Geometric Consistency–Guided LiDAR Odometry with Adaptive Feature Weighting in Unstructured Outdoor Environments

Narayan Longani[1], Gon-Woo Kim[1*]

[1]Department of Intelligent Systems and Robotics,
Chungbuk National University, Cheongju 28644, South Korea.
narayan@chungbuk.ac.kr, gwkim@cbnu.ac.kr

[*]Corresponding author

Abstract

Reliable 3D LiDAR odometry in complex, naturally unstructured, and perceptually challenging environments remains a critical problem for field robots. This paper presents a geometric-consistency-guided LiDAR odometry framework that assesses feature reliability before optimization and uses the resulting information to weight constraints during incremental pose estimation. First, a distance-adaptive neighborhood selection strategy compensates for range-dependent variations in point density. Local geometric analysis then classifies the feature points as stable or unstable according to their geometric consistency. Finally, constraints from stable and unstable features are assigned different weights within a nonlinear least-squares formulation. The framework is evaluated on self-collected agricultural datasets and public datasets covering botanical gardens, university environments in the GRACO dataset, and mines and caves in the DARPA Subterranean dataset. Evaluations using absolute trajectory error, relative pose error, ablation studies, and processing-frequency analyses at the component and system levels demonstrate competitive accuracy, reduced drift in challenging sequences, and computational performance that keeps pace with the input LiDAR stream. Overall, the results confirm the effectiveness of the proposed feature-reliability strategy for LiDAR odometry in diverse and challenging environments.

Keywords: LiDAR Odometry, Field Robotics, Unstructured Environments, Geometric Consistency, Adaptive Windowing, Feature Reliability, Weighted Optimization

1 Introduction

3D Light Detection and Ranging (LiDAR) based odometry estimation plays a crucial role in autonomous navigation, especially in outdoor unstructured environments such as forest trails, construction sites, and off-road terrains. Unlike structured urban areas, these environments present significant challenges due to irregular surfaces, vegetation clutter, and the lack of consistent geometric landmarks [13]. The varying density and quality of geometric features in such scenes often degrade optimization, leading to inaccurate odometry or even failure in degenerate configurations such as long, featureless paths [18].

Figure 1: Visualization of challenging agricultural field environments encountered during data collection. (A–C) scenes represent wheat fields with varying conditions, including waterlogged and uneven terrain, dry and compacted soil with high mechanical vibrations, and tilled loose soil with pronounced surface irregularities; (D–E) scenes depict rice fields characterized by dense vegetation and waterlogged surfaces. These environments present a significant challenge for consistent and accurate LiDAR odometry.

3D LiDAR based odometry is estimated either using the scan-matching methods [1, 21, 25] or the feature matching based methods [8, 15, 26, 31]. To further improve the estimation process, LiDAR is often fused with inertial sensors like IMU [24, 27, 29, 33]. Extracting robust and reliable features from 3D LiDAR scans is inherently challenging due to varying point density with range, occlusions, and environmental complexity. Unlike structured urban scenes, which are rich in geometric information, unstructured environments often exhibit irregular surfaces, sparse geometric constraints, and dynamic elements such as vegetation [18]. These factors can lead to unstable geometric features that degrade the scan-matching process [15]. Loop-closure detection can reduce accumulated drift in full SLAM systems; however, it does not directly address the reliability of local geometric constraints used for frame-to-map pose estimation [10, 11, 24]. These observations motivate the assessment of geometric feature reliability before the corresponding constraints are incorporated into pose optimization.

To address these challenges, we propose a 3D LiDAR odometry framework that achieves robust performance in outdoor unstructured environments, including uneven agricultural fields, as illustrated in Figure 1, as well as challenging environments such as caves and mines represented in the DARPA Subterranean dataset [20]. Through this work, we make the following contributions:

- A distance-adaptive feature extraction strategy that adjusts the local neighborhood according to LiDAR range to compensate for variations in point density.
- A geometric-consistency-guided feature reliability assessment that classifies extracted feature points as stable or unstable based on local geometric characteristics.
- A unified LiDAR odometry formulation using weighted nonlinear least-squares optimization, where stable and unstable feature constraints are assigned different weights to improve estimation stability under varying feature quality.
- Extensive evaluation on self-collected and public datasets demonstrating improved robustness and reduced drift across diverse unstructured and geometrically challenging environments.

The remainder of this paper is organized as follows: Section 2 reviews related LiDAR odometry methods, and Section 3 presents the proposed framework. Section 4 describes the experimental evaluation and discusses the results. Section 5 analyzes the runtime performance, while Section 6 presents the ablation study. Finally, Section 7 concludes the paper.

2 Related Work

Invariance to lighting conditions of LiDAR makes it favourable for odometry and SLAM operation in diverse operating environments[2, 14]. Scan matching starts with Iterative Closest Point (ICP) [1]. The ICP iteratively establishes correspondences and estimates the transformation between point clouds in Euclidean space. This leads to optimization based on different metrics like point-to-point [1] or point-to-plane [21]. Generalized ICP (GICP) [22] and voxelized GICP (VGICP) [12] incorporate local covariance information into the registration objective. Chen et al.[3] employed the NanoGICP for scan-matching using a local map based on the selected keyframes. Dellenbach et al.[4] introduce their work, Continuous-Time ICP (CT-ICP), which performs motion un-distortion of the point cloud followed by the registration process. Vizzo et al. [25] applied the point-to-point error metric with adaptive thresholding and robust kernel.

Recent methods have introduced adaptive weighting and geometry-aware representations to improve registration robustness. Guo et al.[6] proposed adaptive weighted robust ICP (AW-RICP), which learns sparse correspondence weights from registration errors. Correspondence pairs with larger alignment errors receive lower weights or are removed from the registration. However, AW-RICP focuses on pairwise rigid point-set registration and derives reliability from residual magnitude, rather than evaluating the local geometric reliability of features within a sequential LiDAR odometry pipeline. MAD-ICP [5] introduces a robust LiDAR odometry framework based on a PCA-oriented kd-tree. The kd-tree is used to represent local planar patches, estimate surface normals, support point-to-plane correspondence matching, and manage the local map using pose-information criteria. Although MAD-ICP employs PCA and eigenvalue information, these quantities primarily support its data representation, normal estimation, correspondence search, and local-map management. In contrast, the proposed framework uses local eigenvalue structure to classify the geometric stability of individual feature points before weighted optimization. GenZ-ICP[15] classifies scan-to-map correspondences into planar and non-planar groups according to their local geometric structure. Point-to-plane residuals are applied to planar correspondences, while point-to-point residuals are applied to non-planar correspondences. A global adaptive weighting factor is then determined from the relative numbers of the two correspondence groups. This strategy addresses the environment-dependent limitations of relying on a single ICP error metric and improves numerical robustness in degenerate environments. In contrast, the proposed method retains edge and planar residuals and evaluates the reliability of the extracted features before and during scan-to-map optimization. It does not switch between point-to-point and point-to-plane error metrics; instead, it controls the contribution of stable and unstable geometric features within their corresponding edge and planar residuals.

Feature-based methods extract geometric information from LiDAR measurements and use the resulting constraints for incremental pose estimation. These approaches are computationally efficient with improved performance, particularly in structured environments [7, 13, 24, 29]. LOAM[30] introduced LiDAR odometry based on the extraction and matching of edge and planar features. LeGO-LOAM[23] incorporates ground segmentation to improve estimation for ground vehicles, while LIO-SAM[24] combines edge and planar feature matching with inertial preintegration and factor-graph optimization. LeGO-LOAM[23] incorporates ground segmentation to improve estimation for ground vehicles, while LIO-SAM[24] combines edge and planar feature matching with inertial preintegration and factor-graph optimization. F-LOAM[26] reduces the computational cost of LOAM-style feature extraction and optimization. FAST-LIO and FAST-LIO2 [28, 29] employ tightly coupled iterated Kalman filtering and efficient incremental map structures for LiDAR–inertial estimation. Feature-based methods commonly use fixed-size neighbourhoods of adjacent scan points during feature extraction. However, LiDAR point spacing increases with measurement range. A fixed neighbourhood therefore provides different geometric support in near and distant regions. In sparse, range-varying, or vegetation-rich scenes, this can produce unreliable curvature estimates and unstable edge or planar features.

The closely related methods discussed above operate at different stages of the registration pipeline: they modify the registration error metric, derive correspondence weights from alignment errors, or adapt point-cloud representation and local-map management. However, these methods do not jointly address range-dependent feature extraction and feature-level geometric reliability within a scan-to-map LiDAR odometry framework. To address this gap, the proposed method combines a distance-adaptive neighborhood for feature extraction, an eigenvalue-based geometric-consistency assessment for classifying stable and unstable edge and planar features, and stability-guided residual weighting during scan-to-map optimization. Therefore, the contribution is not adaptive weighting alone, but the coupling of range-adaptive feature extraction, feature-level reliability assessment, and weighted optimization within a unified LiDAR odometry framework.

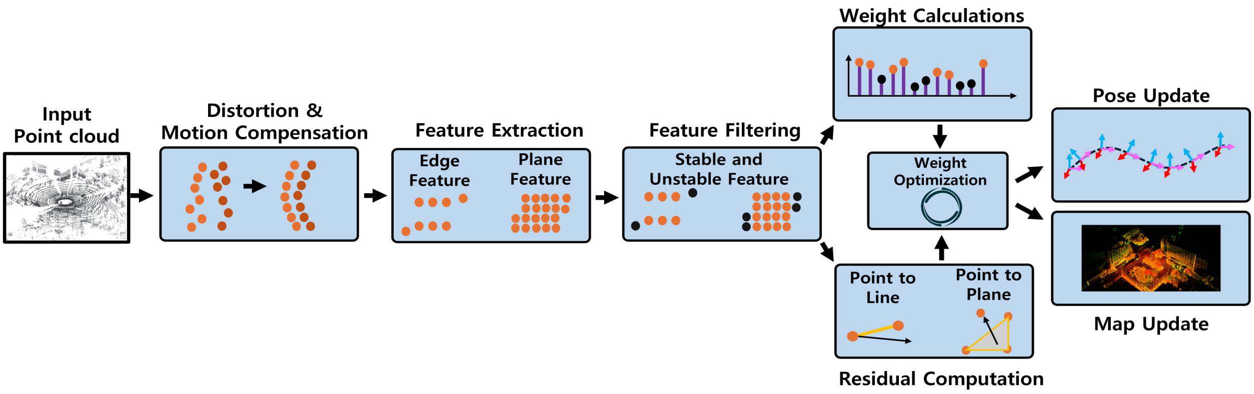

Figure 2: Overview of the proposed Geometric Consistency-Guided and Weighted Feature Points-Based odometry pipeline. The system comprises three main modules: (i) Dynamic-Window based Feature Extraction, which adapts neighbourhood size with range to obtain stable geometric features; (ii) Feature Filtering, which assesses local geometric consistency via eigenvalue analysis; and (iii) Weighted Optimization, which estimates the pose using point-to-edge and point-to-plane residuals modulated by feature stability weights.

3 Geometric-Consistency-Guided 3D LiDAR Odometry

Our method is designed to be both simple and robust, focusing on extracting stable geometric features from sparse point clouds to enable efficient and reliable incremental pose estimation in complex or challenging environments. The proposed framework introduces three main components, illustrated in Figure 2: *a)* dynamic-window-based feature extraction, *b)* geometric-consistency assessment, and *c)* weighted optimization-based pose estimation. The important notations are summarized in Table 1. This hierarchical design ensures that unstructured geometry has limited influence while preserving strong geometric cues, resulting in robust odometry across various environments.

3.1 Preliminaries

The point cloud from the laser sensor is presented as $\mathcal{P}$, which consists of the M parallel lines with N points. The k_{th} laser scan is represented as $\mathcal{P}_k$ and a single point in the scan is represented by $p_k^{m,n}$ where $m \in [1, M]$ and $n \in [1, N]$. The organized scan contains M scan lines with up to N points per line.

Table 1: Nomenclature and Notation Used in the Paper

Symbol	Description
$\mathcal{P}_k$	LiDAR point cloud of the k-th scan
$p_k^{m,n}$	Point at ring m and index n within P_k
$r_{m,n}$	Range of point $\mathbf{p}_{m,n}$
$\mathcal{W}(r_{m,n})$	Distance-adaptive local window size
$\mathcal{W}_{min}$, $\mathcal{W}_{max}$	Minimum and maximum local window size
$c_i^{m,n}$	Curvature or smoothness value of point $p_k^{m,n}$
τ_e, τ_p	Edge and Planar features curvature threshold
$\mathcal{F}^p$, $\mathcal{F}^e$	Sets of edge and planar feature points
$\lambda_1, \lambda_2, \lambda_3$	Eigenvalues of covariance matrix ($\lambda_1 \geq \lambda_2 \geq \lambda_3$)
$\mathcal{L}$	Number of points in local neighbourhood
$\mu, \mathcal{C}$	Mean, & Covariance matrix of local neighbourhood
N_{stable}, N_{total}	Count of stable and total features in a frame
$w_{k,s}^{\mathcal{F}}$, $w_{k,u}^{\mathcal{F}}$	Weights for stable and unstable features points
$\mathbf{T}_k$	Estimated pose at k-th scan

3.2 Dynamic Window-Based LiDAR Feature Extraction

For the latest laser scan $\mathcal{P}_k$, we extract geometric features by calculating the local surface smoothness. Typically, smoothness is computed using a fixed number of neighbouring points; however, in open grounds or in agricultural fields, local surface estimation becomes challenging. We aim to enhance the local surface estimation using the dynamic windows-based approach. Using the proposed approach, the number of neighbour points selected depends on the distance from the laser sensor, as shown in Figure 3. For point $p_i^{m,n}$, the window $\mathcal{W}$ of neighbour points is selected using the eq. (1).

$$\mathcal{W}(r_{m,n}) = \mathcal{W}_{min} + (\mathcal{W}_{max} - \mathcal{W}_{min}) \times \frac{r_{m,n}}{r_{m,n} + \kappa} \tag{1}$$

Whereas $\mathcal{W}_{min}$ is the minimum window size, $\mathcal{W}_{max}$ is the maximum window size, $r_{m,n}$ is the distance from the laser sensor, and κ controls the characteristic transition distance of the adaptive neighborhood. The resulting window size is rounded to the nearest valid integer. The range-dependent term $r_{m,n}/(r_{m,n} + \kappa)$ becomes 0.5 when $r_{m,n} = \kappa$; such that the neighborhood reaches the midpoint between $\mathcal{W}_{min}$ and $\mathcal{W}_{max}$. The same parameter values were used for all

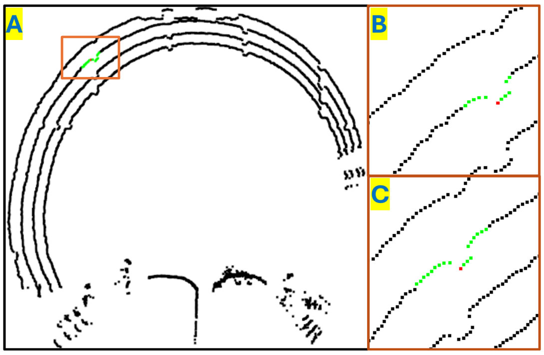

Figure 3: Distance-Adaptive neighbourhood selection in a campus environment. (A) A trimmed top view of the LiDAR point cloud shows the target point (in red) and neighbours (in green), located approximately 46 meters from the sensor. (B) Conventional fixed-window method selects only 5 local neighbours, whereas in (C), Distance-adaptive window selects 9 spatially consistent local neighbours.

evaluated sequences without sequence-specific parameter adjustment. The influence of κ is further examined in Section 4. Now the local surface smoothness $c_i^{m,n}$ for point $p_i^{m,n}$ within the neighborhood $\mathcal{W}$ is then calculated using (2).

$$c_i^{m,n} = \frac{1}{|\mathcal{W}_i^{m,n}|} \sum_{p_i^{(m,j)} \in \mathcal{W}_i^{(m,n)}} \|p_i^{m,j} - p_i^{m,n}\| \tag{2}$$

After this, we will classify the planar features $\mathcal{F}_i^p$ and edge features $\mathcal{F}_i^e$ based on the curvature measure c. For the continuous structures like walls, the c is small, whereas for sharp edges, the value is high. Based on this phenomenon, edge points are selected with higher c values, and planar points with lower values. A point is classified as a planar feature if $c_i < \tau_p$ and as an edge feature if $c_i > \tau_e$, where τ_p & τ_e are the planar and edge curvature thresholds, respectively.

3.3 Feature Point Filtering

Before proceeding to the estimation process, we incorporate a feature-filtering mechanism to enhance performance in complex environments. The distance window-based neighbourhood selection is applied to retain candidate features, some of which may exhibit inconsistent local structure. To assess the reliability of feature points, we analyse their geometric stability using eigenvalues, inspired by [16].

Once the feature cloud $\mathcal{F}_i = \{\mathcal{F}_i^e, \mathcal{F}_i^p\}$ is extracted, it is organized within an incremental voxel grid structure. Then each voxel is examined to assess the consistency of geometric structure. This is achieved by computing the covariance matrix of a point using a local neighbourhood and performing eigenvalue analysis. For a given point $p_i \in \mathcal{F}_i$, a local neighbourhood of window $\mathcal{W} \subset \mathbb{R}^3$ consisting of $\mathcal{L}$ neighbouring points is used to compute the mean μ and covariance $\mathcal{C}$.

$$\mu = \frac{1}{\mathcal{L}} \sum_{p \in \mathcal{W}} p, \quad \mathcal{C} = \frac{1}{\mathcal{L}} \sum_{p \in \mathcal{W}} (p - \mu)(p - \mu)^T \tag{3}$$

The eigenvalues $\{\lambda_1, \lambda_2, \lambda_3\}$ of the covariance matrix, $\mathcal{C}$, reflect the linearity and spread of the point distribution in different directions. Following [16], the stability of a local geometric structure can be quantified by the ratios between consecutive eigenvalues of the covariance matrix. An edge feature is classified as stable when $\lambda_1/\lambda_2 > 6$, while a planar feature is classified as stable when $\lambda_2/\lambda_3 > 6$. Features that do not satisfy the corresponding category-specific criterion are labeled unstable. Figure 4 illustrates this classification in a structured campus environment. Stable planar points (pink) are primarily found on smooth, planar regions such as roads and building facades, offering strong geometric consistency. Stable edge points (blue) are located at sharp edges such as wall corners and signposts, providing high localization reliability. In contrast, unstable planar points (green) typically occur in areas with partial occlusion or reflective surfaces, such as vehicle bodies or trees, where LiDAR returns may be noisy or sparse. Unstable edge points (yellow) tend to appear on geometrically ambiguous structures such as thin poles or highly textured surfaces like parked cars.

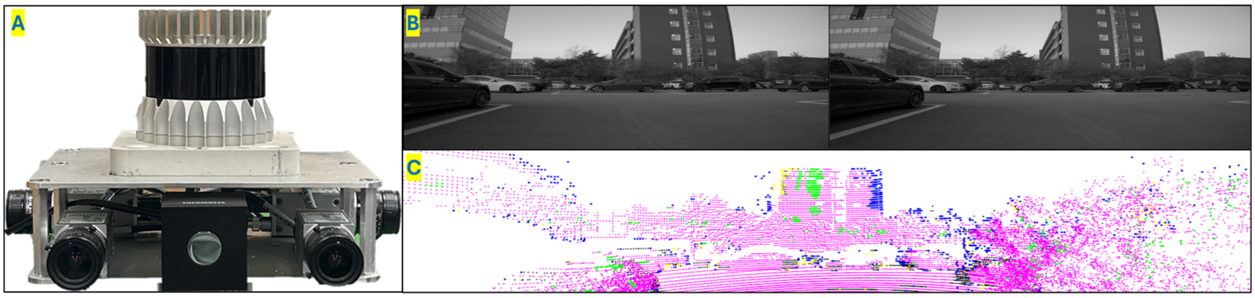

Figure 4: Scene captured using the multi-sensor setup in the university, capturing a road view with cars and buildings. (A) Multi-sensor platform comprising 3D Ouster LiDAR, 4-Basler cameras, Microstrain IMU, and Thermodyne camera. (B) Corresponding visuals captured by Basler cameras at the front. (C) Corresponding feature clouds: stable planar points (pink), unstable planar points (green), stable edge points (blue), and unstable edge points (yellow).

This eigenvalue-based classification enables the system to prioritize stable and geometrically consistent features for scan registration. Also, it significantly enhances the performance in both structured and unstructured environments.

3.4 Distortion Compensation

Given two consecutive laser scans, $\mathcal{P}_{i-1}$ and $\mathcal{P}_i$, the time interval between them is sufficiently small to infer constant velocity motion. This assumption enables initial motion compensation by interpolating the pose during the scan. Once a reliable pose estimate is obtained, the scan is refined by correcting each point's position based on the updated pose. The resulting undistorted points are then used for mapping, improving overall consistency & accuracy. Relative motion between two successive scans is computed as:

$$\zeta_{i-1}^{i} = \log\left(\mathbf{T}_{i-1}^{-1}\mathbf{T}_i\right) \tag{4}$$

Here $\mathbf{T}_i \in \mathrm{SE}(3)$ is the robot pose at i^{th} scan and $\boldsymbol{\zeta}_{i-1}^{i} \in \mathfrak{se}(3)$ is the Lie-algebra representation of the relative motion. The interpolated pose at time τ during the scan duration is:

$$\mathbf{T}_i(\tau) = \mathbf{T}_{i-1}\exp\left(\frac{\tau}{\Delta t}\zeta_{i-1}^{i}\right) \tag{5}$$

where $\tau \in [0, \Delta t]$ is the time elapsed since the start of scan i, and Δt is the total scan duration. Each point $p_i^{m,n}$ in the scan $\mathcal{P}_i$ is undistorted using:

$$\tilde{p}_i^{m,n} = \mathbf{T}_{i-1}^{-1}\mathbf{T}_i(\tau_n)p_i^{m,n}, \qquad p_i^{m,n} \in \mathcal{P}_i. \tag{6}$$

where τ_n is the acquisition time of a point $p^{m,n}$.

3.5 Feature Map

The global feature map $\mathcal{M}_i^G$ consists of the edge feature map $\mathcal{M}_i^e$ and planar feature map $\mathcal{M}_i^p$. The corresponding feature clouds $\{\mathcal{F}_i^e, \mathcal{F}_i^p\}$ are downsampled and transformed into the map frame using the previous pose estimate. This provides the initial alignment $\tilde{\mathcal{F}}_i = \mathbf{T}_{i-1}\mathcal{F}_i$ for scan-to-map correspondence search and subsequent pose refinement. To reduce memory usage, duplicate feature points within the same voxels are removed.

3.6 Feature-Map Correspondences

Now, the latest Feature cloud $\tilde{\mathcal{F}}_i = \{\tilde{\mathcal{F}}_i^p, \tilde{\mathcal{F}}_i^e\}$ is processed to estimate the pose. The estimation is based on the correspondences or alignment generated between the current feature cloud, $\tilde{\mathcal{F}}_i$, and the corresponding feature maps $\mathcal{M}_{i-1} = \{\mathcal{M}_{i-1}^p, \mathcal{M}_{i-1}^e\}$. These correspondences are established through two strategies: point-to-edge and point-to-plane matching. These matching steps yield residuals d_k, which are used in the estimation process.

3.6.1 Point-to-Edge Matching

Now, we will align the edge feature cloud $\tilde{\mathcal{F}}_i^e$ with the edge map cloud $\mathcal{M}_{i-1}^e$ to get point to edge residual, d_{e_k}. For each point $\tilde{p}_{i,k}^e \in \tilde{\mathcal{F}}_i^e$, a local set of nearest neighbouring points is retrieved from the edge map. To verify the local linear structure, the centroid and covariance matrix of the neighbouring points are computed, followed by eigenvalue analysis. The eigenvector corresponding to the largest eigenvalue defines the principal direction $\mathbf{v}$. If the largest eigenvalue is sufficiently larger than the second-largest eigenvalue, the neighbourhood is considered to represent a reliable linear structure.

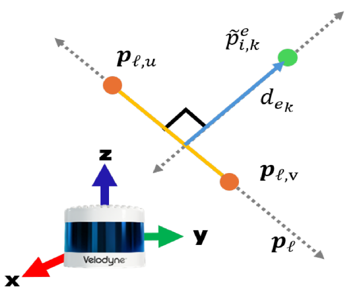

Figure 5: Point-to-edge residual computation. The point $\tilde{p}_{i,k}^e$ in green represents an edge feature from $\tilde{\mathcal{F}}_i^e$. The auxiliary points $\mathbf{p}_{\ell,u}$ and $\mathbf{p}_{\ell,v}$ shown in orange illustrate the fitted local line. The blue arrow indicates the perpendicular distance d_{e_k} from the feature point to the fitted line, representing the point-to-edge residual.

The reference line is therefore defined by its centroid $\mathbf{p}_\ell$ and principal direction $\mathbf{v}$, both estimated from the local edge-map neighbourhood. For geometric representation, two auxiliary points, $\mathbf{p}_{\ell,u}$ and $\mathbf{p}_{\ell,v}$, are constructed along the fitted line as

$$\mathbf{p}_{\ell,u} = \mathbf{p}_\ell + \alpha \mathbf{v}, \qquad \mathbf{p}_{\ell,v} = \mathbf{p}_\ell - \alpha \mathbf{v} \tag{7}$$

where $\alpha > 0$ specifies a small distance from the centroid. The point-to-edge correspondence is represented by the following cross-product residual vector, as illustrated in Figure 5:

$$\mathbf{d}_{e_k} = \mathbf{v} \times \left(\tilde{\mathbf{p}}_{i,k}^{e} - \mathbf{p}_\ell \right) \tag{8}$$

So, the corresponding point-to-edge residual distance is

$$\begin{aligned} \|\mathbf{d}_{e_k}\| &= \left\| \mathbf{v} \times \left(\tilde{\mathbf{p}}_{i,k}^{e} - \mathbf{p}_\ell \right) \right\| \\ &= \frac{\left\| \left(\tilde{\mathbf{p}}_{i,k}^{e} - \mathbf{p}_{\ell,u} \right) \times \left(\tilde{\mathbf{p}}_{i,k}^{e} - \mathbf{p}_{\ell,v} \right) \right\|}{\|\mathbf{p}_{\ell,u} - \mathbf{p}_{\ell,v}\|}. \end{aligned} \tag{9}$$

3.6.2 Point-to-Plane Matching

For the planar feature cloud $\tilde{\mathcal{F}}_i^p$, a similar approach is followed to associate each feature with a local planar region in the planar map $\mathcal{M}_{i-1}^p$ and obtain the point-to-plane residual d_{p_k}. For each point $\tilde{p}_{i,k}^p \in \tilde{\mathcal{F}}_i^p$, a local set of nearest neighbouring points is retrieved from the planar map. The neighbouring map points are used to estimate a local plane through least-squares fitting.

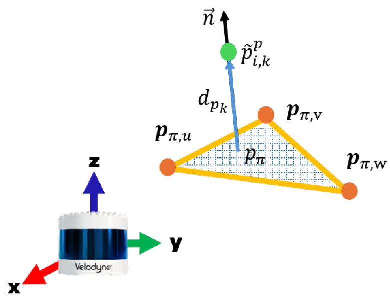

Figure 6: Point-to-plane residual computation. The point $\tilde{p}_{i,k}^p$ in green represents a planar feature from $\tilde{\mathcal{F}}_i^p$. The auxiliary points $\mathbf{p}_{\pi,u}$, $\mathbf{p}_{\pi,v}$, and $\mathbf{p}_{\pi,w}$ shown in orange illustrate the fitted local plane. The blue arrow indicates the perpendicular distance d_{p_k} from the feature point to the fitted plane, representing the point-to-plane residual.

The fitted plane is represented by a unit normal vector $\mathbf{n}$ and plane offset d_π. An anchor point $\mathbf{p}_\pi$ on the fitted plane therefore satisfies $\mathbf{n}^\top \mathbf{p}_\pi + d_\pi = 0$. For an equivalent geometric representation of the fitted plane, let $\mathbf{p}_{\pi,u}$, $\mathbf{p}_{\pi,v}$, and $\mathbf{p}_{\pi,w}$ denote three non-collinear auxiliary points on the fitted plane, such that

$$\mathbf{n}^\top \left(\mathbf{p}_{\pi,r} - \mathbf{p}_\pi \right) = 0, \qquad r \in \{u, v, w\} \tag{10}$$

To ensure geometric consistency, the perpendicular distance of each neighbouring map point $\mathbf{q}_j$ from the fitted plane is evaluated as

$$d_{\pi,j} = \left| \mathbf{n}^\top \left(\mathbf{q}_j - \mathbf{p}_\pi \right) \right| \leq \tau_\pi \tag{11}$$

where τ_π denotes the predefined plane-validation threshold; the correspondence is accepted only when all neighbouring-point deviations satisfy this condition. A unit normal to the same fitted plane, up to its orientation, can equivalently be expressed as

$$\mathbf{n} = \frac{(\mathbf{p}_{\pi,u} - \mathbf{p}_{\pi,v}) \times (\mathbf{p}_{\pi,u} - \mathbf{p}_{\pi,w})}{\|(\mathbf{p}_{\pi,u} - \mathbf{p}_{\pi,v}) \times (\mathbf{p}_{\pi,u} - \mathbf{p}_{\pi,w})\|} \tag{12}$$

The point-to-plane correspondence is represented by projecting the displacement from the fitted plane to the current planar feature onto the unit normal. Figure 6 illustrates the geometric interpretation of the point-to-plane residual, and given as:

$$d_{p_k} = \mathbf{n}^\top \left(\tilde{\mathbf{p}}^p_{i,k} - \mathbf{p}_\pi\right) \tag{13}$$

Similarly, the point-to-plane residual distance will become:

$$\begin{aligned} |d_{p_k}| &= \left|\mathbf{n}^\top \left(\tilde{p}^p_{i,k} - \mathbf{p}_\pi\right)\right| \\ &= \frac{\left|\left(\tilde{p}^p_{i,k} - \mathbf{p}_{\pi,u}\right) \cdot [(\mathbf{p}_{\pi,u} - \mathbf{p}_{\pi,v}) \times (\mathbf{p}_{\pi,u} - \mathbf{p}_{\pi,w})]\right|}{\|(\mathbf{p}_{\pi,u} - \mathbf{p}_{\pi,v}) \times (\mathbf{p}_{\pi,u} - \mathbf{p}_{\pi,w})\|} \end{aligned} \tag{14}$$

3.7 Weighted Optimization

Given the previous pose estimate $\mathbf{T}_{i-1}$ and the geometric residuals $\{d_{e_k}, d_{p_k}\}$ computed from the current scan, the objective is to estimate the current pose $\mathbf{T}_i \in \mathrm{SE}(3)$. The previous pose estimate provides the initial alignment for the current scan-to-map registration. The subsequent pose refinement can be interpreted within a probabilistic framework as a maximum a posteriori (MAP) estimation problem. Specifically, we seek the pose $\mathbf{T}_i$ that maximizes the posterior probability $p(\mathbf{T}_i \mid \tilde{\mathcal{F}}_i)$, where $\tilde{\mathcal{F}}_i$ denotes the current geometric features and $p(\mathbf{T}_i)$ denotes the pose prior.

$$\hat{\mathbf{T}}_i = \arg\max_{\mathbf{T}_i} p(\mathbf{T}_i \mid \tilde{\mathcal{F}}_i) = \arg\max_{\mathbf{T}_i} p(\tilde{\mathcal{F}}_i \mid \mathbf{T}_i)\, p(\mathbf{T}_i) \tag{15}$$

During the scan-to-map refinement, the pose prior is treated as locally non-informative, such that the optimization is primarily determined by the geometric measurement likelihood. Assuming independent measurement noise and Gaussian-like residual distributions, the likelihood term becomes:

$$p(\tilde{\mathcal{F}}_i \mid \mathbf{T}_i) \propto \exp\left(-\frac{1}{2}\sum_k \|\mathbf{d}_k\|^2\right) \tag{16}$$

where $\mathbf{d}_k$ denotes the generic residual associated with the k-th correspondence, corresponding to the vector-valued point-to-edge residual $\mathbf{d}_{e_k}$ or the scalar point-to-plane residual d_{p_k}. Motivated by this probabilistic formulation, geometric reliability is incorporated through feature-dependent weights. The resulting weighted nonlinear least-squares objective is

$$\begin{aligned} \hat{\mathbf{T}}_i &= \arg\min_{\mathbf{T}_i} \sum_k w_k^{\tilde{\mathcal{F}}} \|\mathbf{d}_k\|^2 \\ &= \arg\min_{\mathbf{T}_i} \left(\sum_{\tilde{\mathbf{p}}^p_{i,k} \in \tilde{\mathcal{F}}^p_i} w_k^{\tilde{\mathcal{F}}}\, d^2_{p_k} + \sum_{\tilde{\mathbf{p}}^e_{i,k} \in \tilde{\mathcal{F}}^e_i} w_k^{\tilde{\mathcal{F}}}\, \|\mathbf{d}_{e_k}\|^2 \right) \end{aligned} \tag{17}$$

Here, $w_k^{\tilde{\mathcal{F}}}$ is the weight assigned to the corresponding feature point from the feature cloud $\tilde{\mathcal{F}}$. These weights are computed and applied independently within each feature category. Stable and unstable feature points are first identified using the eigenvalue-based geometric consistency criterion described in Section 3.3. Then to each point, the feature weight $w_k^{\tilde{\mathcal{F}}}$ is assigned as either stable weight $w_s^{\tilde{\mathcal{F}}}$ or the unstable weight $w_u^{\tilde{\mathcal{F}}}$. These weights are determined based on the number of stable feature points, N_{stable}, and the total number of feature points, N_{total}, within the corresponding feature cloud, as described in Section 3.3.

$$w_s^{\tilde{\mathcal{F}}} = \frac{N_{\text{stable}}}{N_{\text{total}} + \varepsilon}, \qquad w_u^{\tilde{\mathcal{F}}} = 1 - w_s^{\tilde{\mathcal{F}}} \tag{18}$$

where $\tilde{\mathcal{F}}_s$ and $\tilde{\mathcal{F}}_u$ denote the stable and unstable feature sets, respectively, and $0 \leq w_s^{\tilde{\mathcal{F}}}, w_u^{\tilde{\mathcal{F}}} \leq 1$. Each feature point is assigned either $w_s^{\tilde{\mathcal{F}}}$ or $w_u^{\tilde{\mathcal{F}}}$ according to its stability label. A small constant ε is added to the denominator in Eq. 18 to avoid division-by-zero cases. This weighting strategy adaptively adjusts the contribution of stable and unstable feature points according to the geometric composition of the current feature cloud. In addition, pose optimization is performed only when a sufficient number of valid features is available in the corresponding feature category, preventing poorly constrained updates under feature-sparse conditions. To solve the optimization problem, we linearize each residual around the current estimate using its Jacobian. For a small pose perturbation $\delta\boldsymbol{\xi} \in \mathbb{R}^6$, with $(\delta\boldsymbol{\xi})^\wedge \in \mathfrak{se}(3)$, the perturbed pose is defined as

$$\mathbf{T}_i \oplus \delta\boldsymbol{\xi} \triangleq \exp\left((\delta\boldsymbol{\xi})^\wedge\right) \mathbf{T}_i. \tag{19}$$

Here, $(\cdot)^\wedge : \mathbb{R}^6 \to \mathfrak{se}(3)$ maps the 6-DoF perturbation vector to its matrix representation in the Lie algebra. Now, using this perturbation model, each residual is locally approximated to first order as

$$\mathbf{d}_k\left(\mathbf{T}_i \oplus \delta\boldsymbol{\xi}\right) \approx \mathbf{d}_k(\mathbf{T}_i) + \mathbf{J}_k \delta\boldsymbol{\xi}. \tag{20}$$

where $\mathbf{J}_k$ is the Jacobian of the k-th residual with respect to the 6-DoF pose perturbation $\delta\boldsymbol{\xi}$. The jacobians for point-to-edge residual and point-to-plane residual are $\mathbf{J}_{e_k}$ and $\mathbf{J}_{p_k}$, respectively:

$$\begin{aligned} \mathbf{J}_{e_k} &= [\mathbf{v}]_\times \begin{bmatrix} \mathbf{I} & -[\tilde{p}^e_{i,k}]_\times \end{bmatrix} \\ \mathbf{J}_{p_k} &= \mathbf{n}^\top \begin{bmatrix} \mathbf{I} & -[\tilde{p}^p_{i,k}]_\times \end{bmatrix} \end{aligned} \tag{21}$$

Here, $[\cdot]_\times$ denotes the skew-symmetric matrix operator for cross product. Substituting the linearized residuals into the cost function yields the following quadratic approximation:

$$\delta\boldsymbol{\xi}^* = \arg\min_{\delta\boldsymbol{\xi}} \sum_k w_k^{\tilde{\mathcal{F}}} \left\|\mathbf{d}_k(\mathbf{T}_i) + \mathbf{J}_k \delta\boldsymbol{\xi}\right\|^2 \tag{22}$$

The expression in eq. 22 defines a quadratic objective in the perturbation $\delta\boldsymbol{\xi}$. Now, the per-residual Jacobians $\mathbf{J}_k$ are stacked to construct the system Jacobian $\mathbf{J}$, while the corresponding residuals d_k are stacked to form the residual vector $\mathbf{r}$. The weighting matrix $\mathbf{W}$ contains geometric-reliability weights $w_k^{\tilde{\mathcal{F}}} \in [0, 1]$; Further substituting, $\mathbf{A} = \sqrt{\mathbf{W}}\mathbf{J}$ & $\mathbf{b} = \sqrt{\mathbf{W}}\mathbf{r}$ the weighted-linear least squares problem becomes

$$\delta\boldsymbol{\xi}^* = \arg\min_{\delta\boldsymbol{\xi}} \|\mathbf{A}\delta\boldsymbol{\xi} + \mathbf{b}\|^2 \tag{23}$$

The linearized cost in (23) leads to the normal equations:

$$\mathbf{H}\delta\boldsymbol{\xi} = -\mathbf{g} \tag{24}$$

where the system matrix (Hessian approximation) $\mathbf{H} = \mathbf{J}^\top\mathbf{W}\mathbf{J}$, and the gradient, $\mathbf{g} = \mathbf{J}^\top\mathbf{W}\mathbf{r}$. Solving this linear system provides the optimal update direction $\delta\boldsymbol{\xi}$ at each iteration. We update the pose using the exponential map:

$$\mathbf{T}_i \leftarrow \exp\left((\delta\boldsymbol{\xi})^\wedge\right) \cdot \mathbf{T}_i \tag{25}$$

The correspondence search, weighted optimization, and pose update are repeated until the convergence criterion is satisfied. This formulation improves robustness under perceptually degraded or feature-sparse conditions while promoting reliable convergence during pose optimization.

4 Experimental Evaluation

The experiments evaluate the accuracy and robustness of the proposed LiDAR odometry framework in naturally unstructured and geometrically challenging environments. The evaluation primarily focuses on outdoor vegetated scenes, while additional public datasets containing botanical gardens, mines, caves, and mixed natural and structural environments are included to examine generalization across different geometric conditions.

The performance is compared with established methodologies, including F-LOAM[26], A-LOAM[30], KISS-ICP[25], and GenZ-ICP[15]. The evaluation was performed on a dedicated computer system featuring an AMD Ryzen 7 5800 8-Core/16-threads Processor running Ubuntu 20.04 with Robot Operating System (ROS). All methods were evaluated on the same input datasets and computing platform to provide a consistent comparison. Publicly available implementations and the default or recommended configurations were used for the baseline methods. Only necessary sensor- and dataset-specific configurations were adjusted, and no method was tuned separately for individual sequences.

We have primarily used the absolute trajectory error (ATE) and relative pose error (RPE) [9] as the primary evaluation metrics. The standard deviation (SD) and average back-end optimization frequency (ABoF) are used as secondary metrics to analyze estimation consistency and computational efficiency. The average back-end optimization frequency is computed as

$$f_{\text{ABoF}} = \frac{1}{N-1}\sum_{i=1}^{N-1}\frac{1}{\Delta t_i} \tag{26}$$

where N denotes the total number of pose updates, and $\Delta t_i = t_{i+1} - t_i$ represents the time interval between two consecutive back-end optimization outputs.

The parameters of the proposed framework were kept fixed across the experiments unless otherwise specified. The minimum and maximum neighborhood sizes were set to $\mathcal{W}_{\min} = 5$ and $\mathcal{W}_{\max} = 15$, respectively, while the transition parameter was fixed at $\kappa = 5$. To examine whether the selected transition parameter introduces strong parameter dependence, a compact sensitivity analysis was conducted on four representative agricultural sequences by varying $\kappa \in \{1, 2.5, 5, 10, 20\}$ while keeping all other parameters unchanged. The corresponding results are summarized in Table 2.

Table 2: Sensitivity analysis of the dynamic-window transition parameter κ using ATE(m) on representative agricultural sequences. Lower values indicate better performance.

Dataset	**Seq.**	**Transition Parameter, κ**				
		1.0	2.5	5.0	10.0	20.0
Rice Field	RF-01	2.0725	2.0559	**2.0552**	2.0606	2.0639
	RF-02	2.4465	2.4458	**2.4421**	2.4426	2.4465
Wheat Field	WF-01	2.9628	2.9564	**2.9277**	2.9712	2.9878
	WF-02	2.9755	2.9041	**2.8740**	2.9553	3.0176

As shown in Table 2, $\kappa = 5$ yields the lowest ATE for all four evaluated sequences. Nevertheless, the variation across the tested settings is relatively small. Compared with $\kappa = 5$, the maximum observed increase in ATE is approximately 5%, while most tested configurations differ by less than 2%. These results indicate relatively limited sensitivity to the transition parameter over the evaluated range and support the use of $\kappa = 5$ as a common choice, without sequence-specific tuning.

The edge curvature threshold was set to 0.1, and the planar curvature threshold was set to 0.05. These values were chosen to balance feature stability and geometric distinctiveness while maintaining consistent performance across multiple datasets. The remaining parameters were kept fixed, except for the environment-specific voxel configurations. For indoor environments, the voxel sizes were set to 0.1 m and 0.2 m for edge and planar features, respectively. For outdoor environments, the corresponding voxel sizes were set to 0.2 m and 0.4 m. The indoor configuration was used for all DARPA Subterranean sequences, whereas the outdoor configuration was used for other datasets. These configurations were fixed for all sequences within their corresponding environmental categories.

4.1 Evaluation on Agri-Field Dataset

We evaluated the proposed framework on self-collected datasets from outdoor agricultural environments, comprising rice and wheat fields, as shown in Figure 1. The details of the self-collected Agri-field datasets are provided in Table 3. These datasets include challenging conditions such as sparse geometric features, unstructured surroundings, and terrain-induced vibrations.

4.1.1 Sensor Setup

The sensor suite, shown in Figure 7, comprises a ZED 2 stereo camera, Velodyne VLP-16 LiDAR, and MicroStrain 3DM-GX5-25 and CV7 IMUs. Additionally, a TDR2000 GPS and a NovAtel PWR 7D GPS were mounted on top of the tractor. In this work, only the Velodyne VLP-16

Table 3: Summary of the Agricultural Field Sequences Used in the Experiments. The Environments Include Rice Field (RF) and Wheat Field (WF) Terrains with Varying Surface Conditions. Scenes of the fields are shown in Figure 1.

Environment	Sequence	Length (m)	Duration (s)
Rice Field	RF-01	203	242.2
	RF-02	210	242.7
	RF-03	158.5	195.5
	RF-04	253	306
	RF-05	211	503
	RF-06	241	537
Wheat Field	WF-01	376	502
	WF-02	291.6	313.5
	WF-03	111.3	143
	WF-04	228.9	278.3

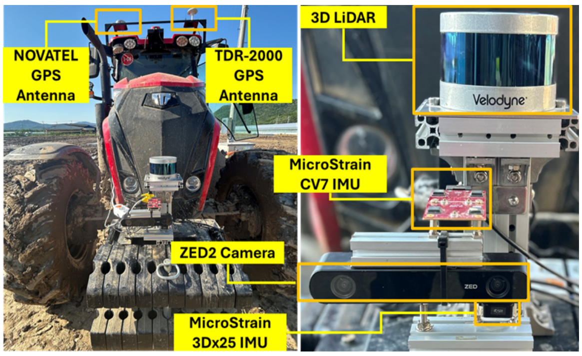

Figure 7: Mini multi-sensor platform rigidly mounted at the front of a tractor for data acquisition in unstructured agricultural fields such as rice and wheat. GPS antennas were mounted on top of the tractor.

measurements were used by the proposed odometry framework, while the NovAtel PWR 7D GPS was used for ground-truth evaluation. The ATE and SD results for the rice and wheat fields are summarized in Table 4, and the odometry results can be visualized in Figure 8a and 8b.

4.1.2 Rice Field

These sequences were collected in waterlogged agricultural terrain with standing water and rice plants of moderate height. The soft, saturated soil forced the tractor to move slowly, while uneven patches occasionally caused it to become momentarily stuck, resulting in significant vibrations and irregular motion.

4.1.3 Wheat Field

These sequences were collected in mixed dry and waterlogged agricultural terrain without standing crops. The combination of hard, dry surfaces and soft saturated sections produced uneven motion

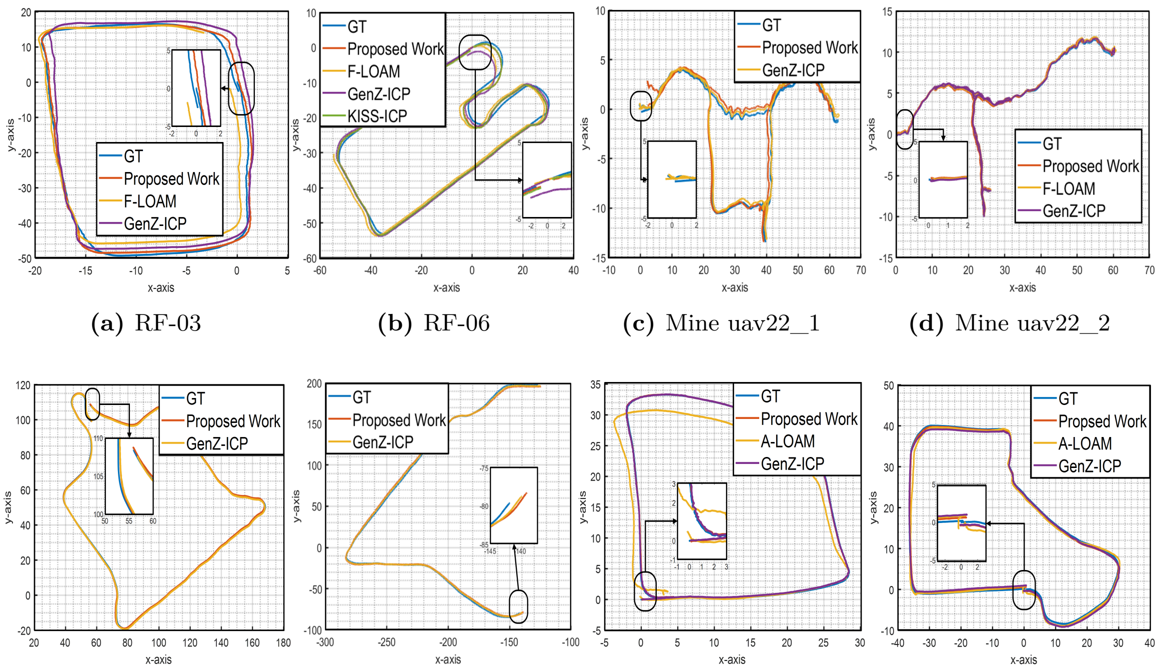

(a) RF-03 **(b)** RF-06 **(c)** Mine uav22_1 **(d)** Mine uav22_2

(e) GRACO Ground 01 **(f)** GRACO Ground 04 **(g)** Botanic Field 1800 **(h)** Botanic Field 1813

Figure 8: Comparison of Proposed Work With Existing Frameworks on Self-Collected Agriculture Dataset and Public dataset as DARPA Subterranean, GRACO, and Botanic Garden dataset. The Starting Point is Presented in Squircle.

and strong vibrations. In several areas, the tractor experienced reduced traction and occasional sticking due to the contrasting ground conditions.

Despite these challenges, the proposed framework maintained accurate and smooth LiDAR odometry, indicating its robustness to varying feature density and terrain-induced motion.

Table 4: ATE and SD-Based Performance Comparison on the Self-Collected Agricultural Dataset Comprising Rice & Wheat Fields. The Comparison is Presented as ATE±SD in Meters (m). The "χ" Indicates Failed odometry (with ATE $\geq$ 11m). The Best ATE is ***Bold and Underlined***; the Second-Best is **bold** only.

Method	Rice Field						Wheat Field			
	RF-01	RF-02	RF-03	RF-04	RF-05	RF-06	WF-01	WF-02	WF-03	WF-04
F-LOAM	χ	χ	10.9260±5.45	χ	χ	χ	χ	χ	**1.0830±0.55**	χ
A-LOAM	10.9930±6.08	χ	χ	χ	χ	χ	χ	χ	2.3770±0.88	8.8090±2.7500
GenZ-ICP	5.8250±3.76	2.5850±0.89	**2.9460±1.17**	**2.3600±0.5**	**6.1180±3.63**	3.1590±1.27	χ	8.7250±3.96	2.1960±0.91	**4.0800±1.69**
KISS-ICP	**2.9750±1.07**	**2.5210±0.81**	3.3240±0.95	2.3620±0.5	6.1820±3.62	**2.7610±1.24**	**2.7680±1.26**	**3.3070±1.10**	1.0840±0.56	4.9730±2.26
Proposed	**2.0552±0.57**	**2.4421±0.45**	**2.437±0.64**	**2.3480±0.5**	**6.033±3.59**	**2.7340±1.24**	**2.9277±0.55**	**2.8740±0.91**	**1.081±0.54**	**3.0447±1.21**

4.2 Evaluation on Public Datasets

To further assess the adaptability of the proposed framework under perceptually degraded, structurally diverse, and sensor-challenging conditions, we extended the evaluation to three public datasets: the DARPA Subterranean dataset, the GRACO dataset, and the Botanic Garden dataset. These datasets represent a broad range of environmental conditions, including underground mines and caves, mixed vegetated and structural environments, and dense natural vegetation. They therefore enable a comprehensive analysis of the proposed framework under diverse real-world conditions.

4.2.1 DARPA Subterranean Dataset

The DARPA Subterranean Challenge (SubT) dataset contains challenging underground environments such as mines, caves, and metro systems. Mine sequences are particularly difficult because dust can degrade the LiDAR measurements [19]. These environments also contain limited geometric features and high localization uncertainty. For evaluation, we have utilized sequences from underground environments to benchmark our framework. The ATE and SD results are summarized in Table 5, and the corresponding visual results are shown in Figure 8c and Figure 8d.

Table 5: ATE and SD-Based Comparison Across DARPA Subterranean Dataset (as the ATE±SD) in Meters (m). The "χ" Indicates Failed LiDAR Odometry (with ATE $\geq$ 11m). The Best ATE is **<u>Bold and Underlined</u>**; the Second-Best is **bold** only.

Sequence	F-LOAM	A-LOAM	GenZ-ICP	KISS-ICP	Proposed
212	**5.012±3.3**	5.02±3.31	6.19±3.32	χ	**<u>5.0±3.3</u>**
213	**1.56±0.93**	1.683±1.043	**<u>1.49±0.94</u>**	2.463±1.11	1.867±1.21
221	8.887±4.15	9.70±5.01	**<u>8.611±5.094</u>**	χ	**8.863±4.99**
222	**5.29±3.3**	5.414±3.6	5.324±3.3	χ	**<u>5.14±3.2</u>**
223	5.22±2.97	χ	**5.212±2.98**	7.4±3.21	**<u>5.211±2.96</u>**

One particularly challenging case, the uav22_1 sequence, spanned approximately 131 meters over 332 seconds in a dust-heavy mine environment. The degraded geometry and dust interference provided a stringent test of system robustness. The resulting point cloud map is shown in Figure 9.

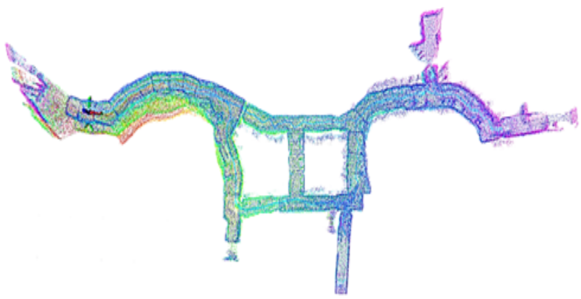

Figure 9: Point cloud map estimated from the uav22_1 sequence of the DARPA Subterranean dataset. The sequence spans approximately 131 meters over 332 seconds in a mine environment, characterized by dust-induced LiDAR degradation.

4.2.2 GRACO Dataset

The GRACO dataset [32] is designed to support research in cooperative localization involving both drones and ground robots. We use the ground-robot sequences, which contain squares, streets, bridges, buildings, and vegetation, to evaluate the adaptability of our framework. The RPE and SD results are summarized in Table 6, and the corresponding visual results are shown in Figure 8e and Figure 8f. The proposed framework achieves the lowest RPE on all evaluated GRACO sequences, indicating strong adaptability to mixed structural and vegetated environments.

4.2.3 Botanic Garden Dataset

The Botanic Garden dataset [17] provides a diverse and complex natural garden environment, ideal for testing proposed systems in unstructured outdoor scenarios. This dataset was collected using a wheeled robot equipped with multiple sensors. The dataset spans over 17 km of trajectories across grasslands, riversides, forested paths, and narrow botanical trails. The Botanic Garden dataset additionally provides measurements from a Livox AVIA MEMS LiDAR, which employs a non-repetitive scanning pattern characteristic of semi-solid-state LiDARs. In this work, however, all experiments were conducted using only the Velodyne VLP-16 data to maintain consistency with the other evaluated datasets and baseline methods. These sequences pose unique challenges due to dense vegetation, terrain irregularities, and repetitive geometry, and weak geometric constraints that can lead to perceptual aliasing. In our experiments, we selected representative sequences with varying vegetation density and terrain complexity to evaluate the proposed system's ability to maintain accurate localization and geometric consistency. The RPE and SD results for selected sequences are summarized in Table 6, and visual results in Figure 8g and 8h. Table 6 provides a rigorous evaluation of the system's performance in unstructured, vegetated environments with variable terrain.

Table 6: RPE (%) and SD-based performance comparison across the GRACO Ground (GO) and Botanic Garden datasets, reported as RPE (%)±SD. The "χ" indicates failed LiDAR odometry (RPE $\geq$ 11%). The best RPE is **<u>bold and underlined</u>**; the second-best is **bold** only.

Seq.	F-LOAM	A-LOAM	GenZ-ICP	KISS-ICP	Proposed
GRACO Ground Dataset					
GO-01	4.149±0.016	1.339±0.006	**0.933±0.006**	χ	**<u>0.502±0.003</u>**
GO-02	2.128±0.012	0.466±0.002	**0.4±0.002**	χ	**<u>0.327±0.002</u>**
GO-03	1.399±0.011	0.702±0.005	**0.33±0.002**	χ	**<u>0.239±0.001</u>**
GO-04	1.883±0.009	0.465±0.002	**0.327±0.002**	χ	**<u>0.240±0.001</u>**
GO-05	3.113±0.011	1.763±0.008	**1.231±0.01**	χ	**<u>0.501±0.003</u>**
Botanic Garden Dataset					
1005-00	χ	**2.272±0.011**	5.434±0.037	χ	**<u>0.944±0.006</u>**
1005-01	χ	3.755±0.019	**3.495±0.02**	χ	**<u>2.998±0.019</u>**
1005-05	χ	**3.326±0.015**	**3.326±0.015**	χ	**<u>2.600±0.016</u>**
1005-07	χ	4.151±0.0197	**4.094±0.0243**	χ	**<u>3.289±0.019</u>**
1006-03	χ	**4.63±0.025**	6.399±0.0368	χ	**<u>3.907±0.025</u>**
1018-00	χ	**<u>1.298±0.008</u>**	1.356±0.008	5.628±0.032	**1.320±0.008**
1018-13	χ	**1.676±0.011**	1.773±0.012	χ	**<u>1.672±0.011</u>**

The comparative results from Tables 4-6 indicate that the benefit of the proposed framework is not limited to mean trajectory accuracy. The complete method remained below the predefined failure threshold in all evaluated sequences and exhibited lower estimation variability in most direct comparisons. Nevertheless, its accuracy was not uniform and could still degrade when the available geometric constraints were weak or insufficient. Overall, these results demonstrate that the proposed framework provides consistent odometry estimation across diverse environmental and geometric conditions while maintaining competitive performance.

5 Run-Time Analysis

5.1 Component-Level Operating Frequency

To examine the computational behavior of the proposed framework, we first evaluate the effective average operating frequency of its major processing stages. The pipeline is divided into two stages: feature extraction and back-end optimization. The feature extraction stage includes distance-adaptive feature extraction and classification of stable or unstable feature points. The back-end optimization stage includes correspondence search, residual construction, and weighted nonlinear optimization. Since correspondence search and optimization are implemented within the same processing module, they are reported together. The average operating frequency of the feature extraction stage is computed in the same manner as the average back-end optimization frequency (ABoF) defined in Section 4. However, instead of using the final optimization output, the output of the feature extraction module is used.

Table 7: Component-level operating frequency (in Hz) of the proposed framework on the self-collected agricultural field datasets. The feature extraction stage includes distance-adaptive feature extraction and classification of stable or unstable features. The back-end optimization stage includes correspondence search, residual construction, and weighted optimization.

Dataset	Seq.	Feature Extraction	Back-End Optimization
Rice Field (15Hz)	RF-01	15.01	15.18
	RF-02	15.03	15.33
	RF-03	15.01	15.38
	RF-04	15.03	15.27
Wheat Field (15Hz)	WF-01	15.00	16.30
	WF-02	15.04	16.36
	WF-03	15.04	15.46
	WF-04	15.03	16.51

Table 7 reports the average component-level operating frequency for the Agri-field datasets. The results demonstrate that both processing stages operate at frequencies comparable to the nominal LiDAR scan rate. This indicates that both modules maintain real-time execution without introducing noticeable module-level processing delay.

5.2 System-Level Back-End Optimization Frequency

After analyzing the average component-level operating frequency, we evaluate the overall runtime performance using the average back-end optimization frequency (ABoF) with baseline methods. The LiDAR scan publication rate serves as a reference for the expected processing rate. As summarized in Table 8, the proposed framework consistently operates at or above the scan rate across all evaluated datasets, indicating the framework processes the incoming measurements without accumulating processing delay. For the Rice Field and Wheat Field datasets, which have an average LiDAR data frequency of 15 Hz, the proposed method achieves an average back-end frequency ranging from approximately 15.18 Hz to 16.46 Hz. Similarly, for the GRACO Ground, Botanic Garden, and DARPA Subterranean datasets, which provide LiDAR measurements at 10 Hz, the average back-end frequency ranges from approximately 10.07 Hz to 11.26 Hz. These results indicate that the additional processing introduced by dynamic windowing, feature classification, and adaptive weighted optimization does not prevent the framework from keeping pace with incoming LiDAR measurements.

Table 8: ABoF (Hz) based Comparison of the Proposed and the SOTA frameworks Across All Datasets. The Dataset Title indicates the LiDAR Scan Rate (in Hz). The Failed Optimization Cases Are Marked With χ. The Best Rate Is **<u>Bold and Underlined</u>**; the Second-Best is **Bold** only.

Method	**Rice Field(15)**			**Wheat Field(15)**			**GRACO Ground(10)**			**Botanic Garden(10)**			**DARPA Subt.(10)**		
	RF-01	RF-02	RF-03	WF-01	WF-02	WF-03	GO-01	GO-02	GO-03	1005-00	1005-01	1005-05	221	222	223
F-LOAM	11.92	12.94	12.86	12.06	12.76	14.82	6.42	7.11	8.33	χ	χ	χ	10.02	10.03	10.05
A-LOAM	14.86	14.85	14.76	15.05	15.1	15.09	10.02	10.01	10	9.94	9.94	9.95	10.03	10.03	10.1
GenZ-ICP	**<u>15.56</u>**	**15.32**	**15.19**	**<u>18.79</u>**	**<u>19.1</u>**	**<u>20.81</u>**	**10.47**	**10.3**	**10.35**	**9.97**	**10.33**	**10.26**	**<u>10.24</u>**	**10.35**	**<u>10.5</u>**
KISS-ICP	15.11	14.78	14.79	15.31	16.07	15	10	10	10	9.92	9.92	9.92	10.01	10.03	10.03
Proposed	**15.18**	**<u>15.33</u>**	**<u>15.38</u>**	**16.3**	**16.37**	**16.46**	**<u>11.26</u>**	**<u>11.08</u>**	**<u>10.82</u>**	**<u>10.77</u>**	**<u>10.94</u>**	**<u>10.63</u>**	**10.07**	**<u>10.66</u>**	**10.16**

The proposed method achieves competitive results, and its performance naturally varies across datasets, reflecting the diversity of geometric and environmental conditions. Although GenZ-ICP [15] achieves a higher back-end optimization frequency on some sequences, the proposed framework maintains an effective processing frequency comparable to or above the nominal LiDAR scan rate across all evaluated datasets.

6 Ablation Studies

To assess the effectiveness of the proposed developments, we conduct an ablation study using four representative variants of the framework. (i) *Base Method (BM):* a standard LOAM-style implementation using a fixed-size neighbourhood and uniform residual weighting. (ii) *Baseline + Weighted Optimization (BWO):* stability-guided feature weighting is added while retaining fixed-window feature extraction. (iii) *Dynamic Windowing (DW):* distance-adaptive windowing is added to BM while retaining uniform residual weighting. (iv) *Dynamic Windowing + Weighted*

Optimization (DW+WO): using the overall proposed system. Experiments are conducted on both the self-collected agricultural datasets and the Botanic Garden dataset. For comparison, we report the ATE as the primary accuracy metric, while the ABoF is reported in parentheses to evaluate computational responsiveness in Table 9. The LiDAR scan rate is indicated with the dataset as the reference for the processing rate.

Table 9: Ablation Study Reporting the ATE (m) and ABoF (Hz) Across the Agri-Field and Botanic Garden Sequences. The LiDAR Scan Rate is Indicated with the Dataset Type. The "χ" Indicates Failed LiDAR Odometry (with ATE $\geq$ 11m). The Best ATE is **Bold and Underlined**; the Second-Best is **Bold** only.

Seq.	BM	BWO	DW	DW+WO
Rice Field Dataset (15Hz)				
RF-01	2.649(15.11)	**2.642(14.81)**	2.651(15.15)	**2.559(15.18)**
RF-02	2.457(15.05)	**2.451(14.81)**	2.456(15.08)	**2.441(15.32)**
RF-03	2.464(15.23)	**2.456(14.82)**	2.463(15.16)	**2.453(15.38)**
RF-04	2.396(15.24)	2.365(14.82)	**2.354(15.18)**	**2.348(15.26)**
Wheat Field Dataset (15Hz)				
WF-01	χ(14.98)	**3.143(14.81)**	χ(14.93)	**3.089(16.3)**
WF-02	2.922(14.98)	2.993(14.80)	**2.91(14.95)**	**2.874(16.36)**
WF-03	1.090(14.97)	1.137(14.81)	**1.082(14.94)**	**1.081(16.46)**
WF-04	4.01(15.05)	**3.939(14.80)**	**3.987(14.96)**	**3.045(16.25)**
Botanic Garden (10Hz)				
1005-00	χ(8.71)	**1.118(10.06)**	χ(8.62)	**2.176(10.77)**
1005-01	χ(8.66)	**8.774(10.00)**	χ(8.71)	**9.125(10.93)**
1005-05	χ(8.48)	χ(9.53)	χ(8.42)	**4.252(10.62)**
1005-07	χ(8.97)	**3.816(9.95)**	χ(8.86)	**7.775(10.6)**

Across the agricultural field sequences, BWO generally improves upon the baseline method, indicating that stability-guided weighting enhances robustness by reducing the influence of unreliable geometric features. DW provides additional refinement through adaptive neighbourhood selection. The complete DW+WO framework consistently achieves the best overall performance while maintaining processing frequencies near or above the nominal LiDAR scan rate.

A different trend is observed on the Botanic Garden dataset. Both BM and DW fail on all four evaluated sequences and operate below the nominal LiDAR rate of 10 Hz. BWO substantially improves robustness, recovering successful trajectories in three sequences and maintaining frequencies close to the scan rate. However, it fails on sequence 1005-05. DW+WO is the only configuration that successfully estimates all four trajectories while maintaining ABoF above 10 Hz. These results indicate that stability-guided weighting is particularly important in this environment, while its combination with dynamic windowing provides the most consistent overall performance.

Overall, the ablation study validates that the proposed developments reduce trajectory error while maintaining consistent processing frequency, thereby improving both accuracy and computational reliability in challenging environments.

7 Conclusion

In this work, we presented a geometric-consistency-guided 3D LiDAR odometry framework for unstructured and geometrically challenging environments. The method combines distance-adaptive feature extraction, eigenvalue-based stability classification, and adaptive feature weighting. Evaluations on both self-collected agri-field datasets and multiple public benchmarks demonstrate that the proposed method achieves competitive accuracy and reduced drift across diverse environmental conditions. The complete framework remained below the predefined failure threshold across all evaluated sequences, although performance can still degrade under extremely sparse or repetitive geometric conditions. High-rate IMU measurements could further improve motion prediction and intra-scan distortion correction, thereby providing more reliable initialization for scan-to-map registration. Since inertial constraints do not directly assess the geometric reliability of individual LiDAR constraints, the proposed feature-reliability mechanism may remain complementary within a LiDAR–Inertial framework; however, this requires dedicated experimental validation. Future work will investigate LiDAR–inertial fusion, improved feature detection, and loop-closure integration for SLAM.

Data Availability

The authors do not have permission to share data.

Conflict of Interest

The authors have no conflicts of interest to declare that are relevant to the content of this article.

Acknowledgments

This research was supported in part by the National Research Foundation (NRF) funded by the Korean government (MSIT) (RS-2024-00421129) and in part by the Innovative Human Resource Development for Local Intellectualization program through the Institute of Information & Communications Technology Planning & Evaluation(IITP) grant funded by the Korean government (MSIT)(IITP-2026-RS-2020-II201462)

References

[1] Besl, P. J. and McKay, N. D. (1992). Method for registration of 3-d shapes. In *Sensor fusion IV: control paradigms and data structures*, volume 1611, pages 586–606. Spie.

[2] Cadena, C., Carlone, L., Carrillo, H., Latif, Y., Scaramuzza, D., Neira, J., Reid, I., and Leonard, J. J. (2016). Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on robotics*, 32(6):1309–1332.

[3] Chen, K., Lopez, B. T., Agha-mohammadi, A.-a., and Mehta, A. (2022). Direct lidar odometry: Fast localization with dense point clouds. *IEEE Robotics and Automation Letters*, 7(2):2000–2007.

[4] Dellenbach, P., Deschaud, J.-E., Jacquet, B., and Goulette, F. (2022). Ct-icp: Real-time elastic lidar odometry with loop closure. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 5580–5586. IEEE.

[5] Ferrari, S., Di Giammarino, L., Brizi, L., and Grisetti, G. (2024). Mad-icp: It is all about matching data–robust and informed lidar odometry. *IEEE Robotics and Automation Letters*.

[6] Guo, Y., Zhao, L., Shi, Y., Zhang, X., Du, S., and Wang, F. (2022). Adaptive weighted robust iterative closest point. *Neurocomputing*, 508:225–241.

[7] He, D., Li, H., and Yin, J. (2025). Ligo: A tightly coupled lidar-inertial-gnss odometry based on a hierarchy fusion framework for global localization with real-time mapping. *IEEE Transactions on Robotics*.

[8] Hyeonjae, G., Dongjae, L., Kwanhyung, S., Seungwook, A., and Ayoung, K. (2023). Gnss-lidar-inertial state estimation system for mapping and autonomous navigation in unstructured environments. *Journal of Robotics Society*, 18(1):72–81.

[9] Kanso, H., Singh, A., El Zarif, E., Almohammed, N., Mounsef, J., Maalouf, N., and Arain, B. (2025). Semantic slam: A comprehensive survey of methods and applications. *Intelligent Systems with Applications*, page 200591.

[10] Khan, E. U. and Kim, G.-W. (2026). Efficient 3D LiDAR pattern feature extraction and descriptor computation for robust loop closure detection. *Intell. Serv. Robot.*, 19(2).

[11] Kim, G. and Kim, A. (2018). Scan context: Egocentric spatial descriptor for place recognition within 3d point cloud map. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4802–4809. IEEE.

[12] Koide, K., Yokozuka, M., Oishi, S., and Banno, A. (2021). Voxelized gicp for fast and accurate 3d point cloud registration. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11054–11059. IEEE.

[13] Le, T., Gjevestad, J. G. O., and From, P. J. (2019). Online 3d mapping and localization system for agricultural robots. *IFAC-PapersOnLine*, 52(30):167–172.

[14] Lee, D., Jung, M., Yang, W., and Kim, A. (2024a). Lidar odometry survey: recent advancements and remaining challenges. *Intelligent Service Robotics*, 17(2):95–118.

[15] Lee, D., Lim, H., and Han, S. (2024b). Genz-icp: Generalizable and degeneracy-robust lidar odometry using an adaptive weighting. *IEEE Robotics and Automation Letters*.

[16] Li, J., Zeng, D., Luo, Q., Luo, X., Chen, C. P., and Yang, C. (2025). Feature assessment and enhanced vertical constraint lidar odometry and mapping on quadruped robot. *IEEE Transactions on Instrumentation and Measurement*.

[17] Liu, Y., Fu, Y., Qin, M., Xu, Y., Xu, B., Chen, F., Goossens, B., Sun, P. Z., Yu, H., Liu, C., et al. (2024). Botanicgarden: A high-quality dataset for robot navigation in unstructured natural environments. *IEEE Robotics and Automation Letters*, 9(3):2798–2805.

[18] Longani, N. and Kim, G.-W. (2025). A robust lidar-inertial multi constraint-based localization for agricultural environments. *IEEE Robotics and Automation Letters*, 10:9638–9645.

[19] Parsons, T., Seo, J., Kim, B., Lee, H., Kim, J.-C., and Cha, M. (2024). Dust de-filtering in lidar applications with conventional and cnn filtering methods. *IEEE Access*, 12:22032–22042.

[20] Rouček, T., Pecka, M., Čížek, P., Petříček, T., Bayer, J., Šalanskỳ, V., Heřt, D., Petrlík, M., Báča, T., Spurnỳ, V., et al. (2020). Darpa subterranean challenge: Multi-robotic exploration of underground environments. In *Modelling and Simulation for Autonomous Systems: 6th International Conference, MESAS 2019, Palermo, Italy, October 29–31, 2019, Revised Selected Papers 6*, pages 274–290. Springer.

[21] Rusinkiewicz, S. and Levoy, M. (2001). Efficient variants of the icp algorithm. In *Proceedings third international conference on 3-D digital imaging and modeling*, pages 145–152. IEEE.

[22] Segal, A., Haehnel, D., and Thrun, S. (2009). Generalized-icp. In *Robotics: science and systems*, volume 2, page 435. Seattle, WA.

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

[24] Shan, T., Englot, B., Meyers, D., Wang, W., Ratti, C., and Rus, D. (2020). Lio-sam: Tightly-coupled lidar inertial odometry via smoothing and mapping. In *2020 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 5135–5142. IEEE.

[25] Vizzo, I., Guadagnino, T., Mersch, B., Wiesmann, L., Behley, J., and 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*, 8(2):1029–1036.

[26] Wang, H., Wang, C., Chen, C.-L., and Xie, L. (2021). F-loam: Fast lidar odometry and mapping. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4390–4396. IEEE.

[27] Wu, Y., Guadagnino, T., Wiesmann, L., Klingbeil, L., Stachniss, C., and Kuhlmann, H. (2024). Lio-ekf: High frequency lidar-inertial odometry using extended kalman filters. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 13741–13747. IEEE.

[28] Xu, W., Cai, Y., He, D., Lin, J., and Zhang, F. (2022). Fast-lio2: Fast direct lidar-inertial odometry. *IEEE Transactions on Robotics*, 38(4):2053–2073.

[29] Xu, W. and Zhang, F. (2021). Fast-lio: A fast, robust lidar-inertial odometry package by tightly-coupled iterated kalman filter. *IEEE Robotics and Automation Letters*, 6(2):3317–3324.

[30] Zhang, J. and Singh, S. (2014). Loam: Lidar odometry and mapping in real-time. *Robotics: Science and systems*, 2(9):1–9.

[31] Zhang, J. and Singh, S. (2017). Low-drift and real-time lidar odometry and mapping. *Autonomous robots*, 41:401–416.

[32] Zhu, Y., Kong, Y., Jie, Y., Xu, S., and Cheng, H. (2023). Graco: A multimodal dataset for ground and aerial cooperative localization and mapping. *IEEE Robotics and Automation Letters*, 8(2):966–973.

[33] Zou, Z., Yuan, C., Xu, W., Li, H., Zhou, S., Xue, K., and Zhang, F. (2024). Lta-om: Long-term association lidar–imu odometry and mapping. *Journal of Field Robotics*, 41(7):2455–2474.