remote sensing

MDPI

Article

Odometry and Mapping for Complex Environment Perception Under Partial-View Sensing

Xinye Dai [1], Dingxi Wang [1], Jin Xing [2,3,*], Zhibo Zhang [3,4], Xiaoxiao Zhang [3,4], Xiao Wang [1], Shiqi Zheng [1], Yusheng Wang [1,5] and Lijian Feng [6]

[1] GNSS Research Center, School of Geodesy and Geomatics, Wuhan University, Luoyu Road No. 129, Wuhan 430079, China; xinyed@whu.edu.cn (X.D.); wangdx@whu.edu.cn (D.W.); kingxiao0902@whu.edu.cn (X.W.); shiqizheng@whu.edu.cn (S.Z.); yushengwhu@whu.edu.cn (Y.W.)
[2] School of Artificial Intelligence, China University of Geosciences Beijing, Beijing 100083, China
[3] State Key Laboratory of CNS/ATM, Beijing 100191, China; zhangzhibo94@buaa.edu.cn (Z.Z.); 10699@buaa.edu.cn (X.Z.)
[4] School of Electronic and Information Engineering, Beihang University, Beijing 100191, China
[5] Shanghai Huace Navigation Technology Ltd., Wuhan 430073, China
[6] Wenzhou Municipal Bureau of Housing and Urban-Rural Development, Wenzhou 325000, China
* Correspondence: jinxing@cugb.edu.cn

Highlights

What are the main findings?

- A LiDAR-only factor-graph SLAM framework is built for high-resolution partial-view LiDARs using complementary ground, planar, edge, and line constraints.
- Degeneracy-aware feature selection and mixed scan-path loop detection improve trajectory consistency in tunnels, urban roads, and opposite-direction revisits.

What is the implication of the main finding?

- Dense forward-looking LiDARs can support reliable mapping when overlap loss and local geometric weakness are explicitly modeled.
- The results offer deployment guidance for mobile mapping and autonomous-driving systems that rely on high-resolution partial-view sensing.

Abstract

Dense partial-view LiDAR observations are attractive for outdoor perception, but limited overlap and viewpoint sensitivity make odometry and mapping less reliable than with spinning LiDARs. Many recent algorithms for this sensing regime are built as LiDAR-inertial odometry frameworks, whose localization and mapping performance can degrade or fail when the IMU state estimation becomes unstable. This paper presents a LiDAR-only framework for complex outdoor scenes using a factor-graph back-end. After denoising and motion compensation, the point cloud is projected onto a range image for ground, planar, edge, and line extraction. Pose estimation is strengthened by degeneracy-aware feature selection, while loop closing combines scan-based and path-based cues to handle partial-view revisits. Experiments in tunnels, urban roads, residential areas, and other challenging scenes show reduced drift and improved mapping consistency for dense limited-FoV LiDAR data.

Keywords: LiDAR SLAM; odometry and mapping; partial-view sensing; limited-FoV LiDAR; complex environments; mobile mapping

Check for updates

Academic Editors: Jorge Delgado García, Yiqing Yao and Johnson Ihyeh Agbinya

Received: 17 April 2026
Revised: 3 July 2026
Accepted: 3 August 2026
Published: 2 September 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Light detection and ranging (LiDAR) is widely used in autonomous driving and field robotics because it delivers direct three-dimensional range measurements and is less affected by ambient illumination than camera-only sensing. Recent automotive LiDAR designs have moved beyond conventional rotating scanners toward sensors with higher angular sampling, larger point rates, and more diverse scanning patterns, including hybrid solid-state and narrow-field-of-view (FoV) devices. Such sensors generate dense, image-like point clouds and therefore provide useful geometric detail for perception, mapping, and multi-sensor fusion [1–3].

This sensing mode also changes the difficulty profile of simultaneous localization and mapping (SLAM). General LiDAR odometry has benefited from efficient LiDAR-inertial and voxelized systems, including FAST-LIO2 [4], Point-LIO [5], and Voxel-SLAM [6]. However, these advances do not remove the specific limitations introduced by hybrid solid-state and other limited-FoV sensors. Their observations cover only part of the surrounding scene, the available features are often directionally biased, and revisit geometry can vary sharply with the vehicle heading. These effects are especially apparent in urban traffic, tunnels, and opposite-direction traversals [7–9]. Related measurement-system studies on sensor-fusion localization, intensity-aided LiDAR SLAM, and multi-LiDAR calibration further show that sensing configuration remains an important factor in localization performance [10–12]. In this study, representative high-resolution forward-looking LiDARs are used to examine this partial-view operating regime.

In this paper, "complex outdoor environments" refer to scenes in which reliable LiDAR registration is challenged by limited scan overlap, uneven or repetitive geometric structures, dynamic objects, or substantial viewpoint changes. Urban built-up areas, tunnels, and bridges represent different combinations of these challenges.

Motivated by this dependency, this paper deliberately studies a LiDAR-only setting. In LIO systems, the estimated IMU state is often used not only as a motion prior, but also as an essential input for motion compensation, scan registration initialization, and back-end optimization. When the IMU state estimation becomes abnormal or divergent, these coupled modules may receive inconsistent motion information and the failure can propagate through the entire SLAM pipeline. We therefore use the LiDAR-only formulation to investigate whether dense limited-FoV LiDAR observations can maintain stable odometry and mapping when inertial state estimation fails. This minimum-sensor setting also provides a clearer baseline for understanding the geometric capability and limitations of dense partial-view LiDARs.

Two practical difficulties motivate the proposed design.

Dense sampling does not guarantee well-conditioned estimation: high-resolution sensors provide many points per frame, but the information carried by these points can be spatially redundant or concentrated in a few structures. Using all measurements directly, or applying a standard feature extractor without considering observability, may increase computation while leaving weak motion directions poorly constrained. This problem is closely related to the feature-selection assumptions inherited from LOAM-style pipelines [13–16].

Partial FoV reduces registration overlap: unlike 360-degree spinning LiDARs, forward-looking hybrid solid-state sensors observe only a limited sector at each scan. The resulting overlap loss makes scan matching more vulnerable to rapid rotation, nearby dynamic objects, and revisits from the reverse direction. Prior studies on FoV variation and heterogeneous LiDAR place recognition indicate that overlap-aware odometry and loop detection remain open issues for such sensors [8,9,17–19].

To address these issues, we propose a LiDAR-only SLAM framework for high-resolution narrow-FoV LiDARs operating in complex outdoor environments. The main contributions are as follows:

(1) A LiDAR-only SLAM framework is formulated for high-resolution narrow-FoV LiDARs, where ground, scan registration, and line constraints are jointly optimized in a factor graph.
(2) A geometry-aware feature extraction and matching strategy is designed with degeneracy-aware feature selection to strengthen estimation in feature-poor and weak-overlap scenes.
(3) The system is evaluated in challenging real-world scenarios, including tunnels, urban roads, and repeated traversals, to analyze its behavior under dense but partial-view sensing.

The remainder of the article is organized as follows. Section 2 describes the sensor configurations and proposed method. Section 3 presents the experimental evaluation. Section 4 discusses scope and limitations, and Section 5 concludes the paper.

1.1. Related Works

This section reviews LiDAR odometry and mapping for solid-state or limited-FoV sensors, followed by geometric constraints and place-recognition methods that are relevant to the proposed framework.

1.2. Limited-FoV LiDAR SLAM

LiDAR SLAM has progressed from classical feature pipelines to tightly coupled, direct, and voxel-based formulations. FAST-LIO2 [4] provides efficient direct LiDAR-inertial odometry for multiple sensor types, including small-FoV solid-state LiDARs, and Point-LIO [5] improves performance under high-frequency measurements and aggressive motion. VoxelMap++ [20] and Voxel-SLAM [6] further demonstrate that compact voxel representations can support accurate and efficient LiDAR(-inertial) estimation. Other work has emphasized flexible software architectures and lightweight deployment [16,21].

When the LiDAR has a hybrid solid-state or otherwise limited-FoV scanning pattern, sensor geometry becomes a dominant design consideration. LOAM-Livox [7] adapts the LOAM framework to irregular, non-repetitive small-FoV scans. Li et al. [22] develop a feature extraction method for Livox Horizon, and Wang et al. [23] report a lightweight SLAM pipeline for Intel L515. AS-LIO [8] later incorporates spatial overlap into the back-end to handle aggressive FoV variation. From a measurement-system perspective, Liu et al. [10] combine binocular vision, 2D LiDAR, and IMU for localization and planar mapping; Dai et al. [11] use intensity information to improve LiDAR SLAM in unstructured environments; and Wei et al. [12] address automatic multi-LiDAR calibration and refinement. Together, these studies suggest that limited-FoV odometry remains sensitive to overlap degradation, vehicle heading, and geometric degeneracy.

Benchmark resources have also expanded. Multi-modal LiDAR datasets [2,3] and heterogeneous place-recognition benchmarks such as HeLiPR [9] make it possible to compare different sensing modalities and revisit conditions. A remaining question is how to design and assess SLAM systems specifically for dense partial-view LiDAR observations rather than for panoramic scans alone.

1.3. Geometric Constraints and Place Recognition

Geometric primitives are often introduced to improve observability when local scene structure is weak. Ground planes [24,25], line features [26,27], and vanishing-point-related structures [28,29] have been widely used in visual SLAM for this purpose. Similar ideas appear in LiDAR SLAM through road-curb constraints [30], wall features [31,32], and explicit ground modeling [33,34]. In outdoor scenes, ground estimates from scan rings [35] or random sample consensus (RANSAC) [36] are particularly useful for limiting vertical drift and roll-pitch errors. These constraints are attractive for limited-FoV LiDARs because partial observations can make the estimator strongly dependent on a small set of local structures.

Loop detection and place recognition face a related visibility problem. Scan Context [37] has been extended by Scan Context++ [17] to improve robustness against lateral and rotational changes. Learning-based descriptors, including OverlapNet [38], OverlapTransformer [18], and RangePlace [19], learn overlap-aware or range-image-based representations for revisit retrieval. Wang et al. [39] also introduce a multi-scale point-cloud feature transformer for LiDAR-SLAM loop closure detection, illustrating the increasing use of discriminative learned descriptors. Even so, a scan-only descriptor may fail when a partial-view sensor revisits a place from the opposite direction. This motivates loop detection that combines local scan evidence with path-level cues.

Overall, existing LiDAR odometry systems provide strong baselines, but dense measurements, limited FoV, and viewpoint-dependent overlap still create sensor-specific failure modes. The method proposed here uses high-resolution point clouds to construct complementary ground, planar, edge, and line constraints, and combines them with mixed scan-path loop detection for LiDAR-only SLAM.

2. Materials and Methods

This section presents a LiDAR-only SLAM framework for dense limited-FoV sensing. While classical point-to-line and point-to-plane residuals are used for local registration, partial-view observations reduce scan overlap, bias feature distribution, and introduce redundant or unreliable correspondences. To address these FoV-induced issues, the proposed method jointly optimizes ground, scan-registration, and line constraints in a factor graph, and uses sensor-specific preprocessing, DOP-inspired feature selection, and mixed scan-path loop detection to improve robustness in weak-overlap and geometrically degenerate scenes.

2.1. Representative High-Resolution LiDARs

The term “image-grade LiDAR” refers to high-resolution LiDARs that generate dense, image-like point clouds. Innovusion Cheetah and Neuvition Titan M1 are used as representative high-resolution limited-FoV sensors in the experiments. Hesai Pandar64, RoboSense RS-Ruby, and Livox Horizon are additionally included to provide context across different LiDAR configurations.

The representative specifications are summarized in Table 1, and example scans are shown in Figure 1. Compared with lower-resolution or wider-view configurations, high-resolution limited-FoV LiDARs preserve fine scene details through dense vertical sampling; distant vehicles and pedestrians remain visible in single scans. This characteristic makes such sensors useful for camera integration and perception modules that benefit from dense local geometry.

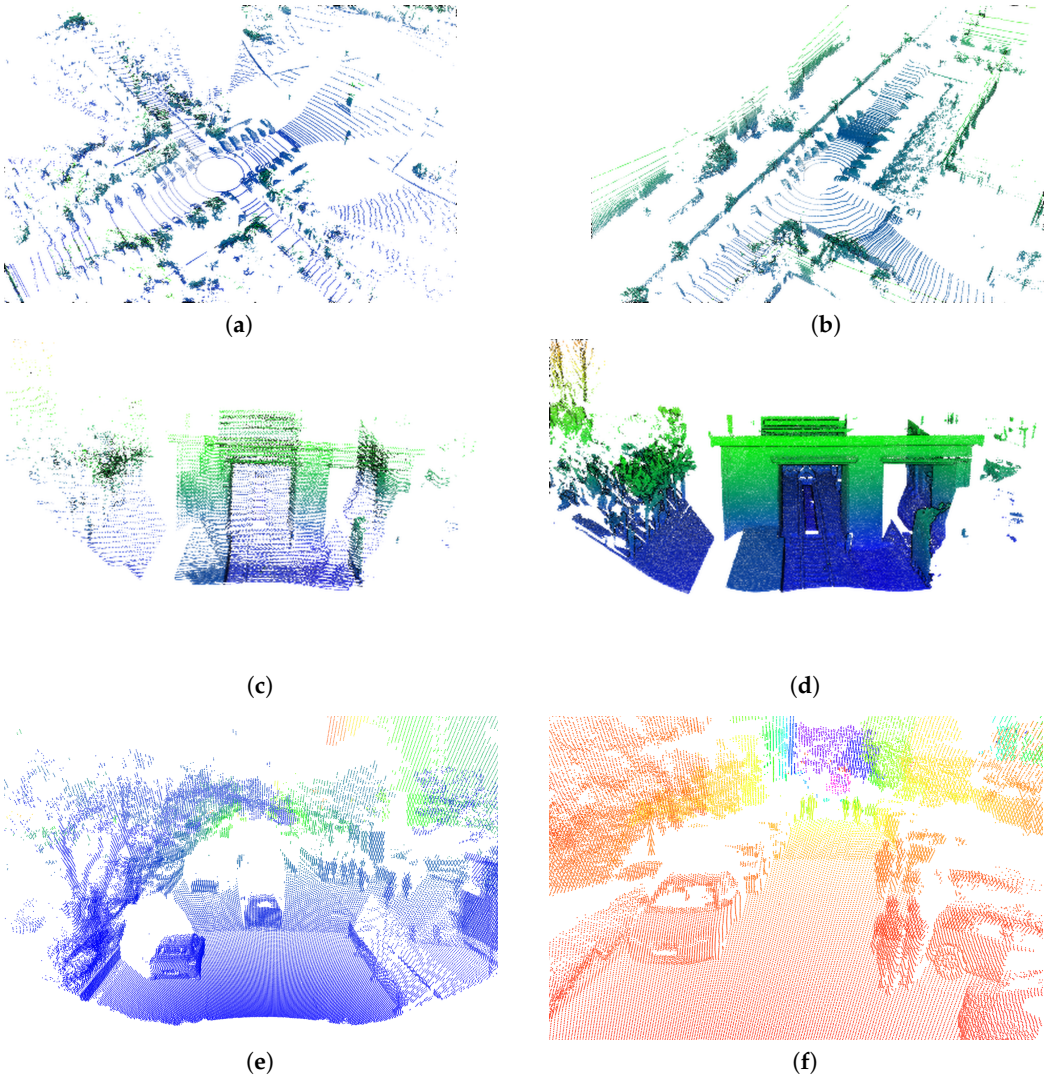


Figure 1. Representative scans of different LiDAR configurations. (**a**,**b**) show single scans of Pandar64 and RS-Ruby. (**c**,**d**) show 0.1 s and 3 s integration scans of Horizon. (**e**,**f**) show single scans of Cheetah and Titan M1.

Table 1. Representative specifications of different LiDAR configurations.

Model	Pandar64	RS-Ruby	Horizon	Cheetah	Titan M1
Type	Mechanical	Hybrid solid-state	Hybrid solid-state	Hybrid solid-state	Hybrid solid-state
Resolution	64	128	Non-repetitive	300	480
Laser wavelength	905 nm	905 nm	905 nm	1550 nm	905 nm
Detection range	200 m (10% reflectivity)	200 m (10% reflectivity)	90 m (10% reflectivity)	200 m (10% reflectivity)	200 m (20% reflectivity)
FoV	$H360°V40°$	$H360°V40°$	$H81.7°V25.1°$	$H100°V40°$	$H45°V25°$
Range precision	2 cm	3 cm	2 cm	2 cm	2 cm
Point rate	1,152,000 points/s	2,304,000 points/s	240,000 points/s	2,142,000 points/s	1,500,000 points/s
Frame rate	10 Hz, 20 Hz	5 Hz, 10 Hz, 20 Hz	10 Hz, 20 Hz, 50 Hz	6–20 Hz	1–30 Hz

2.2. Framework Overview

Figure 2 illustrates the proposed pipeline. Raw point clouds are first denoised and motion-compensated, after which ground points are segmented. The segmented ground points are associated as ground-plane constraints to stabilize roll, pitch, and vertical estimation, as described in Section 2.5. The remaining points are used to extract planar points, edge points, and edge lines. Planar and edge points are obtained through range-image

projection and are used for scan registration in Section 2.6. To improve the conditioning of registration, their point-to-plane and point-to-line correspondences are further screened using a DOP-inspired log-determinant criterion. Depth-continuous edge structures in man-made scenes are extracted as line features and introduced as line constraints in Section 2.7. The back-end is constructed as a factor graph that jointly optimizes ground constraints, LiDAR scan registration constraints, and line constraints. A separate loop-detection thread is used to identify revisited areas.

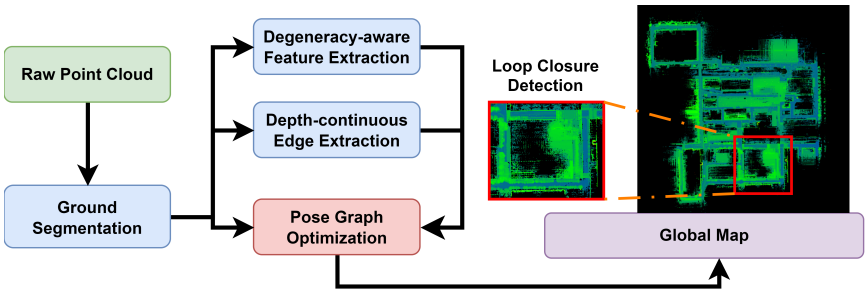


Figure 2. The proposed system overview.

2.3. Methodology

This subsection details the estimation formulation. The objective is to estimate vehicle ego-motion while maintaining a globally consistent map, which can be written as a maximum a posteriori (MAP) problem. For consistency, k denotes the current keyframe index, $\mathcal{Z}_{1:k} = \{\mathbf{z}_1, \mathbf{z}_2, \ldots, \mathbf{z}_k\}$ denotes all measurements up to keyframe k, and $\chi_k = \{\mathbf{x}_0, \mathbf{x}_1, \ldots, \mathbf{x}_k\}$ denotes the corresponding state history. Each state $\mathbf{x}_i$ contains the LiDAR pose at keyframe i, represented by a rotation and translation in $SE(3)$. The symbol $p(\cdot)$ in Equation (1) is a probability density function rather than an optimizable parameter. Given the measurements $\mathcal{Z}_{1:k}$ and the state history χ_k, the MAP estimate is

$$\chi_k^* = \arg\max_{\chi_k} p(\chi_k \mid \mathcal{Z}_{1:k}) \propto p(\chi_0) p(\mathcal{Z}_{1:k} \mid \chi_k). \tag{1}$$

In Equation (1), $p(\chi_k \mid \mathcal{Z}_{1:k})$ is the posterior distribution of the state history after incorporating all measurements, $p(\chi_0)$ is the prior distribution of the initial state, and $p(\mathcal{Z}_{1:k} \mid \chi_k)$ is the measurement likelihood. The proportional relation removes the evidence term $p(\mathcal{Z}_{1:k})$, which is independent of the optimized state. Under the usual conditional-independence assumption and Gaussian measurement-noise model, Equation (1) leads to the following least-squares problem:

$$\chi_k^* = \arg\min_{\chi_k} \sum_{i=1}^{k} \|\mathbf{r}_i(\chi_k)\|^2. \tag{2}$$

where $\mathbf{r}_i(\chi_k)$ denotes the residual induced by measurement $\mathbf{z}_i$, namely the difference between the observed quantity and the quantity predicted from the current state estimate. Minimizing the squared residuals is equivalent to maximizing the posterior when the residual covariance is treated as fixed or absorbed into the residual normalization. The problem is represented by a factor graph with three measurement-factor types. To keep the optimization suitable for real-time operation, keyframes are organized in a sliding window. Let $\mathcal{W}_k$ be the active sliding window and let $\chi_{\mathcal{W}}$ be the states contained in this window. The notations $N_{\mathcal{L}}$, $N_{\mathcal{P}}$, and N_{ℓ} denote the numbers of LiDAR-odometry, ground-plane, and line-feature factors in the window, respectively. The calligraphic subscripts are used

consistently as follows: $\mathcal{L}$ for LiDAR odometry, $\mathcal{P}$ for ground planes, and ℓ for extracted line features. For a window containing N_K keyframes, the states are estimated by minimizing

$$\chi^*_{\mathcal{W}} = \arg\min_{\chi_{\mathcal{W}}} \left\{ \|\mathbf{r}_{\text{pr}}\|^2 + \sum_{i=1}^{N_{\mathcal{L}}} \|\mathbf{r}_{\mathcal{L},i}\|^2 + \sum_{i=1}^{N_{\mathcal{P}}} \|\mathbf{r}_{\mathcal{P},i}\|^2 + \sum_{i=1}^{N_{\ell}} \|\mathbf{r}_{\ell,i}\|^2 \right\}. \tag{3}$$

Here, $\mathbf{r}_{\text{pr}}$ is the prior factor obtained after Schur-complement marginalization [40]; it summarizes information from states that have been removed from the sliding window. The term $\mathbf{r}_{\mathcal{L},i}$ is the ith scan-registration residual constructed from point-to-line and point-to-plane correspondences; $\mathbf{r}_{\mathcal{P},i}$ is the ith ground-plane residual used to constrain roll, pitch, and vertical drift; and $\mathbf{r}_{\ell,i}$ is the ith line-to-line residual used in structured scenes. Equation (3) therefore states that the final pose update is obtained by jointly balancing the prior information and the three types of geometric constraints.

2.4. Preprocessing

A mechanical spinning LiDAR observes the environment by rotating a vertically arranged laser array, so its density and angular resolution are largely determined by the number of laser channels. The limited-FoV high-resolution LiDAR considered here follows a different optical design. As shown in Figure 3, a rotating polygon mirror focuses and steers the emitted beams [41], producing point-cloud density comparable to that of high-resolution rotating sensors within its forward view.

Even after 3D voxel-grid downsampling, outliers may remain in the map and degrade localization. We therefore apply three point-removal rules before feature extraction.

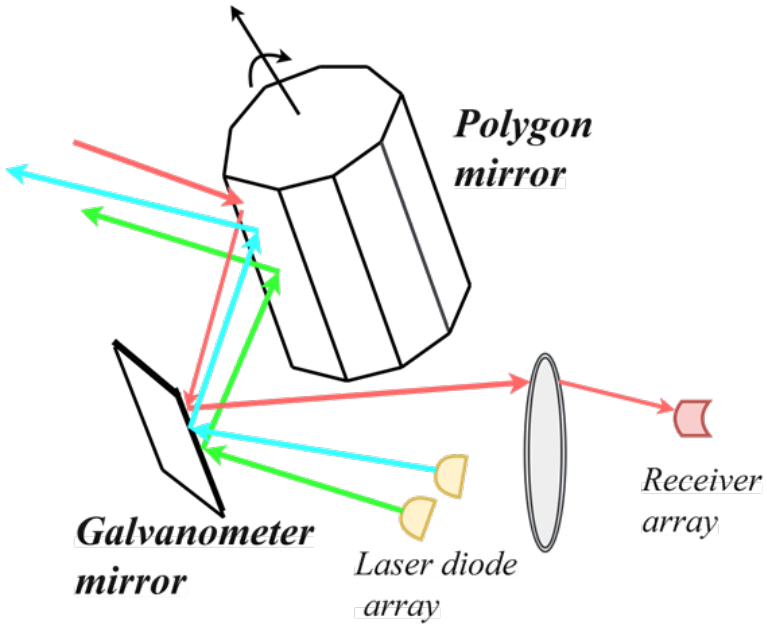


Figure 3. Visual illustration of the scan mechanism. The galvanometer and polygon mirror perform vertical and horizontal scans, respectively.

(1) Points near the FoV boundary. Outliers are concentrated around the scan fringe, so a 5-degree margin on both sides is discarded, leaving a 96-degree horizontal view for the forward-looking sensor used in the experiment.
(2) Points behind foreground objects. Some erroneous returns appear behind visible foreground surfaces; these points are filtered using the empirical criterion reported in [7].

(3) Points below the ground surface. Hybrid solid-state LiDARs may mirror high-intensity objects, for example, projecting an above-ground road sign below the road plane. Points immediately below the detected ground are therefore removed.

Assuming constant angular and linear velocities between two consecutive LiDAR frames, the pose estimated from the previous frame can be used to predict motion and compensate scan distortion in the current frame. Let the estimated transformation from frame $k-1$ to frame k be $\boldsymbol{\xi}_{k-1}^{k}$. For a raw point $\tilde{\mathbf{p}}_i^k$ acquired at timestamp τ_i within frame k, the distortion is corrected point by point as

$$\mathbf{p}_i^k = \mathrm{slerp}(\mathbf{R}(\boldsymbol{\xi}_{k-1}^{k}), \delta_i)\tilde{\mathbf{p}}_i^k + \delta_i \mathbf{t}(\boldsymbol{\xi}_{k-1}^{k}). \tag{4}$$

where $\tilde{\mathbf{p}}_i^k$ is the raw point before motion compensation and $\mathbf{p}_i^k$ is the corresponding motion-compensated point. The functions $\mathbf{R}(\boldsymbol{\xi}_{k-1}^{k})$ and $\mathbf{t}(\boldsymbol{\xi}_{k-1}^{k})$ give the rotation and translation of $\boldsymbol{\xi}_{k-1}^{k}$, and slerp denotes spherical linear interpolation. The interpolation ratio δ_i is defined as

$$\delta_i = \frac{\tau_e - \tau_i}{\tau_e - \tau_b} \tag{5}$$

where τ_i is the acquisition timestamp of point i, and τ_b and τ_e denote the beginning and ending timestamps of the scan acquisition for frame k, respectively. The compensated points are then packed into the frame used for subsequent pose estimation.

2.5. Ground Plane Constraints

Our previous work [42] shows that limited-FoV LiDAR pose estimation is vulnerable to rotational disturbances, which may produce large roll and pitch errors. Once such divergence appears, loop closures alone may not fully correct it. Segmented ground points can constrain roll and pitch effectively [35], but angle-based ground extraction can be unreliable on gently sloped roads because small height changes may be missed and later accumulate as vertical drift. We therefore associate ground planes through their parameters in a shared coordinate frame under a distance metric.

Rough ground points are first selected using thresholds derived from the LiDAR mounting height and orientation, and RANSAC [36] is then used for refinement. The extracted plane $\mathbf{m}$ is represented by its normal vector $\mathbf{n}_{\mathcal{P}}$ and distance scalar $d_{\mathcal{P}}$, namely $\mathbf{m} = [\mathbf{n}_{\mathcal{P}}^T, d_{\mathcal{P}}]^T$. The normal $\mathbf{n}_{\mathcal{P}}$ gives the plane orientation in the current LiDAR frame, and $d_{\mathcal{P}}$ is the signed distance from the frame origin to the plane. The correspondence of each ground point between consecutive scans k and $k+1$ is established as

$$\mathbf{p}_i^{k+1} = \mathbf{T}_k^{k+1}\mathbf{p}_i^k. \tag{6}$$

where $\mathbf{p}_i^{k+1}$ and $\mathbf{p}_i^k$ denote the same point in two consecutive LiDAR frames. The transformation is $\mathbf{T}_k^{k+1} = \left\{\mathbf{R}_k^{k+1}, \mathbf{t}_k^{k+1}\right\}$, with $\mathbf{R}_k^{k+1}$ and $\mathbf{t}_k^{k+1}$ representing rotation and translation. Equation (6) transforms a ground point from frame k into frame $k+1$, so the fitted ground planes from two adjacent frames can be compared in the same coordinate system. The corresponding ground-plane residual is

$$\mathbf{r}_{\mathcal{P},k} = \mathbf{m}_{k+1} - \mathbf{T}_k^{k+1}\mathbf{m}_k. \tag{7}$$

where $\mathbf{m}_k$ and $\mathbf{m}_{k+1}$ are the ground-plane parameters extracted from frames k and $k+1$, respectively. The notation $\mathbf{T}_k^{k+1}\mathbf{m}_k$ denotes the plane from frame k after being transformed into frame $k+1$. Thus, $\mathbf{r}_{\mathcal{P},k}$ measures the mismatch between adjacent ground planes and provides a direct constraint on the relative pose.

2.6. LiDAR Odometry Constraints

The point cloud is next projected into a two-dimensional range image that represents the depth map generated from the LiDAR point cloud. For the kth scan $\mathbf{S}_k = \{\mathbf{p}_1^k, \mathbf{p}_2^k, \ldots, \mathbf{p}_n^k\}$, the projection is determined from the horizontal and vertical angles of each point. Given $\mathbf{p}_i^k = \{x_i, y_i, z_i\}$, the horizontal angle α_i in the x-y plane is

$$\alpha_i = \text{atan2}(y_i, x_i). \tag{8}$$

Here, $\text{atan2}(\cdot, \cdot)$ returns the quadrant-aware azimuth angle. The vertical angle, computed from z_i and the horizontal range, determines the row index of the range image, while α_i determines the column index. This projection preserves the neighborhood relation of adjacent laser measurements and enables roughness-based feature extraction.

$$c_i = \frac{1}{|\mathcal{N}_i| \cdot \|\mathbf{p}_i^k\|} \left\| \sum_{j \in \mathcal{N}_i} (\mathbf{p}_i^k - \mathbf{p}_j^k) \right\|, \tag{9}$$

where c_i denotes the local roughness of point $\mathbf{p}_i^k$, and $\mathcal{N}_i$ is the index set of the continuous neighboring points of $\mathbf{p}_i^k$ in the same row of the range image. The term $|\mathcal{N}_i|$ is the number of neighboring points, and $\|\mathbf{p}_i^k\|$ normalizes the roughness by the point range. To avoid concentrating features in a narrow spatial region, the range image is divided horizontally into equal sub-blocks. Points in each row are sorted and classified by thresholding. Unlike LOAM or LeGO-LOAM, which use a fixed threshold, this work applies distance-dependent thresholds. A larger c_i indicates a sharper local surface and is assigned to an edge feature, whereas a smaller c_i is treated as a surface feature.

LiDAR odometry estimates the relative motion between consecutive scans through point-to-line and point-to-plane scan matching. For each current edge point $\mathbf{p}_i^\varepsilon$, two neighboring points $\mathbf{p}_1^\varepsilon$ and $\mathbf{p}_2^\varepsilon$ are selected from the local map. Separate local maps for edge and planar features are maintained to improve search efficiency. The superscript ε denotes edge features, and the subscript i indexes the current feature point. The point-to-line residual $d_{\varepsilon 2\varepsilon}$ is

$$d_{\varepsilon 2\varepsilon} = \frac{\left\| (\mathbf{p}_i^\varepsilon - \mathbf{p}_1^\varepsilon) \times (\mathbf{p}_i^\varepsilon - \mathbf{p}_2^\varepsilon) \right\|}{\left\| \mathbf{p}_1^\varepsilon - \mathbf{p}_2^\varepsilon \right\|}. \tag{10}$$

Equation (10) computes the perpendicular distance from the current edge point to the line determined by the two matched edge points in the local map. The numerator is the area of the parallelogram formed by the two vectors, and the denominator normalizes it by the length of the reference line segment.

For each planar point $\mathbf{p}_i^\rho$, three nearest points, $\mathbf{p}_1^\rho$, $\mathbf{p}_2^\rho$, and $\mathbf{p}_3^\rho$, are retrieved from the local planar feature map. The superscript ρ denotes planar features. The point-to-plane residual $d_{\rho 2\rho}$ is

$$d_{\rho 2\rho} = \frac{\left| (\mathbf{p}_i^\rho - \mathbf{p}_1^\rho)^T ((\mathbf{p}_1^\rho - \mathbf{p}_2^\rho) \times (\mathbf{p}_1^\rho - \mathbf{p}_3^\rho)) \right|}{\left\| (\mathbf{p}_1^\rho - \mathbf{p}_2^\rho) \times (\mathbf{p}_1^\rho - \mathbf{p}_3^\rho) \right\|}. \tag{11}$$

In Equation (11), the cross product of two vectors on the reference plane gives the plane normal. The numerator is the signed point-to-plane distance scaled by this normal, and the denominator converts it into a Euclidean distance.

Dynamic objects, such as vehicles and pedestrians, are partially suppressed by discarding the largest quarter of residuals from the two correspondence sets. If the current frame

contains N_ε edge correspondences and N_ρ planar correspondences, the LiDAR odometry residual is

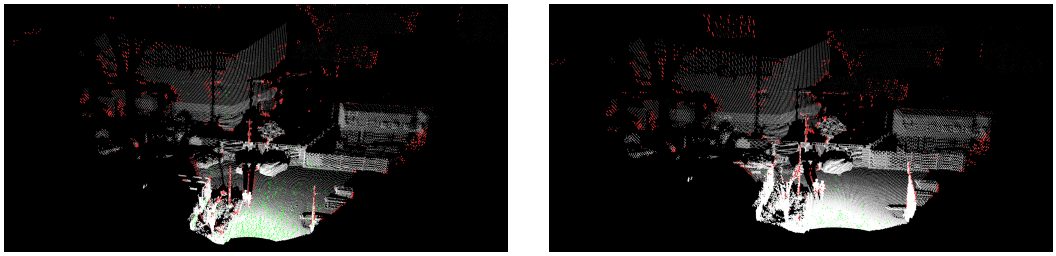
$$\mathbf{r}_{\mathcal{L},k} = \sum_{i=1}^{N_\varepsilon} (d^i_{\varepsilon 2\varepsilon})^2 + \sum_{j=1}^{N_\rho} (d^j_{\rho 2\rho})^2. \tag{12}$$

where $d^i_{\varepsilon 2\varepsilon}$ and $d^j_{\rho 2\rho}$ are the scalar distances from the ith edge correspondence and the jth planar correspondence, respectively. The residual $\mathbf{r}_{\mathcal{L},k}$ is therefore the scan-registration cost contributed by frame k.

Equations (10) and (11) may produce many constraints, increasing the computational cost of real-time optimization. Less informative features should therefore be removed while preserving estimation accuracy. This creates a tradeoff between correspondence count and geometric strength. Inspired by dilution of precision (DOP) in satellite navigation, which quantifies the effect of measurement geometry on position uncertainty, we select features with more favorable geometric distributions. Given the Jacobian matrix $\mathbf{J}$ of the LiDAR odometry residual $\mathbf{r}_{\mathcal{L},k}$, the information matrix is $\mathbf{\Lambda} = \mathbf{J}^T\mathbf{J}$, and the selected feature subset is obtained by

$$\arg\max_{\mathbb{G}_k \subset \mathbb{S}_k} \log\det(\mathbf{\Lambda}). \tag{13}$$

where $\mathbb{G}_k$ denotes the selected feature set and log det is the log-determinant function. The candidate set $\mathbb{S}_k$ contains all edge and planar correspondences available in frame k, and $\mathbb{G}_k \subset \mathbb{S}_k$ is the subset retained for optimization. Maximizing $\log\det(\mathbf{\Lambda})$ favors a subset whose Jacobian spans the pose parameters more evenly; such a subset reduces degeneracy compared with redundant features concentrated in a single direction. The optimization is solved using the stochastic greedy method [15], with an example shown in Figure 4. The point-selection ratio is determined by the current and empirical degeneracy factors, λ_i and λ_{emp} [43]. Specifically, $\mathbb{G}_k$ is set to 70% of $\mathbb{S}_k$ when $\lambda_i \geq \lambda_{emp}$, and to 30% of $\mathbb{S}_k$ when $\lambda_i < \lambda_{emp}$. The 70% value is empirical ratios, selected to balance retaining sufficient geometric constraints for stable optimization and suppressing redundant or poorly distributed features that may aggravate degeneracy. Here, λ_i is the degeneracy indicator estimated from the current frame, and λ_{emp} is an empirical threshold used to decide whether the current geometry is sufficiently well conditioned.



(**a**) (**b**)

Figure 4. Visual illustration of the selected feature results. White points show the background point cloud, red points denote edge features, and green points denote planar features. (**a**) shows the extracted full feature set of 4907 points, whereas (**b**) shows the retained feature set of 3420 points used for pose estimation. Points providing stronger geometric constraints are preserved.

2.7. Line Feature Constraints

Line-assisted visual SLAM has demonstrated robustness in man-made structured scenes [26]. Dense LiDAR point clouds allow similar line cues to be extracted in tunnels and indoor environments. Because junctions between planes often appear as depth-continuous edges and dominate these scenes, we extract such edges as line features. The point cloud is divided into fixed-size voxels, and RANSAC is repeatedly applied to fit planes within

each voxel and identify connected plane pairs. A plane pair is kept only when its included angle falls within a specified threshold, and the intersection of the two planes is treated as a candidate line. Multiple points, including the two endpoints, are sampled from each extracted line, and only lines satisfying the infinite straight-line model [44] are retained.

An infinite straight line is parameterized by a rotation matrix $\mathbf{R}$ and two scalars $a, b \in \mathbb{R}$. In this representation, $\mathbf{R}$ describes the orientation of the line coordinate frame, while a and b describe the two perpendicular offsets of the line in that local frame. The subscripts i and j identify two line observations to be matched. The error operator $\ominus$ between two lines $\mathbf{L}_i$ and $\mathbf{L}_j$ is defined as

$$\mathbf{L}_i \ominus \mathbf{L}_j = \begin{pmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix}^T \mathrm{Log}(\mathbf{R}_i^T \mathbf{R}_j), a_i - a_j, b_i - b_j \end{pmatrix}. \tag{14}$$

The operator $\mathrm{Log}(\cdot)$ maps the relative rotation $\mathbf{R}_i^T \mathbf{R}_j$ from $SO(3)$ to its tangent-space vector. Multiplication by the fixed 3×2 selection matrix keeps the two rotational degrees of freedom relevant to the line orientation, and the last two terms compare the two offset scalars.

The residual between the two lines is then written as

$$\mathbf{r}_{\ell,i} = (\mathbf{T}_{\mathbf{L}_i}^{\mathbf{L}_j} \mathbf{L}_i) \ominus \mathbf{L}_j. \tag{15}$$

where $\mathbf{T}_{\mathbf{L}_i}^{\mathbf{L}_j} = (\mathbf{R}_i^j, \mathbf{t}_i^j)$ describes the transformation between the two lines. Equation (15) first transforms line $\mathbf{L}_i$ into the coordinate frame of $\mathbf{L}_j$ and then evaluates their orientation and offset difference using Equation (14). This residual is inserted into the factor graph as $\mathbf{r}_{\ell,i}$, consistent with the notation used in Equation (3). The derivatives of Equation (15) are computed with the symmetric difference method during optimization.

2.8. Map Association and Loop Detection

The local map contains edge and planar feature maps and, in structured environments, an optional line-feature map. The current scan is associated with the local map by minimizing feature-point distances. An ikd-tree [45] is used for efficient nearest-neighbor search and map updating.

Loop closures are used to reduce accumulated drift in the back-end. For dense partial-view LiDARs, conventional loop-closing methods can be sensitive to moving objects and rotation changes, including histogram-based solutions [7] and Scan Context [37]. We therefore use OverlapNet [38] to detect loop candidates by estimating the overlap ratio between two scans with a deep neural network. Because a limited FoV may prevent scan-based loop detection during opposite-direction revisits, the path-based loop detector from [35] is also incorporated. In this paper, mixed scan-path loop detection denotes this two-source strategy: scan-based candidates are generated from OverlapNet, path-based candidates are generated from the trajectory-pattern detector, and both types of accepted candidates are added as loop-closure constraints in the factor graph.

3. Results

3.1. Experimental Setup

The experiments examine four aspects of the proposed system: the influence of high-resolution narrow-FoV sensing, the mixed scan-path loop-detection strategy, degeneracy-aware feature selection, and the role of geometric constraints in degenerate scenes. The

evaluation is designed as module-level and scene-specific validation for limited-FoV LiDARs, rather than as an exhaustive benchmark against all recent LiDAR-inertial SLAM methods.

Data from the Innovusion Cheetah LiDAR are collected using the autonomous driving platform in Figure 5a. Ground truth is provided by an MPSTNAV M40 GNSS/INS integrated navigation system, with real-time kinematic (RTK) corrections transmitted from our base station. Data from the Neuvition Titan M1 are collected using the unmanned ground vehicle (UGV) in Figure 5b. Its ground truth is obtained from an Xsens MTi-680G GNSS/INS system, with RTK corrections transmitted by a Qianxun SI data terminal unit (DTU). The two LiDARs serve as representative experimental platforms, not as the sole objects of comparison. In the following experiments, the long road, campus, industrial-park, and tunnel sequences are collected with the Cheetah-based autonomous driving platform, whereas the residential-area loop-closing sequence and the colored reconstruction examples are collected with the Titan M1 UGV platform.

The following experiments are organized to support the evaluation from sensor characteristics to system behavior. The multi-LiDAR comparison first examines how sensing configuration influences odometry under the same road condition. The remaining experiments evaluate whether the proposed modules improve mapping consistency and reduce drift in weak or repetitive geometry. Urban and residential scenes contain abundant structures but also dynamic objects and viewpoint-dependent revisits; tunnels provide weak geometric variation; and bridges and noise barriers contain repetitive structures that can cause registration degeneracy.



(**a**) (**b**)

Figure 5. The two platforms for data collection. (**a**) is a Cheetah-based autonomous driving platform; (**b**) is a Titan M1 UGV platform.

The system is implemented in C++ and runs on Ubuntu Linux with ROS. Nonlinear factor-graph optimization is performed using GTSAM [46]. All raw data are processed on a desktop equipped with an Intel i9-11900K CPU (3.5 GHz, 8 cores), 128 GB RAM, and an RTX 3080 Ti GPU with 12 GB memory.

3.2. Comparison Across LiDAR Configurations

To compare odometry behavior under different sensing patterns, six LiDARs are tested on the same driving sequence: Intel RealSense L515, Velodyne VLP-16, Livox Horizon, Hesai Pandar 40P, Ouster OS1-64, and RoboSense RS-Ruby. A-LOAM is used for all sensors, with modifications limited to the front-end. All LiDARs are mounted on the same vehicle, and the extrinsic calibration from each LiDAR to the GNSS antenna is configured beforehand. The vehicle is manually driven on a busy road containing an 850 m tunnel. The trajectories are shown in Figure 6a,b. Figure 6c further provides side-view mapping examples for representative sensors so that vertical drift and map-layer consistency can be visually inspected. The comparison is therefore made from two perspectives: the trajectory

curves in Figure 6a,b quantify horizontal and vertical pose consistency, whereas the point-cloud side views in Figure 6c show how this pose consistency affects the accumulated 3D structure. The L515, whose depth-camera-like design limits its maximum range to 9 m, is not suitable for road-scale observation and gives the weakest result. The horizontal trajectories indicate that denser point clouds can improve robustness under geometric degeneracy. In this sequence, denser forward-looking sensors are less affected by the tunnel than lower-density sensors, and the high-resolution narrow-FoV example maintains relatively small horizontal and vertical deviations after a 2.1 km traversal.

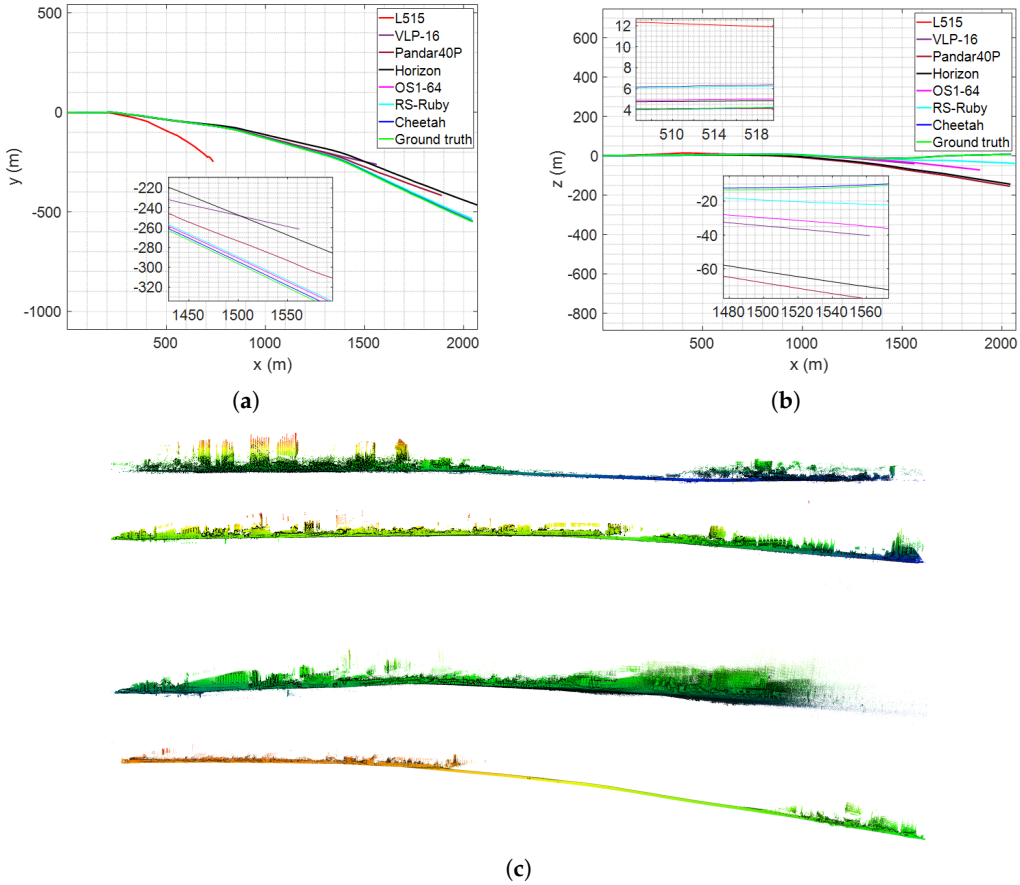


Figure 6. Pose estimation results for different LiDAR configurations. (**a**) shows the trajectories projected onto the horizontal xy plane. (**b**) presents the same trajectories in the xz plane to reveal vertical drift. (**c**) gives side-view point-cloud maps accumulated by four representative sensors; from top to bottom, the maps are from Intel L515, Livox Horizon, RoboSense RS-Ruby, and Innovusion Cheetah, respectively. The side views illustrate the structural consequence of the trajectory errors: severe vertical drift produces tilted or separated map layers, whereas more stable odometry preserves a flatter road surface and more consistent roadside structures.

The six LiDARs differ in field of view, scanning mechanism, point density, and measurement range, which enables an analysis of how sensor configuration influences LiDAR odometry under the same road environment. The results show that dense forward-looking measurements can provide more stable geometric observations in challenging scenes such as tunnels, whereas low-density or short-range sensors are more susceptible to insufficient structural constraints and accumulated drift. These observations support the need to study high-resolution narrow-FoV LiDAR SLAM separately from conventional panoramic LiDAR settings.

3.3. Effectiveness of Loop Closing

This subsection tests whether mixed scan-path loop detection is useful for partial-view revisits, particularly when the same area is traversed from different directions. Ground truth is used to train OverlapNet, with overlap and yaw angles computed directly. All sequences follow the KITTI training preset. The path-based method [35] and Scan Context [37] are selected as representative loop-detection baselines.

In the first experiment, the vehicle circles a tall building three times, as shown in Figure 7a. The route contains many pedestrians, vehicles, and motorcycles. Figure 7b compares the resulting trajectories. Because the second and third laps do not pass through the starting area of the first lap, neither the path-based method nor Scan Context fully detects the revisits between the first and second laps. Since the later laps largely follow the same route, most loop closures occur between the second and third laps. After applying the OverlapNet loop detector, the reconstructed map shows better layer consistency than the map obtained with the single Scan Context loop detector: the green point cloud inside the black rectangle in Figure 7c is visibly stratified, whereas this stratification is largely removed in Figure 7d.

The second experiment studies the cooperation between OverlapNet and path-based loop closure during a more difficult residential-area traversal. The UGV is manually driven for 20 min, and each building is scanned multiple times from both directions. Figure 8a,b compare the maps produced by OverlapNet alone and by the mixed scan-path method. The OverlapNet-only result leaves visible stacking artifacts along building edges, whereas the mixed scan-path detector improves alignment for opposite-direction revisits. This is consistent with limited-FoV sensing, where scan-based loop detection may miss revisits under large viewpoint changes. The absolute trajectory error (ATE) of the two approaches is plotted in Figure 8c using the evopackage. In this sequence, the maximum trajectory error decreases from 7.622 m with OverlapNet alone to 3.644 m with the mixed scan-path loop detector.

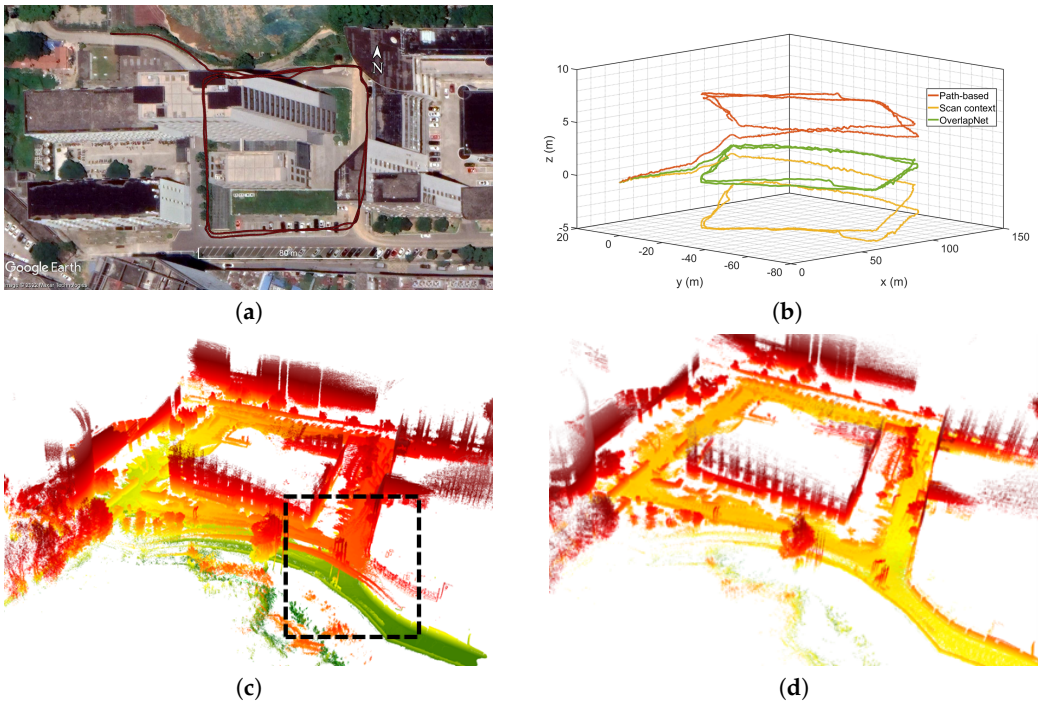


Figure 7. Visualization of different trajectories. (**a**) is the ground truth plotted onto the satellite map. (**b**) compares the estimated trajectories, with the red star showing the start point. (**c**) shows the mapping result using the Scan Context loop detector, where the black rectangle marks revisited areas that were not detected. (**d**) shows the corresponding result using the OverlapNet loop detector, which preserves the revisited structure more completely. The colors in (**c**,**d**) encode height variations.

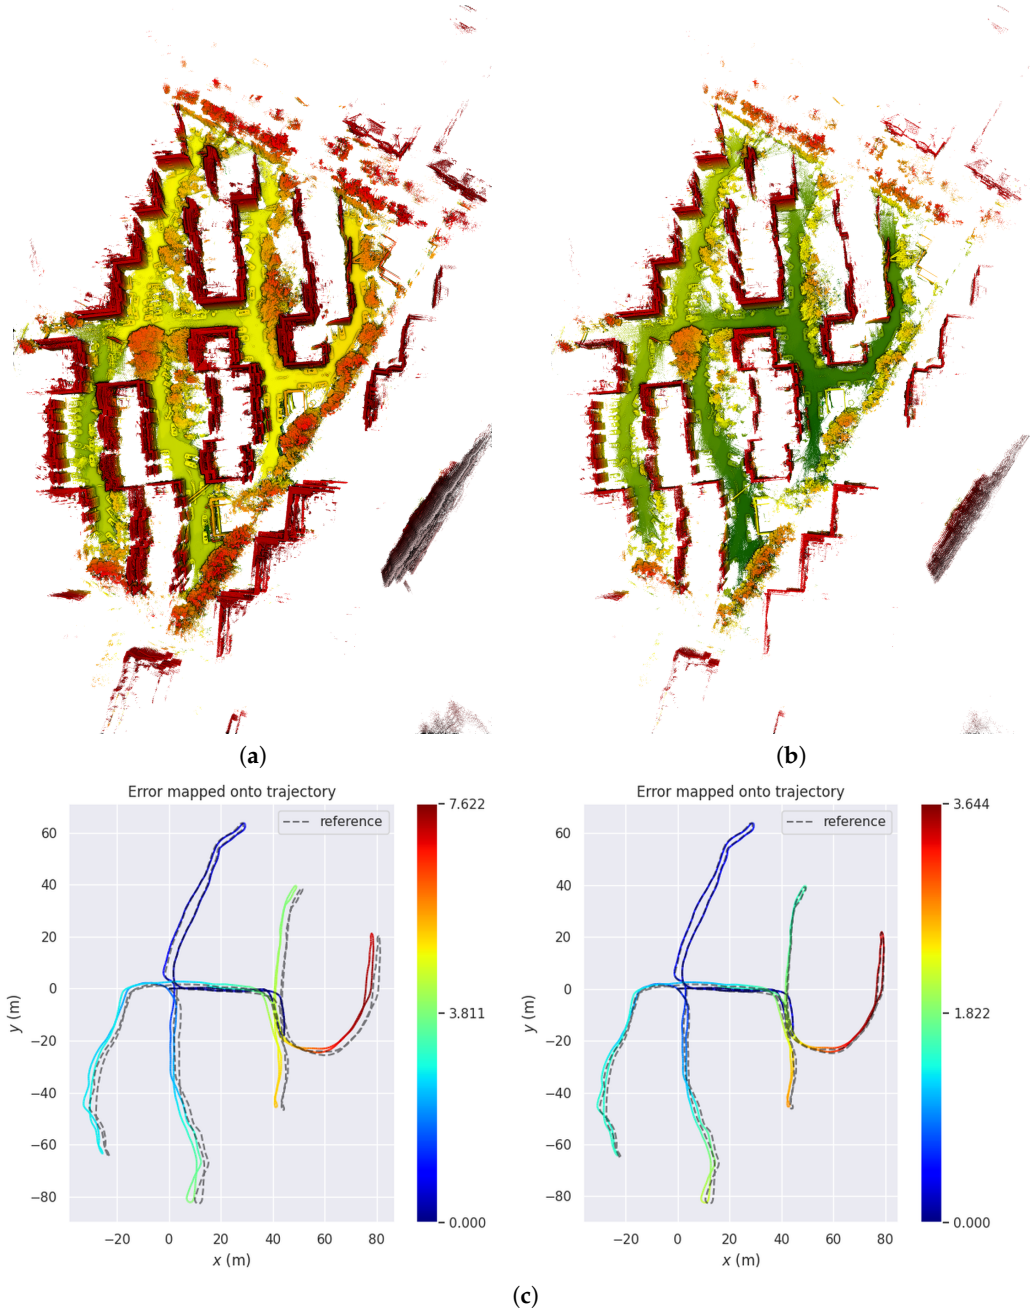


Figure 8. Mapping and trajectory-error comparison of the second loop-closing experiment. (**a**) shows the top-view point-cloud map produced with the OverlapNet loop detector only, with colors coded by height variation. (**b**) shows the corresponding map produced with the proposed mixed scan-path loop detector. The comparison between (**a**) and (**b**) focuses on the alignment of repeated building edges, road boundaries, and courtyard structures after opposite-direction revisits. (**c**) visualizes ATE mapped onto the estimated trajectories; the left panel corresponds to OverlapNet only, the right panel corresponds to the mixed scan-path method, the dashed gray curves denote the ground truth, and warmer colors indicate larger errors.

Figure 8 should be read by first comparing the map geometry in Figure 8a,b and then checking whether the trajectory-error map in Figure 8c supports the same observation. The OverlapNet-only map contains duplicated or blurred building boundaries after opposite-direction revisits, which indicates that some loop constraints are missed or weak. After adding the path-based cue, the building outlines and road boundaries become more compact, and the ATE map shows a lower maximum error. These two observations lead to the same conclusion: the mixed scan-path loop detector improves global consistency in partial-view residential mapping.

Overall, the loop-closing experiments show that scan-based and path-based cues are complementary for partial-view revisits. Combining the two reduces trajectory error and improves building-boundary alignment, confirming that the mixed scan-path loop detector improves global map consistency.

3.4. Effectiveness of Feature Selection

The contribution of degeneracy-aware feature selection is quantified through a campus driving sequence. A-LOAM and F-LOAM are used as LiDAR-only SLAM baselines. The variants w/o FS and w/FS denote the proposed system without and with feature selection, respectively. The vehicle is driven slowly for 406 s. Figure 9a shows the trajectories, and Table 2 reports the error statistics. The inset in Figure 9a enlarges a curved segment where the difference among methods is more visible; the w/FS trajectory follows the ground truth more closely than the w/o FS trajectory and the two LOAM baselines. Figure 9b plots the number of features used for LiDAR odometry. The left curve group in Figure 9b reports the edge-feature numbers, while the right curve group reports the planar-feature numbers. The blue curves are the counts before feature selection, and the orange curves are the counts retained after feature selection. Three abrupt increases appear in the planar-point curve when vehicles are directly ahead of the sensor, since many planar features are extracted from those vehicles. Feature selection suppresses these less informative correspondences and reduces the pose-estimation error. The figure therefore links the mechanism and the result: feature selection reduces redundant or unstable correspondences, especially planar correspondences affected by foreground vehicles, and this reduction corresponds to the lower RMSE and maximum error reported in Table 2.

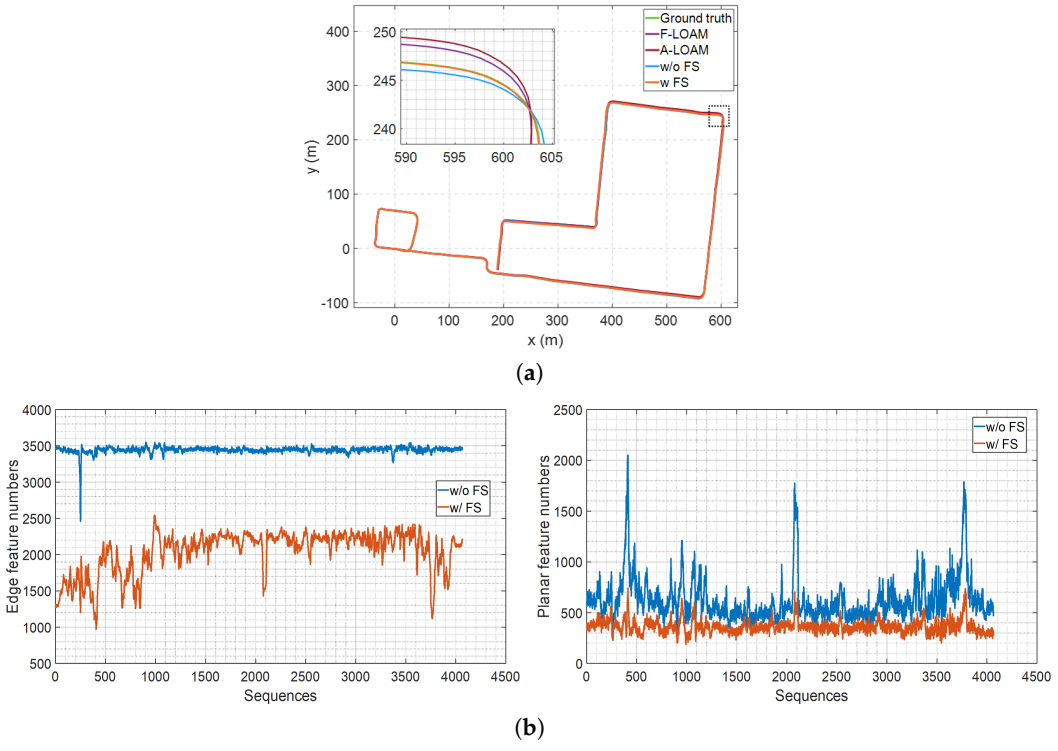


Figure 9. Illustration of the effective feature selection method. (**a**) compares the estimated *xy*-plane trajectories of A-LOAM, F-LOAM, the proposed method without feature selection (w/o FS), and the proposed method with feature selection (w/FS) against the ground truth; the inset enlarges the black dashed region to show local trajectory differences. (**b**) compares the feature numbers used by LiDAR odometry along the sequence. The left plot gives edge-feature counts, and the right plot gives planar-feature counts. The w/o FS curves show all extracted correspondences, whereas the w/FS curves show the retained correspondences after degeneracy-aware selection.

Table 2. Error statistics of various methods.

	A-LOAM	F-LOAM	w/o FS	w/FS
RMSE (m)	4.58	3.71	3.26	2.88
MAX (m)	13.14	9.47	6.57	5.29

3.5. Performance in Degenerate Areas

Performance in degenerate environments is assessed on a city fast road with a long bridge and noise barriers, shown in Figure 10a. These regions contain repetitive structures, such as bridge lamp posts and barrier walls. The trajectories in Figure 10b are compared to examine how different methods behave after traversing degenerate areas. For the error evaluation, the estimated odometry trajectory is compared with the RTK-GNSS/INS ground-truth trajectory. Estimated poses are associated with the ground-truth poses according to their timestamps, transformed into the same vehicle/world coordinate system using the calibrated sensor extrinsics, and aligned by the initial pose before computing the frame-wise absolute trajectory error.

A-LOAM and F-LOAM both produce large heading errors after the large roundabout because passing vehicles are incorrectly registered. With the proposed outlier rejection and constraint design, the estimated trajectory stays closer to the ground truth, with an overall RMSE of 4.23 m.

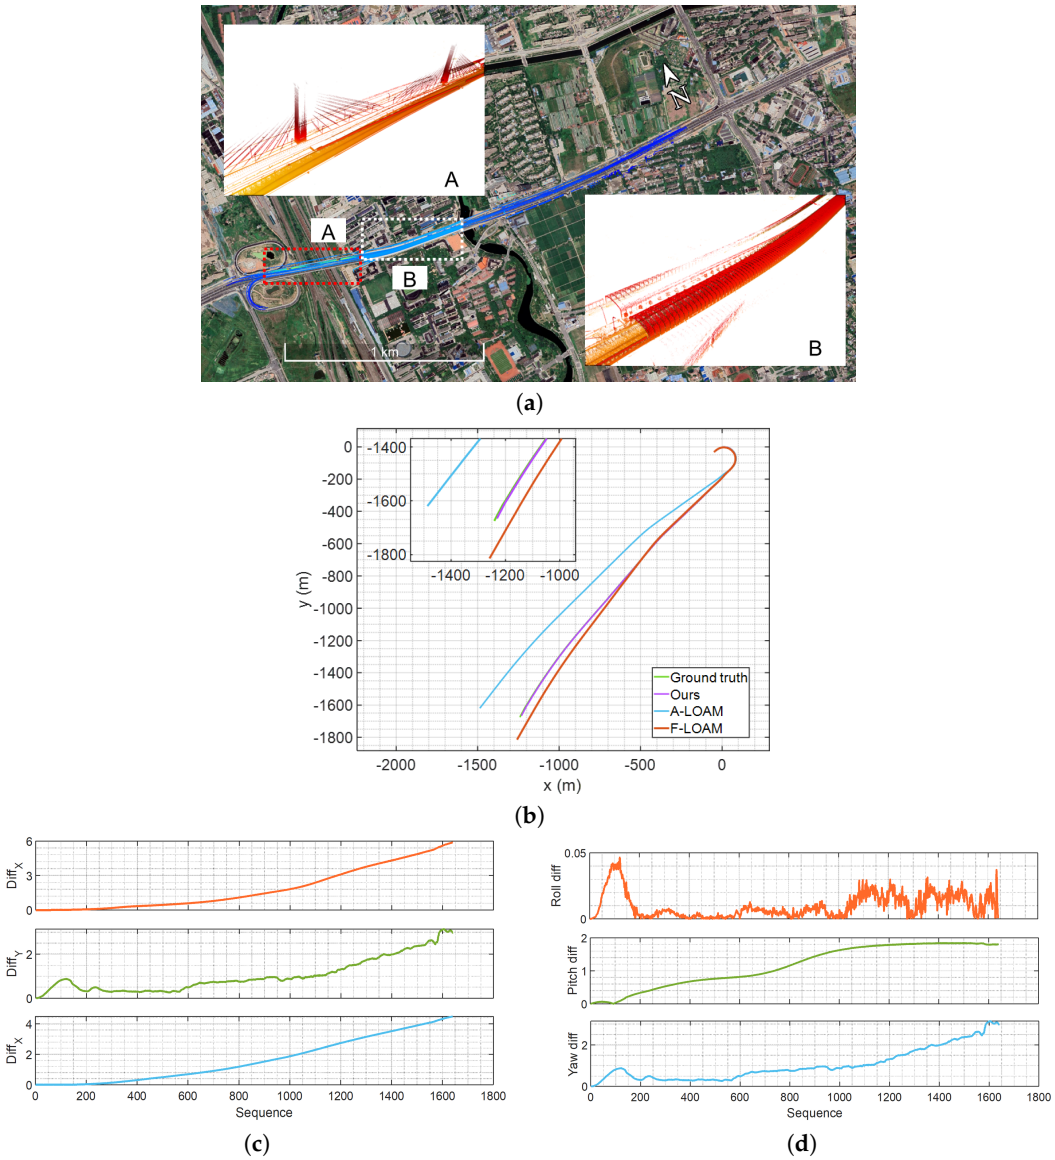
The detailed 6DoF ATE is shown in Figure 10c,d. The roll error remains smaller in magnitude than the errors in the other two attitude directions. These results show that the proposed constraint design remains effective in the presence of repetitive roadside structures and passing vehicles. The trajectory remains close to the ground truth after the degenerate road section, indicating that the system can still provide stable pose estimation when local geometric constraints are weak or potentially ambiguous.

3.6. Ablation Study

This subsection isolates the effect of the additional geometric constraints used in the proposed factor-graph back-end. Here, the additional geometric constraints refer to the ground constraints (GC) and line constraints (LC), which are introduced beyond the standard scan-registration constraints. The first experiment ablates the ground constraints to evaluate their effect on vertical drift suppression, while the second experiment ablates the line constraints to evaluate their contribution in tunnel scenes with weak geometric variation.

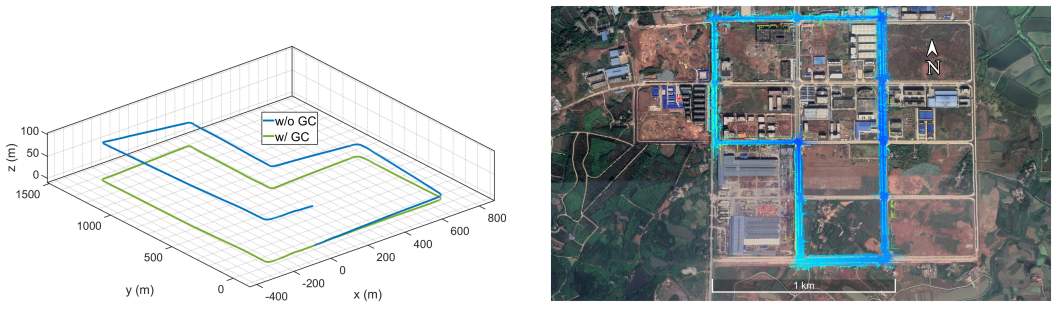
The first experiment is performed in an industrial park, where the vehicle travels for 876 s and eventually returns to the starting point. The variants w/o GC and w/GC denote the proposed system without and with ground constraints, respectively, and loop detection is disabled in both cases. The trajectories are shown in Figure 11a. After a 3.9 km traversal, w/GC returns close to the starting point with an end-to-end error of 1.89 m. By contrast, w/o GC accumulates a large vertical error of 97.6 m. The w/GC map is projected onto the satellite map in Figure 11b, where the alignment indicates that ground constraints help suppress long-term vertical drift.

The second experiment examines line constraints in tunnel scenes. The data are collected in a busy urban environment containing a 1.5 km tunnel, and the vehicle speed exceeds 64 km/h on average. A-LOAM and F-LOAM show clear failures in the tunnel: real-time pose estimation either stops or moves backwards. The variants w/o LC and w/LC denote the proposed method without and with line constraints. Figure 12a plots the ATE along the route. Adding line constraints reduces the maximum ATE by 4.5 m. The map in Figure 12b further indicates that line constraints are useful in long structured corridors with weak lateral variation.



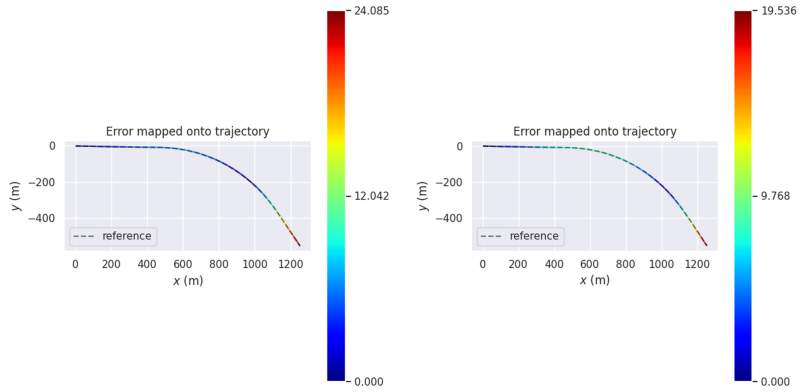
(**a**) (**b**) (**c**) (**d**)

Figure 10. System performance in degenerate areas. (**a**) is the mapping result aligned with the satellite map, with the two insets showing details of the bridge and noise barriers. (**b**) presents the different trajectories of various methods. (**c**) visualizes the frame-wise translational ATE between the estimated odometry and the ground truth along the X, Y, and Z directions in meters. (**d**) shows the corresponding attitude errors in roll, pitch, and yaw in degrees.

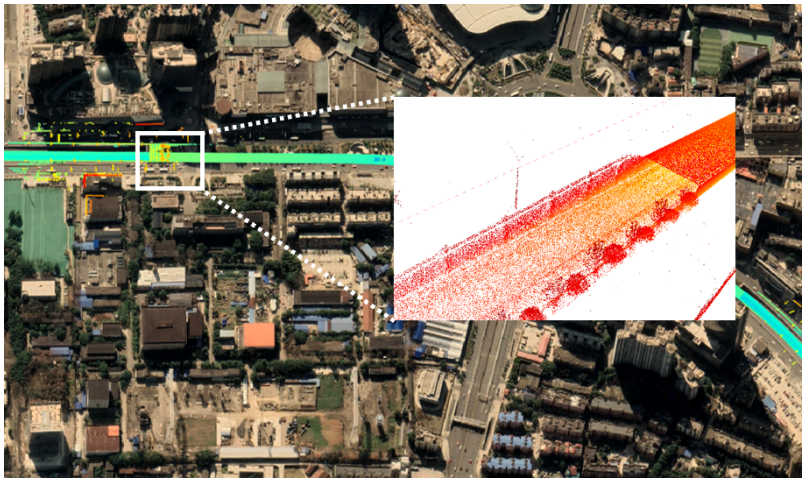


(**a**) (**b**)

Figure 11. Contribution of ground constraints. (**a**) presents the two pose estimation results. (**b**) is the mapping with ground constraints aligned with the satellite map, where the color is coded by height data reflecting the result improvement in vertical map consistency.



(**a**)



(**b**)

Figure 12. Contribution of line constraints in tunnels. (**a**) presents the two ATE curves, where the left inset is w/o LC and the right inset is w/LC. (**b**) is the mapping with line constraints aligned with the satellite map, with the inset showing the mapping at the tunnel entrance and the color coded by height information reflecting the result improvement in consistent scan accumulation and line structures in the local map.

3.7. Run Time Analysis

Table 3 reports the average per-scan runtime of the main modules. The measured timings show that the proposed system can run in real time on the test platform used in this work.

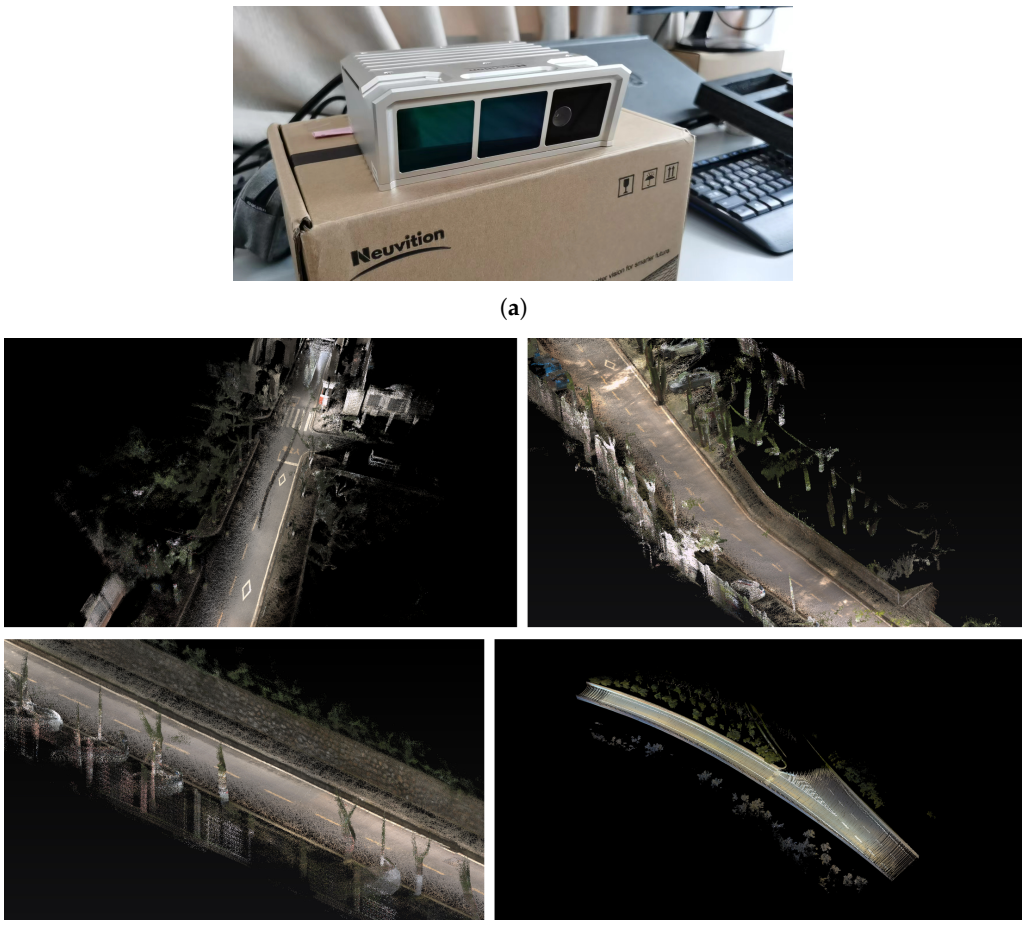
Table 3. Average time consumption of each module [ms].

Module	Submodule	Time Consumption
Feature extraction	Preprocessing	8
	Ground segmentation	13
	Point extraction	12
	Line extraction	14
LiDAR odometry	Feature selection	11
	LiDAR odometry	34
Map optimization	Data association	93

3.8. Colored Mapping Result

Neuvition Titan M1 can generate colored point clouds directly using its internal camera, as shown in Figure 13a. This enables convenient reconstruction of colored 3D

models, as visualized in Figure 13b. The result is included as a qualitative illustration of sensing output rather than as a primary SLAM-accuracy benchmark.



(**a**)

(**b**)

Figure 13. Colored mapping results. (**a**) presents the exterior of the Titan M1 LiDAR. (**b**) shows examples of real-time colored mapping results; the first three are reconstructed in the daytime and the last is generated at night.

Although this experiment is qualitative, the colored reconstructions provide an intuitive assessment of mapping consistency. The clear building facades, road surfaces, and scene details reconstructed in both daytime and nighttime conditions indicate that the estimated poses are sufficiently coherent for dense colored point-cloud accumulation.

4. Discussion

Taken together, the experimental results indicate that the advantage of high-resolution narrow-FoV LiDARs does not come from point density alone, but from using the dense observations with appropriate constraints and data association strategies. In the tested partial-view setting, overlap loss, viewpoint-dependent revisits, and locally repetitive structures are more influential than point density alone. The proposed combination of ground, line, feature-selection, and loop-detection modules is therefore most relevant to scenes where the available geometry is unevenly distributed, such as tunnels, urban roads, and repeated traversals observed from different directions.

The contribution of this work lies in organizing and validating a LiDAR-only system for a specific sensing regime. It does not introduce a new estimation theory; instead, it examines how several geometric components can be combined in a factor graph to address the practical behavior of high-resolution narrow-FoV LiDARs. This positioning is important

because methods designed for panoramic spinning LiDARs may not transfer directly to dense forward-looking sensors. The LiDAR-only design also makes the role of limited-FoV geometry explicit: the observed performance comes from ground, planar, edge, line, and loop constraints extracted from the point cloud, rather than from inertial propagation or image-based correction. This makes the system useful as a baseline for LiDAR-only post-processing, simplified mobile-mapping platforms, and degraded operation when IMU or camera data are not available with sufficient quality.

Several limitations remain. First, the experiments emphasize module ablation and scene-level analysis, and do not constitute a complete benchmark against recent LiDAR-inertial systems such as FAST-LIO2, Point-LIO, or Voxel-SLAM. These systems often rely on inertial coupling or different map representations, whereas this work focuses on LiDAR-only module design under partial-view sensing. Second, the tested sensors still exhibit practical artifacts, including dense outliers, mirrored returns from high-intensity objects, and unstable measurements near reflective surfaces such as water.

5. Conclusions

This article presents a LiDAR-only odometry and mapping system for high-resolution limited-FoV LiDARs in complex outdoor scenes. The method uses a factor graph to combine ground-plane, scan-registration, and line-feature constraints, together with degeneracy-aware feature selection and mixed scan-path loop detection. Experiments on tunnels, urban roads, residential areas, and industrial-park sequences show that these modules help control drift when scan overlap is weak or the local geometry is poorly conditioned.

Future work will extend the evaluation to stronger modern LiDAR-inertial and LiDAR-visual-inertial baselines and will investigate tighter inertial and visual fusion on top of the proposed LiDAR geometric constraints to improve accuracy, efficiency, and deployment stability.

Author Contributions: Conceptualization, X.D. and J.X.; methodology, X.D. and D.W.; software, X.D. and D.W.; validation, X.D., D.W. and X.W.; formal analysis, X.D.; investigation, X.D., D.W., X.W. and S.Z.; resources, J.X., Z.Z., X.Z. and Y.W.; data curation, X.D. and D.W.; writing—original draft preparation, X.D.; writing—review and editing, X.D., D.W., J.X., Z.Z., X.Z., X.W., S.Z., Y.W. and L.F.; visualization, X.D. and D.W.; supervision, J.X. and Y.W.; project administration, J.X.; funding acquisition, J.X. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the Open Project Program of State Key Laboratory of CNS/ATM (Grant No. 2025B51) and the Fundamental Research Funds for the Central Universities (Grant No. 2652025040).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in this study are available on request from the corresponding author.

Conflicts of Interest: Author Yusheng Wang was employed by the company Shanghai Huace Navigation Technology Ltd. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

1. Yuan, C.; Liu, X.; Hong, X.; Zhang, F. Pixel-level Extrinsic Self Calibration of High Resolution LiDAR and Camera in Targetless Environments. *IEEE Robot. Autom. Lett.* **2021**, *6*, 7517–7524. [CrossRef]
2. Li, Q.; Yu, X.; Queralta, J.P.; Westerlund, T. Multi-Modal Lidar Dataset for Benchmarking General-Purpose Localization and Mapping Algorithms. In Proceedings of the 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Kyoto, Japan, 23–27 October 2022.
3. Sier, H.; Li, Q.; Yu, X.; Queralta, J.P.; Zou, Z.; Westerlund, T. A Benchmark for Multi-Modal Lidar SLAM with Ground Truth in GNSS-Denied Environments. *Remote Sens.* **2023**, *15*, 3314. [CrossRef]
4. Xu, W.; Cai, Y.; He, D.; Lin, J.; Zhang, F. FAST-LIO2: Fast Direct LiDAR-inertial Odometry. *IEEE Trans. Robot.* **2021**, *38*, 2053–2073. [CrossRef]
5. He, D.; Xu, W.; Chen, N.; Kong, F.; Yuan, C.; Zhang, F. Point-LIO: Robust High-Bandwidth Light Detection and Ranging Inertial Odometry. *Adv. Intell. Syst.* **2023**, *5*, 2200459. [CrossRef]
6. Liu, Z.; Li, H.; Yuan, C.; Liu, X.; Lin, J.; Li, R.; Zheng, C.; Zhou, B.; Liu, W.; Zhang, F. Voxel-SLAM: A Complete, Accurate, and Versatile LiDAR-Inertial SLAM System. *arXiv* **2024**, arXiv:2410.08935.
7. Lin, J.; Zhang, F. Loam livox: A fast, robust, high-precision LiDAR odometry and mapping package for LiDARs of small FoV. In Proceedings of the 2020 IEEE International Conference on Robotics and Automation (ICRA), Paris, France, 31 May–31 August 2020; pp. 3126–3131.
8. Zhang, T.; Zhang, X.; Liao, Z.; Xia, X.; Li, Y. AS-LIO: Spatial Overlap Guided Adaptive Sliding Window LiDAR-Inertial Odometry for Aggressive FOV Variation. In Proceedings of the 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Abu Dhabi, United Arab Emirates, 14–18 October 2024.
9. Jung, M.; Yang, W.; Lee, D.; Gil, H.; Kim, G.; Kim, A. HeLiPR: Heterogeneous LiDAR Dataset for inter-LiDAR Place Recognition under Spatiotemporal Variations. *Int. J. Robot. Res.* **2023**, *43*, 1867–1883. [CrossRef]
10. Liu, Z.; Li, Z.; Liu, A.; Sun, Y.; Jing, S. Fusion of binocular vision, 2D lidar and IMU for outdoor localization and indoor planar mapping. *Meas. Sci. Technol.* **2022**, *34*, 025203. [CrossRef]
11. Dai, Z.; Zhou, J.; Li, T.; Yao, H.; Sun, S.; Zhu, X. An intensity-enhanced LiDAR SLAM for unstructured environments. *Meas. Sci. Technol.* **2023**, *34*, 125120. [CrossRef]
12. Wei, P.; Yan, G.; Li, Y.; Fang, K.; Cai, X.; Liu, W.; Yang, J. Automatic multi-LiDAR calibration and refinement method. *Meas. Sci. Technol.* **2024**, *35*, 075203. [CrossRef]
13. Marshall, G.F. Risley prism scan patterns. In *Optical Scanning: Design and Application*; SPIE: Bellingham, WA, USA, 1999; Volume 3787, pp. 74–86.
14. Li, W.; Hu, Y.; Han, Y.; Li, X. KFS-LIO: Key-Feature Selection for Lightweight Lidar Inertial Odometry. In Proceedings of the 2021 IEEE International Conference on Robotics and Automation (ICRA), Xi'an, China, 30 May–5 June 2021; pp. 5042–5048.
15. Jiao, J.; Zhu, Y.; Ye, H.; Huang, H.; Yun, P.; Jiang, L.; Wang, L.; Liu, M. Greedy-based feature selection for efficient lidar slam. In Proceedings of the 2021 IEEE International Conference on Robotics and Automation (ICRA), Xi'an, China, 30 May 2021–5 June 2021; pp. 5222–5228.
16. Yi, S.; Lyu, Y.; Hua, L.; Pan, Q.; Zhao, C. Light-LOAM: A Lightweight LiDAR Odometry and Mapping based on Graph-Matching. *IEEE Robot. Autom. Lett.* **2024**, *9*, 3219–3226. [CrossRef]
17. Kim, G.; Choi, S.; Kim, A. Scan Context++: Structural Place Recognition Robust to Rotation and Lateral Variations in Urban Environments. *IEEE Trans. Robot.* **2021**, *38*, 2336–2356. [CrossRef]
18. Ma, J.; Zhang, J.; Xu, J.; Ai, R.; Gu, W.; Chen, X. OverlapTransformer: An Efficient and Rotation-Invariant Transformer Network for LiDAR-Based Place Recognition. *IEEE Robot. Autom. Lett.* **2022**, *7*, 6958-6965. [CrossRef]
19. Li, J.; Liu, Q.; Wang, B.; Liu, H.; Han, Y. RangePlace: A Hierarchical Range Image Transformer for LiDAR-Based Place Recognition. *IEEE Trans. Intell. Veh.* **2024**, 1–11. [CrossRef]
20. Yuan, Y.; Wu, C.; You, Y.; Kong, X.; Zhang, Y.; Li, Q. VoxelMap++: Mergeable Voxel Mapping Method for Online LiDAR(-inertial) Odometry. *IEEE Robot. Autom. Lett.* **2023**, *9*, 427–434. [CrossRef]
21. Blanco-Claraco, J.L. A Flexible Framework for Accurate LiDAR Odometry, Map Manipulation, and Localization. *Int. J. Robot. Res.* **2024**, *44*, 1553–1599. [CrossRef]
22. Li, K.; Li, M.; Hanebeck, U.D. Towards high-performance solid-state-lidar-inertial odometry and mapping. *IEEE Robot. Autom. Lett.* **2021**, *6*, 5167–5174. [CrossRef]
23. Wang, H.; Wang, C.; Xie, L. Lightweight 3-D localization and mapping for solid-state LiDAR. *IEEE Robot. Autom. Lett.* **2021**, *6*, 1801–1807. [CrossRef]
24. Lee, G.H.; Fraundorfer, F.; Pollefeys, M. MAV visual SLAM with plane constraint. In Proceedings of the 2011 IEEE International Conference on Robotics and Automation, Shanghai, China, 9–13 May 2011; pp. 3139–3144.
25. Li, B.; Zou, D.; Sartori, D.; Pei, L.; Yu, W. Textslam: Visual slam with planar text features. In Proceedings of the 2020 IEEE International Conference on Robotics and Automation (ICRA), Paris, France, 31 May–31 August 2020; pp. 2102–2108.

26. Zhou, H.; Zou, D.; Pei, L.; Ying, R.; Liu, P.; Yu, W. StructSLAM: Visual SLAM With Building Structure Lines. *IEEE Trans. Veh. Technol.* **2015**, *64*, 1364–1375. [CrossRef]
27. Fu, Q.; Wang, J.; Yu, H.; Ali, I.; Guo, F.; He, Y.; Zhang, H. PL-VINS: Real-Time Monocular Visual-Inertial SLAM with Point and Line Features. *arXiv* **2020**, arXiv:2009.07462.
28. Lim, H.; Jeon, J.; Myung, H. UV-SLAM: Unconstrained Line-based SLAM Using Vanishing Points for Structural Mapping. *IEEE Robot. Autom. Lett.* **2022**, *7*, 1518–1525. [CrossRef]
29. Wang, P.; Fang, Z.; Zhao, S.; Chen, Y.; Zhou, M.; An, S. Vanishing Point Aided LiDAR-Visual-Inertial Estimator. In Proceedings of the 2021 IEEE International Conference on Robotics and Automation (ICRA), Xi'an, China, 30 May–5 June 2021; pp. 13120–13126.
30. Wang, Z.; Fang, J.; Dai, X.; Zhang, H.; Vlacic, L. Intelligent vehicle self-localization based on double-layer features and multilayer LIDAR. *IEEE Trans. Intell. Veh.* **2020**, *5*, 616–625. [CrossRef]
31. Zhou, L.; Koppel, D.; Kaess, M. LiDAR SLAM with Plane Adjustment for Indoor Environment. *IEEE Robot. Autom. Lett.* **2021**, *6*, 7073–7080. [CrossRef]
32. Zhou, L.; Wang, S.; Kaess, M. pi-LSAM: LiDAR smoothing and mapping with planes. In Proceedings of the 2021 IEEE International Conference on Robotics and Automation (ICRA), Xi'an, China, 30 May–5 June 2021; pp. 5751–5757.
33. Su, Y.; Wang, T.; Shao, S.; Yao, C.; Wang, Z. GR-LOAM: LiDAR-based sensor fusion SLAM for ground robots on complex terrain. *Robot. Auton. Syst.* **2021**, *140*, 103759. [CrossRef]
34. Wei, X.; Lv, J.; Sun, J.; Pu, S. Ground-SLAM: Ground Constrained LiDAR SLAM for Structured Multi-Floor Environments. *arXiv* **2021**, arXiv:2103.03713.
35. Shan, T.; Englot, B. Lego-loam: Lightweight and ground-optimized lidar odometry and mapping on variable terrain. In Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Madrid, Spain, 1–5 October 2018; pp. 4758–4765.
36. Fischler, M.A.; Bolles, R.C. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM* **1981**, *24*, 381–395.
37. Kim, G.; Kim, A. Scan context: Egocentric spatial descriptor for place recognition within 3d point cloud map. In Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Madrid, Spain, 1–5 October 2018; pp. 4802–4809.
38. Chen, X.; Läbe, T.; Milioto, A.; Röhling, T.; Vysotska, O.; Haag, A.; Behley, J.; Stachniss, C. OverlapNet: Loop closing for LiDAR-based SLAM. *arXiv* **2021**, arXiv:2105.11344.
39. Wang, S.; Zheng, D.; Li, Y. LiDAR-SLAM loop closure detection based on multi-scale point cloud feature transformer. *Meas. Sci. Technol.* **2023**, *35*, 036305. [CrossRef]
40. Sibley, G.; Matthies, L.; Sukhatme, G. Sliding window filter with application to planetary landing. *J. Field Robot.* **2010**, *27*, 587–608. [CrossRef]
41. Stutz, G.E. Polygonal Scanners: Components, Performance, and Design. In *Handbook of Optical and Laser Scanning*; CRC Press: Boca Raton, FL, USA, 2018; p. 247.
42. Wang, Y.; Lou, Y.; Song, W.; Tu, Z. A Tightly-Coupled Framework for Large-Scale Map Construction with Multiple Non-Repetitive Scanning LiDARs. *IEEE Sens. J.* **2022**, *22*, 3626–3636. [CrossRef]
43. Zhang, J.; Kaess, M.; Singh, S. On degeneracy of optimization-based state estimation problems, 2016. In Proceedings of the 2016 IEEE International Conference on Robotics and Automation (ICRA), Stockholm, Sweden, 16–21 May 2016; pp. 809–816.
44. Taylor, C.J.; Kriegman, D.J. Minimization on the Lie group SO (3) and related manifolds . *Yale Univ.* **1994**, *16*, 6.
45. Cai, Y.; Xu, W.; Zhang, F. ikd-Tree: An Incremental KD Tree for Robotic Applications. *arXiv* **2021**, arXiv:2102.10808.
46. Dellaert, F. *Factor Graphs and GTSAM: A Hands-on Introduction*; Georgia Institute of Technology: Atlanta, GA, USA, 2012.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.