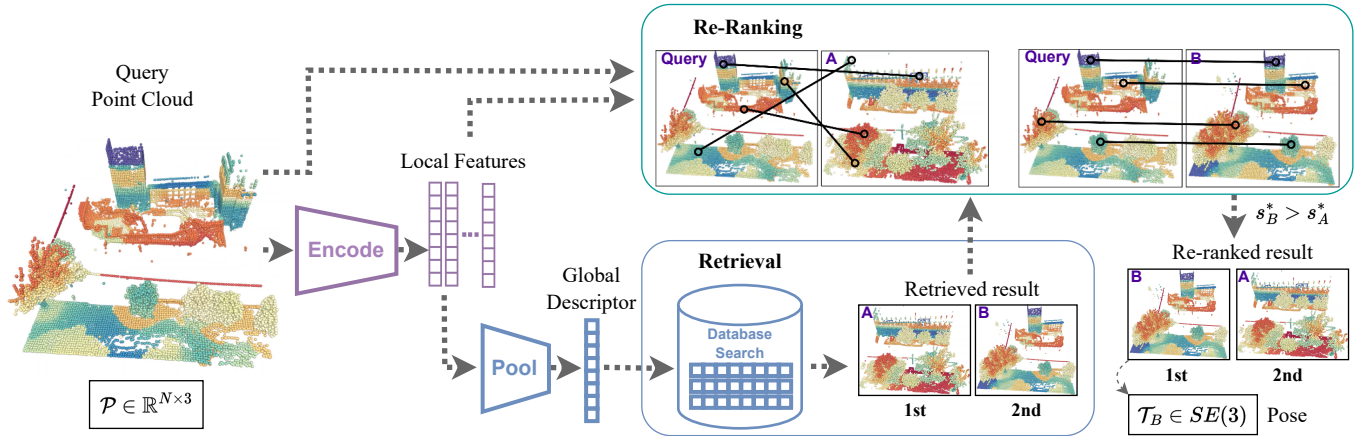


Fig. 1: Overview of our SpectralGV method. Local features are extracted from the query point cloud using an encoder and then aggregated to form a global descriptor which is used in the ‘Retrieval’ block to find the top-k place candidates ( $k=2$  in the diagram). The local features and point cloud are used for re-ranking the retrieved candidates in the ‘Re-Ranking’ module which estimates a matching confidence based on the spatial consistency of point correspondences (black lines). In this example, retrieved point cloud ‘B’ is more consistent with the query and is therefore ranked higher than ‘A’ after re-ranking. The 6DoF pose estimate  $\mathcal{T}_B$  is obtained by aligning the query point cloud to the top re-ranked candidate. Point clouds are coloured based on the t-SNE embeddings of local features to better visualize the point correspondences.

while being the only geometric verification method with sub-linear time complexity, enabling real-time deployment.

## II. RELATED WORK

Point cloud retrieval methods have mainly been proposed under the two topics of place recognition and metric localization, which we discuss in II-A and II-B respectively. We review re-ranking methods used in other domains in II-C and discuss them in relation to our proposal.

### A. Place Recognition

Large-scale place recognition is formulated as a retrieval problem where methods encode point clouds to a compact global descriptor to be used for querying a database of previously visited places. Global descriptors can be categorised as handcrafted [10], [11], hybrid [12], and end-to-end learning [13]–[17]. While handcrafted methods such as ScanContext [10] and hybrid methods such as Locus [12] still act as strong baselines, end-to-end learning methods have demonstrated superior performance [18], [19]. End-to-end learning methods define a neural network to map point clouds to global descriptors. PointNetVLAD [13] pioneered the end-to-end trainable global descriptor by combining PointNet [20] and NetVLAD [21]. Works such as LPD-Net [14] and MinkLoc3D [15] have addressed the limitations of PointNetVLAD. Recently, LoGG3D-Net [18] demonstrated the benefits of using joint constraints on the local and global embeddings during the training. In this paper, we demonstrate how the recall of these methods can be improved by a large margin by our re-ranking method.

### B. Metric localization

DH3D [4] pioneered correspondence-based metric localization using LiDAR data by proposing the unification of

global place recognition with local 6DoF refinement. They used a simultaneous detect-and-describe formulation to extract local features, NetVLAD to generate a global descriptor, and RANSAC [22] to estimate the 6DoF pose. LCDNet [5] proposed the use of a differentiable relative pose head based on Unbalanced Optimal Transport (UOT) [23] to match local features. EgoNN [6] extended the MinkLoc3D architecture for the task of metric localization with the addition of keypoint detection and saliency prediction modules. Global descriptors formed by GeM [7] pooling were used for retrieval and RANSAC was used for pose estimation.

Our proposed SpectralGV can be readily integrated into all the above metric localization methods [4]–[6] as well as place recognition methods such as [18] which produce discriminative local features. We demonstrate how this integration results in an improvement of both recall and pose estimation for all architectures across all datasets, while having an insignificant effect on their runtime.

### C. Re-ranking

In many retrieval problems, re-ranking methods take the initial retrieved candidates and re-order them such that correct candidates will be ranked higher. These methods can be categorised as those which use only the global descriptors and those which use both global and local embeddings. Re-ranking methods which use only global descriptors can be classified under Query Expansion (QE) [7] methods that use first-order neighbour information and methods that use higher-order neighbour information such as k-reciprocal [24], [25] and region-diffusion [26]. QE methods take the nearest neighbours of the query and generate an updated query which is used to retrieve a new set of candidates. The updated query is an aggregation of the initial top-candidates (*e.g.*, average-QE,

alpha-QE [7]). Higher-order methods consider the neighbours-of-neighbours of the query and are non-trivial to extend to large-scale scenarios [8]. In this paper, we demonstrate how such methods which only use global information are not robust in the point cloud domain.

As opposed to general information retrieval tasks (*e.g.*, text-retrieval), when re-ranking in the context of localization (using either images or point clouds), the geometric cues present in the scene can be used to aid the re-ranking. For this, the local features can be used in addition to the global descriptors to verify geometric consistency between the query and the retrieved candidates. In this paper, we study how such geometric verification can be performed on point clouds by adapting methods from other fields and also introduce a more efficient method for this purpose.

Geometric verification based re-ranking has become popular for visual place recognition [27]–[29]. DELG [27] used RANSAC [22] based geometric verification on the local features for re-ranking. PatchNetVLAD [28] compared RANSAC based re-ranking with a novel rapid-spatial-scoring mechanism for efficient geometric verification on images. In this paper, we study how RANSAC and newer geometric estimators [30], [31] can be adapted for performing geometric verification on point clouds, and how they are inefficient when re-ranking a large number of potential candidates. We propose a geometric verification method with sublinear time complexity to address this drawback.

### III. PROPOSED METHOD

This section describes the re-ranking based metric localization pipeline depicted in Fig. 1 and presents our SpectralGV method for efficient geometric verification.

#### A. Problem Formulation

Our approach to metric localization is split into 3 sub-tasks: retrieval, re-ranking and 6DoF pose estimation.

**Retrieval:** The sub-task of point cloud retrieval is illustrated in the ‘Retrieval’ module of Fig. 1 and formulated as follows. Given a point cloud  $\mathcal{P} \in \mathbb{R}^{N \times 3}$  representing varying number of  $N$  3D points, a mapping function  $\Phi$  such that

$$\Phi : \mathcal{P} \rightarrow (\mathcal{F}, g) \quad (1)$$

is learnt to represent the point cloud  $\mathcal{P}$  as a set of local features  $\mathcal{F} \in \mathbb{R}^{N \times d'}$ , and also as a fixed-size vector global descriptor  $g \in \mathbb{R}^d$ . We refer the reader to [18] for a detailed description of how  $\Phi$  can be learnt to produce discriminative local and global embeddings. During retrieval, the global descriptor  $g$  of the query point cloud  $\mathcal{P}$  is matched with a database of global descriptors of previously visited places to obtain the list of top-k retrieval candidates,  $\mathcal{L}_R = \{\mathcal{P}_{R_1}, \dots, \mathcal{P}_{R_k}\}$ , where  $\mathcal{L}_R$  is ordered based on similarity of retrieved global descriptors with  $g$ . If not employing re-ranking,  $\mathcal{P}_{R_1}$  is taken as the solution for place recognition.

**Re-Ranking:** In the second sub-task depicted in the ‘Re-Ranking’ module of Fig. 1, geometric verification based re-ranking techniques aim to calculate a geometric ‘fitness score’ ( $s$ ) between the query and each retrieved candidate, where

the score measures the spatial consistency between the query and each candidate using point correspondences estimated by matching the local features  $\mathcal{F}$ .

The initial top-k retrieval candidates are then re-ranked by sorting based on descending fitness score to obtain,

$$\mathcal{L}_{RR} = \{\mathcal{P}_{RR_1}, \dots, \mathcal{P}_{RR_k}\}, \quad (2)$$

where the list  $\mathcal{L}_{RR}$  is a permutation/re-ordering of the list  $\mathcal{L}_R$ . A successful re-ranking method will ensure that,

$$\mathcal{D}(x, x_{RR_i}) \leq \mathcal{D}(x, x_{R_i}) \quad \forall i, \quad (3)$$

for all queries where  $x, x_{R_i}, x_{RR_i}$  are the ground-truth geolocations of  $\mathcal{P}, \mathcal{P}_{R_i}, \mathcal{P}_{RR_i}$  respectively, and  $\mathcal{D}$  represents Euclidean distance. We note that re-ranking methods which are not robust will occasionally violate this inequality (due to failure cases such as noisy inputs, degenerate scenes). In some instances, they will reduce retrieval performance after re-ranking. We empirically demonstrate how all geometric verification methods increase performance after re-ranking and how our SpectralGV shows the biggest improvement while being the only geometric verification method with sub-linear time complexity, hence enabling its real-time use.

**Pose Estimation:** During the third sub-task, the pose estimate is obtained by registering the query  $\mathcal{P}$  with the top-1 re-ranked candidate  $\mathcal{P}_{RR_1}$  to obtain their relative transformation:  $\mathcal{T}_{RR_1} \in SE(3)$ . This transformation can be estimated using any point cloud registration technique. Typically, for LiDAR-based metric-localization, since local features and global descriptors are extracted concurrently, correspondence-based registration techniques using robust estimators such as RANSAC are preferred over alternatives.

#### B. Spectral Geometric Verification

This paper focuses on the sub-task of re-ranking, and the proposed method can be readily incorporated with existing architectures which support the formulation presented in Eq.1 - global descriptor based retrieval and local feature correspondence based registration. Baseline geometric verification methods first register the query with each of the retrieved point clouds, calculate the point inlier-ratio after registration, and then sort the retrievals based on the Registered-Inlier-Ratio (RIR). This type of re-ranking is inefficient as it requires registering the query with each retrieval. In this paper, we introduce a spectral fitness score as an alternative to the RIR for geometric verification based re-ranking.

Inspired by the Spectral Matching method introduced in [9], we define an un-directed graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  between the query  $\mathcal{P}$  and each retrieved candidate  $\mathcal{P}_R \in \mathcal{L}_R$ . Vertices  $\mathcal{V} = \{1, \dots, n\}$  represent the  $n$  elements in the set of local point correspondences  $\mathcal{C}$  between the two point clouds obtained by nearest-neighbour search on their local features. Here  $c_i \in \mathcal{C}$  is denoted as  $c_i = (x_i, y_i)$ , where  $x_i \in \mathcal{P}, y_i \in \mathcal{P}_R$  are the corresponding points from the two point clouds. The edges  $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$  represent the spatial compatibility score of these correspondences (defined later). The motivation here is that this graph will form clusters of nodes which have high spatial consistency with each other, and the largest cluster of this

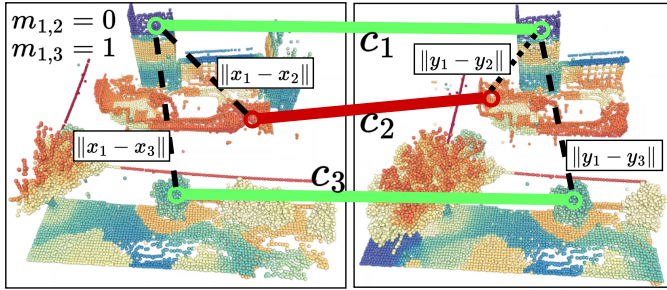


Fig. 2: Visualization of values in the adjacency matrix  $M$ . The colored lines represent inter-pointcloud correspondences. Out of the 3 depicted correspondences, the green ones ( $c_1, c_3$ ) are inliers and the red ( $c_2$ ) are outliers. The compatibility of these correspondences are represented by the  $m_{i,j}$  values in the top left of this figure. Inlier correspondences will preserve their spatial consistency (low  $d_{1,3}$ ) and form strong edges in  $\mathcal{G}$  (high  $m_{1,3}$ ), while outliers will not form strong edges with the inliers (low  $m_{1,2}$ ).

graph is statistically likely to be the inlier correspondences, as it is highly unlikely that outlier correspondences will be spatially consistent with the inliers or with each other.

The adjacency matrix of this graph represented as  $M \in \mathbb{R}^{n \times n}$  is a symmetric matrix of all positive elements where each element  $m_{i,j}$  represents a spatial compatibility score between the correspondences  $c_i, c_j$ . Specifically,  $m_{i,j}$  measures how well the pairwise geometry of two points in the query point cloud are preserved after matching them with two points in the retrieved point cloud. This is defined as,

$$m_{i,j} = [1 - \frac{d_{ij}^2}{d_{thr}^2}]_+, d_{ij} = |\|x_i - x_j\| - \|y_i - y_j\||, \quad (4)$$

where  $[\cdot]_+ = \max(\cdot, 0)$  ensures a non-negative value for  $m_{i,j}$ , and  $d_{thr}$  controls the sensitivity to the length difference. As depicted in Fig. 2, correspondence pairs with length difference (of dashed lines) larger than  $d_{thr}$  (e.g.  $c_1, c_2$ ) are considered incompatible and get a zero value in  $M$ , while pairs which preserve their length relationships (e.g.  $c_1, c_3$ ) obtain high positive scores.

The task of finding the inlier correspondences is formulated as finding the cluster  $\hat{C} \subseteq \mathcal{C}$  of correspondences that maximize the inter-cluster spatial compatibility score,

$$s = \sum_{c_i, c_j \in \hat{C}} m_{i,j}, \quad (5)$$

subject to the constraint that no pairs of incompatible correspondences are present in  $\hat{C}$ . By representing the cluster  $\hat{C}$  using a binary indicator vector  $v$  such that  $v(c_i) = 1$  if  $c_i \in \hat{C}$  and zero otherwise, we can re-write the inter-cluster score in Eq.5 as  $s = v^T M v$ . Given this formulation the indicator vector that obtains the maximum inter-cluster score can be computed as,

$$v^* = \arg \max_v (v^T M v). \quad (6)$$

As demonstrated in [9], the compatibility constraint and the integer constraint on  $v$  can be relaxed and it's norm can be

fixed to 1 since only the relative values of it's elements are important. Then, Eq. 6 can be represented as,

$$v^* = \arg \max_v \left( \frac{v^T M v}{v^T v} \right) \quad (7)$$

where  $v^*$  can be approximated using the principal eigenvector of  $M$  according to the Rayleigh quotient. We note that obtaining an exact solution for  $v^*$  accounts to finding the inlier correspondences given initial noisy correspondences and this problem does not have a closed form solution. While the methods of Spectral Matching [9] and its recent variants [32] have used  $v^*$  to solve the registration problem, in this paper, we explore a different direction. We interpret the correspondence compatibility graph as a graphical representation of the geometric compatibility of the two point clouds, and thus a scalar-value summary of this compatibility can be obtained via the optimal inter-cluster score. With this interpretation, we demonstrate (for the first time) that a fitness score for geometric verification based re-ranking can be obtained without registration of the two point clouds.

The maximum inter-cluster score (*i.e.*, the sum of spatial compatibility scores  $m_{i,j}$  of the approximated inlier correspondences) is obtained using:

$$s^* = v^{*T} M v^*, \quad (8)$$

and can be interpreted as a scalar value summarizing the geometric consistency between two point clouds and therefore can be used as a fitness score for geometric verification based re-ranking. Additionally, as demonstrated in [9], the approximation of  $v^*$  in this manner is statistically robust against outlier correspondences. Hence, we postulate that this robustness of  $v^*$  naturally translates to the score  $s^*$  which results in a robust geometric verification method enabling the preservation of the inequality in 3. This robustness for the task of re-ranking substantiated by our extensive experimental results.

Using this interpretation of  $s^*$ ,  $\mathcal{L}_{RR}$  can be obtained by sorting  $\mathcal{L}_R$  based on descending  $s^*$  calculated between the query  $\mathcal{P}$  and each  $\mathcal{P}_R \in \mathcal{L}_R$ , using Eq.2. This process is illustrated in Fig. 1 for the case of 2 candidates where the initial top-1 candidate 'A' is replaced by the 2nd candidate 'B' after re-ranking because  $s_B^* > s_A^*$ . Additionally, we provide an implementation that parallelizes this computation across all retrieved candidates, enabling robust large-scale point cloud re-ranking in realtime.

## IV. EXPERIMENTS

We conduct extensive experiments to demonstrate that SpectralGV (1) outperforms other re-ranking methods, and (2) improves performance for all architectures and datasets.

### A. Experiment setup

We integrate SpectralGV with the architectures EgoNN [6], LCDNet [5] and LoGG3D-Net [18]. For EgoNN we use the pre-trained model released by the authors which was

Train/Test Split	Length	Step	Number of scans
<i>Train split:</i>			
MulRan [33]: Sejong (0.1/02)	19 km	0.2 m	35,871
SouthBay [34]: excl. Sunnyvale	37 km	1.0 m	72,706
<i>Test split (Easy):</i>			
MulRan [33]: Sejong (01/02)	4 km	0.2 m	(database / query) 3,764 / 3,453
SouthBay [34]: Sunnyvale	38 km	1.0 m	48,959 / 16,999
KITTI [35]: Sequence 00	4 km	0.1 m	1,620 / 621
<i>Test split (Hard):</i>			
MulRan [33]: DCC (01/02)	4.9 km	10.0 m	469 / 307
ALITA [36] (val:5)	1.4 km	3-20 m	107 / 84
KITTI-360 [37]: Sequence 09	10.5 km	3.0 m	766 / 398

TABLE I: Details of training and evaluation sets.

trained in the same train set as used in this paper. LoGG3D-Net and LCDNet are trained on a cluster of P100-16GB GPUs. LCDNet was trained for 150 epochs with batch size of 4 using identical hyper-parameters from [5]. We use the power-iteration method to efficiently approximate the leading eigenvector  $v^*$  in Eq. 7.

We evaluate on five public LiDAR datasets, using an identical training set as in [6]: MulRan [33] and Apollo-SouthBay [34] datasets. We define 2 evaluation settings (‘easy’ and ‘hard’) where the easy setting consists of the test sequences proposed in [6], and for the hard setting we introduce additional sequences (unseen during training) in which current methods struggle to generalize: MulRan DCC [33], KITTI-360 [35] and the ALITA [36] datasets. In MulRan DCC, sequences 01 and 02 are used as database and query sets. In KITTI-360, sequence 09 is used with the first 300 seconds of the sequence as a database and the rest as queries. In ALITA, we evaluate on the data released at the ICRA 2022 UGV Challenge and use the validation sequence 5. Details on the number of scans and sampling step lengths of each dataset are given in Table. I.

### B. Evaluation Criteria

For place recognition, we compute the similarity between the global descriptors of each query point cloud with global descriptors from the database set and get the top-k retrieval candidates. We report Recall@k (Rk) (for k = 1 and 5) metrics for 5 and 20m revisit thresholds. We also use the Mean Reciprocal Rank (MRR) as it serves as a scalar summary of the Recall@k curve defined as  $MRR = \frac{1}{|Q|} \sum_{q \in Q} \frac{1}{r_q}$  where  $r_q$  is the rank of the first true-positive retrieval for a given query  $q$ . We report both metrics as percentages.

We first retrieve the top candidates for each query and calculate the baseline R1, R5 and MRR results (presented in the 1st row of table II and the top-halves of tables III and IV). Then we re-rank the  $n_{topk}$  retrieval candidates and recalculate R1, R5 and MRR. We use  $n_{topk} = 20$  for the evaluations in the bottom halves of Table. III and IV.

For metric localization, the 6DoF pose estimate of each method is compared to the ground truth. The success rate (S) is the percentage of pose estimates within 2 m and  $5^\circ$  threshold from the ground truth. The mean Relative Rotation Error (RRE) and mean Relative Translation Error (RTE) metrics are as defined in [32] and are calculated for all queries (not only the successful queries as done in [6]).




Fig. 3: Analysis of re-ranking methods with varying number of candidates used for re-ranking:  $n_{topk}$ . SpectralGV shows consistently increasing Recall@1 (left plot) while maintaining sub-linear time complexity (right plot).

### C. Results

1) *Re-ranking evaluation:* We explore different methods for re-ranking of point cloud retrievals. We evaluate on the ‘hard’ evaluation sequences and report Recall@1, MRR metrics and the time taken ‘t’ (in ms) for re-ranking in Table. II using the EgoNN architecture. We present results under two settings: using 2 and 20  $n_{topk}$  candidates for re-ranking.

We group the re-ranking methods into 3 categories. The first category (group (a) in Tab. II) only uses the global descriptor. We evaluate two Query Expansion (QE) methods: Average-QE and Alpha-QE [7]. We observe that Average-QE does not produce an improvement in any setting. This is because majority of the candidate descriptors in the initial retrieval are incorrect matches, hence averaging these in to the new query simply produces a noisy query. Alpha-QE shows some improvements in Recall@1 when  $n_{topk} = 2$ , but has degrading performance with increasing  $n_{topk}$ . While global-only methods are efficient in point cloud re-ranking, they lack robustness and perform worse than the baseline.

The second category consists of geometric verification (GV) methods that re-rank based on the registered-inlier-ratio (RIR). We test 3 robust geometric estimators to perform the registration: RANSAC [22], FGR [30] and TEASER++ [31]. For RANSAC and FGR we use the efficient implementations provided in [38]. We observe that all RIR methods improve performance in all settings. RANSAC obtains the best results in the RIR category (group (b) in Tab. II). The final category of re-ranking methods (group (c) in Tab. II) consists of geometric verification methods that do not perform registration but directly compute a spatial consistency (SC) score. SpectralGV falls into this category and we observe that it outperforms all other re-ranking methods in Tab. II.

All geometric verification methods show increasing recall with increasing  $n_{topk}$ . RIR methods significantly increase runtime with increasing  $n_{topk}$  (increasing by an order of magnitude in all cases when increasing  $n_{topk}$  from 2 to 20) but SpectralGV maintains roughly constant computation time. In the  $n_{topk}=20$  setting, our method obtains the best recall overall while also being more efficient than other GV methods as shown in the final row of Tab. II operating at 3-5 ms while other GV methods take 50-300 ms. The effect of  $n_{topk}$  on re-ranking performance and on efficiency for the best performing

	$n_{topk}$	MulRan DCC					ALITA					KITTI-360				
		5m		20m		t (ms)	5m		20m		t (ms)	5m		20m		t (ms)
		R1	MRR	R1	MRR		R1	MRR	R1	MRR		R1	MRR	R1	MRR	
Baseline	-	67.1	77.2	90.2	92.3	-	79.8	87.3	91.7	94.9	-	86.9	90.1	88.2	91.4	-
<i>(a) Global only:</i>																
Average-QE	2	67.4	76.7	90.2	91.8	0.2	79.8	86.5	91.7	94.3	0.1	86.9	89.3	88.2	90.8	0.2
	20	35.2	47.6	54.7	65.0	0.2	13.1	27.4	26.2	44.0	0.1	46.2	59.4	49.7	63.8	0.2
Alpha-QE [7]	2	67.4	76.7	90.6	92.0	0.2	81.0	87.1	92.9	94.9	0.1	86.7	89.2	87.9	90.6	0.2
	20	34.5	47.4	54.7	65.0	0.3	13.1	28.4	26.2	45.0	0.2	47.0	59.9	50.8	64.4	0.3
<i>(b) GV: RIR</i>																
RANSAC [22]	2	71.0	79.0	92.8	93.6	2.8	90.5	92.7	95.2	96.7	4.1	90.5	91.9	92.5	93.5	2.8
	20	<u>72.0</u>	<u>80.9</u>	<b>95.8</b>	<b>96.0</b>	63.6	<b>100</b>	<b>100</b>	<b>100</b>	<b>100</b>	60.4	<u>95.0</u>	<u>96.4</u>	<u>98.5</u>	<u>98.6</u>	56.3
FGR [30]	2	70.4	78.7	91.2	92.8	79.7	89.3	92.1	95.2	96.7	54.8	90.2	91.7	92.2	93.4	74.4
	20	70.7	78.9	92.2	93.2	175.7	97.6	98.8	<b>100</b>	<b>100</b>	126.7	93.7	95.8	98.0	98.3	171.3
TEASER++ [31]	2	70.0	78.5	92.5	93.4	26.3	90.5	92.7	95.2	96.7	30.1	89.9	91.6	92.2	93.4	35.1
	20	71.3	80.6	93.8	94.8	261.3	98.8	99.4	<b>100</b>	<b>100</b>	272.8	94.0	95.8	96.7	97.6	320.5
<i>(c) GV: SC</i>																
<b>SpectralGV (Ours)</b>	2	71.7	79.3	92.8	93.6	2.7	90.5	92.7	95.2	96.7	5.3	90.7	92.0	92.5	93.5	2.4
	20	<b>73.6</b>	<b>82.2</b>	<u>95.1</u>	<u>95.7</u>	3.0	<b>100</b>	<b>100</b>	<b>100</b>	<b>100</b>	5.7	<b>95.7</b>	<b>96.9</b>	<b>98.7</b>	<b>98.8</b>	2.8

TABLE II: Comparison of re-ranking methods. The first row shows the baseline retrieval results without re-ranking. Methods in subsection (a) perform re-ranking using only the global descriptors. Methods in subsections (b) and (c) perform Geometric Verification (GV) based re-ranking where (b) computes the Registered-Inlier-Ratio (RIR) and (c) directly computes a Spatial Consistency (SC) score.  $n_{topk}$ : number of top-k retrieval candidates used during re-ranking.

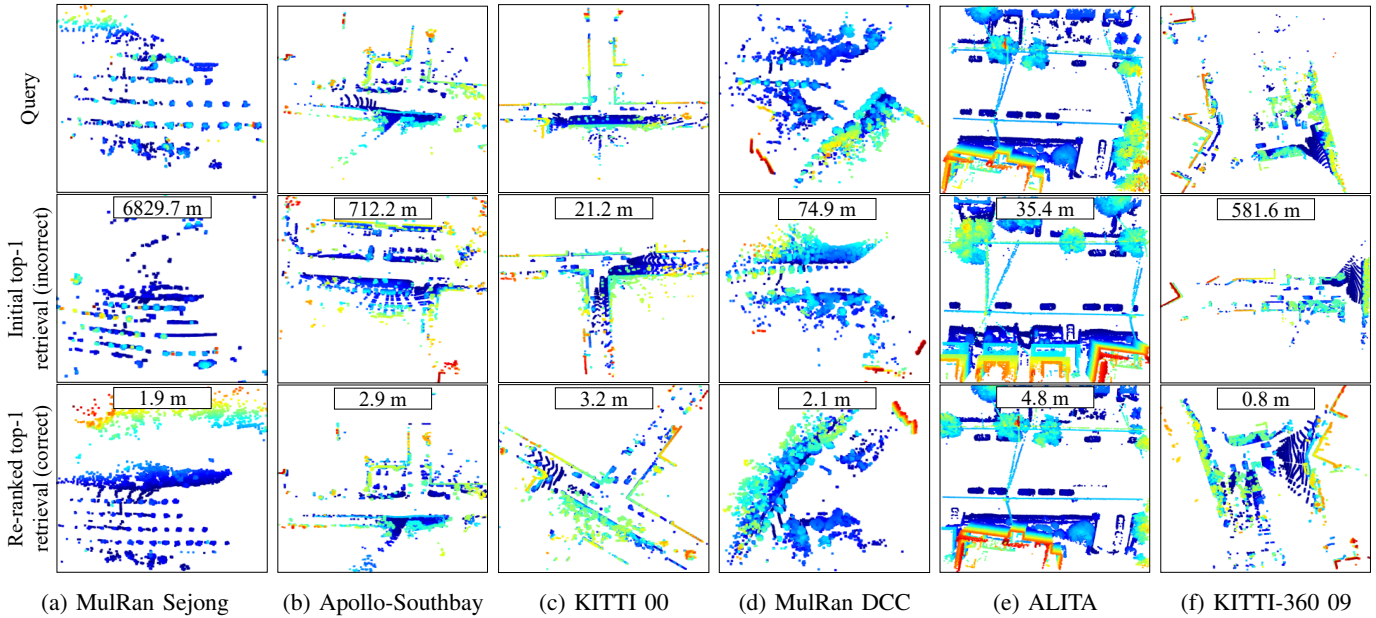


Fig. 4: Visualizations of retrieval corrections by SpectralGV on the five datasets using the EgoNN architecture. Our re-ranking method is able to filter out incorrect initial candidates which are highly structurally similar to the query and correct them by accurately identifying the correct candidate. The distance of each retrieval to the query is noted in the text box.

method in each category is presented in Fig. 3 using EgoNN on KITTI-360 09 under the 5 m revisit criteria. Only GV methods show consistently increasing Recall@1. While RIR methods have linearly increasing runtime and become impractical at large  $n_{topk}$  values, SpectralGV maintains almost constant runtime while also outperforming RIR methods in recall.

We conduct a qualitative evaluation of re-ranking results on the 5 datasets. Fig. 4 shows point clouds corresponding to the query, initial top-1 retrieval and re-ranked top-1 retrievals in cases where an incorrect initial candidate is corrected after re-ranking. In all instances, the incorrect initial candidate point cloud is highly structurally similar to the query with only

minor inconsistencies visible upon close inspection. Our re-ranking method is able to handle such complex cases and accurately identify the correct candidate.

As demonstrated by our extensive experiments, re-ranking enables better use of test-time knowledge. Current retrieval methods still struggle to discriminate between complex scenes which have largely similar structural elements as they only rely on global embeddings. These ambiguities can be addressed at test time using re-ranking of the top retrieved candidates as the correct candidate is more likely to be within the top-k retrievals than top-1.

	MulRan Sejong				Apollo-SouthBay				KITTI				Average			
	5 m		20 m		5 m		20 m		5 m		20 m		5 m		20 m	
	R1	R5	R1	R5	R1	R5	R1	R5	R1	R5	R1	R5	R1	R5	R1	R5
MinkLoc3D* [15]	82.3	94.4	92.1	96.6	77.2	93.8	95.0	98.3	95.7	96.9	97.7	98.0	85.0	95.0	94.9	97.6
DiSCO* [39]	94.0	97.5	95.8	98.1	95.1	96.6	95.4	97.0	90.7	91.3	92.3	94.5	93.2	95.1	94.5	96.5
DH3D* [4]	32.4	56.2	58.9	75.5	25.3	49.6	50.4	70.7	75.5	91.1	86.8	96.1	44.4	65.6	65.3	80.7
ScanContext [10]	86.1	89.6	88.5	91.5	91.0	91.0	92.1	92.4	96.3	97.4	96.3	97.6	91.1	92.6	92.3	93.8
Locus [12]	67.0	75.8	70.6	74.6	85.5	95.5	94.9	98.1	99.0	99.7	99.8	<b>100</b>	83.8	90.3	88.4	90.9
EgoNN [6]	98.3	<b>99.9</b>	<b>99.6</b>	99.9	95.7	97.7	96.3	98.2	97.4	98.2	97.9	98.7	97.1	98.6	97.9	98.9
LCDNet [5]	63.1	82.0	85.8	92.1	64.7	81.7	87.4	92.1	96.9	99.4	99.5	99.8	74.9	87.7	90.9	94.6
LoGG3D-Net [18]	97.7	99.7	98.8	99.8	95.2	98.4	98.2	<u>99.2</u>	<b>99.7</b>	<b>99.8</b>	<b>100</b>	<b>100</b>	97.5	<u>99.3</u>	99.0	<u>99.6</u>
EgoNN + SGV	<b>99.1</b>	<b>99.9</b>	<b>99.6</b>	<b>100</b>	<u>98.3</u>	<u>98.7</u>	<u>99.1</u>	<u>99.2</u>	99.2	99.4	99.7	99.7	98.8	<u>99.3</u>	99.4	<u>99.6</u>
LCDNet + SGV	90.8	92.4	94.3	95.8	90.0	90.8	95.2	95.6	<b>99.7</b>	99.7	<b>100</b>	<b>100</b>	93.5	94.3	96.5	97.1
LoGG3D-Net + SGV	<u>99.0</u>	<b>99.9</b>	99.5	<b>100</b>	<b>98.6</b>	<b>99.3</b>	<b>99.6</b>	<b>99.6</b>	<b>99.7</b>	<b>99.8</b>	<b>100</b>	<b>100</b>	<b>99.1</b>	<b>99.6</b>	<b>99.7</b>	<b>99.8</b>

TABLE III: Place recognition on the ‘easy’ evaluation set. (\*) Results are obtained from [6]. ‘R1’: Recall@1, ‘R5’: Recall@5.

	MulRan DCC				ALITA				KITTI-360						
	5 m		20 m		5 m		20 m		5 m		20 m				
	R1	MRR	$\Delta$ MRR	R1	MRR	R1	MRR	R1	MRR	R1	MRR	R1	MRR		
ScanContext [10]	72.0	80.5	-	93.2	94.0	57.1	66.5	-	59.5	69.4	93.7	94.9	-	95.2	95.9
Locus [12]	46.3	55.6	-	71.0	77.5	95.2	97.6	-	<b>100</b>	<b>100</b>	92.7	95.1	-	97.0	98.1
EgoNN [6]	67.1	77.2	-	90.2	92.3	79.8	87.3	-	91.7	94.9	86.9	90.1	-	88.2	91.4
LCDNet [5]	57.7	71.6	-	90.9	93.8	90.5	94.6	-	96.4	97.8	85.9	89.7	-	95.7	96.6
LoGG3D-Net [18]	66.8	76.7	-	91.2	92.9	15.5	28.0	-	26.2	40.0	93.2	95.5	-	97.5	98.1
EgoNN + SGV	<u>73.6</u>	<b>82.2</b>	5.0	95.1	95.7	<b>100</b>	<b>100</b>	12.7	<b>100</b>	<b>100</b>	<u>95.7</u>	<u>96.9</u>	6.8	98.7	98.8
LCDNet + SGV	72.3	83.0	11.4	<b>99.3</b>	<b>99.4</b>	<b>100</b>	<b>100</b>	5.4	<b>100</b>	<b>100</b>	<u>95.7</u>	96.2	6.5	<u>99.0</u>	<u>99.1</u>
LoGG3D-Net + SGV	<b>73.9</b>	<u>82.0</u>	5.3	<u>96.1</u>	<u>96.5</u>	60.7	63.2	35.2	75.0	77.3	<b>96.7</b>	<b>98.1</b>	2.6	<b>99.5</b>	<b>99.5</b>

TABLE IV: Place recognition on the unseen ‘hard’ evaluation set.  $\Delta$ MRR: increase in MMR after re-ranking in 5 m threshold.

2) *Place recognition evaluation:* We compare place recognition performance using the handcrafted method ScanContext [10], hybrid method Locus [12] and end-to-end learning methods MinkLoc3D [15], DiSCO [39], DH3D [4], EgoNN [6], LCDNet [5] and LoGG3D-Net [18]. Table III shows place recognition results on the ‘easy’ evaluation sets while Table IV shows the corresponding results in the ‘hard’ evaluation sequences. On the ‘easy’ evaluation set, we see that LoGG3D-Net outperforms the previous state-of-the-art EgoNN on average performance across all evaluation metrics. For the methods which also output local features (EgoNN, LCDNet and LoGG3D-Net), we also evaluate performance after re-ranking. The results on the ‘easy’ set are already saturated prior to re-ranking but we still see improvements for all methods after re-ranking using SGV. LoGG3D-Net + SGV obtains the best performance overall on the ‘easy’ datasets.

On the ‘hard’ evaluation setting all methods do worse as not only are these sequences unseen during training, but they also contain difficult evaluation scenarios such as reverse/orthogonal revisits and lane offsets during revisits. Therefore, the performance gain after re-ranking is much more pronounced on the ‘hard’ evaluation set. We observe that LoGG3D-Net which fails to generalize to the ALITA dataset<sup>1</sup> shows an improvement of 32.5 MRR after re-ranking which is more than double the initial MRR while also almost quadrupling R1 after re-ranking. The results after re-ranking using our SGV are better for all methods in all datasets. This demonstrates that our method improves retrieval performance on all datasets while being architecture agnostic.

<sup>1</sup>The ALITA dataset is out-of-distribution as it contains accumulated point clouds as opposed to the raw LiDAR scans in the training datasets.

	S (%) ( $\uparrow$ )	RTE (m) ( $\downarrow$ )	RRE ( $^\circ$ ) ( $\downarrow$ )
EgoNN [6]	88.2	42.45	13.36
LCDNet [5]	94.0	10.40	6.37
LoGG3D-Net [18]	97.0	9.51	3.64
EgoNN + SGV	98.7	4.32	1.44
LCDNet + SGV	98.5	<b>0.79</b>	2.35
LoGG3D-Net + SGV	<b>99.5</b>	2.88	<b>1.07</b>

TABLE V: Results of 6DoF metric localization on the unseen KITTI-360 09 dataset with and without re-ranking.

3) *Metric localization evaluation:* We evaluate LCDNet [5], EgoNN [6] and LoGG3D-Net [18]. All methods use RANSAC for pose estimation. Although LCDNet uses unbalanced optimal transport for transformation estimation during training, RANSAC is used during inference [5].

Table V shows 6DoF pose estimation results. We observe that prior to re-ranking, all 3 methods have large RTE (close to or greater than 10m), limiting their reliability for re-localization applications. After re-ranking using SpectralGV (with  $n_{topk} = 20$ ), all 3 methods show significant performance improvements. For EgoNN, the RTE and RRE have reduced by 89.9% and 89.2% respectively. The RTE of LCDNet and LoGG3D-Net have reduced by 92.4% and 69.4% respectively. All methods now have RTE < 5m and RRE < 3 $^\circ$ , making them more reliable for re-localization.

SpectralGV can be used to refine the retrieval candidates such that the top-1 retrieval used for registration is closer to the the query. For example, in Fig. 4c, we see that SpectralGV has replaced the initial candidate which is 21 m away from the query with a candidate which is only 3 m away from the query. This will result in an easier registration task and hence enable more accurate metric localization results.

## V. CONCLUSION

In this paper we explore re-ranking of point cloud retrievals for large-scale place recognition and metric localization. We study multiple methods for re-ranking point cloud retrieval and introduce an efficient spectral method for geometric verification based re-ranking which does not require registering the query to each candidate. Our method can be readily integrated with current metric localization methods without any changes to their architecture. We show consistent and significant performance improvements in 5 different large-scale datasets in all evaluation settings. We empirically demonstrate how, in the context of point cloud retrieval, all geometric verification methods are robust, and Spectral Geometric Verification obtains the best performance improvement after re-ranking while being the only method efficient for real-time robotic application.

## REFERENCES

- [1] P. Yin, S. Zhao, I. Cisneros, A. Abuduweili, G. Huang, M. Milford, C. Liu, H. Choset, and S. Scherer, "General place recognition survey: Towards the real-world autonomy age," *arXiv preprint arXiv:2209.04497*, 2022.
- [2] F. Pomerleau, F. Colas, R. Siegwart *et al.*, "A review of point cloud registration algorithms for mobile robotics," *Foundations and Trends® in Robotics*, vol. 4, no. 1, pp. 1–104, 2015.
- [3] C. Park, S. Kim, P. Moghadam, J. Guo, S. Sridharan, and C. Fookes, "Robust photogeometric localization over time for map-centric loop closure," *IEEE Robotics and Automation Letters*, 2019.
- [4] J. Du, R. Wang, and D. Cremers, "DH3D: Deep hierarchical 3d descriptors for robust large-scale 6dof relocalization," in *European Conference on Computer Vision*, 2020.
- [5] D. Cattaneo, M. Vaghi, and A. Valada, "LCDNet: Deep Loop Closure Detection and Point Cloud Registration for LiDAR SLAM," *IEEE Transactions on Robotics*, 2022.
- [6] J. Komorowski, M. Wysoczanska, and T. Trzcinski, "EgoNN: Egocentric neural network for point cloud based 6dof relocalization at the city scale," *IEEE Robotics and Automation Letters*, 2022.
- [7] F. Radenović, G. Toliás, and O. Chum, "Fine-tuning cnn image retrieval with no human annotation," *IEEE transactions on pattern analysis and machine intelligence*, vol. 41, no. 7, pp. 1655–1668, 2018.
- [8] F. Tan, J. Yuan, and V. Ordonez, "Instance-level image retrieval using reranking transformers," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 12 105–12 115.
- [9] M. Leordeanu and M. Hebert, "A spectral technique for correspondence problems using pairwise constraints," in *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, 2005.
- [10] 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*, 2018.
- [11] G. Kim, S. Choi, and A. Kim, "Scan context++: Structural place recognition robust to rotation and lateral variations in urban environments," *IEEE Transactions on Robotics*, vol. 38, no. 3, pp. 1856–1874, 2022.
- [12] K. Vidanapathirana, P. Moghadam, B. Harwood, M. Zhao, S. Sridharan, and C. Fookes, "Locus: Lidar-based place recognition using spatiotemporal higher-order pooling," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [13] M. A. Uy and G. H. Lee, "PointNetVLAD: Deep point cloud based retrieval for large-scale place recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018.
- [14] Z. Liu, S. Zhou, C. Suo, P. Yin, W. Chen, H. Wang, H. Li, and Y.-H. Liu, "Lpd-net: 3d point cloud learning for large-scale place recognition and environment analysis," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 2831–2840.
- [15] J. Komorowski, "MinkLoc3D: Point Cloud Based Large-Scale Place Recognition," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2021.
- [16] 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*, 2022.
- [17] J. Knights, P. Moghadam, M. Ramezani, S. Sridharan, and C. Fookes, "InCloud: Incremental learning for point cloud place recognition," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022.
- [18] K. Vidanapathirana, M. Ramezani, P. Moghadam, S. Sridharan, and C. Fookes, "LoGG3D-Net: Locally guided global descriptor learning for 3d place recognition," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2022.
- [19] J. Knights, K. Vidanapathirana, M. Ramezani, S. Sridharan, C. Fookes, and P. Moghadam, "Wild-Places: A Large-Scale Dataset for Lidar Place Recognition in Unstructured Natural Environments," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2023.
- [20] 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*, 2017.
- [21] 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*, 2016.
- [22] 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*, 1981.
- [23] L. Chizat, G. Peyré, B. Schmitzer, and F.-X. Vialard, "Scaling algorithms for unbalanced optimal transport problems," *Mathematics of Computation*, vol. 87, no. 314, pp. 2563–2609, 2018.
- [24] D. Qin, S. Gammeter, L. Bossard, T. Quack, and L. Van Gool, "Hello neighbor: Accurate object retrieval with k-reciprocal nearest neighbors," in *CVPR 2011*. IEEE, 2011, pp. 777–784.
- [25] Z. Zhong, L. Zheng, D. Cao, and S. Li, "Re-ranking person re-identification with k-reciprocal encoding," in *IEEE conference on computer vision and pattern recognition*, 2017.
- [26] A. Iscen, G. Toliás, Y. Avrithis, T. Furon, and O. Chum, "Efficient diffusion on region manifolds: Recovering small objects with compact cnn representations," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2077–2086.
- [27] B. Cao, A. Araujo, and J. Sim, "Unifying deep local and global features for image search," in *European Conference on Computer Vision*. Springer, 2020, pp. 726–743.
- [28] S. Hausler, S. Garg, M. Xu, M. Milford, and T. Fischer, "Patch-NetVLAD: Multi-scale fusion of locally-global descriptors for place recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 14 141–14 152.
- [29] S. Lee, H. Seong, S. Lee, and E. Kim, "Correlation verification for image retrieval," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 5374–5384.
- [30] Q.-Y. Zhou, J. Park, and V. Koltun, "Fast global registration," in *European conference on computer vision*, 2016.
- [31] H. Yang, J. Shi, and L. Carlone, "Teaser: Fast and certifiable point cloud registration," *IEEE Transactions on Robotics*, vol. 37, no. 2, pp. 314–333, 2020.
- [32] X. Bai, Z. Luo, L. Zhou, H. Chen, L. Li, Z. Hu, H. Fu, and C.-L. Tai, "PointDSC: Robust point cloud registration using deep spatial consistency," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 15 859–15 869.
- [33] G. Kim, Y. S. Park, Y. Cho, J. Jeong, and A. Kim, "MulRan: Multimodal Range Dataset for Urban Place Recognition," in *2020 IEEE International Conference on Robotics and Automation*, 2020.
- [34] W. Lu, Y. Zhou, G. Wan, S. Hou, and S. Song, "L3-Net: Towards learning based lidar localization for autonomous driving," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- [35] 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.
- [36] P. Yin, S. Zhao, R. Ge, I. Cisneros, R. Fu, J. Zhang, H. Choset, and S. Scherer, "ALITA: A large-scale incremental dataset for long-term autonomy," *arXiv preprint arXiv:2205.10737*, 2022.
- [37] Y. Liao, J. Xie, and A. Geiger, "KITTI-360: A novel dataset and benchmarks for urban scene understanding in 2d and 3d," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [38] Q.-Y. Zhou, J. Park, and V. Koltun, "Open3d: A modern library for 3d data processing," *arXiv preprint arXiv:1801.09847*, 2018.
- [39] X. Xu, H. Yin, Z. Chen, Y. Li, Y. Wang, and R. Xiong, "Disco: Differentiable scan context with orientation," *IEEE Robotics and Automation Letters*, 2021.