Improving Map Consistency in Graph-Based LiDAR SLAM Through Information-Aware Odometry and Retroactive Loop Closure

Saurabh Gupta Niklas Trekel Louis Wiesmann Cyrill Stachniss

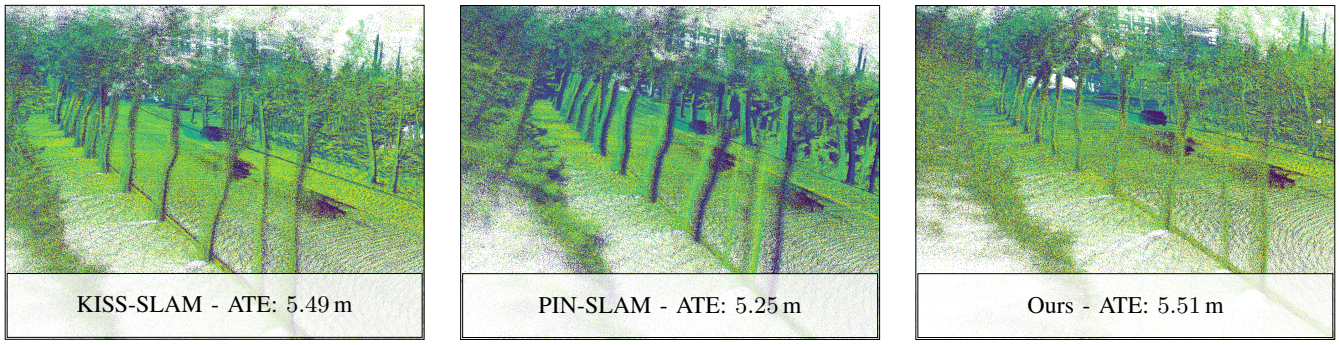

Fig. 1: **Trajectory accuracy does not necessarily imply consistent map quality.** KISS-SLAM [13], PIN-SLAM [28], and our approach achieve comparable absolute trajectory error (ATE) values, yet have substantial differences in local geometric consistency. Our method better preserves the alignment of repeated observations, reducing duplicated structures in the reconstructed map. The points in all images are colored based on their timestamp, with the contrast between points reflecting temporally separated revisits.

***Abstract*— High-quality maps are fundamental for robotics tasks such as navigation and planning. Although modern graph-based LiDAR SLAM systems achieve good trajectory accuracies, a low trajectory error alone does not guarantee geometrically consistent maps, particularly at revisit locations where missed loop closures and residual drift can produce local misalignments. In this work, we address the problem of jointly improving global trajectory estimation and local map quality in 3D LiDAR SLAM. We first propose a framework to efficiently estimate geometry-dependent information matrices for ICP, enabling principled weighting of odometry constraints in a pose graph. We then introduce a hierarchical loop-closure module that decouples place recognition from geometric registration, together with a retroactive loop-closure module that exploits the optimized pose graph to recover missed loop closures. We also propose an evaluation protocol to measure map consistency at revisit locations. We evaluate our SLAM system on several datasets against state-of-the-art LiDAR SLAM systems. Experimental results demonstrate global trajectory accuracies on par with or better than existing methods while consistently improving local geometric map consistency at revisit locations. These results suggest that coupling uncertainty-aware odometry with geometry-guided loop-closure refinement leads to more accurate trajectories and higher-quality maps.**

I. Introduction

Mobile robots require accurate maps of their environments for reliable localization and planning. Three-dimensional LiDARs have become a popular sensing modality because of their accurate range measurements over large outdoor environments. Consequently, LiDAR-based simultaneous localization and mapping (SLAM) has become a fundamental component of autonomous robotic systems.

Graph-based SLAM systems [2], [11], [22] combine sensor odometry, loop closure detection, and pose graph optimization to produce globally consistent trajectories. Although these systems achieve good trajectory accuracies, conventional trajectory metrics do not fully capture the geometric consistency of the reconstructed map, see Fig. 1. Small residual pose errors that accumulate over long trajectories, or weakly constrained revisited regions can still result in duplicated surfaces and locally inconsistent geometry despite low global trajectory error.

This limitation arises because pose graph optimization can only exploit the constraints available in the graph. If individual loop closures are missed, revisited regions remain weakly constrained and may still exhibit duplicated surfaces or local geometric inconsistencies even after optimization. Improving map quality therefore requires generating additional geometric constraints in revisited regions.

We address this problem by treating trajectory estimation and map quality as complementary objectives within the SLAM framework. We first improve trajectory estimates through efficient geometry-aware information estimation for iterative closest point (ICP) registration, enabling principled weighting of odometry constraints in pose graph optimization. In parallel, we also focus on improving map quality through a hierarchical loop-closure framework followed by a retroactive geometric refinement stage that recovers revisits missed by the appearance-based front-end.

The main contribution of this paper is a novel and effective 3D LiDAR SLAM approach that improves global trajectory estimates and local map quality by coupling real-time infor-

All authors are with the University of Bonn, Center for Robotics. Cyrill Stachniss is additionally with the Lamarr Institute for Machine Learning and Artificial Intelligence. This work has partially been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy, EXC-2070 – 390732324 – PhenoRob, by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under STA 1051/5-1 within the FOR 5351 (AID4Crops), and by the German Federal Ministry of Research, Technology and Space (BMFTR) under the Robotics Institute Germany (RIG).

mation estimation for point-to-point ICP, hierarchical loop closure estimation, and geometry-based retroactive loop closure detection within a pose-graph optimization framework.

First, we introduce a computationally efficient method for estimating geometry-dependent information matrices for point-to-point ICP in real time, making information-weighted pose graph optimization practical without modifying the underlying registration algorithm. Second, we introduce a hierarchical multi-resolution loop-closure framework that decouples place recognition from geometric registration by performing them at different spatial resolutions. Large local maps enhances the robustness of detected loops, while smaller submaps achieve better registration accuracy. Third, we propose a geometry-based retroactive loop-closure detector that closes the feedback loop between the back-end and the front-end by using the optimized pose graph to recover valid revisits missed during appearance-based place recognition. Finally, because conventional trajectory metrics do not quantify local geometric consistency, we introduce an evaluation protocol that explicitly measures map quality at revisit locations alongside global trajectory accuracy.

In summary, we make two key claims: Our approach first provides global trajectory accuracy that is on par with or better than state-of-the-art SLAM systems and second, it improves the local map quality in revisited areas, even in scenes with perceptual aliasing. These claims are backed up by the paper and our experimental evaluation.

II. Related Work

The majority of today's 3D LiDAR SLAM systems typically follow the graph-based SLAM paradigm [11], [22], consisting of an odometry front-end, loop closure detection, and pose graph optimization. Existing systems differ primarily in their scan registration and map representation, including point-to-point ICP [13], multi-metric registration [27], [34], continuous-time methods [8], surfel maps [1], neural representations [9], [28], or semantic front-ends [7], [32]. Despite these differences, they all ultimately rely on pose graph optimization to reduce accumulated drift after revisits.

Pose graph optimization assumes appropriately weighted measurement constraints typically expressed through the information matrix of the measurement. While graph optimization supports adaptive and anisotropic information matrices, many systems still rely on fixed isotropic weights as estimating uncertainty for ICP remains computationally expensive. Although analytical approaches can estimate the ICP information matrix from the Gauss-Newton Hessian [5], [6], these approximations become inaccurate for point-to-point or other multi-metric ICP objectives because correspondence changes violate the known-correspondence assumption underlying the Gauss-Newton approximation [3]. Sampling-based approaches estimate uncertainty through repeated perturbation and registration trials [21], [23], [25], but they are computationally expensive. Real-time information estimation for ICP therefore remains an obstacle to fully exploiting uncertainty-aware pose graph optimization. We propose a linear regression estimator for the information matrix, fit directly to the samples drawn from the loss landscape, compatible with real-time ICP algorithms.

Another important component of a SLAM front-end is a loop-closure detector. Reliable loop closures require both robust place recognition and accurate geometric verification. Practical systems must remain robust under perceptual aliasing while avoiding catastrophic false-positive loop closures [4]. Scan-based loop-closure detectors are used by several SLAM algorithms [19], [20]. Scan-based methods provide accurate geometric alignment but operate on limited spatial context, making them more susceptible to perceptual aliasing. Practical systems therefore use conservative thresholds, which reduce false positives at the expense of missed loop closures.

An alternative is to perform loop closure using aggregated local maps or submaps instead of individual scans [15], [24], [33]. Aggregating scans leads to descriptor robustness and naturally produces sparse pose-graph constraints. However, larger submaps also accumulate odometry drift, degrading the accuracy of geometric registration. In contrast, scan-level registration is not affected by accumulated drift but provides less spatial context for place recognition. Our framework combines the strengths of both representations through a map hierarchy in the front-end, using large local maps for robust place recognition and smaller submaps for accurate geometric alignment. Existing hierarchical SLAM systems, on the other hand, employ hierarchy primarily in the back-end to improve optimization scalability and efficiency [12].

Conventional graph-based SLAM pipelines treat the front-end and back-end as sequential modules: the front-end generates loop-closure constraints, and the back-end optimizes the pose graph. Since front-end loop closure detection is inherently imperfect, considerable work has focused on making pose graph optimization robust to false-positive constraints through robust kernels [30], or max-mixture models [26].

Our work addresses the complementary problem. Rather than increasing the back-end's tolerance to incorrect loop closures, we deliberately employ a conservative front-end to avoid false positives and then use the optimized graph geometry to recover valid loop closures that were initially missed, particularly in perceptually aliased environments. This creates a feedback loop between the front-end and the back-end, and still allows for the use of a robust kernel. Overall, our framework improves both global trajectory accuracy and map quality through information-aware odometry, hierarchical loop-closure front-end, and geometry-guided back-end loop closures.

III. Our Approach

In this section, we describe the main components of our SLAM pipeline, as illustrated in Fig. 2. Given a sequence of LiDAR scans, we compute the incremental odometry estimates along with an approximate uncertainty estimate for each incremental frame-to-map alignment in Sec. III-A and Sec. III-B. Then, we introduce the hierarchical front-end setup for loop closure detection in Sec. III-C. We introduce the pose-graph back-end in Sec. III-D. Finally, in Sec. III-E

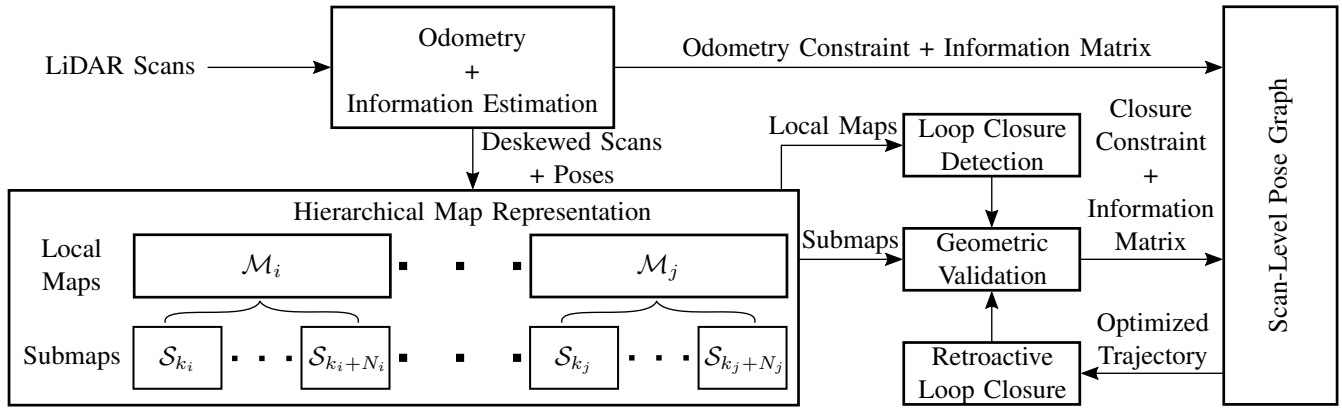

Fig. 2: **Overview of our LiDAR SLAM approach.** The front-end computes LiDAR odometry and information matrix estimates. Registered scans are organized into a hierarchical map representation with large local maps for place recognition and smaller submaps for geometric validation. Odometry and loop-closure constraints are both optimized in a scan-level pose graph. The optimized trajectory is fed back to a retroactive loop-closure module, which recovers loop closures missed by the appearance-based detector.

we elaborate on the retroactive loop-closure detector in the back-end.

A. LiDAR Odometry

We use KISS-ICP [31] as our LiDAR odometry front-end; no contribution is claimed for Sec. III-A. Given a point cloud $\mathcal{P}_t = \{\boldsymbol{p}_i \mid \boldsymbol{p}_i \in \mathbb{R}^3\}$ at time t, we create a deskewed and downsampled $\mathcal{P}_t^* \subset \mathcal{P}_t$. We initialize the pose estimate as ${}^o\hat{\mathsf{T}}_t = {}^o\mathsf{T}_{t-1}\Delta\mathsf{T}_t$, where ${}^o\mathsf{T}_{t-1} \in \mathbb{SE}(3)$ is the previous pose estimate, $\Delta\mathsf{T}_t \in \mathbb{SE}(3)$ is the constant-velocity motion prediction and o refers to the odometry frame.

We register $\mathcal{P}_t^*$ to a voxel local map $\mathcal{Q} = \{\boldsymbol{q}_j \mid \boldsymbol{q}_j \in \mathbb{R}^3\}$ by minimizing the point-to-point ICP objective

$$\chi(\mathsf{T}) = \sum_{(i,j)\in\mathcal{C}(\mathsf{T})} \left\|\mathsf{T}\boldsymbol{p}_i - \boldsymbol{q}_j\right\|_2^2, \tag{1}$$

where $\mathcal{C}$ are the nearest-neighbor matches between $\mathcal{P}_t^*$ and $\mathcal{Q}$ for a given pose $\mathsf{T} \in \mathbb{SE}(3)$. Starting from ${}^o\hat{\mathsf{T}}_t$, we solve

$$\delta\boldsymbol{\xi} = \underset{\delta\boldsymbol{\xi}}{\operatorname{argmin}}\, \chi({}^o\hat{\mathsf{T}}_t \boxplus \delta\boldsymbol{\xi}), \tag{2}$$

where $\delta\boldsymbol{\xi} \in \mathbb{R}^6$ is a pose increment and $\boxplus$ applies this increment to the current estimate. Upon convergence, we obtain the pose ${}^o\mathsf{T}_t$ and insert the scan into the local map.

B. Local Information Matrix Approximation

KISS-ICP provides the frame-to-map transformation but no uncertainty information. We propose a computationally efficient method to approximate the local information matrix of the transformation. We focus on modelling the local geometry of the ICP objective around the converged pose ${}^o\mathsf{T}_t$.

Instead of relying on a fixed-correspondence Gauss-Newton approximation or repeated registration trials, we directly estimate a local Hessian approximation around ${}^o\mathsf{T}_t$. We model the point-to-point ICP objective with a second-order Taylor-series expansion

$$\chi({}^o\mathsf{T}_t \boxplus \delta\boldsymbol{\xi}_m) \approx \chi({}^o\mathsf{T}_t) + \delta\boldsymbol{\xi}_m^\mathsf{T}\boldsymbol{g} + \frac{1}{2}\delta\boldsymbol{\xi}_m^\mathsf{T} H \delta\boldsymbol{\xi}_m, \tag{3}$$

where $\boldsymbol{g} = \nabla\chi({}^o\mathsf{T}_t)$ and $H = \nabla^2\chi({}^o\mathsf{T}_t)$ denote the local gradient and Hessian, respectively. We retain the gradient term to account for imperfect convergence. Subtracting the constant term $\chi({}^o\mathsf{T}_t)$ yields

$$b_m = \chi({}^o\mathsf{T}_t \boxplus \delta\boldsymbol{\xi}_m) - \chi({}^o\mathsf{T}_t) \tag{4}$$

$$\approx \delta\boldsymbol{\xi}_m^\mathsf{T}\boldsymbol{g} + \frac{1}{2}\delta\boldsymbol{\xi}_m^\mathsf{T} H \delta\boldsymbol{\xi}_m. \tag{5}$$

We sample the ICP objective around ${}^o\mathsf{T}_t$ using perturbations $\mathcal{U} = \{\delta\boldsymbol{\xi}_m \mid \delta\boldsymbol{\xi}_m \in \mathbb{R}^6\}$ on the tangent space of $\mathbb{SE}(3)$. We empirically choose perturbations in $[-0.1, 0.1]\,\mathrm{m}$ for translation and $[-0.01, 0.01]\,\mathrm{rad}$ for rotation, large enough to induce correspondence changes while staying inside the convergence basin. For each sample, we recompute correspondences and evaluate $\chi({}^o\mathsf{T}_t \boxplus \delta\boldsymbol{\xi}_m)$.

Since the coefficients of $\boldsymbol{g}$ and H appear linearly in Eq. (5), we obtain a linear system of the form $A\boldsymbol{x} = \boldsymbol{b}$. To reduce computational cost, we approximate translational and rotational sensitivities independently near convergence and neglect second-order translation-rotation coupling terms. This yields the block diagonal approximation

$$H = \begin{bmatrix} H_{\mathrm{tr}} & 0 \\ 0 & H_{\mathrm{ro}} \end{bmatrix}, \tag{6}$$

where $H_{\mathrm{tr}} \in \mathbb{R}^{3\times3}$ and $H_{\mathrm{ro}} \in \mathbb{R}^{3\times3}$ denote the translational and rotational Hessian blocks, respectively. We can therefore solve two independent regression problems

$$A_{\mathrm{tr}}\boldsymbol{x}_{\mathrm{tr}} = \boldsymbol{b}_{\mathrm{tr}} \qquad \text{and} \qquad A_{\mathrm{ro}}\boldsymbol{x}_{\mathrm{ro}} = \boldsymbol{b}_{\mathrm{ro}}, \tag{7}$$

where $A \in \mathbb{R}^{M\times9}$, $\boldsymbol{x} \in \mathbb{R}^9$, $\boldsymbol{b} \in \mathbb{R}^M$, and M denotes the number of perturbations.

For a translational perturbation $\delta\boldsymbol{\xi}_{\mathrm{tr},m} = [x_m\ y_m\ z_m]$, the corresponding row of the design matrix A_{tr} is

$$A_{\mathrm{tr},m} = \left[x_m, y_m, z_m, \tfrac{x_m^2}{2}, \tfrac{y_m^2}{2}, \tfrac{z_m^2}{2}, x_m y_m, y_m z_m, x_m z_m\right], \tag{8}$$

with unknown parameter vector

$$\boldsymbol{x}_{\mathrm{tr}}^\mathsf{T} = [g_x, g_y, g_z, h_{xx}, h_{yy}, h_{zz}, h_{xy}, h_{yz}, h_{xz}]. \tag{9}$$

Here, $[g_x\ g_y\ g_z]$ denote the translational gradient $\boldsymbol{g}_{\mathrm{tr}}^\mathsf{T}$ and h_{ij} the six coefficients of the symmetric Hessian block H_{tr}.

We construct H_{ro} analogously using rotational perturbations on the tangent space of $\mathbb{SO}(3)$. We then combine H_{tr} and H_{ro} as per Eq. (6) to obtain H. We project H onto a positive-definite cone through eigendecomposition, where we clamp all eigenvalues below 10^{-3} to 10^{-3} to obtain H^{+}.

Following the standard interpretation of curvature as information in nonlinear least-squares estimation, we use $\Lambda_{\text{odo}} = H^{+}$ as the information matrix for the odometry constraint in the pose graph optimization. Note that this information estimation framework can also be used with other ICP objectives without loss of generality.

C. Hierarchical Front-End

We maintain a hierarchical map representation in the front-end, as illustrated in Fig. 2. Large local maps provide sufficient spatial context for reliable place recognition. Each local map consists of multiple smaller submaps that accumulate less odometry drift and are therefore used for geometric registration. We generate both representations online using a displacement-based splitting criterion.

We create a submap whenever the platform displaces more than $\tau_{\text{sub}} = 25\,\text{m}$. A submap is represented as $\mathcal{S}_k = ({}^{m}\mathsf{T}_k^{\text{sub}}, \mathcal{V}_k^{\text{sub}}, \Gamma_k^{\text{sub}})$, where ${}^{m}\mathsf{T}_k^{\text{sub}}$ denotes the keypose corresponding to the first scan, $\mathcal{V}_k^{\text{sub}}$ is a voxel grid with resolution ν_{sub} and $\Gamma_k^{\text{sub}} = ({}^{t_k^s}\mathsf{T}_{t_k^s}, \ldots, {}^{t_k^s}\mathsf{T}_{t_k^e})$ is the local trajectory segment, where t_k^s and t_k^e denote the timestamps of the first and the last scan of the kth submap, respectively.

In parallel, we create local maps using a larger displacement threshold $\tau_{\text{loc}} = 100\,\text{m}$. To maintain an unambiguous association between submaps and local maps, every local-map split simultaneously triggers a submap split. This ensures that each submap belongs to exactly one local map. A local map is represented as $\mathcal{M}_n = ({}^{m}\mathsf{T}_n^{\text{loc}},\ \mathcal{V}_n^{\text{loc}}, \mathcal{I}_n)$, where ${}^{m}\mathsf{T}_n^{\text{loc}}$ denotes the keypose corresponding to the first scan, $\mathcal{V}_n^{\text{loc}}$ is a voxel grid with resolution ν_{loc} and $\mathcal{I}_n = \{k_n, \ldots, k_n + N_n\}$ are the indices of its constituent submaps. We also define the local-map trajectory by concatenating the trajectories of its constituent submaps,

$$\Gamma_n^{\text{loc}} = \text{concat}\left({}^{m}\mathsf{T}_{k_n}^{\text{sub}}\Gamma_{k_n}^{\text{sub}}, \ldots, {}^{m}\mathsf{T}_{k_n+N_n}^{\text{sub}}\Gamma_{k_n+N_n}^{\text{sub}}\right), \quad (10)$$

where transforming a trajectory by a pose denotes left multiplication of every pose in the trajectory.

The place-recognition stage, based on the method of Gupta et al. [14], identifies two overlapping local maps and provides an approximate relative pose. For a query local map $\mathcal{M}_i$, we search for the local map $\mathcal{M}_j$ with the largest number of inlier feature matches exceeding τ_{inliers}. We accept such a pair (i, j) of local maps as a loop closure candidate and get an initial alignment ${}^{j}\hat{\mathsf{T}}_i \in \mathbb{SE}(3)$ between their keyposes ${}^{m}\mathsf{T}_i^{\text{loc}}$ and ${}^{m}\mathsf{T}_j^{\text{loc}}$ from the feature matches.

Since each local map contains multiple submaps within the hierarchical structure, we next identify the pairs of submaps that observe the same physical region. Only those pairs are used to generate loop-closure constraints. We identify such pairs using trajectory proximity instead of voxel map-based overlap because it is computationally cheaper yet provides a reliable estimate of spatial overlap between submaps.

For each source submap with index $k \in \mathcal{I}_i$ and target submap with index $r \in \mathcal{I}_j$, we express their keyposes relative to their parent local-map coordinate frame as

$$ {}^{i}\tilde{\mathsf{T}}_k^{\text{sub}} = ({}^{m}\mathsf{T}_i^{\text{loc}})^{-1}\,{}^{m}\mathsf{T}_k^{\text{sub}} \quad \text{and} \quad {}^{j}\tilde{\mathsf{T}}_r^{\text{sub}} = ({}^{m}\mathsf{T}_j^{\text{loc}})^{-1}\,{}^{m}\mathsf{T}_r^{\text{sub}}. \quad (11)$$

Using the initial loop closure estimate ${}^{j}\hat{\mathsf{T}}_i$, we transform the source submaps into the target local-map frame as ${}^{j}\hat{\mathsf{T}}_k^{\text{sub}} = {}^{j}\hat{\mathsf{T}}_i\,{}^{i}\tilde{\mathsf{T}}_k^{\text{sub}}$. We then transform the source and target submap trajectories as

$$\hat{\Gamma}_k^{\text{sub}} = {}^{j}\hat{\mathsf{T}}_k^{\text{sub}}\Gamma_k^{\text{sub}} \quad \text{and} \quad \tilde{\Gamma}_r^{\text{sub}} = {}^{j}\tilde{\mathsf{T}}_r^{\text{sub}}\Gamma_r^{\text{sub}}. \quad (12)$$

This expresses both submap trajectories in a common coordinate frame of the target local map j, enabling direct geometric comparison.

We identify corresponding submaps by comparing these transformed trajectories using the Hausdorff distance

$$\begin{aligned} d_H(\Gamma_1, \Gamma_2) = \max\Bigg(&\max_{\boldsymbol{t}_a \in \Gamma_1}\left(\min_{\boldsymbol{t}_b \in \Gamma_2} d(\boldsymbol{t}_a, \boldsymbol{t}_b)\right), \\ &\max_{\boldsymbol{t}_b \in \Gamma_2}\left(\min_{\boldsymbol{t}_a \in \Gamma_1} d(\boldsymbol{t}_a, \boldsymbol{t}_b)\right)\Bigg), \end{aligned} \quad (13)$$

where $d(\boldsymbol{t}_a, \boldsymbol{t}_b)$ is the L2 norm between the translational components $\boldsymbol{t}_a$ and $\boldsymbol{t}_b$ of the corresponding trajectories.

For each source submap $k \in \mathcal{I}_i$, we identify the nearest target submap using Eq. (13) as

$$r^*(k) = \operatorname*{argmin}_{r \in \mathcal{I}_j} d_{\text{H}}\left(\hat{\Gamma}_k^{\text{sub}}, \tilde{\Gamma}_r^{\text{sub}}\right), \quad (14)$$

and accept the pair $(k, r^*(k))$ whenever

$$d_{\text{H}}\left(\hat{\Gamma}_k^{\text{sub}}, \tilde{\Gamma}_{r^*(k)}^{\text{sub}}\right) \le \tau_{\text{H}}, \quad (15)$$

where τ_{H} denotes the Hausdorff distance threshold, yielding the candidate correspondence set

$$\mathcal{Z}_{ij} = \left\{(k, r^*(k)) \mid k \in \mathcal{I}_i, d_{\text{H}}\left(\hat{\Gamma}_k^{\text{sub}}, \tilde{\Gamma}_{r^*(k)}^{\text{sub}}\right) \le \tau_{\text{H}}\right\}. \quad (16)$$

Finally, for every candidate pair $(k, r) \in \mathcal{Z}_{ij}$, we perform fine registration between $\mathcal{V}_k^{\text{sub}}$ and $\mathcal{V}_r^{\text{sub}}$. We initialize the fine registration with

$${}^{r}\hat{\mathsf{T}}_k = ({}^{j}\tilde{\mathsf{T}}_r^{\text{sub}})^{-1}\,{}^{j}\hat{\mathsf{T}}_k^{\text{sub}}, \quad (17)$$

yielding a refined relative pose ${}^{r}\mathsf{T}_k$, the associated information matrix Λ_{lc}, and a registration fitness score s_{rk}, where larger fitness scores indicate better alignment. If the registration fitness score s_{rk} exceeds the acceptance threshold τ_{f}, we insert the loop-closure constraint into the pose graph.

D. Pose Graph Back-End

We maintain a pose graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ at the back-end of our pipeline, where $\mathcal{V}$ is the set of nodes and $\mathcal{E} = \mathcal{E}_{\text{odo}} \cup \mathcal{E}_{\text{lc}}$ consists of sequential odometry constraints ($\mathcal{E}_{\text{odo}}$) and loop-closure constraints ($\mathcal{E}_{\text{lc}}$). We employ a scan-level pose graph representation to enable smooth deformation of the estimated global trajectory, with errors distributed over each scan in the graph upon the insertion of a loop closure. As we detect

loop closure at most once per local map (every $100\,\mathrm{m}$), pose graph optimization is performed infrequently and therefore does not affect the runtime of our SLAM algorithm.

Each incoming LiDAR scan is a node in the pose graph. For every scan acquired at time t belonging to submap k, we obtain its pose in the global map frame ${}^{m}\mathsf{T}_t = {}^{m}\mathsf{T}_k^{\text{sub}}\mathsf{T}_{t_k^s \to t}$, where $t \in [t_k^s, t_k^e]$. We connect consecutive scan nodes by relative odometry constraints ($\mathcal{E}_{\text{odo}}$) along with their corresponding information matrices Λ_{odo}. After detecting and validating loop closures between submaps belonging to local maps $\mathcal{M}_i$ and $\mathcal{M}_j$, we add a loop-closure constraint ($\mathcal{E}_{\text{lc}}$) between the nodes corresponding to the first scan in the submaps with the corresponding information matrix Λ_{lc}.

After a pose graph optimization, we propagate the optimized poses back to the hierarchical map representation. For every submap, we update its keypose with the optimized pose of the node corresponding to its first scan and recompute the complete trajectory segment from the subsequent nodes.

E. Geometry-Based Retroactive Loop Closure Detection

The front-end loop-closure detector used in this work [14] favors precision over recall through a self-similarity feature-pruning stage in order to avoid false positives, particularly in perceptually aliased environments. While this conservative strategy improves robustness and avoids wrong loop closures, it inevitably misses valid loop closures, leaving portions of the global trajectory weakly constrained, potentially producing locally imperfect maps.

However, once the robust loop closures are incorporated, the pose graph optimization redistributes the accumulated drift, improving the global alignment of revisited regions. Trajectories corresponding to the same physical location, previously separated in the estimated map due to accumulated drift are now in closer geometric proximity. Exploiting this observation, we perform a retroactive loop closure based on geometric proximity, recovering loop closures from areas where appearance-based loop closure acts conservatively.

In particular, we revisit previously created local maps and identify candidate loop closures in the local vicinity by comparing their optimized trajectories. We could naïvely revisit every local map after each optimization, but this would be computationally expensive. Instead, we revisit only the contiguous sequence of local maps that lie between the current local map and the previous local map to have established a successful appearance-based loop closure. These maps have not previously benefited from loop-closure corrections, making them the most promising candidates for recovering additional loop closures after optimization.

We compute the Hausdorff distance, see Eq. (13), between the optimized trajectories of each candidate local-map pair. We retain a local-map pair (i, j) if $d_{\text{H}}\left(\Gamma_i^{\text{loc}}, \Gamma_j^{\text{loc}}\right) \leq \tau_{\text{H}}$. This coarse geometric search substantially reduces the number of submap pairs that require fine registration.

We then pass the retained pairs to the hierarchical submap association module described in Sec. III-C. The optimized pose graph provides an accurate initial estimate of the relative pose between nearby submaps as

$$ {}^{r}\hat{\mathsf{T}}_k = \left({}^{m}\mathsf{T}_r^{\text{sub}}\right)^{-1}\, {}^{m}\mathsf{T}_k^{\text{sub}}, \tag{18} $$

which we use directly to initialize the fine registration.

Consequently, the retroactive loop-closure detector exploits the geometric consistency of the optimized pose graph, recovering loop closures that the front-end intentionally rejected or failed to detect. This creates a feedback loop between the front-end and the back-end by using the optimized graph to generate new front-end loop closure associations.

IV. Experimental Evaluation

The main focus of this work is an effective 3D LiDAR SLAM approach that improves global trajectory estimation and local map quality through information estimation-aware odometry, appearance-based front-end loop closure, and geometry-based retroactive loop closure.

We evaluate our system on several large-scale benchmark datasets and compare it against state-of-the-art LiDAR SLAM systems. The experiments also support our key claims, that our system (i) provides global trajectory accuracy that is on par with or better than state-of-the-art SLAM systems; (ii) improves the local map quality in revisited areas, even in scenes with perceptual aliasing.

A. Experimental Setup

We evaluate our LiDAR SLAM approach on the MulRan [18], HeLiPR [17], and Apollo [16] datasets. The HeLiPR dataset is particularly challenging, with long sequences in urban scenarios, driving through highways, narrow alleys, and perceptually aliased environments. It also provides data from multiple LiDAR sensors, allowing us to evaluate the robustness of SLAM algorithms to different scanning patterns and fields of view. Additionally, we also report performance metrics on the Newer College dataset [29], which is captured with a handheld LiDAR device.

We compare our approach against several state-of-the-art LiDAR SLAM systems, including CT-ICP [8], MULLS [27], PIN-SLAM [28], and KISS-SLAM [13]. For all baselines, we use the authors' publicly available implementations, with the recommended dataset-specific configurations.

Unless stated otherwise, all experiments use the same parameter set across all datasets: $\nu_{\text{sub}} = 0.5$ m, $\nu_{\text{loc}} = 0.5$ m, $\tau_{\text{inliers}} = 5$, $\tau_{\text{H}} = 100$ m, and $\tau_{\text{f}} = 0.25$. We perform no dataset- or sensor-specific parameter tuning.

B. Pose Accuracy Evaluation

We first evaluate global trajectory accuracy using the Absolute trajectory error (ATE) and Relative pose error (RPE) following the KITTI metric [10]. Since the HeLiPR and MulRan datasets contain three recordings per environment, we report the average metric over the corresponding sequences.

We present the quantitative results in Tab. Ia, Tab. Ib and Tab. II. Overall, our method achieves global trajectory accuracy that is competitive with state-of-the-art baseline approaches across all datasets with the same set of system parameters. On the HeLiPR dataset Tab. Ia, we obtain the

TABLE I: Quantitative results for global trajectory accuracy. We report the ATE in meters and RPE in percentage as [m] / [%]. The best and second-best methods are in **bold** and underline, respectively. We provide ablation studies at the bottom of each table.

(a) HeLiPR dataset

Method	Bridge Aeva	Bridge Avia	Bridge Ouster	Roundabout Aeva	Roundabout Avia	Roundabout Ouster	Town Aeva	Town Avia	Town Ouster
MULLS	356.06 / 9.17	321.87 / 4.06	52.65 / 0.68	19.08 / 1.43	16.39 / 1.06	2.65 / 0.24	39.82 / 2.85	14.93 / 1.35	4.45 / 0.27
CT-ICP	- / -	41.47 / 0.36	- / -	10.04 / 0.72	3.36 / 0.25	1.81 / 0.16	65.29 / 1.37	63.72 / 2.02	- / -
PIN-SLAM	- / -	365.72 / 2.15	19.23 / **0.13**	- / -	7.02 / 0.51	1.47 / **0.11**	41.19 / 1.03	11.40 / 0.64	2.55 / **0.12**
KISS-SLAM	98.61 / 1.57	148.88 / 1.82	19.47 / 0.28	6.06 / 0.47	3.84 / **0.23**	**1.18** / 0.17	14.44 / 0.85	12.01 / 0.46	1.99 / 0.21
Ours	**16.83 / 0.47**	**20.17 / 0.25**	**14.18** / 0.22	**4.67 / 0.46**	**3.34 / 0.23**	1.21 / 0.18	**10.50 / 0.47**	**7.57 / 0.38**	**1.89** / 0.20
Ablation Study: no-info → without uncertainty estimation module; no-feedback → without the retroactive loop closing									
no-info	107.34 / 1.46	27.03 / 0.39	18.18 / 0.29	6.01 / 0.44	3.93 / 0.25	1.31 / 0.20	11.78 / 0.83	8.70 / 0.59	1.94 / 0.22
no-feedback	16.37 / 0.47	21.33 / 0.25	13.80 / 0.21	4.43 / 0.39	3.89 / 0.22	1.14 / 0.18	13.47 / 0.48	8.31 / 0.38	1.75 / 0.20

(b) Apollo and MulRan dataset

Method	Apollo BTS	Apollo CP	Apollo H237	Apollo MAVE	Apollo SB	MulRan KAIST	MulRan Riverside	MulRan Sejong
MULLS	104.14 / 0.27	47.03 / 0.90	354.21 / 0.39	182.59 / **0.16**	- / -	34.21 / 0.60	66.60 / 2.27	1082.55 / 2.53
CT-ICP	10.81 / 0.27	0.97 / 0.30	258.87 / 0.62	61.39 / 0.33	667.28 / 1.60	2.76 / 0.50	8.37 / 0.58	- / -
PIN-SLAM	6.10 / 0.24	**0.64 / 0.16**	97.11 / 1.66	7.06 / 0.20	2.23 / 0.25	**2.42 / 0.32**	7.55 / 0.46	782.96 / 7.68
KISS-SLAM	**3.74** / 0.23	0.80 / 0.22	30.16 / 0.40	9.08 / 0.26	2.51 / 0.25	2.96 / 0.33	8.34 / 0.41	178.35 / **0.25**
Ours	4.22 / **0.22**	0.81 / 0.19	**26.50 / 0.29**	**6.17** / 0.19	**1.61 / 0.17**	2.67 / **0.32**	**7.22 / 0.35**	**169.14** / 0.28
Ablation Study: no-info → without uncertainty estimation module; no-feedback → without the retroactive loop closing								
no-info	9.45 / 0.23	1.02 / 0.22	88.25 / 0.32	19.38 / 0.19	22.13 / 0.45	2.79 / 0.31	10.04 / 0.63	173.44 / 0.28
no-feedback	3.83 / 0.22	- / -	27.31 / 0.30	5.33 / 0.20	1.94 / 0.19	2.64 / 0.32	7.25 / 0.35	169.16 / 0.28

TABLE II: Quantitative results on the Newer College dataset. We report the ATE in meters and RPE in percentage as [m] / [%]. The best and second-best methods are in **bold** and underline, respectively. We provide ablation studies at the bottom of the table.

Method	01-short-sequence	02-long-sequence
MULLS	0.47 / **0.17**	8.47 / 4633.62
CT-ICP	0.63 / 0.42	25.06 / 3.01
PIN-SLAM	0.42 / 0.33	**0.31 / 0.22**
KISS-SLAM	**0.30** / 0.26	1.58 / 2.14
Ours	0.34 / 0.38	0.42 / 0.46
Ablation Study		
no-info	0.36 / 0.43	0.42 / 0.48
no-feedback	0.34 / 0.38	0.42 / 0.46

lowest ATE score on the majority of the sequences and also consistently maintain competitive RPE values.

We observe the largest improvements on the challenging Bridge sequences across all three LiDARs (Aeva, Avia, and Ouster), which contain long traversals through structurally repetitive environments that are prone to perceptual aliasing. These results suggest that the conservative loop-closure strategy together with geometry-aware constraint weighting increases robustness in such scenarios.

On the Apollo and MulRan datasets Tab. Ib, our approach obtains the best ATE on five of the eight sequences. On the Newer College dataset Tab. II, our method's performance matches the strongest baselines despite the significantly different motion profile of the handheld sensing platform. These results indicate that our approach generalizes well across different LiDAR sensors and operating environments.

The ablation study further illustrates the roles of the individual components. Removing geometry-aware information weighting consistently degrades global trajectory accuracy across most sequences, demonstrating the benefit of information-weighted pose graph optimization. Disabling retroactive loop closure has only a minor effect on ATE, indicating that its primary contribution is not improved localization accuracy but improved local map consistency, which we evaluate in the following experiment. An exception is the narrow field-of-view Avia sensor, where retroactive loop closures yield more accurate trajectories. The reduced field-of-view makes appearance-based place recognition more difficult, allowing the geometry-guided recovery stage to successfully identify additional valid loop closures.

Overall, these quantitative results support our first claim that our SLAM system provides on par with or better than state-of-the-art global trajectory accuracy.

C. Revisit Map Consistency Evaluation

Trajectory-based metrics alone do not fully capture the quality of the reconstructed map. We therefore complement the conventional evaluation with a quantitative assessment of local map consistency at revisit locations. Our protocol measures the geometric consistency between independently reconstructed maps of the same physical region acquired during different traversals. The results suggest that our approach enhances the local map accuracy at revisit locations.

We first identify revisit locations from the reference trajectory by projecting it onto a coarse grid in the xy-plane. Within each occupied cell, we extract contiguous trajectory segments. We retain cells containing multiple such segments as revisit regions, and store the associated trajectory segments with the corresponding scan indices.

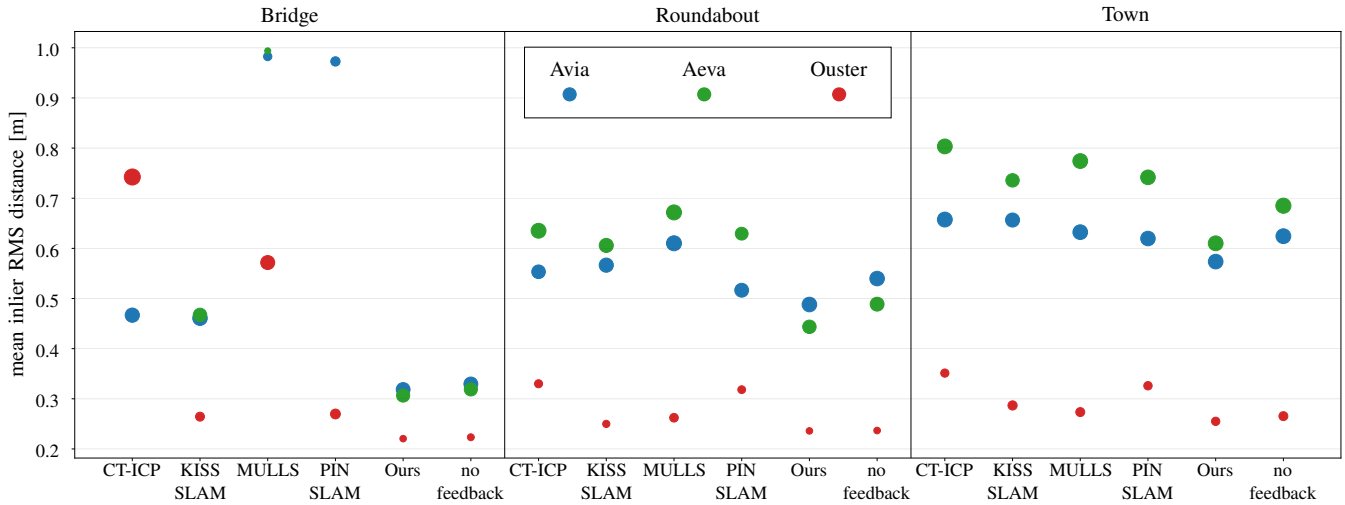

Fig. 3: Map quality evaluation on the HeLiPR dataset averaged over the three recording instances in each scenario across all three LiDARs. We show a scatter plot of the average inlier RMS scores for the nearest-neighbor distances between local maps at revisit locations. The size of the markers represent the magnitude of the standard deviation. The inlier threshold is set to 1.0 m.

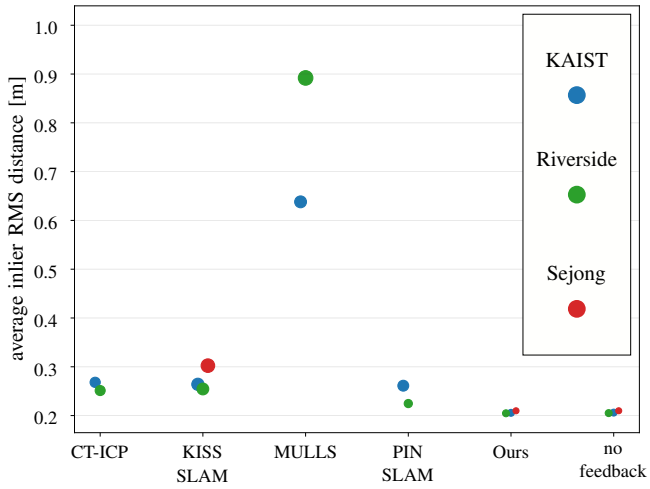

Fig. 4: Map quality evaluation on the MulRan dataset averaged over the three recording instances in each scenario. We show a scatter plot of the average inlier RMS scores for the nearest-neighbor distances between local maps at revisit locations. The size of the markers represent the magnitude of the standard deviation. The inlier threshold is set to 1.0 m.

To evaluate each SLAM baseline, we use the stored scan indices to construct a local map for every trajectory segment within a revisit cell based on the estimated global poses, yielding multiple local maps representing independent observations of the same region. We compare the local maps directly without any additional registration. We compute the root mean square (RMS) distances between nearest-neighbor point correspondences between every pair of local maps within each revisit cell. We aggregate the RMS values across all revisit cells and report their mean and standard deviation as a measure of map consistency.

We first discuss the map quality evaluation results for the HeLiPR dataset. Fig. 3 shows that our method occupies the lower regions of the plots, with lower mean RMS error and therefore better geometric agreement between independent revisits. This trend is consistent across all environments and all three different LiDAR sensors. The ablation study in the last column shows the results when we disable the retroactive loop closure feedback. We see a slight increase in RMS distances across all sequences and sensors, confirming its contribution to improving local map consistency.

The Bridge sequences remain the most challenging scenarios because they contain long traversals through perceptually aliased environments. Nevertheless, our method achieves substantially lower RMS errors than all baselines. Combined with the trajectory evaluation, these results indicate that our approach remains robust in perceptually aliased environments, and simultaneously provides more consistent local reconstructions. The evaluation also reveals a clear dependence on LiDAR sensor characteristics. Across all methods, the spinning Ouster sensor (red) consistently achieves lower RMS errors and lower variance than the Aeva (green) and Avia (blue) sensors, suggesting that a wider field-of-view and denser scan pattern provides better reconstruction consistency.

We observe a similar performance trend on the MulRan dataset (Fig. 4), where our method again achieves the lowest mean RMS error across all recording sessions. Since these recordings are not affected significantly by perceptual aliasing, our front-end loop-closure module does not miss many loops for the latter retroactive loop-closure stage to recover. Therefore, the ablation study (last column) does not show a perceivable difference in map quality. This observation is consistent with the ATE scores on these sequences as reported in Tab. Ib.

Furthermore, we are the only baseline alongside KISS-SLAM to succeed on the Sejong sequences which are recorded on a highway with a little overlap in the trajectory. We obtain a better map quality in comparison to KISS-SLAM to complement the best ATE score. This demonstrates that the improvements generalize across different datasets without any parameter tuning.

Interestingly, several methods exhibit comparable ATE

values but have noticeably different map consistency scores. This highlights that trajectory accuracy alone is insufficient to characterize SLAM performance and motivates the proposed revisit-based evaluation.

In summary, our proposed SLAM approach achieves competitive global trajectory accuracy while improving local map consistency at revisit locations. The results demonstrate that information-aware pose graph optimization along with hierarchical and retroactive loop closure improves both trajectory estimation and map quality.

V. Conclusion

In this work, we presented a novel and effective 3D LiDAR SLAM approach that improves global trajectory estimation and local map quality through geometry-aware odometry information estimation, hierarchical multi-resolution loop closure, and geometry-based retroactive loop closure detection. Experimental results obtained on several challenging benchmark datasets demonstrate competitive trajectory accuracy together with consistently better map quality at revisit locations compared with state-of-the-art SLAM systems.

References

[1] 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.

[2] J.L. Blanco-Claraco. A flexible framework for accurate lidar odometry, map manipulation, and localization. *Intl. Journal of Robotics Research (IJRR)*, 0(0):02783649251316881, 2025.

[3] S. Bonnabel, M. Barczyk, and F. Goulette. On the Covariance of ICP-Based Scan-Matching Techniques. In *Proc. of the IEEE American Control Conf. (ACC)*, pages 5498–5503, 2016.

[4] M. Bosse, P. Newman, J. Leonard, and S. Teller. Simultaneous Localization and Map Building in Large-Scale Cyclic Environments Using the Atlas Framework. *Intl. Journal of Robotics Research (IJRR)*, 23(12):1113–1139, 2004.

[5] M. Brossard, S. Bonnabel, and A. Barrau. A New Approach to 3D ICP Covariance Estimation. *IEEE Robotics and Automation Letters (RA-L)*, 5(2):744–751, 2020.

[6] A. Censi. An Accurate Closed-Form Estimate of ICP's Covariance. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 3167–3172, 2007.

[7] X. Chen, A. Milioto, E. Palazzolo, P. Giguère, J. Behley, and C. Stachniss. SuMa++: Efficient LiDAR-based Semantic SLAM. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2019.

[8] P. Dellenbach, J. Deschaud, B. Jacquet, and F. Goulette. CT-ICP Real-Time Elastic LiDAR Odometry with Loop Closure. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2022.

[9] J. Deng, Q. Wu, X. Chen, S. Xia, Z. Sun, G. Liu, W. Yu, and L. Pei. NeRF-LOAM: Neural Implicit Representation for Large-Scale Incremental LiDAR Odometry and Mapping. In *Proc. of the IEEE/CVF Intl. Conf. on Computer Vision (ICCV)*, 2023.

[10] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun. Vision meets Robotics: The KITTI Dataset. *Intl. Journal of Robotics Research (IJRR)*, 32(11):1231–1237, 2013.

[11] 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.

[12] G. Grisetti, R. Kümmerle, C. Stachniss, U. Frese, and C. Hertzberg. Hierarchical Optimization on Manifolds for Online 2D and 3D Mapping. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2010.

[13] T. Guadagnino, B. Mersch, S. Gupta, I. Vizzo, G. Grisetti, and C. Stachniss. KISS-SLAM: A Simple, Robust, and Accurate 3D LiDAR SLAM System With Enhanced Generalization Capabilities. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 5363–5370, 2025.

[14] S. Gupta, T. Guadagnino, B. Mersch, N. Trekel, M. Malladi, and C. Stachniss. Efficiently Closing Loops in LiDAR-Based SLAM Using Point Cloud Density Maps. *Intl. Journal of Robotics Research (IJRR)*, 2026.

[15] 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.

[16] X. Huang, X. Cheng, Q. Geng, B. Cao, D. Zhou, P. Wang, Y. Lin, and R. Yang. The ApolloScape Dataset for Autonomous Driving. In *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition Workshops*, 2018.

[17] 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.

[18] G. Kim, Y. Park, Y. Cho, J. Jeong, and A. Kim. Mulran: Multimodal range dataset for urban place recognition. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2020.

[19] G. Kim and A. Kim. Scan Context: Egocentric Spatial Descriptor for Place Recognition within 3D Point Cloud Map. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2018.

[20] H. Kim, J. Choi, T. Sim, G. Kim, and Y. Cho. Narrowing Your FOV With SOLiD: Spatially Organized and Lightweight Global Descriptor for FOV-Constrained LiDAR Place Recognition. *IEEE Robotics and Automation Letters (RA-L)*, 9(11):9645–9652, 2024.

[21] D. Landry, F. Pomerleau, and P. Giguère. CELLO-3D: Estimating the Covariance of ICP in the Real World. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 8190–8196, 2019.

[22] F. Lu and E. Milios. Globally Consistent Range Scan Alignment for Environment Mapping. *Autonomous Robots*, 4:333–349, 1997.

[23] S. Ma, H. Zhang, and M. Toussaint. SVN-ICP: Uncertainty Estimation of ICP-Based LiDAR Odometry Using Stein Variational Newton. *IEEE Robotics and Automation Letters (RA-L)*, 10(11):12181–12188, 2025.

[24] M. Magnusson, H. Andreasson, A. Nuechter, Achim, and J. Lilienthal. Appearance-Based Loop Detection from 3D Laser Data Using the Normal Distributions Transform. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2009.

[25] J. Nieto, T. Bailey, and E. Nebot. Recursive Scan-Matching SLAM. *Journal on Robotics and Autonomous Systems (RAS)*, 55(1):39–49, 2007.

[26] E. Olson and P. Agarwal. Inference on Networks of Mixtures for Robust Robot Mapping. *Intl. Journal of Robotics Research (IJRR)*, 32:826–840, 06 2013.

[27] Y. Pan, P. Xiao, Y. He, Z. Shao, and Z. Li. MULLS: Versatile LiDAR SLAM Via Multi-Metric Linear Least Square. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2021.

[28] Y. Pan, X. Zhong, L. Wiesmann, T. Posewsky, J. Behley, and C. Stachniss. PIN-SLAM: LiDAR SLAM Using a Point-Based Implicit Neural Representation for Achieving Global Map Consistency. *IEEE Trans. on Robotics (TRO)*, 40:4045–4064, 2024.

[29] M. Ramezani, Y. Wang, M. Camurri, D. Wisth, M. Mattamala, and M. Fallon. The newer college dataset: Handheld lidar, inertial and vision with ground truth. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2020.

[30] N. Sünderhauf and P. Protzel. Towards a Robust Back-End for Pose Graph SLAM. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2012.

[31] 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.

[32] N. Wang, H. Lu, Z. Zheng, Y.H. Liu, and X. Chen. Leveraging Semantic Graphs for Efficient and Robust LiDAR SLAM. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2025.

[33] 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.

[34] J. Zhang and S. Singh. LOAM: Lidar Odometry and Mapping in Real-time. In *Proc. of Robotics: Science and Systems (RSS)*, 2014.