

# Loop Closure via Maximal Cliques in 3D LiDAR-Based SLAM

Javier Laserna<sup>1</sup>   Saurabh Gupta<sup>2</sup>   Oscar Martinez Mozos<sup>1</sup>   Cyrill Stachniss<sup>2</sup>   Pablo San Segundo<sup>1</sup>




Fig. 1: Successful loop closure association between two local LiDAR submaps in a complex urban environment. The red and green points represent keypoints from the query and reference maps, respectively. The left image shows the result of our approach, while the right image shows the result of a RANSAC-based method.

**Abstract**—Reliable loop closure detection remains a critical challenge in 3D LiDAR-based SLAM, especially under sensor noise, environmental ambiguity, and viewpoint variation conditions. RANSAC is often used in the context of loop closures for geometric model fitting in the presence of outliers. However, this approach may fail, leading to map inconsistency. We introduce a novel deterministic algorithm, CliReg, for loop closure validation that replaces RANSAC verification with a maximal clique search over a compatibility graph of feature correspondences. This formulation avoids random sampling and increases robustness in the presence of noise and outliers. We integrated our approach into a real-time pipeline employing binary 3D descriptors and a Hamming distance embedding binary search tree-based matching. We evaluated it on multiple real-world datasets featuring diverse LiDAR sensors. The results demonstrate that our proposed technique consistently achieves a lower pose error and more reliable loop closures than RANSAC, especially in sparse or ambiguous conditions. Additional experiments on 2D projection-based maps confirm its generality across spatial domains, making our approach a robust and efficient alternative for loop closure detection.

## I. INTRODUCTION

Loop closure detection is an essential component of simultaneous localization and mapping (SLAM) systems [2], [3], [8], allowing robots and autonomous vehicles to identify locations they have previously visited and correct the global drift in trajectory obtained from sequential odometry sources. A globally consistent trajectory, in return, allows such robots to maintain a consistent map representation of their environment, enabling downstream autonomy tasks such as global re-localization and path planning.

Identifying loop closures can be difficult in large outdoor 3D spaces where LiDAR point clouds are sparse, noisy, and

undergo spatio-temporal changes due to dynamic scenes. Recent state-of-the-art methods for loop closure detection [5], [10], [12], [14], [41] rely on computing local and/or global feature descriptors from LiDAR point clouds, narrowing the search space of loop closure candidates through a feature matching process. Subsequently, they apply geometric verification of these candidate loop closures to find consistent 3D pose constraints for the pose-graph optimization in SLAM. Such geometric verification can be challenging due to outlier correspondences, especially in purely feature-based correspondence sets generated by typical loop closure methods [5], [10], [18], leading to incorrect loop closures. Already, a few wrong closures can substantially affect the overall performance of the SLAM system.

Among the various strategies for geometric verification of loop closure candidates, outlier rejection techniques like RANSAC [6] are widely used, as they iteratively sample minimal sets of correspondences to fit a transformation model while handling outliers. However, this approach can be sensitive to outlier correspondences and may fail to find the correct loop closure in challenging environments. Moreover, such a sampling-based approach can be computationally expensive in the presence of a large number of outlier correspondences.

The main contribution of this paper is a loop closure validation based on a strong combinatorial optimization backend. We propose replacing the RANSAC-based verification process with a maximal clique search algorithm CliReg [16] to verify candidate loop closures. Our approach aims to find the largest mutually consistent set of feature correspondences with a graph-theoretic approach to achieve geometric consistency without relying on random sampling. Our method provides a principled alternative to traditional consensus-based verification with improved robustness and

<sup>1</sup>The authors are with the ETSIDI, Universidad Polit cnica de Madrid, 28012 Madrid, Spain. <sup>2</sup>The authors are with the Center for Robotics, University of Bonn, Germany. Cyrill Stachniss is additionally with the Lamarr Institute for Machine Learning and Artificial Intelligence, Germany.

precision. As shown in Fig. 1, our approach produces a consistent set of feature matches between local maps. The resulting geometric transformation is directly usable as a loop closure constraint in pose-graph optimization. We propose a complete loop closure detection and validation pipeline using binary 3D descriptors and an efficient binary search tree-based database for appearance matching. We evaluate the method on challenging 3D LiDAR datasets and show that it achieves better pose accuracy than RANSAC-based geometric verification while being computationally efficient for real-time SLAM. We demonstrate that our approach successfully validates loop closures even when RANSAC fails to find consistent inlier correspondences. Our method improves alignment accuracy and robustness without introducing runtime overheads that hinder real-time use. These characteristics make it a compelling alternative for loop closure validation in LiDAR-based SLAM systems.

## II. RELATED WORK

The precision and robustness of loop closure detection within any SLAM pipeline directly influence mapping quality and localization accuracy. Traditional loop closure detection methods for LiDAR SLAM rely on geometric feature-based algorithms and corresponding descriptors such as SHOT [27], FPFH [25], NARF [35], as well as point feature-based place recognition [34]. Although effective in controlled environments, these techniques often suffer from reduced recall under viewpoint changes, sensor noise, or environmental variability. More recent pipelines leverage global place recognition or voting schemes [4], [5], but they typically depend on robust feature filtering for precise loop closure validation. Efforts to improve robustness across diverse sensor setups have also led to new benchmarks such as HeLiPR [13] which support place recognition evaluation in heterogeneous LiDAR configurations.

Robust loop closure detection plays a central role in maintaining global map consistency and reducing drift in SLAM systems, especially under large-scale 3D LiDAR-based scenarios. The literature proposes a wide range of approaches, from geometric descriptor matching to learned place recognition, each with distinct robustness, efficiency, and generalization trade-offs.

RANSAC-based techniques [6], [40] remain the most widely used solution for outlier rejection within geometric validation of candidate loop closures. Despite their popularity, they rely on random sampling and often fail under high outlier ratios or sparse correspondence sets [23]. Alternative robust verification methods such as pairwise consistent measurement have also been explored in multi-robot SLAM [19]. This motivates the use of deterministic alternatives. Earlier work on deterministic methods includes branch-and-bound strategies for data association, such as the joint compatibility branch and bound algorithm by Neira and Tardós [20]. Olson et al. [21] proposed a real-time correlative scan matching strategy, influencing early SLAM pipelines.

Clique-based methods have recently been introduced to generate robust data associations for point cloud registration

as an alternative to sampling-based methods. In particular, maximal clique search can extract mutually consistent correspondences under rigid-body constraints. Recent research has focused on combinatorial optimization and graph-based methods [33], particularly maximal clique enumeration techniques. While earlier works such as [28], [29], [30], [31] studied the theoretical aspects of clique enumeration and introduced graph pruning strategies for general-purpose applications, their focus was not on 3D registration. The CliReg algorithm [16] applies maximal clique search specifically for 3D point cloud registration, demonstrating significant gains in robustness and accuracy.

Efficient indexing and retrieval of binary descriptors are also crucial to the clique-based correspondence graph generation methods. HBST [32] enables scalable feature lookup using binary trees, and its use with binary descriptors like ORB [24] and B-SHOT [22] has shown promise in place recognition.

Learning-based solutions [36], [39] offer high retrieval recall but suffer from training dependency and limited generalization. While these models can perform well on curated datasets, their robustness in diverse and previously unseen real-world conditions remains a challenge [17]. Earlier works explored binary vocabularies [7] and nonlinear embedding for nearest neighbor classification [26].

In summary, the state of the art reveals a gap between robustness and efficiency. Our approach bridges this by integrating binary 3D descriptors with deterministic clique-based filtering. Unlike prior methods, our approach looks for mutual consistency among feature matches, reducing false positives and improving loop closure reliability without sacrificing runtime performance.

## III. CLIREG-BASED LOOP CLOSURES

We present a loop closure validation pipeline that replaces RANSAC with a clique-based approach for geometric verification. As illustrated in Fig. 2, our pipeline comprises three main stages: (1) feature extraction and encoding using binary descriptors, (2) construction of a correspondence graph based on descriptor matches, and (3) pose estimation through maximal clique search and least-squares alignment. Our approach identifies the largest subset of geometrically consistent correspondences without relying on random sampling.

### A. Preliminaries

Let  $\mathcal{M}$  and  $\mathcal{Q}$  denote two local maps composed of 3D point clouds or their 2D projections. From both maps, we extract a set of feature keypoints along with their associated descriptors. Based on descriptor matching, we then establish a tentative set of point-to-point correspondences  $\mathcal{C} = \mathbf{m}_i, \mathbf{q}_i$ , where  $\mathbf{m}_i \in \mathcal{M}$  and  $\mathbf{q}_i \in \mathcal{Q}$ .

Assuming a rigid body transformation model, we expect inlier correspondences to satisfy the following relation:

$$\mathbf{q}_i = \mathbf{R}\mathbf{m}_i + \mathbf{t} + \boldsymbol{\epsilon}_i, \quad i \in \{1, \dots, |\mathcal{C}|\}, \quad (1)$$

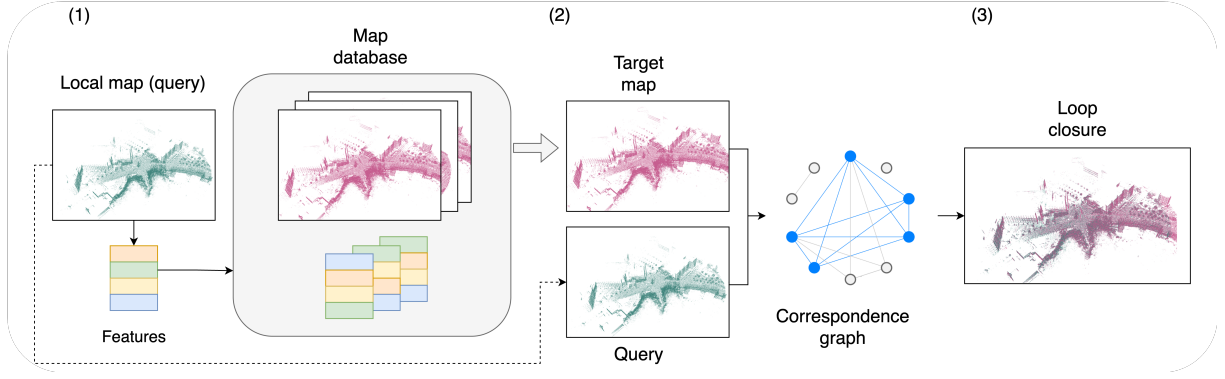


Fig. 2: An overview of our loop closure detection and validation pipeline. It consists of three stages: (1) Feature Extraction and Encoding, (2) Correspondence Graph generation, and (3) Pose Estimation via maximal clique search. This integration enables efficient and robust loop closure detection across 2D and 3D feature representations.

where  $R \in \text{SO}(n)$  is a rotation matrix,  $t \in \mathbb{R}^n$  is a translation vector,  $\epsilon_i \in \mathbb{R}^n$  is an additive noise term, and  $n \in \{2, 3\}$  denotes the spatial dimension (depending on whether 3D point clouds or 2D image projections are used).

The validation of matches aims to find the largest subset of correspondences consistent with a single rigid body transformation.

### B. Clique-Based Registration

We construct a correspondence graph where each node represents a candidate correspondence, and edges connect mutually consistent pairs (i.e., pairs satisfying the rigidity constraint in Eq. (1) within some tolerance). We then employ a branch-and-bound search to identify the maximum clique of mutually consistent correspondences, denoted by  $C^* \subseteq C$ .

Given the inlier set  $C^*$ , we estimate the optimal transformation parameters  $(R^*, t^*)$  by solving the following least-squares problem:

$$R^*, t^* = \underset{R, t}{\operatorname{argmin}} \sum_{(m_i, q_i) \in C^*} \|q_i - Rm_i - t\|^2. \quad (2)$$

Researchers have extensively documented a closed-form solution for Eq. (2) in the literature [1], [11], [38]. If  $|C^*|$  exceeds a predefined inlier threshold, we accept the estimated transformation as a valid loop closure and directly integrate it as a constraint in pose graph optimization.

Unlike sampling-based methods such as RANSAC, CliReg provides a deterministic and globally optimal inlier set under the rigid body transformation, making it particularly robust in scenarios with high outlier rates.

### C. Feature Extraction and Encoding

For each local map, we extract keypoints and compute descriptors depending on the dimensionality. In 3D, we detect keypoints using intrinsic shape signatures (ISS) algorithm [42] directly on voxelized LiDAR point clouds. We describe them using SHOT descriptors [27], and then apply a median thresholding strategy to binarize them into compact B-SHOT descriptors [22].

In the 2D setting, we follow Gupta et al. [10], first projecting the 3D LiDAR point clouds into bird's eye




Fig. 3: Example of a correspondence graph  $G$  with nodes representing feature matches and edges denoting geometric consistency. The highlighted 5-clique forms the basis for  $\text{SE}(3)$  transformation.

view (BEV) density images and extracting ORB descriptors [24], which provide binary descriptors suitable for fast matching. Regardless of the representation, we organize all binary descriptors using an HBST data structure [32] to enable real-time insertion and nearest-neighbor lookup via Hamming distance metric.

### D. Descriptor Database and Correspondence Graph Construction

From the set of tentative matches  $C = \{m_i, q_i\}$ , we build a correspondence graph  $G = (V, E)$  where each vertex represents a descriptor match between  $\mathcal{M}$  and  $\mathcal{Q}$ . An edge connects two nodes if their associated correspondences are mutually compatible, meaning the rigid transformation preserves their pairwise distances. Formally, two correspondences  $(m_i, q_i)$  and  $(m_j, q_j)$ , we consider them consistent if the following condition is met:

$$\|m_i - m_j\| - \|q_i - q_j\| < \epsilon. \quad (3)$$

This constraint preserves pairwise distances under a rigid-body model, enforcing geometric consistency. The threshold  $\epsilon$  is set relative to the voxel resolution. A clique thus represents mutually consistent correspondences explainable by a single transformation.

Fig. 3 shows an example of such a correspondence graph, where the blue subgraph represents a maximal clique selected to estimate the rigid body transformation.

#### IV. EXPERIMENTAL EVALUATION

We evaluated the effectiveness of our CliReg-based loop closure validation pipeline in the Bridge01, Bridge02 and Roundabout01 sequences of the HeLiPR dataset [13], which include realistic loop closures in urban semi-structured environments. We captured these sequences using three different LiDAR sensors: Aeva (solid-state), Avia (MEMS), and Ouster (mechanical), to represent a wide range of sensor characteristics.

Our evaluation focusses on three key aspects: (i) loop closure validation quality (number of geometrically consistent inliers), (ii) accuracy of pose estimation, and (iii) computational efficiency. We compare our CliReg-based verification with a RANSAC-based baseline under identical conditions, using 3D and 2D features.

##### A. Metrics and Experimental Setup

To ensure a fair comparison, we configure a standard set of parameters across all experiments. In both the 2D and 3D pipelines, appearance-based descriptor matching uses a Hamming distance threshold of 50 bits. A loop closure is accepted only if a minimum number of mutually consistent inliers is found: 10 in the 2D case and 5 in the 3D case.

To compute the absolute pose error (APE) with and without loop closures (denoted APE w and APE w/o respectively), we follow a consistent evaluation pipeline. We begin by incorporating all detected loop closures into a pose graph optimization using g2o [15], initializing the graph with the trajectory estimated by the KISS-ICP algorithm [37], and the loop closure constraints obtained from our loop closure validation algorithm. After optimization, we compare the corrected poses with the ground-truth trajectory available for each sensor and scene. Finally, we compute the APE using the evo Python toolbox [9], a widely adopted tool for evaluating odometry and SLAM performance metrics.

##### B. 2D and 3D Pipeline Description

**2D pipeline:** We adopt a BEV density map representation of the environment, with each local map discretised into a grid using a resolution of 0.5 m. We discard low-density regions by applying a threshold of 5 % relative to the maximum cell occupancy. Local maps are formed incrementally until the platform travels a fixed distance, of 100 m, which defines the spatial extent of each local representation. This setup favours compact, viewpoint-agnostic descriptors, enabling efficient matching using HBST. A similar representation for loop closure has been proposed by Gupta et al. [10].

**3D pipeline:** We employ voxelization with a 0.5 m resolution to downsample the point cloud, ensuring uniform spatial coverage. Keypoints are extracted using ISS with scale-adaptive parameters: a salient radius of 3 m, a nonmaximum suppression radius of 2 m, and an eigenvalue threshold of 0.975. Around these keypoints, SHOT descriptors are computed within a support radius of 2 m and binarized to obtain B-SHOT signatures. Geometric consistency is enforced using a pairwise distance between the aligned features with tolerance set to 1 m. This configuration balances descriptiveness

and computational efficiency, making it suitable for real-time SLAM deployments.

**RANSAC configuration:** For the 3D experiments, we employ RANSAC in which the number of iterations is explicitly fixed to 10,000 (denoted as RANSAC-10K). This decision is motivated by the observation that, when using the number of iterations derived from the classical probabilistic formulation (assuming an inlier ratio of 0.3 and a success probability of 0.999), RANSAC may fail to detect any loop closures across all evaluated scenes. The same result was observed even when the number of iterations was fixed to 1,000. Therefore, we perform 10,000 RANSAC iterations in all 3D experiments to ensure a meaningful baseline for comparison.

##### C. Evaluation Results

We summarize the results for the 3D pipeline in Tab. I. Our approach consistently yields geometrically more reliable loop closures than RANSAC, particularly in challenging cases where RANSAC fails to detect any closure. Although our approach typically validates fewer matches, the identified inliers exhibit higher mutual consistency and produce a lower absolute pose error (APE).

In particular, RANSAC often fails entirely in sequences such as Aeva-Bridge01 or Ouster-Bridge02, while our approach reliably detects loop closures with 50–130 inliers. These results illustrate the robustness of our clique-based method under viewpoint and structural changes.

We also evaluate the 2D pipeline on the Roundabout01 sequence to analyze CliReg’s performance in dynamic and structurally complex environments. Tab. II presents the results for the 2D projection-based pipeline. Here, both RANSAC and our approach produce comparable F1 scores and APE, but our approach achieves this with a run time more than 10 times faster in all cases. This reinforces CliReg’s suitability for real-time applications in scenarios where descriptor dimensionality is reduced.

##### D. Runtime Performance and Variability

Across all 3D experiments, our approach operates within 2.8 ms to 6.3 ms per match, with a standard deviation below 1.3 ms. Although this is higher than 2D runtimes, it remains practical for use in real-time SLAM systems, especially considering the robustness gains.

##### E. Case Study: Detailed Failure and Success Example

In the Bridge01-Aeva sequence, RANSAC does not detect a valid loop closure. In contrast, our approach identifies a set of 61 inliers forming a mutually consistent clique, leading to an accurate pose estimate with an APE of 27.40 m, down from 94.67 m without loop closures. This demonstrates the ability of our approach to validate challenging loop closures that heuristic methods often miss.

##### F. Failure Modes in 2D Scenarios

In some 2D cases like Bridge01-Aeva, we observe an increase in APE after loop closure insertion for both RANSAC

TABLE I: **Loop closure performance using 3D pipeline across sensors and scenes with our method and RANSAC.** Metrics include number of inliers, runtime, and APE with and without loop closures.

Scene	Sensor	Algorithm	Inliers	Mean Time (ms)	APE w (m)	APE w/o (m)
Bridge01	Aeva	RANSAC-10K	-	-	-	94.67
		Ours	<b>61</b>	<b>6.25</b>	<b>27.40</b>	
	Avia	RANSAC-10K	6	324.17	79.05	167.83
		Ours	<b>24</b>	<b>2.83</b>	<b>33.98</b>	
	Ouster	RANSAC-10K	6	317.05	50.67	157.37
		Ours	<b>136</b>	<b>4.92</b>	<b>18.80</b>	
Bridge02	Aeva	RANSAC-10K	-	-	-	120.93
		Ours	<b>57</b>	<b>4.11</b>	<b>108.94</b>	
	Avia	RANSAC-10K	21	291.31	<b>16.48</b>	76.09
		Ours	<b>70</b>	<b>3.61</b>	19.00	
	Ouster	RANSAC-10K	-	-	-	53.95
		Ours	<b>71</b>	<b>4.52</b>	<b>68.25</b>	

TABLE II: **Loop closure performance using 2D pipeline across sensors and scenes with our method and RANSAC.** Metrics include number of inliers, runtime, precision, recall, F1 score, and APE with and without loop closures.

Scene	Sensor	Algorithm	Inliers	Mean Time (ms)	Precision	Recall	F1 Score	APE w (m)	APE w/o (m)
Bridge01	Aeva	RANSAC	<b>2985</b>	2.26	<b>0.9908</b>	0.0692	0.1292	<b>331.00</b>	94.67
		Ours	2927	<b>0.20</b>	0.9898	<b>0.0693</b>	<b>0.1295</b>	362.31	
	Avia	RANSAC	<b>1675</b>	2.13	<b>0.9877</b>	<b>0.0466</b>	<b>0.0890</b>	408.30	167.83
		Ours	1624	<b>0.13</b>	0.9869	0.0463	0.0883	<b>366.65</b>	
	Ouster	RANSAC	<b>4128</b>	2.24	0.9835	0.0737	0.1371	<b>332.80</b>	157.37
		Ours	4022	<b>0.23</b>	<b>0.9843</b>	<b>0.0737</b>	<b>0.1372</b>	380.01	
Bridge02	Aeva	RANSAC	<b>2059</b>	2.05	<b>0.9891</b>	0.0611	0.1151	277.00	120.93
		Ours	2011	<b>0.20</b>	0.9888	<b>0.0612</b>	<b>0.1152</b>	<b>276.49</b>	
	Avia	RANSAC	<b>1018</b>	2.05	<b>0.9667</b>	<b>0.0426</b>	<b>0.0816</b>	263.74	76.09
		Ours	995	<b>0.11</b>	0.9667	0.0422	0.0809	<b>255.88</b>	
	Ouster	RANSAC	<b>2200</b>	1.82	<b>0.9954</b>	<b>0.0589</b>	0.1113	257.90	53.95
		Ours	2120	<b>0.22</b>	<b>0.9954</b>	<b>0.0589</b>	<b>0.1113</b>	<b>204.26</b>	
Roundabout01	Aeva	RANSAC	<b>539</b>	2.00	<b>0.9972</b>	<b>0.0376</b>	<b>0.0725</b>	3.5217	21.7571
		Ours	518	<b>0.13</b>	0.9971	<b>0.0376</b>	<b>0.0725</b>	<b>3.5208</b>	
	Avia	RANSAC	174	1.95	0.9468	<b>0.0120</b>	<b>0.0233</b>	<b>2.2475</b>	21.9412
		Ours	<b>181</b>	<b>0.14</b>	<b>0.9469</b>	<b>0.0120</b>	<b>0.0233</b>	2.2584	
	Ouster	RANSAC	1130	1.97	0.8757	0.0245	0.0477	<b>1.5481</b>	11.7912
		Ours	1132	<b>0.21</b>	<b>0.8753</b>	<b>0.0246</b>	<b>0.0479</b>	1.5670	

and our method, an effect not seen in 3D. We attribute this to repetitive structures in bridge scenes which, when represented in BEV, may lead to globally inconsistent constraints despite locally correct matches. This suggests a limitation of the 2D representation due to spatial information loss, rather than a flaw in the verification algorithm.

### G. Summary

Our results confirm that: (i) our approach reliably validates loop closures even where RANSAC fails, (ii) produces a lower or comparable APE in both 2D and 3D, (iii) and executes efficiently enough for practical deployment.

Future work will explore temporal and semantic consistency, as well as faster clique search and hybrid descriptors, to enhance robustness and efficiency.

## V. CONCLUSION

In this work, we introduced a loop closure validation algorithm that formulates geometric verification as a maximal clique search over feature correspondences. Our approach optimizes for mutual consistency among matches and eliminates the reliance on random sampling. Using binary descriptors, we integrated it into a full 3D LiDAR-based loop closure pipeline and validated its performance on challenging real-world datasets. Our results show improved robustness and comparable or better accuracy than RANSAC, even in sparse or noisy conditions. Despite solving an NP-hard problem, determining loop closures via searching cliques in a correspondence graph maintains real-time applicability in typical mobile robotics setups. These characteristics make our approach a reliable and practical alternative for loop

closure detection in SLAM.

Despite these encouraging results, there is still room for improvement. Future work will focus on accelerating the clique search step through parallel strategies or approximate solutions, integrating learning-based descriptors to enhance matching robustness, and extending it to exploit temporal consistency across consecutive map segments.

## REFERENCES

- [1] K. Arun, T. Huang, and S. Blostein. Least-Squares Fitting of Two 3-D Point Sets. *IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, 9(5):698–700, 1987.
- [2] J. Behley and C. Stachniss. Efficient Surfel-Based SLAM using 3D Laser Range Data in Urban Environments. In *Proc. of Robotics: Science and Systems (RSS)*, 2018.
- [3] J.L. Blanco-Claraco. A flexible framework for accurate LiDAR odometry, map manipulation, and localization. *Intl. Journal of Robotics Research (IJRR)*, 0(0), 2025.
- [4] M. Bosse and R. Zlot. Place Recognition Using Keypoint Voting in Large 3D Lidar Datasets. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2013.
- [5] R. Dubé, D. Dugas, E. Stumm, J. Nieto, R. Siegwart, and C. Lerma. SegMatch: Segment Based Place Recognition in 3D Point Clouds. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2017.
- [6] M. Fischler and R. Bolles. Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography. *Communications of the ACM*, 24(6):381–395, 1981.
- [7] D. Galvez-Lopez and J.D. Tardos. Bags of Binary Words for Fast Place Recognition in Image Sequences. *IEEE Trans. on Robotics (TRO)*, 28(5):1188–1197, 2012.
- [8] G. Grisetti, R. Kümmerle, C. Stachniss, and W. Burgard. A tutorial on graph-based SLAM. *IEEE Trans. on Intelligent Transportation Systems Magazine*, 2(4):31–43, 2010.
- [9] M. Grupp. evo: Python package for the evaluation of odometry and SLAM. <https://github.com/MichaelGrupp/evo>, 2017.
- [10] S. Gupta, T. Guadagnino, B. Mersch, I. Vizzo, and C. Stachniss. Effectively Detecting Loop Closures using Point Cloud Density Maps. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2024.
- [11] B.K. Horn. Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America A*, 4:629–642, 1987.
- [12] B. Jiang and S. Shen. Contour Context Abstract Structural Distribution for 3D LiDAR Loop Detection and Metric Pose Estimation. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2023.
- [13] M. Jung, W. Yang, D. Lee, H. Gil, G. Kim, and A. Kim. HeLiPR: Heterogeneous LiDAR Dataset for inter-LiDAR Place Recognition under Spatiotemporal Variations. *Intl. Journal of Robotics Research (IJRR)*, 12(43):1867–1883, 2024.
- [14] G. Kim, S. Choi, and A. Kim. Scan Context++: Structural Place Recognition Robust to Rotation and Lateral Variations in Urban Environments. *IEEE Trans. on Robotics (TRO)*, 38(2):21–27, 2021.
- [15] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard. g2o: A general framework for graph optimization. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2011.
- [16] J. Laserna, P. San Segundo, and D. Álvarez. Clireg: Clique-based robust point cloud registration. *IEEE Trans. on Robotics (TRO)*, 41:1898–1917, 2025.
- [17] S. Lowry, N. Sunderhauf, P. Newman, J. Leonard, D. Cox, P. Corke, and M. Milford. Visual Place Recognition: A Survey. *IEEE Trans. on Robotics (TRO)*, 32(1):1–19, 2016.
- [18] L. Luo, S. Cao, B. Han, H. Shen, and J. Li. BVMatch: Lidar-Based Place Recognition Using Bird’s-Eye View Images. *IEEE Robotics and Automation Letters (RA-L)*, 6(3):6076–6083, 2021.
- [19] J. Mangelson, D. Dominic, R. Eustice, and R. Vasudevan. Pairwise Consistent Measurement Set Maximization for Robust Multi-Robot Map Merging. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2018.
- [20] J. Neira and J. Tardos. Data Association in Stochastic Mapping using the Joint Compatibility Test. *IEEE Trans. on Robotics (TRO)*, 17:890–897, 2001.
- [21] E. Olson. Real-Time Correlative Scan Matching. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2009.
- [22] S.M. Prakhya, B. Liu, and W. Lin. B-SHOT: A binary feature descriptor for fast and efficient keypoint matching on 3D point clouds. In *Proc. of the IEEE/RISJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2015.
- [23] R. Raguram, J.M. Frahm, and M. Pollefeys. A Comparative Analysis of RANSAC Techniques Leading to Adaptive Real-Time Random Sample Consensus. In *Proc. of the Europ. Conf. on Computer Vision (ECCV)*, 2008.
- [24] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski. Orb: an efficient alternative to sift or surf. In *Proc. of the IEEE Intl. Conf. on Computer Vision (ICCV)*, 2011.
- [25] R. Rusu, N. Blodow, and M. Beetz. Fast point feature histograms (fpfh) for 3d registration. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2009.
- [26] R. Salakhutdinov and G. Hinton. Learning a Nonlinear Embedding by Preserving Class Neighbourhood Structure. In *Proc. of the Intl. Conf. on Artificial Intelligence and Statistics*, 2007.
- [27] S. Salti, F. Tombari, and L. Stefano. SHOT: Unique Signatures of Histograms for Surface and Texture Description. *Journal of Computer Vision and Image Understanding (CVIU)*, 125:251–264, 2014.
- [28] P. San Segundo, F. Furini, and J. Artieda. A new branch-and-bound algorithm for the maximum weighted clique problem. *Computers & Operations Research*, 110:18–33, 2019.
- [29] P. San Segundo, A. Lopez, and P.M. Pardalos. A new exact maximum clique algorithm for large and massive sparse graphs. *Computers & Operations Research*, 66:81–94, 2016.
- [30] P. San Segundo, A. Nikolaev, M. Batsyn, and P.M. Pardalos. Improved infra-chromatic bound for exact maximum clique search. *Informatica*, 27(2):463–487, 2016.
- [31] P. San Segundo, A. Nikolaev, and M. Batsyn. Infra-chromatic bound for exact maximum clique search. *Computers & Operations Research*, 64:293–303, 2015.
- [32] D. Schlegel and G. Grisetti. HBST: A Hamming Distance Embedding Binary Search Tree for Visual Place Recognition. *IEEE Robotics and Automation Letters (RA-L)*, 3(4):3741–3748, 2018.
- [33] J. Shi, H. Yang, and L. Carlone. ROBIN: a Graph-Theoretic Approach to Reject Outliers in Robust Estimation using Invariants. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2021.
- [34] B. Steder, G. Grisetti, and W. Burgard. Robust Place Recognition for 3D Range Data Based on Point Features. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2010.
- [35] B. Steder, R. Rusu, K. Konolige, and W. Burgard. NARF: 3D range image features for object recognition. In *Proc. of the IROS Workshop on Defining and Solving Realistic Perception Problems in Personal Robotics*, 2010.
- [36] A. Uy and G. Lee. PointNetVLAD: Deep point cloud based retrieval for large-scale place recognition. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [37] I. Vizzo, T. Guadagnino, B. Mersch, L. Wiesmann, J. Behley, and C. Stachniss. KISS-ICP: In Defense of Point-to-Point ICP – Simple, Accurate, and Robust Registration If Done the Right Way. *IEEE Robotics and Automation Letters (RA-L)*, 8(2):1029–1036, 2023.
- [38] M.W. Walker, L. Shao, and R.A. Volz. Estimating 3-d location parameters using dual number quaternions. *Computer Vision, Graphics, and Image Processing*, 54:1–14, 1991.
- [39] T. Xu, Y.C. Guo, Z. Li, G. Yu, Y.K. Lai, and S.H. Zhang. TransLoc3D: Point Cloud based Large-scale Place Recognition using Adaptive Receptive Fields. *Communications in Information and Systems*, 23:57–83, 2023.
- [40] J. Yang, H. Li, D. Campbell, and Y. Jia. Go-ICP: A Globally Optimal Solution to 3D ICP Point-Set Registration. *IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, 38(11):2241–2254, 2015.
- [41] C. Yuan, J. Lin, Z. Zou, X. Hong, and F. Zhang. STD: Stable Triangle Descriptor for 3D place recognition. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2023.
- [42] Y. Zhong. Intrinsic Shape Signatures: A Shape Descriptor for 3D Object Recognition. In *Proc. of the Intl. Conf. on Computer Vision Workshops*, 2009.