Leveraging LiDAR Reflectance Images for Sparse Odometry

Simone Marmaglio[1,2], Mattia Savardi[3], Nam Nguyen Hoang[1,2], Matteo Sgrenzaroli[2], Giorgio Vassena[1,2], Alberto Signoroni[3]

[1]Università degli Studi di Brescia, Department of Information Engineering, Via Branze 38, 25123 Brescia, Italy
(simone.marmaglio; nam.nguyenhoang; giorgio.vassena)@unibs.it
[2]GEXCEL srl, Via Branze 45, 25123 Brescia, Italy.
(matteo.sgrenzaroli; giorgio.vassena)@gexcel.it
[3]Università degli Studi di Brescia, Department of Medical and Surgical Specialities,
Radiological Sciences and Public Health, Viale Europa 11, 25123 Brescia, Italy
(mattia.savardi; alberto.signoroni)@unibs.it

Technical Commission II

Keywords: Reflectance, Odometry, SLAM, LiDAR, Feature Detection, RANSAC, IMU Integration

Abstract

Accurate odometry is essential for Simultaneous Localization and Mapping (SLAM), yet traditional methods relying on geometric features struggle in feature-poor environments such as tunnels. Our work addresses this issue by leveraging LiDAR reflectance data to develop a robust odometry technique. This approach generates reflectance images, extracts 3D keypoints, and employs an IMU-based outlier detection process eventually refined by RANSAC algorithm. Unlike geometry-based methods, our solution can operate in highly symmetrical environments, producing consistent trajectories even where geometric methods fail. Our evaluations highlight the capacity of our approach to maintain trajectory coherence in GNSS-denied and geometrically degenerate scenarios. This robustness underscores its potential for reliable navigation and mapping when traditional SLAM solutions are inadequate.

1. Introduction

An odometer is an essential element in Simultaneous Localization and Mapping (SLAM) algorithms that not only tracks a system's position within an environment, but also offers motion estimates for building and updating the map of that environment. Accurate odometry is therefore vital for effective SLAM systems, as errors in motion estimation can lead to significant inaccuracies in estimated system locations and, consequently, in generated maps, potentially causing the entire SLAM algorithm to fail. Traditional SLAM techniques with Light Detection and Ranging (LiDAR) often rely on odometers that use geometry-based point cloud registration, such as Iterative Closest Point (ICP) (Besl and McKay, 1992), which leverage local geometric properties to detect lines and planes for precise pose computation (Zhang and Singh, 2014).

Although effective in structured environments, these methods face significant challenges in geometrically degenerate spaces. In such environments, the absence of sufficient physical structures reduces the optimization constraints provided by geometric features, leading to increased odometry drift and potential failure of the SLAM system (Tuna et al., 2024). For example, the absence of frontal and posterior walls in tunnels hampers the accurate estimation of the point cloud pose (Zhang et al., 2016). To tackle this issue, a common approach consists in augmenting the point-to-plane ICP method with data from an Inertial Measurement Unit (IMU) within a Kalman filter framework. This approach, as used in LIO-SAM (Shan et al., 2020) and FAST-LIO (Xu and Zhang, 2021), is designed for motion compensation and provides a reliable initial guess for the ICP point-to-plane optimization algorithm (Chen and Medioni, 1991). However, while this approach can offer improved robustness, it may suffer from increased error propagation if the ICP algorithm fails to converge properly. Incorrect point matches due to a lack of distinctive features can lead to erroneous pose estimates. These errors, when incorporated into a Kalman filter, can result in incorrect state updates, degrading the system's accuracy over time and potentially leading to the failure of the SLAM algorithm.

A promising approach to overcoming the limitations of geometric methods is to leverage complementary information, such as the reflectance or intensity of each point in the cloud. Intensity values depend on reflectance, distance, and the angle of incidence. This additional information, provided by modern 3D LiDARs alongside range measurements, can enhance pose estimation, as demonstrated in various studies (Wang et al., 2021, Park et al., 2020). For rotating multilayer LiDARs, this signal can be projected into a dense reflectance image, effectively enabling the LiDAR to function as an active camera without the need for external illumination. This offers significant advantages over conventional cameras, as each pixel in the reflectance image corresponds directly to a 3D point in the point cloud, eliminating the need for triangulation or scale estimation. The resulting reflectance image captures texture information about the environment, which is particularly useful in geometrically degenerate scenarios. By mapping LiDAR returns onto a reflectance image and extracting image-based features, we can determine their 3D locations in successive LiDAR scans, enabling the computation of the rigid transformation without requiring an initial pose estimate.

However, reflectance images are not easy to handle due to the strong anisotropy typical of LiDAR configurations, where a limited number of rows results in images with a high imbalance between horizontal and vertical resolution. This issue is significantly less pronounced in modern 128-line configurations, but the challenge remains in finding methods and solutions to effectively leverage reflectance to support odometry and SLAM in the more commonly used configurations, such as those with 16 or 32 lines.

This study addresses the limitations of conventional geometry-based SLAM methods by proposing an odometry technique that relies solely on LiDAR reflectance. The key contributions of this work are as follows:

- Reflectance-based keypoint extraction and matching. We present a novel method that uses only the reflectance component of LiDAR scans to extract and match keypoints. This approach serves as an alternative to geometry-based methods and is particularly effective in environments with sparse geometric features.

- Robust 3D outlier detection. We implement a two-stage outlier detection process that combines IMU data for initial filtering with the RANSAC (Random Sample Consensus) (Fischler and Bolles, 1987) algorithm. This significantly reduces incorrect correspondences, enhancing the robustness of pose estimation.
- Feature tracking for drift reduction. To mitigate drift during consecutive scan alignments, we introduce a feature-tracking mechanism that aligns new LiDAR scans with a dynamically updated tracking map. This ensures mid-term accuracy, even when using low-resolution LiDAR sensors.
- Practical validation. We validate our method using a commercial LiDAR acquisition system and demonstrate its reliability with sensors having as few as 32 lines.

2. Related Work

2.1 Geometry-Based LiDAR Odometry

LiDAR-based Simultaneous Localization and Mapping (SLAM) systems have been extensively studied, leading to the development of various methods for odometry and mapping. Among these, the LOAM framework stands out for its real-time pose estimation and map-building capabilities using edge and planar features from LiDAR point clouds (Zhang and Singh, 2014). While effective in structured environments, LOAM and its derivatives, such as LeGO-LOAM (Shan and Englot, 2018), face challenges in featureless environments due to the lack of sufficient geometric features for reliable pose estimation. To overcome these limitations, researchers have proposed improvements like robust error metrics and advanced data association techniques. For example, KISS-ICP (Vizzo et al., 2023) enhances performance in unstructured environments by directly registering downsampled voxel point clouds using a robust point-to-point ICP with an adaptive data association threshold.

Furthermore, algorithms like LIO-SAM (Shan et al., 2020) leverage IMU data to improve robustness in environments with limited geometric features, while FAST-LIO (Xu and Zhang, 2021) employs a Kalman filter for efficient scan-to-map alignment using raw point clouds instead of extracted features. The integration of inertial measurements in LiDAR-Inertial Odometry (LIO) methods significantly boosts performance by enabling accurate point cloud undistortion due to ego-motion and providing a reliable initial guess for registration. A notable advancement in this field is X-ICP (Tuna et al., 2024), which introduces a localizability-aware constrained optimization approach. By analyzing alignment strength along principal directions, it mitigates degeneracy issues in weakly constrained environments.

2.2 Reflectance and Intensity Integration

Integrating intensity or reflectance information from LiDAR sensors into odometry and mapping has proven effective in overcoming the limitations of purely geometric methods. This integration improves SLAM performance, particularly in geometrically deficient environments, by providing texture-like details that aid in feature identification.

Methods such as Intensity-SLAM (Wang et al., 2021) and I-LOAM (Park et al., 2020) incorporate intensity information into a weighted ICP framework, enhancing LiDAR odometry performance in large-scale environments. More recent approaches, like RI-LIO (Zhang et al., 2023) and COIN-LIO (Pfreundschuh et al., 2024), improve robustness by integrating photometric error minimization into an Iterative Kalman Filter. Guadagnino et al. (Guadagnino et al., 2022) leverage SuperPoint, a CNN-based keypoint detector and descriptor, for real-time LiDAR odometry. This self-supervised neural network extracts and matches salient points from intensity images, enabling robust odometry without relying on hand-crafted features. Similarly, Du and Beltrame (Du and Beltrame, 2023) propose a comprehensive SLAM pipeline that combines intensity information with geometric features to enhance localization and mapping in environments prone to geometric degeneracy. Both methods focus on intensity-based feature extraction and outlier rejection. Guadagnino et al. utilize RANSAC and robust optimization techniques, while Du and Beltrame validate correspondences through ORB feature extraction and matching scores. However, neither approach leverages IMU data for improved feature matching robustness.

Our method distinguishes itself by incorporating IMU data to filter out unaligned correspondences before applying RANSAC. This preprocessing step reduces the number of RANSAC iterations needed, enhancing the system's robustness and efficiency. The IMU-based filtering ensures that only well-aligned correspondences proceed to the next stage, streamlining the overall process.

This approach is particularly important because we use a tracking map to maintain feature continuity by searching for previously detected features within a specified radius in the current image. While this technique effectively captures intermittently visible features, especially in low-resolution reflectivity images, it also generates a significant number of outliers. The abundance of outliers arises from the inclusion of features that may not consistently match across frames. To handle this, IMU filtering plays a crucial role in discarding the majority of these outliers, after which RANSAC is applied to further refine correspondences and ensure robust odometry.

Unlike most existing systems designed for high vertical resolution LiDARs (64 or 128-line sensors), our system is optimized for reliability even with lower vertical resolution sensors (e.g., 32-line LiDAR). By combining robust feature detection using SIFT with IMU-based filtering, we significantly enhance spatial localization accuracy and overall odometry robustness.

3. Method

Our method employs a feature-based approach that utilizes LiDAR sensor reflectivity data to generate reflectance images, from which features are extracted. Each LiDAR scan is aligned with a tracking point cloud to enable continuous feature tracking and minimize cumulative drift. To enhance odometry robustness, features from the newly acquired LiDAR scan are matched with those in a pre-existing map. However, this process introduces a significant number of outliers. To mitigate this issue, we apply a two-stage outlier detection strategy to efficiently filter incorrect correspondences, followed by robust point-to-plane optimization to refine the alignment.

The following subsections provide a detailed explanation of our method, including reflectance image generation, 3D keypoint detection from LiDAR scans, two-stage outlier detection, pose estimation, and feature tracking. Each component plays a critical role in improving the robustness and reliability of the odometry system.

3.1 Image Generation

To generate a reflectance image, we leverage the point cloud's grid structure provided by the LiDAR encoder. Unlike conventional projection methods, which can lead to the loss of points, our approach retains all data points. The resulting image dimensions correspond to the number of LiDAR beams for height and the number of points captured in a full rotational sweep for width.

Due to the uneven distribution of reflectivity values, we apply histogram equalization to enhance the image. Analysis reveals that reflectivity values above 127 constitute less than 1% of the data (Figure 1). To address this imbalance, we clip values exceeding 127 to 255 and stretch the 0–127 range to span the entire interval from 0 to 255. This adjustment significantly improves the visual quality of the reflectance image, enhancing contrast and detail representation (Figures 2).

Our equalization technique is compatible with various LiDAR systems, including Ouster, Hesai, and Velodyne.

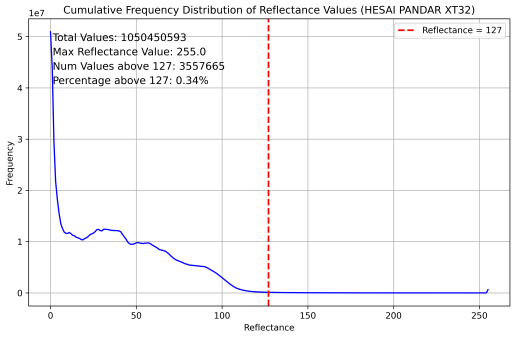

Figure 1. This histogram illustrates the frequency distribution of reflectance values over an entire acquisition comprising 17,000 scans across various types of environments.

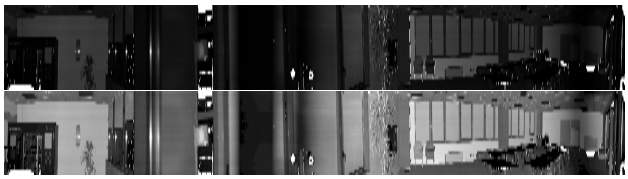

Figure 2. Reflectance images from the HESAI PandarXT 32 sensor. **Top** raw reflectance, **bottom** equalized reflectance.

3.2 3D keypoint detection

One advantage of representing a point cloud within a grid data structure is its natural interpretation as a four-channel image. The first three channels correspond to the 3D coordinates x, y, z, while the fourth channel represents the reflectance value. This format allows us to apply conventional feature detectors, such as SIFT (Lowe, 2004), within the image domain to extract features and determine their corresponding 3D positions. The result is a sparse 3D point cloud where each point has an associated 2D image location, 3D coordinates, and a feature descriptor.

A key challenge with this approach is accurately mapping the 2D features detected in the image to their corresponding 3D points in the point cloud. This difficulty arises from the limited resolution of the image and the inherent sparsity of the point cloud. Although each pixel in the image has a one-to-one correspondence with a 3D point in the cloud, SIFT outputs feature coordinates in floating-point format. As a result, a single feature may correspond to up to four nearby pixels, each linked to a different 3D point, creating ambiguity in the matching process.

To address this issue, we employ bilinear interpolation to compute a synthetic point. This synthetic point is derived by combining the information from the four candidate 3D points associated with the nearest 2D pixels. The weights assigned to each 3D point are based on the relative areas between the SIFT feature location and the corresponding pixel centers. This weighted interpolation provides an approximate 3D position that balances contributions from the surrounding candidate points, ensuring a more accurate representation of the detected feature in 3D space.

3.3 Outlier Detection and Pose Estimation

To compute a robust pose for each point cloud C_t, accurate data association with the tracking point cloud is crucial. To achieve this, we first filter correspondences using IMU data. Specifically, gyroscope readings are utilized to calculate the relative rotation between the previous point cloud C_{t-1} and the current one C_t. By integrating the angular velocity data, we align the current scan with the tracking point cloud, enabling the effective identification and removal of misaligned correspondences as outliers.

To integrate the gyroscope data, we incrementally compute the rotation matrix $\mathbf{R}_k$ using the angular velocity vector ω_k at each IMU timestamp k, as follows:

$$\mathbf{R}_k = \mathbf{R}_{k-1}\mathbf{Exp}(\omega_k \delta t) \tag{1}$$

where $\mathbf{R}_{k-1}$ represents the orientation at the previous IMU timestamp $k-1$, ω_k is the angular velocity at timestamp k, and δt is the time interval between consecutive IMU measurements. For each point cloud C_t, acquired over a time interval $[t_0, t_n]$, we compute the relative rotation $\Delta\mathbf{R}$ by interpolating between the initial and final orientations, $\mathbf{R}_{t_0}$ and $\mathbf{R}_{t_n}$, derived from IMU readings.

For a given time $t \in [k_1, k_2]$, where k_1 and k_2 represent two consecutive IMU timestamps that include t, the orientation $\mathbf{R}_t$ is interpolated between the known orientations $\mathbf{R}_{k_1}$ and $\mathbf{R}_{k_2}$ as follows:

$$\mathbf{R}_t = \mathbf{R}_{k_1} \cdot \mathbf{Exp}\left(\frac{t - t_{k_1}}{t_{k_2} - t_{k_1}} \mathbf{Log}\left(\mathbf{R}_{k_1}^{-1}\mathbf{R}_{k_2}\right)\right) \tag{2}$$

where:

- $\mathbf{Exp} : \mathbb{R}^3 \to SO(3)$ computes a rotation matrix from an axis-angle representation.
- $\mathbf{Log} : SO(3) \to \mathbb{R}^3$ extracts the axis-angle representation.

This interpolation allows us to compute the rotation at timestamps t_0 and t_n.

The relative rotation $\Delta\mathbf{R}$ is then obtained by inverting $\mathbf{R}_{t_0}$ and multiplying it by $\mathbf{R}_{t_n}$.

Next, we establish correspondences between C_t and the tracking point cloud using the SIFT descriptor to match keypoints. We evaluate the quality of each correspondence using cosine similarity. For each pair of keypoints, we compute the cosine of the angle between their corresponding ray vectors, given by:

$$\cos(\theta_{ij}) = \frac{\mathbf{p}_i \cdot \mathbf{q}_j}{\|\mathbf{p}_i\|\|\mathbf{q}_j\|} \tag{3}$$

where $\mathbf{p}_i$ and $\mathbf{q}_j$ are the vectors of the corresponding points. Cosine similarity values close to 1 indicate strong alignment, while values below a threshold suggest outliers. To determine this threshold, we analyze the angular error of the IMU using ground truth data from the from the acquisition system (see Sec.4). Given the IMU-estimated rotation $\mathbf{R}_{\text{IMU}}$ and the ground truth rotation $\mathbf{R}_{\text{GT}}$, the rotation error $\mathbf{R}_{\text{error}}$ is calculated as:

$$\mathbf{R}_{\text{error}} = \mathbf{R}_{\text{GT}} \cdot \mathbf{R}_{\text{IMU}}^{-1}. \tag{4}$$

The angular error θ_{error} is derived from the logarithmic map of the rotation error:

$$\omega_{\text{error}} = \mathbf{Log}(\mathbf{R}_{\text{error}}), \quad \theta_{\text{error}} = \|\omega_{\text{error}}\|. \tag{5}$$

Using the error distribution, we set a threshold at $\theta_{thr} = \mu + 3\sigma$ (where μ is the mean and σ is the standard deviation). The corresponding cosine threshold $\gamma = \cos(\theta_{thr})$ defines the similarity criterion to filter outliers.

To further refine the correspondences, we employ the RANSAC (Fischler and Bolles, 1987) algorithm for robust alignment. In each iteration, we randomly sample three correspondences to compute the pose transformation for C_t using the Umeyama algorithm. The transformation with the highest inlier count, defined by a proximity threshold τ, is selected.

Once we determine the inlier correspondences, we refine the relative pose by minimizing the point-to-plane error using non-linear ICP with a robust cost function. We employ the Geman-McClure robust function ρ to reduce the influence of outliers. The objective function is given by:

$$\min_{\mathbf{R},\mathbf{t}} \sum_{i=1}^{N} \rho\left(||\mathbf{n}_j^\top \left(\mathbf{R}\mathbf{p}_i + \mathbf{t} - \mathbf{q}_j\right)||^2\right), \tag{6}$$

where N is the number of inlier correspondences between the current point cloud and the tracking point cloud, $\mathbf{R}$ and $\mathbf{t}$ are the rotation matrix and translation vector that define the relative pose transformation we seek to refine, $\mathbf{p}_i$ represents a point in the current point cloud C_t, $\mathbf{q}_j$ is the corresponding point in the tracking point cloud, $\mathbf{n}_j$ is the normal vector at the point $\mathbf{q}_i$ in the tracking point cloud, $\rho(\cdot)$ is the Geman-McClure robust function, which helps to reduce the influence of outliers by diminishing the contribution of large residuals.

3.4 Feature Tracking

The low resolution of the point cloud, combined with the 2D floating-point coordinates provided by SIFT, poses a challenge in accurately determining the corresponding 3D positions of features. To address this issue, we use interpolation to estimate feature locations. However, these estimates may still have inaccuracies, especially for features detected at greater distances from the sensor. To mitigate this, we adopt a temporal tracking approach, which facilitates the refinement and convergence of their spatial positions through successive observations. After computing the pose of the point cloud, we use all the detected features for updating the tracking cloud. For each correspondence, we check whether the detected features in the current point cloud are close to their tracked counterparts. If they fall within a defined threshold, the features detected in C_i are used to update the spatial positions of the tracked features, improving accuracy and consistency over time. Any features that do not match closely with existing tracked points are added to the tracking cloud as new features, representing potentially new detections that have not been previously observed.

Additionally, if a feature goes undetected for k consecutive clouds, it is removed from the tracking cloud. This approach accommodates cases where features may go undetected in a few scans due to low resolution or occlusions. By allowing brief detection gaps, we reduce the likelihood of prematurely discarding relevant features. This adaptive tracking process ensures that the tracking point cloud remains robust and dynamically updates, retaining newly detected features while preserving a consistent set of tracked points essential for reliable pose estimation and mapping.

4. Data Collection

4.1 Gexcel Heron System

The Gexcel HERON system, developed in collaboration with the Joint Research Center of the European Commission, is a SLAM-based mobile mapping platform designed for high-precision 3D scanning in a variety of environments, both indoor and outdoor. The platform can be equipped with one or two HESAI Pandar XT32 3D LiDAR sensors, a commercial-grade IMU, and a camera (see Fig. 3).

Figure 3. HERON mobile mapping system

The HERON system's odometry leverages IMU data, capturing both angular velocities and linear accelerations, which are critical for motion compensation during point-cloud acquisition. By integrating IMU readings with geometric ICP through an extended Kalman filter, the system can effectively align point clouds, even during sensor movement.

Equipped with low vertical resolution 32-line LiDAR sensors, the HERON system is particularly well-suited for testing the robustness of the reflectance-based odometry proposed in this study. Data is collected in environments with pre-existing high-precision, millimeter-level maps generated by a static terrestrial laser scanner. These static maps allow the HERON system to operate in tracking mode, aligning each LiDAR scan precisely to produce an accurate trajectory. This setup provides an ideal benchmark for evaluating the performance of the reflectance-based odometry across complex and varied environments.

4.2 Data Acquisition

Two datasets were collected for the purpose of conducting real-world experiments: Underground Parking and Tunnel. The Underground Parking dataset was captured using two Hesai Pandar XT32 mid-range LiDAR sensors and an Xsens IMU. In this setup, a secondary LiDAR was mounted obliquely, allowing the HERON system to effectively capture both ground and ceiling surfaces. This configuration enabled the evaluation of the odometry system using only the horizontal LiDAR (Lite) and the combination of both sensors (Twin).

The Tunnel dataset was designed to test odometry performance in a challenging, feature-poor environment. It involved traversing an 800-meter-long tunnel characterized by repetitive structural patterns and limited geometric features, posing a rigorous test for geometry-based odometry methods. To establish a high-accuracy reference map for this dataset, the tunnel was surveyed using a FARO Focus static laser scanner (TLS). Multiple overlapping TLS scans were captured along the tunnel and aligned using ICP to create a cohesive reference model, as shown in Fig. 4. Additionally, TLS scans were conducted at both the entrance and exit of the tunnel and georeferenced to align with an absolute reference frame. This setup ensured that the HERON system operated in tracking mode with precise alignment to the static map, providing a reliable ground-truth trajectory for performance evaluation.

5. Experiments

In this section, we present the experimental evaluation of our odometry algorithm across various environments and datasets, using distinct LiDAR sensors and acquisition methods. The

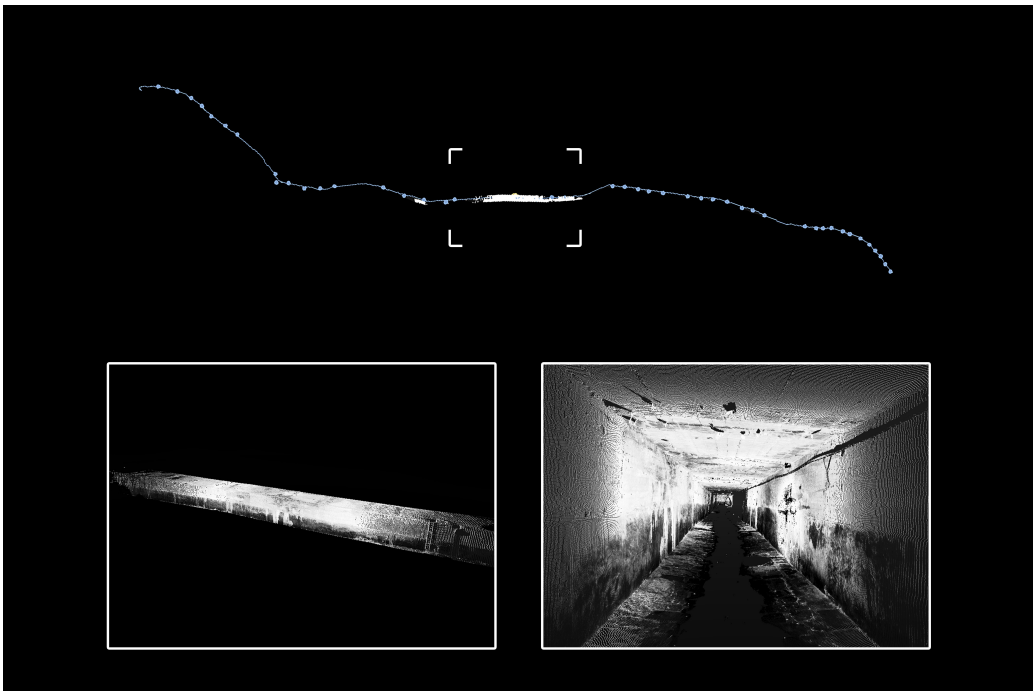

Figure 4. The trajectory acquired using the HERON odometry system aligned with TLS static scans (blue points). The highlighted segment corresponds to a region where HERON's geometric odometry fails to compute the pose accurately due to the regular tunnel geometry, which lacks distinctive features.

primary data source is the Gexcel HERON mobile mapping system, supplemented by tests on two public datasets: the Newer College Dataset (Zhang et al., 2021) and the ENWIDE Dataset (Pfreundschuh et al., 2024), both captured with the Ouster-128 sensor.

5.1 Performance on Gexcel HERON Dataset

In the Underground Parking dataset, which is tested with both Lite and Twin configurations, significant improvements were observed when using both LiDARs (Twin configuration) compared to a single horizontal LiDAR (Lite configuration). The oblique LiDAR provided additional geometric constraints by capturing both ground and ceiling features, effectively reducing vertical drift. This improvement is evident in the Absolute Trajectory Error (ATE) metrics: translational ATE decreased from 1.926 (Lite) to 0.697 (Twin), while rotational ATE dropped from 0.158 to 0.114 rad. Similarly, the mean translational and rotational errors improved from 2.737% to 1.531% and from 0.085 to 0.064 degrees per meter, respectively, underscoring the impact of enhanced geometric constraints on pose estimation accuracy (see Table 1).

In the Tunnel dataset, the reflectance-based odometer exhibited higher translation errors due to a sparse feature tracking map, which failed to provide sufficient constraints for accurate positioning. Small incremental errors accumulated over time, challenging positional accuracy in this feature-poor environment. Notably, in the water-drainage tunnel, the presence of steps introduced abrupt changes in the LiDAR's field of view (FoV). These rapid FoV changes hindered the odometer's ability to establish robust correspondences between tracked features and keypoints in successive LiDAR scans, leading to increased pose estimation errors.

This analysis suggests that a hybrid strategy — combining a dense local map and geometric features to maintain alignment with the pre-built map, while leveraging reflectance-based features to constrain the ICP optimization - could enhance robustness. Such a strategy would address the limitations of relying solely on geometric or reflectance-based methods.

Despite these challenges, the reflectance-based odometer successfully completed the Tunnel dataset trajectory using a 32-line sensor without critical failures. This outcome demonstrates the system's resilience in maintaining functionality under challenging environmental conditions and with low-resolution sensors. For a qualitative comparison of trajectories, see Fig. 5.

5.2 Performance on Newer College Dataset

The Newer College Dataset provides a diverse set of scenarios specifically designed to assess odometry robustness in complex environments. The dataset is captured using a handheld device equipped with synchronized LiDAR and IMU, mirroring the configuration of the HERON system. A high-precision ground-truth trajectory is obtained by aligning each LiDAR scan with a reference map generated using a Leica BLK360 laser scanner. This alignment, performed with the Iterative Closest Point (ICP) algorithm, ensures minimal drift and offers a reliable reference path for evaluating odometry performance.

The errors presented in Table 2 for the Newer College Dataset emphasize the critical role of reflectance features in odometry performance. These findings are consistent with those from the Gexcel HERON Dataset, showing that higher vertical resolution LiDARs, with more lines, not only capture a greater number of points but also produce more detailed reflectance features. These features act as robust constraints for point cloud alignment, significantly improving pose estimation accuracy (see Fig. 6).

In the Cloister scenario, the Absolute Trajectory Error (ATE) for translation is 1.056 m, with a rotational error of 0.032 rad. The relative translation error stands at 1.024%, while the average rotational error is 0.014°/m. These results underscore the benefit of abundant and consistent reflectance features, which enable better alignment and reduce drift. The structured environment of the Cloister provides reliable reflectance cues that the odometry system leverages for accurate trajectory estimation.

Performance on the Math Institute dataset surpasses that of the Cloister dataset, largely due to the open and well-structured environment. This setup ensures that reflectance-based features

Dataset	ATE Tra. (*m*)	ATE Rot. (*rad*)	Avg. Tra. (%)	Avg. Rot. (*deg/m*)
Tunnel	18.384	0.258	17.558	0.048
Underground Lite	1.926	0.158	2.737	0.085
Underground Twin	0.697	0.114	1.531	0.064

Table 1. We report the relative translational error and the relative rotational error using the KITTI (Geiger et al., 2013) metrics, calculated between the ground truth trajectory and the estimated trajectory. Additionally, we show the absolute trajectory error for translation in meters and for rotation in radians.

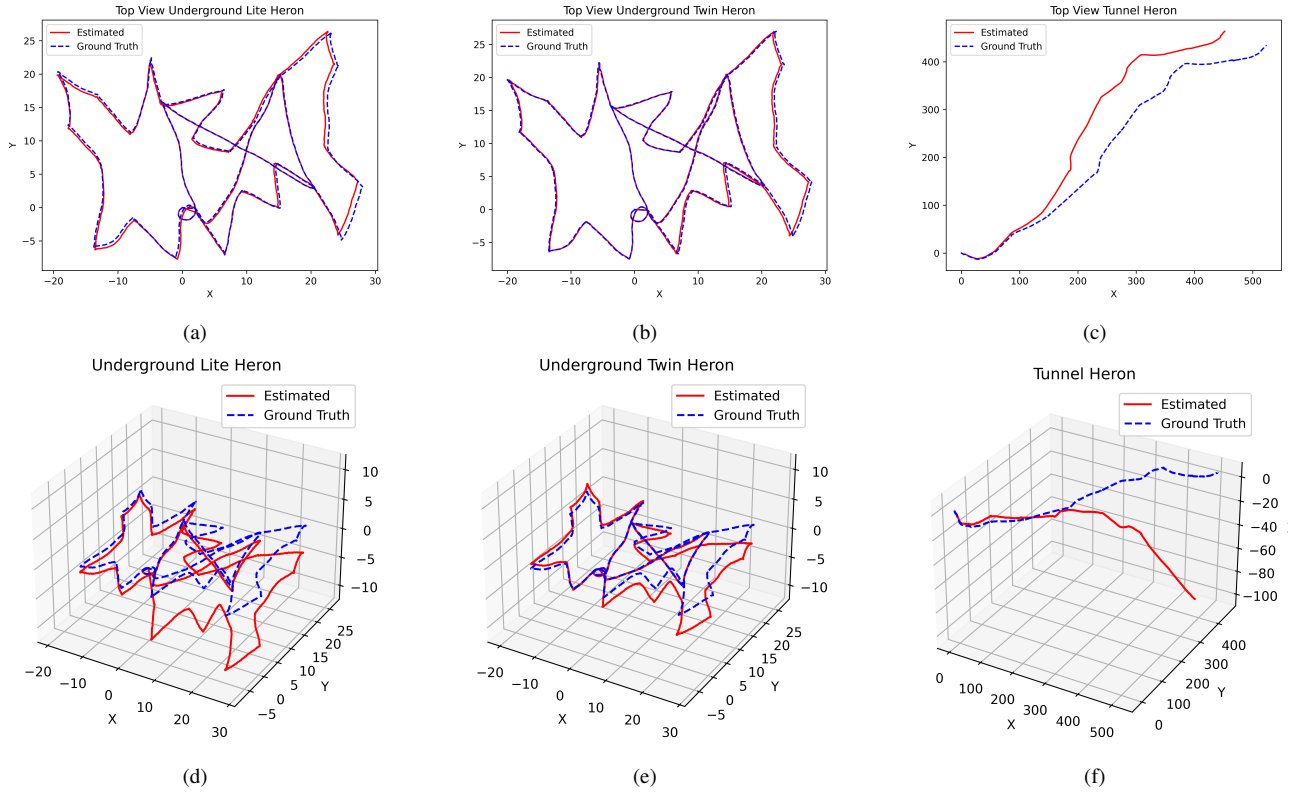

Figure 5. Comparison of ground truth and estimated trajectories for the Gexcel HERON dataset in different configurations: (a) Underground Lite Heron, (b) Underground Twin Heron, (c) Tunnel Heron, (d) Top view Underground Lite Heron, (e) Top view Underground Twin Heron, (f) Top view Tunnel Heron. The ground truth trajectory is shown in blue, and the estimated trajectory is depicted in dashed orange.

remain visible and consistently trackable over extended periods, enhancing odometry robustness and resulting in higher accuracy and stability.

Notably, these trajectories were derived from uncalibrated reflectance images due to the Ouster LiDAR operating with outdated firmware, which added complexity to pose estimation.

5.3 Performance on ENWIDE Dataset

The ENWIDE Dataset, specifically designed to tackle odometry challenges in GNSS-denied environments, features diverse sequences recorded in both indoor and outdoor settings, with a focus on geometrically degenerate environments. Data is collected using a handheld Ouster OS0 128 (Rev D) LiDAR, operating at 10 Hz, coupled with an integrated IMU capturing data at 100 Hz.

High-accuracy ground truth positioning is provided by a Leica MS60 Total Station. Positional data is obtained via a prism mounted on the LiDAR, offering an approximate positional accuracy of 3 cm. This setup is synchronized with the sensor data to establish a reliable reference trajectory.

Despite the challenges posed by sparse prior maps and complex environments, our method consistently estimates a reliable trajectory. This consistency is crucial in GNSS-denied environments, where relative positioning accuracy suffices for applications such as navigation and mapping. However, the ground truth in the ENWIDE dataset provides only positional data, limiting a comprehensive assessment of rotational or full pose accuracy. Consequently, the analysis focuses on the positional consistency of the estimated trajectory.

Figure 7 compares the ground truth trajectory with the trajectory estimated by our reflectance-based odometry. Although some drift is evident, the method accurately captures the overall path structure and maintains consistency even in geometrically degenerate sections, such as the tunnel. Notably, in the tunnel's highly symmetrical environment, traditional odometry methods relying on geometric features fail to generate a trajectory. In contrast, our method not only produces a trajectory but ensures its consistency with the ground truth.

While the sparse prior map affects numerical performance, the ability of our method to estimate a consistent trajectory underscores the potential of reflectance-based odometry in GNSS-denied environments, especially when reliable georeferencing

Dataset	ATE Tra. (*m*)	ATE Rot. (*rad*)	Avg. Tra. (%)	Avg. Rot. (*deg/m*)
Cloister	1.056	0.032	1.024	0.014
Math-Easy	0.312	0.024	0.673	0.013

Table 2. We report the relative translational error and the relative rotational error using the KITTI metrics. Additionally, we show the absolute trajectory error for translation in m and for rotation in rad.

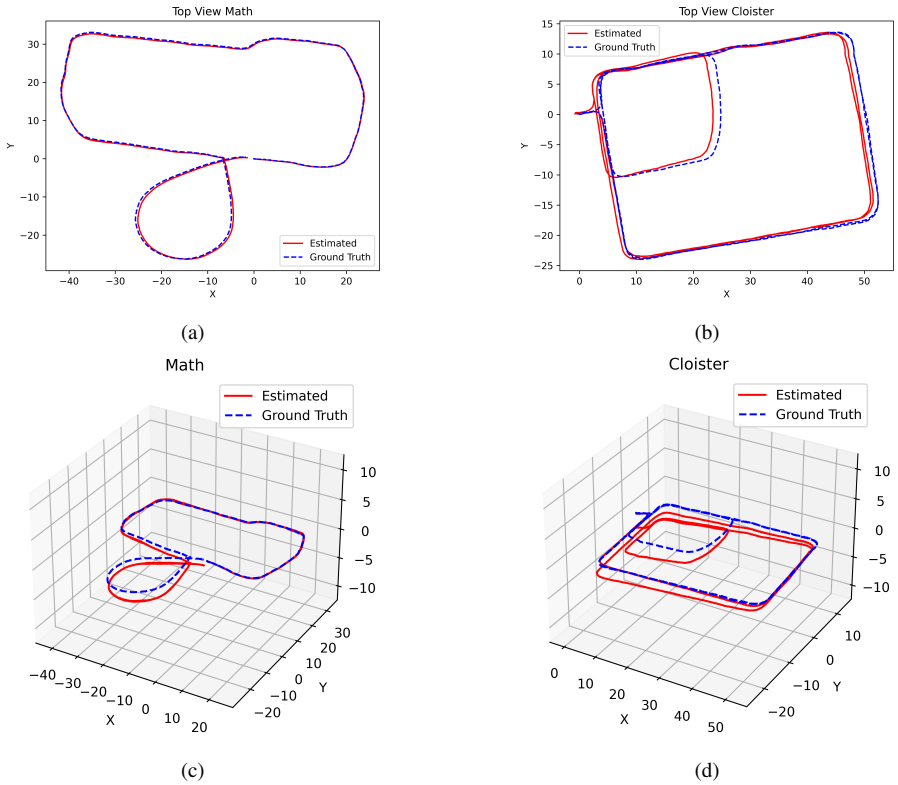

Figure 6. Comparison of ground truth and estimated trajectories for the Newer College dataset: (a) Top view Math, (b) Top view Cloister, (c) Perspective view Math, (d) Perspective view Cloister. The ground truth trajectory is shown in blue, and the estimated trajectory is depicted in dashed orange.

is unavailable. This result highlights the robustness of the proposed approach, which maintains trajectory coherence under challenging conditions. Moreover, in scenarios where traditional methods struggle, such as highly symmetrical environments, our approach delivers reliable results, confirming its suitability for complex real-world applications.

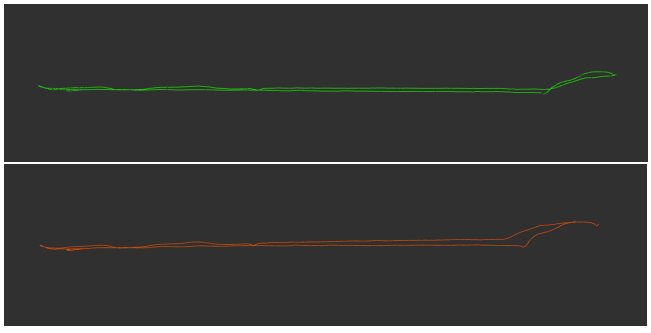

Figure 7. Comparison of trajectories for the ENWIDE Tunnel dataset: (a) Ground truth trajectory, (b) Estimated trajectory.

6. Conclusion

In this paper, we introduce a novel LiDAR-based odometry method that leverages sensor reflectivity data to enhance the robustness and accuracy of pose estimation. By representing the point cloud as a reflectance image, applying SIFT for keypoint detection, and incorporating feature tracking, our approach effectively reduces cumulative drift and improves alignment accuracy. A two-stage outlier detection strategy — combining IMU data and cosine similarity filtering, followed by RANSAC-based refinement — plays a pivotal role in filtering incorrect correspondences, enabling robust point-to-plane optimization.

Experimental evaluations on diverse datasets, including Gexcel HERON, Newer College, and ENWIDE, demonstrate the adaptability of our method in various environments. Despite challenges posed by cumulative errors in sparse feature maps, the system reliably maintains a consistent trajectory, even in GNSS-denied settings and geometrically degenerate scenarios.

Future work will focus on refining the feature matching process, particularly in environments with sparse or inconsistent reflectance features. We also aim to enhance reflectance-based odometry by integrating geometric features to mitigate drift caused by tracking cloud sparsity. Additionally, combining local map-based and reflectance-based odometry—using reflectance-derived features as constraints could further improve the robustness of SLAM systems. This hybrid approach would enhance alignment between new scans and previously built maps, leading to a more resilient and reliable pose estimation framework.

References

Besl, P., McKay, N. D., 1992. A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2), 239-256.

Chen, Y., Medioni, G., 1991. Object modeling by registration of multiple range images. *Proceedings. 1991 IEEE International Conference on Robotics and Automation*, 2724–2729 vol.3.

Du, W., Beltrame, G., 2023. Real-time simultaneous localization and mapping with lidar intensity. *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 4164–4170.

Fischler, M. A., Bolles, R. C., 1987. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. M. A. Fischler, O. Firschein (eds), *Readings in Computer Vision*, Morgan Kaufmann, San Francisco (CA), 726–740.

Geiger, A., Lenz, P., Stiller, C., Urtasun, R., 2013. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11), 1231–1237.

Guadagnino, T., Chen, X., Sodano, M., Behley, J., Grisetti, G., Stachniss, C., 2022. Fast Sparse LiDAR Odometry Using Self-Supervised Feature Selection on Intensity Images. *IEEE Robotics and Automation Letters*, 7(3), 7597-7604.

Lowe, D. G., 2004. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60, 91–110.

Park, Y. S., Jang, H., Kim, A., 2020. I-loam: Intensity enhanced lidar odometry and mapping. *2020 17th International Conference on Ubiquitous Robots (UR)*, 455–458.

Pfreundschuh, P., Oleynikova, H., Cadena, C., Siegwart, R., Andersson, O., 2024. Coin-lio: Complementary intensity-augmented lidar inertial odometry. *2024 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 1730–1737.

Shan, T., Englot, B., 2018. Lego-loam: Lightweight and ground-optimized lidar odometry and mapping on variable terrain. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 4758–4765.

Shan, T., Englot, B., Meyers, D., Wang, W., Ratti, C., Daniela, R., 2020. Lio-sam: Tightly-coupled lidar inertial odometry via smoothing and mapping. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 5135–5142.

Tuna, T., Nubert, J., Nava, Y., Khattak, S., Hutter, M., 2024. X-ICP: Localizability-Aware LiDAR Registration for Robust Localization in Extreme Environments. *IEEE Transactions on Robotics*, 40, 452-471.

Vizzo, I., Guadagnino, T., Mersch, B., Wiesmann, L., Behley, J., Stachniss, C., 2023. 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.

Wang, H., Wang, C., Xie, L., 2021. Intensity-SLAM: Intensity Assisted Localization and Mapping for Large Scale Environment. *IEEE Robotics and Automation Letters*, PP, 1-1.

Xu, W., Zhang, F., 2021. FAST-LIO: A Fast, Robust LiDAR-Inertial Odometry Package by Tightly-Coupled Iterated Kalman Filter. *IEEE Robotics and Automation Letters*, 6, 3317-3324.

Zhang, J., Kaess, M., Singh, S., 2016. On degeneracy of optimization-based state estimation problems. *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 809–816.

Zhang, J., Singh, S., 2014. LOAM : Lidar Odometry and Mapping in real-time. *Robotics: Science and Systems Conference (RSS)*, 109-111.

Zhang, L., Camurri, M., Wisth, D., Fallon, M., 2021. Multi-camera LiDAR inertial extension to the newer college dataset. *arXiv preprint arXiv:2112.08854*.

Zhang, Y., Tian, Y., Wang, W., Yang, G., Li, Z., Jing, F., Tan, M., 2023. RI-LIO: Reflectivity Image Assisted Tightly-Coupled LiDAR-Inertial Odometry. *IEEE Robotics and Automation Letters*, 8(3), 1802-1809.