applied sciences

MDPI

Article

HERO-LO: Heterogeneous-Odometry-Enhanced Robust LiDAR Odometry with Observable-Subspace Correction for Tunnel-like Environments

Dongjie Wu [1,2], Mingyan Liang [1] and Xianyi Chen [3,*]

[1] School of Marine Engineering, Jimei University, Xiamen 361021, China; wudongjiechn@jmu.edu.cn (D.W.); 202512858018@jmu.edu.cn (M.L.)
[2] School of Aerospace Engineering, Xiamen University, Xiamen 361102, China
[3] School of Mechanical and Automotive Engineering, Xiamen University of Technology, Xiamen 361024, China
* Correspondence: chenxy@xmut.edu.cn

Abstract

Tunnel-like environments remain challenging for Light Detection and Ranging (LiDAR) odometry because repetitive geometry and weak axial constraints can lead to degeneracy, accumulated drift, and unstable local pose estimates. This paper presents HERO-LO, a heterogeneous-odometry-enhanced robust LiDAR odometry framework with observable-subspace correction for tunnel-like environments. The proposed method detects LiDAR degeneracy online in the tangent space of $SO(3) \times \mathbb{R}^3$ by interpreting the Hessian eigenstructure from a state-estimation uncertainty perspective. When the LiDAR geometry is well constrained, HERO-LO uses scan-to-scan LiDAR odometry as the motion prior and applies LiDAR matching corrections normally. When tunnel-like degeneracy is detected, a heterogeneous odometry source, such as kinematic odometry, attitude-and-heading-reference-system (AHRS)-aided kinematic odometry, or visual-inertial odometry, provides the motion prior, while LiDAR corrections are applied only in the observable subspace and the prior is preserved in degenerate directions. The correction law is further interpreted as a projection-gain observer, and a Linear Matrix Inequality (LMI)-based Lyapunov analysis establishes its per-direction bounded error propagation and stability under mode switching. Experiments on a real tunnel dataset show that, within the degenerate tunnel segment, the default HERO-LO configuration with an encoder-based kinematic odometry (KO) prior reduces the local anomaly rate to 0.6% and achieves an endpoint horizontal error of approximately 3.6 m, comparable to the best open-source LiDAR odometry (LO), LiDAR-inertial odometry (LIO), and LiDAR-visual-inertial odometry (LVIO) methods given the approximately 2.5 m circular error probable (CEP) uncertainty of the single-point Global Navigation Satellite System (GNSS) reference. Runtime analysis shows that HERO-LO processes each LiDAR frame in approximately 9.5 ms on average, satisfying the real-time requirement of the 20 Hz LiDAR. Ablation results further confirm that robust tunnel traversal requires both a stable heterogeneous prior and observable-subspace correction.

Keywords: LiDAR odometry; tunnel-like environments; degeneracy; observable-subspace correction; heterogeneous odometry prior; observer design; LMI; local trajectory stability

Check for updates

Academic Editor: Luigi Fortuna

Received: 17 July 2026
Revised: 17 October 2026
Accepted: 18 October 2026
Published: 23 October 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

Accurate and robust ego-motion estimation is a fundamental capability for autonomous vehicles, mobile robots, and intelligent transportation systems [1,2]. Among var-

ious sensing modalities, Light Detection and Ranging (LiDAR) has been widely used for odometry and mapping because it provides direct 3D geometric measurements and is relatively insensitive to illumination changes compared with cameras. In recent years, LiDAR odometry and LiDAR SLAM (simultaneous localization and mapping) have achieved remarkable progress in urban driving, indoor navigation, underground inspection, and field robotics applications [3,4]. Most existing LiDAR odometry methods estimate the relative motion between consecutive scans or between the current scan and a local map by minimizing geometric residuals, such as point-to-line, point-to-plane, or point-to-distribution distances [5–8]. When sufficient geometric structures are available, these methods can provide accurate and stable pose estimation.

However, the performance of LiDAR scan matching strongly depends on the geometric constraints provided by the environment. In tunnel-like and corridor-like scenes, the surrounding structures are often highly repetitive and provide insufficient constraints along certain motion directions [9]. For example, a road tunnel may provide strong constraints in the lateral and vertical directions through the walls and ground, but weak constraints along the tunnel axis. In such cases, the LiDAR registration problem becomes ill-conditioned, and the resulting pose update may contain unreliable components in the degenerate directions [10]. Directly applying the full LiDAR matching increment can, therefore, introduce erroneous corrections, leading to large accumulated drift, trajectory distortion, or even tracking failure. This issue is particularly critical in long tunnels and mountainous road environments, where high-precision Global Navigation Satellite System (GNSS) is usually unavailable and external absolute positioning cannot be reliably used.

A straightforward way to improve robustness is to introduce additional odometry sources, such as wheel or encoder odometry, visual-inertial odometry (VIO), inertial odometry, or other platform-dependent motion estimates [10–12]. These heterogeneous odometry sources can provide useful motion priors when LiDAR geometry becomes weak. Nevertheless, simply replacing the initial guess of LiDAR registration with an external odometry prediction is not sufficient to solve the degeneracy problem [13]. Even with a better initial guess, the final pose may still be corrupted if unreliable LiDAR updates in the degenerate directions are fully applied. Conversely, suppressing degenerate LiDAR updates without a sufficiently reliable motion prior cannot recover the unobservable motion components. Therefore, robust pose estimation in tunnel-like environments requires both a reliable motion prior and a direction-aware LiDAR correction strategy [14].

The key observation of this paper is that LiDAR degeneracy is usually direction-dependent rather than globally invalid. Even in a degenerate tunnel scene, LiDAR measurements may still provide reliable constraints in some directions while being unreliable in others. Therefore, LiDAR should not be completely discarded when degeneracy occurs. Instead, the LiDAR matching result should be used only in the observable directions, while the pose components in the degenerate directions should be maintained by a heterogeneous odometry prior. Based on this idea, this paper proposes HERO-LO, a Heterogeneous-Odometry-Enhanced Robust LiDAR Odometry framework with observable-subspace correction for tunnel-like environments. HERO-LO is designed to exploit the high local accuracy of LiDAR scan matching in well-constrained directions and the motion continuity of heterogeneous odometry priors in LiDAR-degenerate directions.

HERO-LO adopts a feature-based LiDAR odometry front end and supports different types of odometry priors through a unified six-dimensional (6D) pose interface. In non-degenerate environments, the system can use a scan-to-scan estimate as the motion prior and mainly rely on LiDAR matching for accurate pose estimation. When a tunnel-like geometric degeneracy is detected online, HERO-LO switches to a non-LiDAR heterogeneous odometry prior, such as encoder odometry or VIO, to maintain the pose estimate in weakly

constrained directions. After LiDAR feature matching, the optimization increment is not directly applied in full. Instead, the local observability of the LiDAR registration problem is analyzed, and the LiDAR increment is projected onto the observable subspace. The final pose is obtained by applying only the observable LiDAR correction to the prior pose, while the degenerate components remain governed by the heterogeneous odometry prediction.

To identify the observable and degenerate directions, HERO-LO revisits the classical Hessian-based degeneracy detection strategy used in LOAM-type methods. In the traditional Zhang-style degeneracy handling approach, the eigenvalues of the Hessian matrix are analyzed in a local linearized space, and directions associated with small eigenvalues are regarded as degenerate [9]. This paper further interprets this strategy from the perspective of state estimation uncertainty. Under the Gaussian approximation of nonlinear least squares, the inverse of the Hessian matrix is proportional to the posterior covariance of the pose perturbation [15]. Therefore, small Hessian eigenvalues correspond to large principal axes of the covariance ellipsoid, indicating high state uncertainty and poor observability in the corresponding directions. Based on this interpretation, HERO-LO measures the degree of degeneracy using the principal-axis lengths of the covariance 1σ ellipsoid in the tangent space of $SO(3) \times \mathbb{R}^3$. This provides a more intuitive uncertainty-aware explanation of LiDAR degeneracy and enables pose-manifold-consistent observable-subspace correction.

The proposed framework is evaluated using a real-world road-tunnel dataset collected by a ground mobile robot platform. To comprehensively evaluate the effectiveness and robustness of HERO-LO, we compare it with representative open-source LiDAR-only, LiDAR-inertial, and LiDAR-visual-inertial odometry methods, including A-LOAM [16], LeGO-LOAM [17], F-LOAM [18], CT-ICP [19], DLO [6], KISS-ICP [7], GenZ-ICP [20], LIO-SAM [21], FAST-LIO2 [5], and LVI-SAM [22]. In addition, several ablation variants are designed to analyze the influence of the heterogeneous odometry prior and the degeneracy handling strategy. The experimental results show that HERO-LO can maintain stable trajectory estimation in tunnel-degenerate environments and clearly improves over the variants without degeneracy handling.

The main contributions of this paper are summarized as follows:

- A heterogeneous-odometry-enhanced feature-based LiDAR odometry framework, HERO-LO, is proposed for tunnel-like degenerate environments. The framework uses scan-to-scan LiDAR odometry as the motion prior in non-degenerate scenes, and preserves a non-LiDAR heterogeneous odometry prior in degenerate directions when tunnel-like degeneracy is detected. Different priors, such as encoder-based kinematic odometry (KO), attitude and heading reference system (AHRS)-aided kinematic odometry (KAO), and visual-inertial odometry (VIO), are integrated through a unified 6D pose interface.
- An uncertainty-aware degeneracy detection and observable-subspace correction strategy is formulated on $SO(3) \times \mathbb{R}^3$. The Hessian eigenstructure is interpreted from the perspective of posterior pose uncertainty, where small Hessian eigenvalues correspond to long principal axes of the covariance 1σ ellipsoid. LiDAR matching increments are then projected onto the observable subspace to avoid unreliable corrections in degenerate directions. The correction law is further interpreted as a projection-gain observer, and its per-direction error propagation and Lyapunov stability are established through a Linear Matrix Inequality (LMI)-based analysis.
- A real-world road-tunnel dataset is collected and used for evaluation against representative open-source LiDAR odometry (LO), LiDAR-inertial odometry (LIO), and LiDAR-visual-inertial odometry (LVIO) methods on a unified computing platform.

The comparison and ablation results show that HERO-LO improves local stability in tunnel-degenerate scenes, and that robust performance requires both a stable heterogeneous odometry prior and observable-subspace correction.

The remainder of this paper is organized as follows. Section 2 reviews related work on LiDAR odometry, multi-sensor fusion, degeneracy handling, and tunnel-like environments. Section 3 presents the HERO-LO framework and its main modules. Section 4 reports the experimental setup, comparisons, ablation studies, and runtime analysis. Section 5 discusses the effectiveness, deployability, and limitations of HERO-LO. Section 6 concludes this paper.

2. Related Work

2.1. LiDAR Odometry and LiDAR SLAM

LiDAR odometry estimates sensor motion by registering consecutive point clouds or by aligning the current scan with a local or global map. Classical scan registration methods are mainly derived from iterative closest point (ICP) [23] and its variants, such as point-to-plane ICP [24] and generalized ICP (GICP) [25]. These methods directly minimize geometric distances between corresponding points, planes, or local distributions. They are simple and effective in geometrically rich environments, but their performance can degrade when the surrounding structure provides insufficient constraints in some directions.

LOAM [26–28] is one of the most influential LiDAR odometry and mapping frameworks. It extracts edge and planar features from LiDAR scans and estimates motion by minimizing point-to-line and point-to-plane residuals. To achieve real-time performance, LOAM separates high-frequency odometry and low-frequency mapping. Many subsequent methods follow this feature-based design. LeGO-LOAM [17] improves efficiency for ground vehicles by using ground segmentation and a two-step optimization strategy. F-LOAM [18] further simplifies and accelerates the LOAM pipeline. These feature-based methods are computationally efficient and work well in structured environments, but their accuracy depends on the quality and distribution of extracted geometric features. In tunnel-like or corridor-like environments, edge and planar features may be abundant but directionally uninformative, which can still lead to an ill-conditioned registration problem.

Recent LiDAR and LiDAR-inertial methods have also explored dense or direct registration strategies. CT-ICP [19] formulates LiDAR odometry in continuous time and elastically compensates scan distortion during scan-to-map registration, achieving strong performance on multiple datasets. DLO [6] uses dense point clouds with an efficient keyframe-based submapping strategy and a lightweight GICP solver. KISS-ICP [7] shows that point-to-point ICP, when combined with adaptive thresholding, robust kernels, motion compensation, and appropriate downsampling, can achieve competitive performance with a simple pipeline. FAST-LIO2 [5] directly registers raw LiDAR points to an incremental map maintained by an ikd-tree in a tightly coupled LiDAR-inertial framework. More recently, GenZ-ICP [20] adaptively combines point-to-plane and point-to-point error metrics according to local geometric characteristics, improving robustness in corridor-like degenerative cases.

These methods have significantly advanced LiDAR odometry in terms of accuracy, efficiency, and generality. Nevertheless, tunnel-like environments remain challenging because the local LiDAR registration problem can become weakly constrained in certain pose directions. In a long corridor or tunnel, the surrounding geometry often provides strong constraints in directions perpendicular to the tunnel axis, while the constraint along the direction of travel or other symmetry-related directions can be much weaker. General-purpose LO and LIO systems may still remain stable in such scenarios by benefiting from robust scan matching, inertial prediction, local mapping, or multi-sensor constraints, as also

observed in our experiments. However, many of them do not explicitly decompose the pose update into LiDAR-observable and LiDAR-degenerate components. While GenZ-ICP improves the conditioning of LiDAR registration by adaptively combining different ICP error metrics, HERO-LO addresses degeneracy from a complementary perspective: it preserves a heterogeneous odometry prior in weakly observable directions and applies LiDAR corrections only in the observable subspace. Therefore, explicitly identifying and treating weakly observable directions is a desirable mechanism for improving the interpretability and robustness of tunnel-oriented LiDAR odometry.

2.2. Multi-Sensor Fusion for Robust Odometry

Multi-sensor fusion is widely used to improve odometry robustness under challenging motion or sensing conditions [29–32]. IMU (inertial measurement unit) measurements are commonly fused with LiDAR to compensate motion distortion, provide short-term motion prediction, and improve robustness during aggressive motion. FAST-LIO [33] tightly couples LiDAR feature points and IMU measurements within an iterated Kalman filter and introduces an efficient Kalman gain computation whose complexity depends on the state dimension rather than the measurement dimension. FAST-LIO2 [5] further removes the feature extraction module and directly registers raw points to the map. LIO-SAM [21] formulates LiDAR-inertial odometry as a factor graph with IMU preintegration, LiDAR odometry factors, optional Global Positioning System (GPS) factors, and loop closure factors. By using keyframes and local submaps, it achieves real-time smoothing and mapping.

Visual and visual-inertial information can also be fused with LiDAR to improve robustness. LVI-SAM [22] consists of a visual-inertial subsystem and a LiDAR-inertial subsystem. The visual-inertial subsystem provides initial guesses for LiDAR scan matching, while the LiDAR subsystem provides depth information for visual features and helps visual initialization. Such LiDAR-visual-inertial systems can improve robustness in texture-poor or structure-poor environments by allowing different sensing modalities to compensate for each other.

Although tightly coupled LIO and LVIO methods can significantly improve motion prediction and reduce drift, LiDAR degeneracy may still exist at the local registration level when the scene provides weak geometric constraints in certain pose directions [13]. In such cases, inertial, visual, wheel, or encoder measurements can help the overall estimator remain stable, as also observed in our experiments. However, many existing systems do not explicitly separate LiDAR-observable and LiDAR-degenerate directions during the registration update. Therefore, using an auxiliary odometry source only as a motion prior or initial guess does not provide an explicit mechanism to prevent weakly constrained LiDAR residuals from affecting the final pose increment. Under severe or long-lasting degeneracy, the optimized LiDAR correction may still contain unreliable components along poorly observable directions [34].

This observation motivates a direction-aware update strategy. Instead of applying the full LiDAR registration increment to all pose dimensions, LiDAR corrections can be applied only in well-constrained directions, while the components in weakly observable directions are preserved from a reliable heterogeneous odometry prior. Different from methods that improve LiDAR registration robustness by modifying residual metrics or adaptive weighting, HERO-LO focuses on preserving an external motion prior in the LiDAR-degenerate subspace. Specifically, HERO-LO introduces a unified 6D pose interface for heterogeneous odometry priors. In non-degenerate scenes, it can use scan-to-scan LiDAR odometry as the motion prior and rely mainly on LiDAR registration. When tunnel-

like degeneracy is detected, it switches to a non-LiDAR odometry prior, such as encoder odometry or VIO, and applies LiDAR matching corrections only in the observable subspace.

2.3. Degeneracy Detection and Handling in LiDAR Odometry

Degeneracy in optimization-based state estimation has been studied from the perspective of constraint geometry. Zhang et al. [9,35] proposed an online degeneracy analysis method for optimization-based state estimation. Their method analyzes the eigenvalues and eigenvectors of the Hessian matrix of the linearized least-squares problem. Small eigenvalues indicate weakly constrained directions, and solution remapping is used to combine the optimization solution in well-conditioned directions with a prediction in degenerate directions. This idea has influenced many LiDAR odometry and SLAM systems because it provides a practical way to detect and mitigate local degeneracy with low computational cost.

Several later works further investigate localizability or observability in LiDAR registration. LION [36] introduces an observability-aware LiDAR-inertial navigation framework for vision-denied environments and uses an observability metric to detect geometrically ill-conditioned cases. AdaLIO [13] detects degenerate indoor environments and adaptively adjusts system parameters to improve robustness. X-ICP [14] proposes a fine-grained localizability detection module and a localizability-aware constrained ICP optimization strategy. It classifies optimization directions into localizable, partially localizable, and non-localizable categories, and then constrains the ICP update accordingly. These methods demonstrate that explicitly reasoning about localizability is beneficial for improving the robustness and interpretability of LiDAR registration in geometrically challenging environments.

Another line of work attempts to reduce degeneracy by modifying the registration formulation itself. GenZ-ICP [20] revisits point-to-plane and point-to-point error metrics and adaptively weights them according to local planarity. In corridor-like scenes, the additional point-to-point components can help prevent severely ill-posed point-to-plane optimization. This strategy improves robustness by changing the residual composition rather than directly projecting the pose update. However, it does not explicitly preserve an external motion prior in the unobservable subspace.

The degeneracy problem is especially relevant in tunnel-like environments. A tunnel or long corridor often provides strong geometric constraints perpendicular to the tunnel axis, while the constraints along the direction of travel or other symmetry-related directions can be much weaker. Consequently, the Hessian matrix of local LiDAR registration may become ill-conditioned, and the estimated update can be sensitive to measurement noise, correspondence errors, or local geometric ambiguities in weakly constrained directions. In severe or long-lasting degeneracy, directly applying the full LiDAR increment may introduce longitudinal drift, local trajectory distortion, or occasional unstable updates. Existing degeneracy-aware methods provide valuable tools for identifying weak directions. However, their handling mechanisms are often designed within the LiDAR registration process itself, such as solution remapping, constrained optimization, adaptive parameter tuning, or residual reweighting. In contrast, HERO-LO focuses on preserving a reliable heterogeneous odometry prior in the LiDAR-degenerate subspace.

HERO-LO is closely related to the Hessian-based degeneracy handling strategy of Zhang et al. [9], but provides a different interpretation and implementation. Instead of treating small Hessian eigenvalues only as numerical indicators, HERO-LO interprets them from the viewpoint of state estimation uncertainty. Under the local Gaussian approximation of nonlinear least squares, the Hessian can be regarded as an approximate information matrix of the pose perturbation, and its inverse or pseudo-inverse is related to the corresponding covariance. Thus, small Hessian eigenvalues correspond to large uncertainty

along the associated eigen-directions, indicating poor local observability. HERO-LO defines these degenerate directions in the tangent space of $SO(3) \times \mathbb{R}^3$ and projects the LiDAR optimization increment onto the observable subspace. The degenerate components are not corrected by potentially unreliable LiDAR matching; instead, they are maintained by the heterogeneous odometry prior. This design combines LiDAR matching in well-constrained directions with external odometry prediction in weakly observable directions, providing an explicit and interpretable mechanism for long corridor and tunnel scenarios.

From a control-theoretic viewpoint, the direction-wise suppression of unreliable corrections is an observer-design problem. In degenerate environments, the registration provides reliable information only in a subspace of the pose tangent space, and the question is how to design a correction law whose estimation-error dynamics remain well behaved under this structural loss of observability. A classical route is to certify observer-error stability through LMIs [37]. Arena et al. [38] design Luenberger observers for piecewise-linear (PWL) chaotic systems by imposing, for each linear region, an LMI with a common Lyapunov certificate that simultaneously stabilises all matched-region error dynamics (their Equations (8) and (11)), and they show that the design problem is feasible only if the system is observable—and infeasible for unobservable (identical) subsystems. HERO-LO exhibits the same structural ingredient, namely a finite set of operating modes with mode-dependent error dynamics, and Section 3.7 carries out the corresponding discrete-time LMI analysis, with an essential difference in scope: HERO-LO does not attempt to reconstruct the unobservable subspace from LiDAR data but hands it over to the heterogeneous odometry prior.

3. Methodology

This section presents the proposed HERO-LO framework. HERO-LO is designed for robust LiDAR odometry in tunnel-like or corridor-like environments, where LiDAR registration may be well constrained in some pose directions but weakly constrained in others. The key idea is to avoid applying the full LiDAR registration increment when local geometric degeneracy is detected. Instead, the LiDAR increment is projected onto the observable subspace, while the components in the degenerate subspace are preserved from a heterogeneous odometry prior.

3.1. System Overview

The overall architecture of HERO-LO is shown conceptually in Figure 1. The system consists of four main modules: feature extraction, scan-to-scan odometry, heterogeneous odometry prior generation, and scan-to-submap registration with degeneracy-aware observable-subspace correction.

Given an incoming LiDAR scan $\mathcal{P}_k$ at time t_k, the feature extraction module extracts edge and planar feature points following the feature-based design of LOAM [26] and LeGO-LOAM [17]. These features are used by both the scan-to-scan odometry module and the scan-to-submap registration module. The scan-to-scan module estimates a short-term LiDAR-only motion prediction by matching the current scan with the previous scan. In normal non-degenerate scenes, this scan-to-scan estimate is used as the motion prior for scan-to-submap registration.

In parallel, HERO-LO provides a unified 6D pose interface for heterogeneous odometry sources, such as wheel or encoder odometry, visual-inertial odometry, or other onboard odometry estimators. After temporal synchronization and spatial alignment, the heterogeneous odometry module predicts the current LiDAR pose from the previous estimated pose and the relative motion measured by the external odometry source.

The scan-to-submap module performs feature-based LiDAR registration against a local submap. During this process, HERO-LO analyzes the Hessian matrix of the local least-squares problem to identify weakly constrained directions. When the environment is non-degenerate, the full LiDAR registration increment is applied. When tunnel-like degeneracy is detected, HERO-LO applies the LiDAR correction only in the observable subspace and preserves the heterogeneous odometry prior in the degenerate subspace.

Compared with conventional feature-based LiDAR odometry pipelines, HERO-LO does not only use heterogeneous odometry as an initial guess. Instead, it explicitly prevents unreliable LiDAR corrections from changing the pose components that are weakly constrained by the current LiDAR geometry.

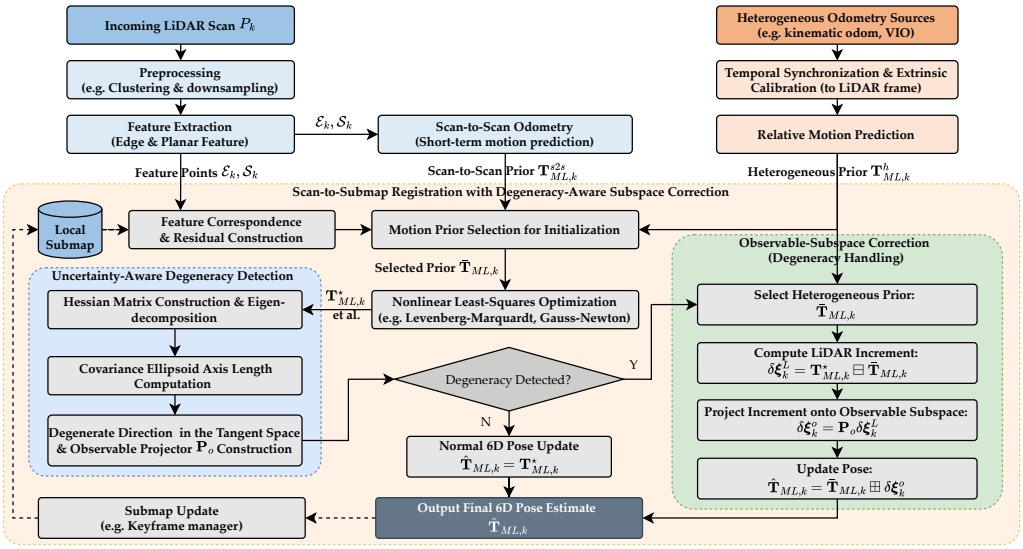


Figure 1. System overview of HERO-LO. The arrows indicate the data flow between the modules: the start of an arrow is the output of the preceding module and its end is the input of the following module. Solid lines denote data flow present in every frame, whereas dashed lines denote data flow present only when a trigger condition is met (e.g., when a keyframe is detected). Abbreviations: LiDAR, light detection and ranging; VIO, visual-inertial odometry; odom, odometry; 6D, six-dimensional.

3.2. Notation and Pose Perturbation on $SO(3) \times \mathbb{R}^3$

A 6D pose is represented as

$$\mathbf{T} = (\mathbf{R}, \mathbf{t}) \in SO(3) \times \mathbb{R}^3, \tag{1}$$

where $\mathbf{R} \in SO(3)$ is the rotation matrix and $\mathbf{t} \in \mathbb{R}^3$ is the translation vector. The corresponding homogeneous transformation is also denoted by $\mathbf{T}$ when no ambiguity exists.

For local optimization, we use a 6D perturbation vector

$$\delta\boldsymbol{\xi} = \begin{bmatrix} \delta\boldsymbol{\theta}^{\top} & \delta\mathbf{p}^{\top} \end{bmatrix}^{\top} \in \mathbb{R}^6, \tag{2}$$

where $\delta\boldsymbol{\theta} \in \mathbb{R}^3$ is the rotational perturbation and $\delta\mathbf{p} \in \mathbb{R}^3$ is the translational perturbation. The pose update operator is defined as

$$\mathbf{T} \boxplus \delta\boldsymbol{\xi} = \big(\mathbf{R}\,\mathrm{Exp}(\delta\boldsymbol{\theta}^{\wedge}), \mathbf{t} + \mathbf{R}\delta\mathbf{p}\big), \tag{3}$$

where $(\cdot)^{\wedge}$ maps a vector in $\mathbb{R}^3$ to its corresponding skew-symmetric matrix, and $\mathrm{Exp}(\cdot)$ denotes the exponential map from $\mathfrak{so}(3)$ to $SO(3)$. Here, $\boxplus$ denotes the pose retraction ("boxplus") operator, which perturbs the pose $\mathbf{T}$ by the tangent-space vector $\delta\boldsymbol{\xi}$.

The inverse operation between two poses is defined as

$$\mathbf{T}_j \boxminus \mathbf{T}_i = \begin{bmatrix} \mathrm{Log}(\mathbf{R}_i^\top \mathbf{R}_j)^\vee \\ \mathbf{R}_i^\top (\mathbf{t}_j - \mathbf{t}_i) \end{bmatrix}, \tag{4}$$

where $\mathrm{Log}(\cdot)$ is the logarithm map from $SO(3)$ to $\mathfrak{so}(3)$, and $(\cdot)^\vee$ is the inverse operation of $(\cdot)^\wedge$. This formulation allows the degeneracy directions to be defined in the tangent space of $SO(3) \times \mathbb{R}^3$. Here, $\boxminus$ denotes the corresponding inverse ("boxminus") operator.

3.3. Heterogeneous Odometry Prior

3.3.1. Unified 6D Pose Interface

HERO-LO assumes that each auxiliary odometry source can provide or be converted into a time-stamped 6D pose sequence,

$$\mathcal{O}^h = \left\{ \mathbf{T}_0^h, \mathbf{T}_1^h, \ldots, \mathbf{T}_k^h \right\}, \tag{5}$$

where the superscript h denotes a heterogeneous odometry source, such as encoder odometry or VIO. For odometry sources that originally provide lower-dimensional motion, such as planar wheel odometry, the measurement can be embedded into a 6D pose representation using available attitude information or reasonable platform motion assumptions. HERO-LO itself only requires a unified 6D pose input and does not depend on the internal implementation of the auxiliary odometry estimator.

Let $\mathbf{T}_{O_h H,k}$ denote the pose of the heterogeneous odometry sensor frame H in its own odometry frame O_h at time t_k. Let $\mathbf{T}_{HL}$ be the calibrated extrinsic transformation from the LiDAR frame L to the heterogeneous sensor frame H. The corresponding relative LiDAR motion predicted from the heterogeneous odometry between two consecutive timestamps is computed as

$$\Delta\mathbf{T}_{L,k-1,k}^h = \mathbf{T}_{HL}^{-1} \left(\mathbf{T}_{O_h H,k-1}^{-1} \mathbf{T}_{O_h H,k} \right) \mathbf{T}_{HL}. \tag{6}$$

This relative-motion formulation avoids the need to directly align the global coordinate frame of the heterogeneous odometry with the LiDAR map frame. Only the extrinsic calibration between the heterogeneous sensor and the LiDAR is required.

3.3.2. Prior Pose Prediction

Given the previous HERO-LO pose estimate $\hat{\mathbf{T}}_{ML,k-1}$ in the LiDAR map frame M, the heterogeneous odometry prior for the current scan is predicted as

$$\bar{\mathbf{T}}_{ML,k}^h = \hat{\mathbf{T}}_{ML,k-1} \Delta\mathbf{T}_{L,k-1,k}^h. \tag{7}$$

Similarly, the scan-to-scan LiDAR odometry module provides a LiDAR-only relative motion estimate $\Delta\mathbf{T}_{L,k-1,k}^{s2s}$, and the corresponding scan-to-scan prior is

$$\bar{\mathbf{T}}_{ML,k}^{s2s} = \hat{\mathbf{T}}_{ML,k-1} \Delta\mathbf{T}_{L,k-1,k}^{s2s}. \tag{8}$$

Therefore, HERO-LO can select different prior predictions depending on the degeneracy condition:

$$\bar{\mathbf{T}}_{ML,k} = \begin{cases} \bar{\mathbf{T}}_{ML,k}^{s2s}, & \text{non-degenerate mode,} \\ \bar{\mathbf{T}}_{ML,k}^{h}, & \text{degenerate mode and valid heterogeneous prior available.} \end{cases} \tag{9}$$

If the current registration is detected as degenerate, the observable-subspace correction described in Section 3.6 is applied using the selected prior.

3.3.3. Non-Degenerate and Degenerate Working Modes

HERO-LO has two working modes.

In non-degenerate scenes, the LiDAR geometry provides sufficient constraints for all or most pose degrees of freedom. In this case, HERO-LO uses the scan-to-scan LiDAR odometry as the motion prior and allows scan-to-submap registration to update the full 6D pose. This mode avoids unnecessary dependence on external odometry and preserves the high accuracy of LiDAR matching in geometrically rich environments.

In degenerate tunnel-like scenes, the LiDAR registration problem can become weakly constrained along the tunnel axis or other symmetry-related directions. In this case, HERO-LO switches to a non-LiDAR odometry prior, such as encoder odometry or VIO, and only applies LiDAR corrections in the observable subspace. The weakly observable pose components are maintained by the heterogeneous prior.

In the implementation, the working mode of the current frame is determined by the degeneracy detection performed on the Hessian of the previous frame, i.e., with a one-frame lag (50 ms at the 20 Hz LiDAR rate). No hysteresis or multi-frame confirmation is applied: the detection is evaluated independently for every frame. Because the motion prior is anchored to the previous HERO-LO pose estimate (Equation (7)), this one-frame lag and the per-frame mode changes do not produce pose resets or discontinuities; only the source of the degenerate-direction component switches between the scan-to-scan and the heterogeneous prior.

3.4. Feature-Based LiDAR Registration

3.4.1. Point Cloud Preprocessing and Feature Extraction

HERO-LO adopts a feature-based LiDAR registration structure similar to LOAM [26] and LeGO-LOAM [17]. For each incoming scan, the raw point cloud is first organized according to the LiDAR scan pattern. Ground points and unreliable points can be filtered according to the sensor configuration and platform requirements. Then, a local smoothness measure is computed for each point, and feature points are selected based on their geometric characteristics.

The extracted features are divided into two sets:

$$\mathcal{E}_k = \{\mathbf{p}^e_{k,i}\}, \qquad \mathcal{S}_k = \{\mathbf{p}^s_{k,j}\}, \tag{10}$$

where $\mathcal{E}_k$ denotes edge feature points and $\mathcal{S}_k$ denotes planar feature points. Edge features are used to construct point-to-line residuals, while planar features are used to construct point-to-plane residuals.

The same feature extraction results are shared by the scan-to-scan odometry module and the scan-to-submap registration module. This design reduces computational redundancy and keeps the overall pipeline lightweight.

3.4.2. Feature Correspondence and Residual Construction

Given an initial pose prediction $\bar{\mathbf{T}}_{ML,k}$, each feature point in the current LiDAR frame is transformed into the map frame:

$$\mathbf{p}^M_{k,i} = \mathbf{R}_{ML,k}\mathbf{p}^L_{k,i} + \mathbf{t}_{ML,k}. \tag{11}$$

For an edge feature point $\mathbf{p}^e_{k,i}$, two corresponding edge points $\mathbf{q}_1$ and $\mathbf{q}_2$ are searched from the target point cloud or the local submap. The point-to-line residual is defined as

$$\mathbf{r}^e_i(\mathbf{T}_{ML,k}) = \frac{\left(\mathbf{p}^M_{k,i} - \mathbf{q}_1\right) \times \left(\mathbf{p}^M_{k,i} - \mathbf{q}_2\right)}{\|\mathbf{q}_1 - \mathbf{q}_2\|}. \tag{12}$$

For a planar feature point $\mathbf{p}_{k,j}^{s}$, several neighboring planar points are found from the target point cloud or submap, and a local plane is estimated. Let the plane be represented by the unit normal vector $\mathbf{n}_j$ and a point $\mathbf{q}_j$ on the plane. The point-to-plane residual is

$$r_j^s(\mathbf{T}_{ML,k}) = \mathbf{n}_j^\top \left(\mathbf{p}_{k,j}^{M} - \mathbf{q}_j \right). \tag{13}$$

The scan-to-scan and scan-to-submap registration problems share the same residual form. Their difference lies in the target point cloud: scan-to-scan registration uses the previous scan as the target, while scan-to-submap registration uses a local submap constructed from recent keyframes.

3.4.3. Nonlinear Least-Squares Optimization

The LiDAR registration problem is formulated as a nonlinear least-squares problem:

$$\mathbf{T}_{ML,k}^{\star} = \arg\min_{\mathbf{T}_{ML,k}} \left(\sum_{i \in \mathcal{E}_k} \rho\left(\left\| \mathbf{r}_i^e(\mathbf{T}_{ML,k}) \right\|_{\mathbf{\Sigma}_e^{-1}}^2 \right) + \sum_{j \in \mathcal{S}_k} \rho\left(\left\| r_j^s(\mathbf{T}_{ML,k}) \right\|_{\sigma_s^{-2}}^2 \right) \right), \tag{14}$$

where the superscript $\star$ denotes the optimal solution of the minimization problem, $\rho(\cdot)$ is a robust loss function, and $\mathbf{\Sigma}_e$ and σ_s^2 denote the residual noise models for edge and planar constraints, respectively.

Using the perturbation model in Equation (3), the residuals are linearized around the current estimate:

$$\mathbf{r}(\mathbf{T} \boxplus \delta\boldsymbol{\xi}) \approx \mathbf{r}(\mathbf{T}) + \mathbf{J}\delta\boldsymbol{\xi}. \tag{15}$$

The corresponding normal equation is

$$\mathbf{H}\delta\boldsymbol{\xi} = -\mathbf{g}, \tag{16}$$

where

$$\mathbf{H} = \mathbf{J}^\top \mathbf{W} \mathbf{J}, \qquad \mathbf{g} = \mathbf{J}^\top \mathbf{W} \mathbf{r}. \tag{17}$$

Here, $\mathbf{W}$ contains residual weighting, noise weighting, and robust-kernel weighting. The Hessian matrix $\mathbf{H}$ is also used for local observability analysis.

In the implementation, the robust loss $\rho(\cdot)$ is a Huber kernel with $\delta = 0.08$ m, applied to the residual norm. The residual noise model is estimated online and adaptively from the local geometry (a fit through the five nearest points; see the Supplementary Materials, Table S6).

The scan-to-submap registration solves the normal equation with a Gauss–Newton optimiser that performs three data-association rounds, each followed by 20 iterations; the total of 60 iterations is fixed and no convergence criterion is evaluated, so the execution is deterministic for identical inputs (Section 4.1.4).

The linearisation of the residuals around the current estimate is well conditioned in the present scenario: the inter-frame interval is 50 ms and the robot speed does not exceed 1.5 m/s, so the inter-frame motion is on the order of centimetres of translation and a small rotation, and the linearised model closely follows the true residual surface; the Hessian is then used as an approximate information matrix under the local Gaussian model of the least-squares solution (Section 3.5). Outside this regime—large inter-frame motions or a strongly wrong initial prior—the linearised observability analysis, and, therefore, the detection, can become unreliable. HERO-LO is intended for the quasi-static, small-motion regime of the mobile-robot applications it targets.

3.5. Uncertainty-Aware Degeneracy Detection

3.5.1. Hessian-Based Local Observability Analysis

The proposed degeneracy detection is closely related to the Hessian-based analysis of optimization-based state estimation problems [9]. For the linearized least-squares problem in Equation (16), the Hessian matrix describes how strongly the residuals constrain the local pose perturbation.

HERO-LO performs eigenvalue decomposition on the Hessian matrix:

$$\mathbf{H} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{\top}, \quad \mathbf{\Lambda} = \mathrm{diag}(\lambda_1, \lambda_2, \ldots, \lambda_6), \quad \mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_6]. \tag{18}$$

where λ_i and $\mathbf{v}_i$ are the i-th eigenvalue and eigenvector, respectively. A small eigenvalue λ_i indicates that the registration residuals provide weak information along the direction $\mathbf{v}_i$ in the local tangent space.

3.5.2. Covariance Ellipsoid Interpretation

Different from directly treating the Hessian eigenvalues as purely numerical indicators, HERO-LO interprets degeneracy from the perspective of state estimation uncertainty. Under the local Gaussian approximation of nonlinear least squares, the Hessian matrix can be regarded as an approximate information matrix of the pose perturbation. Therefore, the covariance of the perturbation can be approximately related to the inverse Hessian as

$$\mathbf{\Sigma}_{\xi} \approx \sigma_r^2 \mathbf{H}^{\dagger}, \tag{19}$$

where σ_r^2 denotes the residual noise scale and $\mathbf{H}^{\dagger}$ denotes the inverse or pseudo-inverse of $\mathbf{H}$. When numerical regularization is used, $\mathbf{H}^{\dagger}$ can be replaced by $(\mathbf{H} + \epsilon\mathbf{I})^{-1}$ with a small $\epsilon > 0$.

According to this interpretation, the principal axes of the covariance ellipsoid are aligned with the eigenvectors of $\mathbf{H}$. The corresponding covariance-inspired uncertainty length along direction $\mathbf{v}_i$ can be written as

$$a_i = \frac{\sigma_r}{\sqrt{\lambda_i + \epsilon}}. \tag{20}$$

In HERO-LO, this covariance ellipsoid interpretation is used only for online degeneracy detection and observable-subspace construction, rather than for reporting an absolutely calibrated pose covariance. Therefore, to avoid additional online residual noise-scale estimation, we set $\sigma_r^2 = 1$ in the implementation. The normalized uncertainty indicator used by HERO-LO is then

$$a_i = \frac{1}{\sqrt{\lambda_i + \epsilon}}. \tag{21}$$

This simplification does not affect the identified eigenspaces because σ_r is only a global scale factor applied to all covariance axes. It does not change the Hessian eigenvectors or the relative ordering of the eigenvalues. Its influence is absorbed into the degeneracy threshold. Thus, a small Hessian eigenvalue still corresponds to a large normalized uncertainty length, indicating that the pose perturbation is weakly constrained along the associated direction. HERO-LO regards such directions as LiDAR-degenerate directions and prevents unreliable LiDAR corrections from being applied along them.

3.5.3. Degenerate Direction Identification in the Tangent Space

A direction $\mathbf{v}_i$ is classified as observable if its uncertainty length is below a predefined threshold τ_a:

$$\mathbf{v}_i \in \mathcal{O} \quad \text{if} \quad a_i \leq \tau_a. \tag{22}$$

Otherwise, it is classified as degenerate:

$$\mathbf{v}_i \in \mathcal{D} \quad \text{if} \quad a_i > \tau_a. \tag{23}$$

Equivalently, one may use an eigenvalue threshold τ_λ:

$$\lambda_i < \tau_\lambda \quad \Rightarrow \quad \mathbf{v}_i \in \mathcal{D}. \tag{24}$$

The covariance-ellipsoid interpretation provides a more intuitive explanation of the threshold selection: the threshold corresponds to the maximum acceptable local pose uncertainty along a principal direction. Following the threshold-determination procedure of Zhang et al. [9], the threshold used in this work, $\tau_a = 0.0315$, was calibrated on the tunnel sequence of the present experiments: the scalar indicator $\max_i a_i$ was computed for a set of frames manually labelled as well-conditioned or degenerate according to the known environment and the observed estimation behaviour, the distributions of the two groups were compared, and the threshold was set at the midpoint of the gap between them. Equivalently, this corresponds to an eigenvalue threshold $\tau_\lambda \approx 1.01 \times 10^3$ on the unnormalised Hessian. With this value, no direction is flagged in the non-degenerate intervals (0–83 s and 145–256 s), while a single direction (the forward translation direction, see Section 4.1.2) is flagged in essentially all frames of the degenerate segment (approximately 83–145 s). A systematic analysis of the influence of the threshold value is beyond the scope of this work and is left to a future dedicated study on degeneracy detection.

3.6. Observable-Subspace LiDAR Correction

3.6.1. Observable and Degenerate Subspace Construction

Let $\mathbf{D}_o$ be a diagonal selection matrix:

$$\mathbf{D}_o = \mathrm{diag}(d_1, d_2, \ldots, d_6), \tag{25}$$

where

$$d_i = \begin{cases} 1, & \mathbf{v}_i \in \mathcal{O}, \\ 0, & \mathbf{v}_i \in \mathcal{D}. \end{cases} \tag{26}$$

The projection matrix onto the LiDAR-observable subspace is then

$$\mathbf{P}_o = \mathbf{V}\mathbf{D}_o\mathbf{V}^\top. \tag{27}$$

The projection matrix onto the degenerate subspace is

$$\mathbf{P}_d = \mathbf{I}_6 - \mathbf{P}_o. \tag{28}$$

If all directions are observable, then $\mathbf{P}_o = \mathbf{I}_6$ and HERO-LO reduces to conventional full LiDAR registration. If one or more directions are degenerate, the LiDAR correction is only applied through $\mathbf{P}_o$.

The construction is not restricted to a single degenerate direction: when several directions are flagged simultaneously, $\mathbf{D}_o$ assigns zero to each of them and the projector handles a degenerate subspace of arbitrary rank. At the opposite extreme, if all six directions were degenerate ($\mathbf{P}_o \to \mathbf{0}$), the update would reduce to a pure prior prediction and the performance would be entirely governed by the quality of the heterogeneous odometry source. In the present dataset this extreme does not occur: at most one direction (the forward translation direction) is degenerate, and only between approximately 83 s and 145 s (Section 4.1.2). This extreme case makes explicit the requirement, stated in the

Limitations, that a reasonably stable heterogeneous prior is needed whenever the LiDAR geometry provides no constraint at all.

3.6.2. LiDAR Increment Projection

Let $\bar{\mathbf{T}}_{ML,k}$ be the selected prior pose from Equation (9), and let $\mathbf{T}^{\star}_{ML,k}$ be the full LiDAR registration result obtained by solving Equation (14). The LiDAR registration increment relative to the selected prior is

$$\delta\boldsymbol{\xi}^{L}_{k} = \mathbf{T}^{\star}_{ML,k} \boxminus \bar{\mathbf{T}}_{ML,k}. \tag{29}$$

The observable component of this LiDAR increment is computed as

$$\delta\boldsymbol{\xi}^{o}_{k} = \mathbf{P}_{o}\delta\boldsymbol{\xi}^{L}_{k}. \tag{30}$$

The degenerate component is

$$\delta\boldsymbol{\xi}^{d}_{k} = \mathbf{P}_{d}\delta\boldsymbol{\xi}^{L}_{k}. \tag{31}$$

HERO-LO discards $\delta\boldsymbol{\xi}^{d}_{k}$ because it is not sufficiently constrained by the LiDAR measurements. This prevents weakly constrained LiDAR residuals from introducing unreliable pose corrections along degenerate directions.

3.6.3. Prior-Preserved Pose Update

The final HERO-LO pose estimate is obtained by applying only the observable LiDAR correction to the selected prior:

$$\hat{\mathbf{T}}_{ML,k} = \bar{\mathbf{T}}_{ML,k} \boxplus \delta\boldsymbol{\xi}^{o}_{k}. \tag{32}$$

Equation (32) is the core update rule of HERO-LO. It has two important properties.

First, in well-constrained directions, HERO-LO uses LiDAR scan-to-submap matching to correct the motion prior. This allows the system to exploit accurate geometric constraints when they are available.

Second, in weakly observable directions, HERO-LO keeps the prior estimate instead of applying the potentially unreliable LiDAR increment. Therefore, the final pose estimate combines LiDAR-based correction in observable directions and heterogeneous odometry prediction in degenerate directions.

From an error point of view, the update rule in Equation (32) can be read as a per-direction selection of the information source. Let the full LiDAR increment carry a registration error, and let the selected prior carry its own prediction error, both expressed in the tangent space. After the update, the error along an observable direction is bounded by the observable component of the LiDAR registration error, while the error along a degenerate direction is exactly the component of the prior error in that direction, because the LiDAR component is discarded. Each pose component is, thus, carried by the information source that is most reliable in that direction under the current geometry, which is the quantitative basis of the effectiveness of the update rule. This per-direction selection is made precise by the observer analysis of Section 3.7 (Equations (38) and (39)).

The projection also has a price: the subspace is estimated from the same Hessian that drives the detection, so for frames whose eigenvalues lie close to the threshold, part of the reliable LiDAR correction may be discarded along a mis-estimated direction. This cost is small in the well-separated eigenvalue regime, and it is visible in the ablation study as the slightly larger endpoint error of HERO-LO-S2S-OSC (78.6 m) compared with HERO-LO-S2S-NoOSC (77.6 m), in which the retained scan-to-scan prior is itself contaminated in the degenerate direction (Section 4.4.1).

Finally, because the selected prior is anchored to the previous HERO-LO pose estimate (Equation (7)), a change of the working mode within one frame does not introduce a pose reset or a discontinuity: the estimate always continues from the previous estimate, and only the source of the degenerate-direction component changes.

3.7. Observer Interpretation and LMI-Based Analysis of the Subspace Correction

This subsection interprets the prior-preserved update of Equation (32) as an observer with a projection gain and analyses its estimation-error dynamics. The purpose is interpretive: no component of the algorithm is changed, and no optimisation problem is solved online. The analysis follows the classical route of certifying observer-error stability through Lyapunov matrix inequalities [37] and, in particular, the design philosophy of Arena et al. [38], who stabilise the error dynamics of a Luenberger observer over the linear regions of a piecewise-linear system by a common Lyapunov certificate expressed as an LMI, and who show that the corresponding design problem is infeasible when the system is unobservable.

The analysis is performed to first order in the tangent space of $SO(3) \times \mathbb{R}^3$, in the quasi-static small-motion regime already declared in Section 3.4: the $\boxplus/\boxminus$ operations of Equations (3) and (4) reduce to vector addition and subtraction in a common tangent-space coordinate system, the tangent-space transport between consecutive frames is the identity to first order, and second-order remainder terms (of the order of the squared centimetre-level per-frame increments) are neglected. Let $\mathbf{x}_k \in \mathbb{R}^6$ denote the true pose state at frame k in these coordinates, $\hat{\mathbf{x}}_k$ the HERO-LO estimate (the tangent-space image of Equation (32)), $\bar{\mathbf{x}}_k$ the selected prior of Equation (9), and $\mathbf{y}_k$ the full LiDAR registration result of Equation (14), so that $\mathbf{y}_k - \bar{\mathbf{x}}_k = \delta\boldsymbol{\xi}_k^L$ (Equation (29)). The prior is the open-loop propagation of the previous estimate by the odometry increment $\mathbf{u}_k$ of the selected source (Equations (7) and (8)); the true state obeys the same propagation up to the increment error $\mathbf{w}_k$ of that source; and the registration provides a full-state measurement candidate corrupted by the registration error $\boldsymbol{\ell}_k$:

$$\bar{\mathbf{x}}_k = \hat{\mathbf{x}}_{k-1} + \mathbf{u}_k, \qquad \mathbf{x}_k = \mathbf{x}_{k-1} + \mathbf{u}_k + \mathbf{w}_k, \qquad \mathbf{y}_k = \mathbf{x}_k + \boldsymbol{\ell}_k. \tag{33}$$

The directional quality of the two sources is quantified by the partition of Section 3.5: along observable directions the registration is well conditioned, so $\mathbf{P}_{o,k}\boldsymbol{\ell}_k$ is small, whereas along degenerate directions it is uncontrolled—which is precisely the defining property of $\mathcal{D}$. The projectors are re-estimated at every frame. The working mode, and hence the prior source selected by Equation (9), is determined by the degeneracy detection performed on the Hessian of the previous frame (one-frame lag, Section 3.3). All statements below refer to the projector actually applied at each frame, whatever the frame whose Hessian it is computed from.

In these coordinates, Equation (32) reads

$$\hat{\mathbf{x}}_k = \bar{\mathbf{x}}_k + \mathbf{P}_{o,k}(\mathbf{y}_k - \bar{\mathbf{x}}_k) = \mathbf{P}_{d,k}\,\bar{\mathbf{x}}_k + \mathbf{P}_{o,k}\,\mathbf{y}_k, \tag{34}$$

which has the structure of a Luenberger observer with output matrix $\mathbf{C} = \mathbf{I}_6$: $\bar{\mathbf{x}}_k$ is the open-loop prediction, $\mathbf{y}_k - \bar{\mathbf{x}}_k$ is the innovation, and the observer gain is the orthogonal projection

$$\mathbf{L}_k = \mathbf{P}_{o,k} = \mathbf{V}_k\mathbf{D}_{o,k}\mathbf{V}_k^\top. \tag{35}$$

Equivalently, the gain factorises as $\mathbf{B}_k\mathbf{C}_k$ with $\mathbf{B}_k = \mathbf{V}_{o,k}$ (the matrix whose columns span the observable subspace) and $\mathbf{C}_k = \mathbf{V}_{o,k}^\top$: the registration may be read as a measurement of the observable components of the pose, and exactly those components are injected. Because $\mathbf{D}_{o,k}$ contains only zeros and ones, every tangent-space direction receives either full

gain (observable directions) or zero gain (degenerate directions): the observable/degenerate split is a direction-wise gain selection with an explicit semantic, not a heuristic clipping. In the non-degenerate mode, $\mathbf{P}_{o,k} = \mathbf{I}_6$ and Equation (34) reduces to $\hat{\mathbf{x}}_k = \mathbf{y}_k$, the full registration update of Algorithm 1.

Substituting Equation (33) into Equation (34), the estimation error $\mathbf{e}_k := \mathbf{x}_k - \hat{\mathbf{x}}_k$ evolves as

$$\mathbf{e}_k = \mathbf{A}_{cl,k}\,\mathbf{e}_{k-1} + \mathbf{P}_{d,k}\mathbf{w}_k - \mathbf{P}_{o,k}\boldsymbol{\ell}_k, \qquad \mathbf{A}_{cl,k} := \mathbf{I}_6 - \mathbf{P}_{o,k} = \mathbf{P}_{d,k}, \tag{36}$$

where the memory term originates from the prior anchoring: by the first two relations of Equation (33), $\mathbf{x}_k - \bar{\mathbf{x}}_k = \mathbf{e}_{k-1} + \mathbf{w}_k$, and $\mathbf{P}_{o,k}\mathbf{P}_{d,k} = \mathbf{0}$ is used. Three properties of this correction are now established.

Algorithm 1 HERO-LO

Require: LiDAR scans $\{\mathcal{P}_k\}$, heterogeneous odometry poses $\{\mathbf{T}_k^h\}$, extrinsic calibration $\mathbf{T}_{HL}$
Ensure: LiDAR odometry trajectory $\{\hat{\mathbf{T}}_{ML,k}\}$

```
Initialize map frame, local submap, and initial pose T̂_{ML,0}
for each incoming scan P_k do
    Extract edge features E_k and planar features S_k
    Estimate scan-to-scan relative motion ΔT^{s2s}_{L,k-1,k}
    Compute scan-to-scan prior T̄^{s2s}_{ML,k} using Equation (8)
    if heterogeneous odometry is available then
        Synchronize heterogeneous odometry to time t_k
        Compute ΔT^h_{L,k-1,k} using Equation (6)
        Compute heterogeneous prior T̄^h_{ML,k} using Equation (7)
    end if
    Select prior T̄_{ML,k} according to the current working mode
    Register current scan features to the local submap by solving Equation (14)
    Obtain full LiDAR registration result T*_{ML,k}
    Construct Hessian matrix H = J^⊤ W J
    Perform eigenvalue decomposition H = V Λ V^⊤
    Identify observable and degenerate directions using Equations (22) and (23)
    Construct observable projector P_o using Equation (27)
    if degeneracy is detected then
        Compute LiDAR increment δξ^L_k using Equation (29)
        Project LiDAR increment onto observable subspace using Equation (30)
        Update pose using prior-preserved correction in Equation (32)
    else
        Set T̂_{ML,k} = T*_{ML,k}
    end if
    Update local submap and keyframes
    Update working mode for the next frame
end for
```

Proposition 1 (Optimality of the projection update)**.** *Let* $\bar{\mathbf{H}}_k := \mathbf{P}_{o,k}\mathbf{H}_k\mathbf{P}_{o,k}$*, where* $\mathbf{H}_k$ *is the Hessian of Equation* (16)*;* $\bar{\mathbf{H}}_k$ *coincides with* $\mathbf{H}_k$ *on the observable subspace (where all eigenvalues exceed the detection threshold) and vanishes on the degenerate subspace. The HERO-LO update of Equation* (34) *is the unique minimiser of*

$$\min_{\mathbf{x}} (\mathbf{y}_k - \mathbf{x})^\top \bar{\mathbf{H}}_k (\mathbf{y}_k - \mathbf{x}) \quad s.t. \quad \mathbf{P}_{d,k}\,\mathbf{x} = \mathbf{P}_{d,k}\,\bar{\mathbf{x}}_k, \tag{37}$$

i.e., among all poses that keep the degenerate components fixed at the prior, the update is the estimate that best fits the registration result in the directions where the LiDAR is informative.

Proof. Every feasible point has the form $\mathbf{x} = \mathbf{P}_{d,k}\bar{\mathbf{x}}_k + \mathbf{P}_{o,k}\mathbf{z}$ with $\mathbf{z}$ free in the observable subspace. Since $\bar{\mathbf{H}}_k\mathbf{P}_{d,k} = \mathbf{0}$, the objective of Equation (37) reduces to

$(\mathbf{P}_{o,k}\mathbf{y}_k - \mathbf{z})^\top \bar{\mathbf{H}}_k(\mathbf{P}_{o,k}\mathbf{y}_k - \mathbf{z})$, a positive-definite quadratic in $\mathbf{z}$ on the observable subspace. It is minimised uniquely at $\mathbf{z} = \mathbf{P}_{o,k}\mathbf{y}_k$, which yields $\mathbf{x} = \mathbf{P}_{d,k}\bar{\mathbf{x}}_k + \mathbf{P}_{o,k}\mathbf{y}_k = \hat{\mathbf{x}}_k$. □

Remark 1 (Bayesian reading of the hard gain)**.** *Under a Gaussian model with LiDAR information* $\mathbf{H}_k$ *and prior information* $\mathbf{\Lambda}_k = \mathbf{V}_k \operatorname{diag}(\mu_i)\mathbf{V}_k^\top$ *in the same eigenbasis (available by construction, because the partition is defined in the eigenbasis of* $\mathbf{H}_k$*), the maximum a posteriori (MAP) estimate is* $\hat{\mathbf{x}} = \bar{\mathbf{x}}_k + (\mathbf{H}_k + \mathbf{\Lambda}_k)^{-1}\mathbf{H}_k(\mathbf{y}_k - \bar{\mathbf{x}}_k)$*. Letting* $\mu_i \to 0$ *along observable directions and* $\mu_i \to \infty$ *along degenerate directions, the gain converges to* $\mathbf{P}_{o,k}$*. The hard* $0/1$ *gain is, thus, the limit of a continuous Bayesian design that trusts the LiDAR completely along observable directions and the prior completely along degenerate ones—the quantitative counterpart of the per-direction selection of the information source discussed in Section 3.6.*

Proposition 2 (Per-direction error propagation)**.** *For every frame* k*,*

$$\mathbf{P}_{o,k}\,\mathbf{e}_k = -\mathbf{P}_{o,k}\boldsymbol{\ell}_k, \qquad \mathbf{P}_{d,k}\,\mathbf{e}_k = \mathbf{P}_{d,k}\,\mathbf{e}_{k-1} + \mathbf{P}_{d,k}\mathbf{w}_k. \tag{38}$$

If the observable/degenerate partition is constant over frames $1, \ldots, k$*, then*

$$\|\mathbf{e}_k\| \le \|\mathbf{P}_o\boldsymbol{\ell}_k\| + \|\mathbf{P}_d\mathbf{e}_0\| + \sum_{j=1}^{k}\|\mathbf{P}_d\mathbf{w}_j\|. \tag{39}$$

Proof. The first relation of Equation (38) follows by left-multiplying $\mathbf{e}_k = \mathbf{x}_k - \mathbf{P}_{d,k}\bar{\mathbf{x}}_k - \mathbf{P}_{o,k}\mathbf{y}_k$ with $\mathbf{P}_{o,k}$ (using $\mathbf{P}_{o,k}\mathbf{P}_{d,k} = \mathbf{0}$, $\mathbf{P}_{o,k}^2 = \mathbf{P}_{o,k}$) and $\mathbf{y}_k - \mathbf{x}_k = \boldsymbol{\ell}_k$; the second by left-multiplying with $\mathbf{P}_{d,k}$ and using $\mathbf{x}_k - \bar{\mathbf{x}}_k = \mathbf{e}_{k-1} + \mathbf{w}_k$. Under a constant partition, the second relation iterates to $\mathbf{P}_d\mathbf{e}_k = \mathbf{P}_d\mathbf{e}_0 + \sum_{j=1}^{k}\mathbf{P}_d\mathbf{w}_j$, and the orthogonal decomposition $\|\mathbf{e}_k\|^2 = \|\mathbf{P}_o\mathbf{e}_k\|^2 + \|\mathbf{P}_d\mathbf{e}_k\|^2$ with the triangle inequality yields Equation (39). □

Equation (38) is the precise form of the per-direction selection of the information source: the observable component of the error is refreshed by the LiDAR at every frame and cannot accumulate, while the degenerate component propagates neutrally and accumulates only the increment errors of the prior. This is the sense in which the correction provides bounded error propagation; it is not a uniform bound—under a persistent prior bias the degenerate-direction error grows linearly, which is the mechanism behind the standalone-prior endpoint errors of Section 4.4.4 and the reason a reasonably stable prior is required (Section 5.6).

Remark 2 (Changing partitions)**.** *When the partition changes between frames, Equation (38) still holds per frame, but unrolling the recursion produces additional terms of the form* $\mathbf{P}_{d,k}\mathbf{P}_{o,k-1}\boldsymbol{\ell}_{k-1}$*: the previous frame's LiDAR refresh error can leak into the current frame's degenerate subspace. Such terms are non-zero only at frames where the partition changes, and each is bounded by the per-frame observable-direction registration accuracy. In the present dataset the leakage is negligible, because the degenerate direction is essentially constant (the forward translation direction) throughout the degenerate segment (Section 4.1.2).*

Proposition 3 (LMI certificate with a common Lyapunov function)**.** *For every observable/degenerate partition, the closed-loop matrix* $\mathbf{A}_{cl} = \mathbf{P}_d$ *of Equation (36) satisfies the discrete-time Lyapunov inequality*

$$\mathbf{A}_{cl}^\top\mathbf{P}\,\mathbf{A}_{cl} - \mathbf{P} \preceq 0, \qquad \mathbf{P} \succ 0, \tag{40}$$

whose left-hand side is affine in the decision variable $\mathbf{P}$ *and is, therefore, an LMI, with the certificate* $\mathbf{P} = \mathbf{I}_6$*, which is independent of the partition and hence common to all partitions and to arbitrary mode switches between frames.*

Proof. The matrix $\mathbf{P}_d$ is a symmetric idempotent orthogonal projection (Section 3.6), hence $\mathbf{P}_d^\top \mathbf{P}_d = \mathbf{P}_d^2 = \mathbf{P}_d$. With $\mathbf{P} = \mathbf{I}_6$, the left-hand side of Equation (40) becomes $\mathbf{P}_d - \mathbf{I}_6 = -\mathbf{P}_o \preceq 0$, because $\mathbf{P}_o$ is positive semidefinite. □

Three consequences follow. First, feasibility of Equation (40) certifies that all eigenvalues of $\mathbf{A}_{cl}$ lie in the closed unit disc: zero on the observable subspace, one on the degenerate subspace. Second, with the common quadratic function $V(\mathbf{e}) = \mathbf{e}^\top \mathbf{e}$ (in the chosen tangent-space coordinates; the rotational and translational components are not unit-normalised, Section 3.8), the noise-free error dynamics satisfy, for every frame and every partition actually applied,

$$V(\mathbf{e}_k) - V(\mathbf{e}_{k-1}) = -\left\|\mathbf{P}_{o,k}\mathbf{e}_{k-1}\right\|^2 \leq 0, \tag{41}$$

so, in the absence of source errors, the correction never amplifies the error, and mode switches between consecutive frames cannot destabilise the estimator. If, in addition, the degenerate direction itself changes between frames, the composition of two projections is a contraction: when the degenerate subspace is 1D with unit direction $\mathbf{v}_{d,k}$, one has $\left\|\mathbf{P}_{d,k+1}\mathbf{P}_{d,k}\right\| = \left|\mathbf{v}_{d,k+1}^\top \mathbf{v}_{d,k}\right|$, which is strictly below one unless the degenerate direction is unchanged. Only the persistently degenerate component of the error is not attenuated. Third, the inequality in Equation (40) cannot be made strict whenever the degenerate subspace is non-empty: for any $\mathbf{P} \succ 0$ and any non-zero $\mathbf{z}$ in $\mathcal{D}$, $\mathbf{z}^\top\left(\mathbf{A}_{cl}^\top \mathbf{P}\mathbf{A}_{cl} - \mathbf{P}\right)\mathbf{z} = \mathbf{z}^\top\mathbf{P}\mathbf{z} - \mathbf{z}^\top\mathbf{P}\mathbf{z} = 0$. A strict (exponentially stable) certificate is, therefore, infeasible exactly along the unobservable directions—the discrete-time counterpart of the result of Arena et al. [38] that their LMI-based observer design is unfeasible when the system is unobservable.

The LMI of Equation (40) certifies the non-amplification mechanism of the correction; by itself it does not distinguish HERO-LO from a full-space correction. A full-space correction ($\mathbf{P}_o = \mathbf{I}_6$ at every frame) has $\mathbf{A}_{cl} = \mathbf{0}$ and satisfies Equation (40) strictly. The distinction lies in the directional structure of the disturbance term of Equation (36): a full-space correction accepts the registration error $\boldsymbol{\ell}_k$ in all directions, including its uncontrolled degenerate components—the accumulation mechanism behind the large endpoint errors of the full-space feature-based methods in Section 4.3.1—whereas the projection gain of Equation (35) admits only the observable component $\mathbf{P}_{o,k}\boldsymbol{\ell}_k$ into the estimate and leaves the degenerate subspace to the prior, whose contribution is governed by Equation (39). The stability certificate and the per-direction error bound are, thus, complementary: the former establishes that the correction cannot amplify the error under any mode sequence; the latter establishes why excluding the degenerate components of the registration error matters.

The relation to Arena et al. [38] is one of design philosophy rather than a transfer of their synchronization result. They design Luenberger observers for piecewise-linear chaotic systems by imposing, for each linear region i, the continuous-time LMI $(\mathbf{A}_i - \mathbf{K}\mathbf{C})^\top \mathbf{P} + \mathbf{P}(\mathbf{A}_i - \mathbf{K}\mathbf{C}) \prec 0$ with a common certificate (their Equations (8) and (11)), recovering the gain as $\mathbf{K} = \mathbf{P}^{-1}\mathbf{Q}$; the problem is feasible only if the pair $(\mathbf{A}, \mathbf{C})$ is observable, and infeasible for unobservable (identical) subsystems. Three differences delimit the transfer. Their certificate is a synthesis tool (the gain is obtained by solving the LMI), whereas here the gain is fixed by the degeneracy detector and the LMI is a verification certificate—constructive ($\mathbf{P} = \mathbf{I}_6$) and never solved online, so the runtime results of Section 4.5 are unaffected. Their inequality is strict and certifies asymptotic stability of the matched-region error dynamics (a necessary condition in their own analysis, complemented numerically and experimentally), whereas Equation (40) is non-strict: HERO-LO deliberately leaves the unobservable subspace unreconstructed, with closed-loop eigenvalue one, and hands it over to the heterogeneous prior—which is precisely why a strict certificate is infeasible. Finally, their setting is continuous-time synchronization of chaotic circuits,

whereas the present system is a discrete, per-frame estimation recursion whose certificate holds for every partition and every mode switch.

In summary, within the first-order tangent-space framework stated above and in the implementation coordinates of the present system, the analysis establishes: (i) the optimality of the projection update given the detected information structure (Proposition 1); (ii) per-direction, non-amplifying error propagation, with the observable component refreshed by the per-frame LiDAR accuracy and the degenerate component accumulating the prior increment errors (Proposition 2); and (iii) Lyapunov stability of the correction under arbitrary mode switching, with one-step convergence of the observable component in the noise-free case and critical stability of the persistently degenerate component (Proposition 3). It does not claim asymptotic convergence of the full pose error—which is impossible along directions that the LiDAR does not observe—nor uniformly bounded long-term error, which is governed by the quality of the heterogeneous prior.

3.8. Implementation Details

The complete HERO-LO pipeline is summarized in Algorithm 1. Several practical details are important for stable real-time implementation.

First, the heterogeneous odometry poses are temporally synchronized with LiDAR timestamps. If the odometry source has a higher output frequency than the LiDAR, interpolation is used to obtain the pose at the LiDAR scan time. If the odometry source has a lower output frequency, the nearest valid odometry interval is used only when the time difference is within an acceptable threshold.

Second, all heterogeneous odometry sources are converted into relative LiDAR-frame motion increments through the calibrated extrinsic transformation in Equation (6). This makes the interface independent of the specific odometry modality. Therefore, HERO-LO can use encoder odometry, VIO, or other 6D odometry estimates without modifying the LiDAR registration module.

Third, the local submap is maintained using selected keyframes. A new keyframe is inserted when the relative motion from the latest keyframe exceeds 2.0 m of translation or 1.0° of rotation. The submap is a count window rather than a time- or space-bounded window: it contains the point clouds of the most recent 20 keyframes, and it is rebuilt from these keyframes (0.6 m voxel grid) on every keyframe insertion, which keeps the computational cost of the scan-to-submap registration bounded.

Fourth, the degeneracy detection threshold is applied in the tangent space of $SO(3) \times \mathbb{R}^3$. Although rotational and translational components carry different units, the present implementation does not normalise the Hessian before the eigenvalue analysis, following the common practice in the literature (e.g., Zhang et al. [9]); the single threshold τ_a is applied directly to the normalised uncertainty indicators a_i. Unit normalisation or per-component thresholds are possible extensions whose influence is not explored here and is left to future work. In the experiments, the same detection and projection mechanism and the same parameter set are used for all compared HERO-LO variants to ensure a fair ablation study.

Fifth, the principal parameters of the HERO-LO implementation used in all experiments are summarised in the Supplementary Materials (Table S6). In brief: the degeneracy threshold is $\tau_a = 0.0315$, calibrated as described in Section 3.5; the regularisation is $\epsilon = 10^{-8}$ and $\sigma_r^2 = 1$ in the detection stage; the robust loss is a Huber kernel with $\delta = 0.08$ m; the scan-to-submap registration performs three data-association rounds with 20 Gauss–Newton iterations each (60 iterations per frame, no convergence test), which, together with the absence of RANSAC and random initialisation, makes the execution fully deterministic (Section 4.1.4); the residual noise model is estimated online and adaptively from the local geometry (a fit through the five nearest points); keyframes are inserted at 2.0 m or 1.0° of

relative motion, and the submap contains the 20 most recent keyframes (0.6 m voxel grid); and no unit normalisation of the Hessian is applied before the eigenvalue analysis (a single threshold is applied directly to a_i).

Finally, HERO-LO does not rely on GNSS or loop closure for online odometry estimation. GNSS is only used as an external reference for evaluation in the experiments.

4. Experiments and Analysis

This section evaluates the proposed HERO-LO framework on a real-world tunnel dataset collected using a ground mobile robot. The experiments are designed to answer the following questions: (1) whether HERO-LO can maintain stable odometry estimation in tunnel-like LiDAR-degenerate environments; (2) whether HERO-LO is competitive with representative LiDAR-only, LiDAR-inertial, and LiDAR-visual-inertial odometry systems; (3) whether the observable-subspace correction and the heterogeneous odometry prior are both necessary for robust tunnel traversal; and (4) whether HERO-LO can be deployed with different heterogeneous odometry sources.

Unless otherwise specified, HERO-LO denotes the default configuration, namely HERO-LO-S2S/KO-OSC. In this configuration, scan-to-scan (S2S) LiDAR odometry is used as the motion prior in non-degenerate segments, encoder-based kinematic odometry (KO) is used as the heterogeneous prior when LiDAR degeneracy is detected, and observable-subspace correction (OSC) is enabled.

4.1. Experimental Setup

4.1.1. Data Collection Platform

The data collection platform is an AGILEX SCOUT 2.0 ground mobile robot (AgileX Robotics, Dongguan, China) equipped with a 16-beam LiDAR (RoboSense RS-LiDAR-16, RoboSense, Shenzhen, China), wheel encoders, a stereo camera with an embedded IMU, an AHRS, and a GNSS/RTK (real-time kinematic) receiver. The LiDAR provides 3D point clouds at 20 Hz. The encoder-based kinematic odometry is built into the robot chassis. It estimates the forward linear velocity and yaw angular velocity from wheel encoders at 20 Hz, and can also output odometry data conforming to the unified 6D pose interface. The stereo camera is a STEREOLABS ZED2 (Stereolabs, San Francisco, CA, USA), which provides stereo images at 20 Hz and IMU measurements at 200 Hz. The AHRS is an XSENS MTi-30 (Xsens, Enschede, The Netherlands), which provides attitude measurements at 400 Hz. The GNSS/RTK receiver provides single-point positioning measurements at 4 Hz.

Table 1 summarizes the sensor and computing configuration used in the experiments. During data collection, an ARK-3510L-00A1E onboard computer (Advantech, Taipei, Taiwan, China) equipped with an Intel i7-3610QM quad-core CPU and 16 GB RAM was used to record all sensor data. All compared algorithms were evaluated offline by replaying the recorded dataset on a separate test computer equipped with an Intel Core i9-13900K CPU and 64 GB RAM. This offline replay setting ensures that all methods are tested using identical sensor data and the same evaluation platform. A companion video (Video S1 in the Supplementary Materials) visualizes the dynamic evolution of the estimated trajectories and point-cloud maps of A-LOAM, LeGO-LOAM, F-LOAM, DLO, KISS-ICP, GenZ-ICP, LIO-SAM, FAST-LIO2, LVI-SAM, HERO-LO, and its variants on the tunnel dataset.

Table 1. Sensor and computing configuration used in the experiments.

Item	Model/Source	Output/Configuration
LiDAR	ROBOSENSE RS-LiDAR-16	3D point cloud, 20 Hz
Encoder-based odometry	AGILEX SCOUT 2.0 chassis	Linear velocity and yaw rate, 20 Hz; 6D odometry output
Stereo camera and IMU	STEREOLABS ZED2	Stereo images, 20 Hz; IMU, 200 Hz
AHRS	XSENS MTi-30	3D attitude, 400 Hz
GNSS/RTK	GNSS/RTK receiver	Single-point positioning data, 4 Hz
Data recording computer	ARK-3510L-00A1E	Intel i7-3610QM, 16 GB RAM
Evaluation computer	Desktop workstation	Intel Core i9-13900K, 64 GB RAM

4.1.2. Tunnel Dataset

A real-world tunnel dataset was collected on a public mountain road using the above mobile robot platform. The robot began the sequence stationary at an open-air site before the tunnel, moving forward continuously between 22 s and 231 s; it entered the tunnel at about 59 s and left it at about 144 s, so the in-tunnel traversal lasted about 85 s. It stopped at about 231 s and remained stationary until the end of the sequence (about 256 s in total). The robot never reversed, and the road surface was approximately flat asphalt.

Figure 2 and Table 2 show an overview of the tunnel dataset. The background is a Google Earth screenshot, on which the approximate trajectory of the mobile robot during data collection is marked in yellow. The tunnel is located on a mountain road. Two onboard camera images, captured at approximately 60 s and 100 s, are overlaid to show representative scenes along the route. The tunnel has a nearly straight longitudinal structure, and its cross-section is similar to a semi-circular shape. Such a structure leads to repeated geometric patterns and weak constraints along the tunnel direction, making it a typical corridor-like LiDAR-degenerate environment. Consistent with this geometry, the online degeneracy detector of HERO-LO flagged approximately 1220 of the approximately 1240 frames between 83 s and 145 s; this degenerate segment overlaps the latter part of the in-tunnel interval and corresponds to the weakly constrained central portion of the tunnel. The 85–135 s evaluation window used for the local stability metrics (Section 4.2) is the central 50 s of this degenerate segment. The detector flags a single degenerate direction in this segment—the forward translation direction of the LiDAR frame, i.e., the tunnel axis—and no degenerate direction outside this interval, as recorded by the per-frame Hessian eigenvalue logs of the original runs.

Table 2. Overview of the tunnel dataset.

Item	Description
Total duration	Approximately 256 s
Initial standstill	0–22 s
Continuous forward motion	22–231 s
Tunnel entry/exit	Approximately 59 s/144 s
In-tunnel duration	Approximately 85 s
Degenerate segment (online detection)	Approximately 83–145 s
Stop at the endpoint	Approximately 231 s; stationary 231–256 s
Road condition	Flat asphalt road
Reference positioning	Single-point GNSS outside the tunnel, 4 Hz
GNSS horizontal accuracy	Approximately 2.5 m CEP

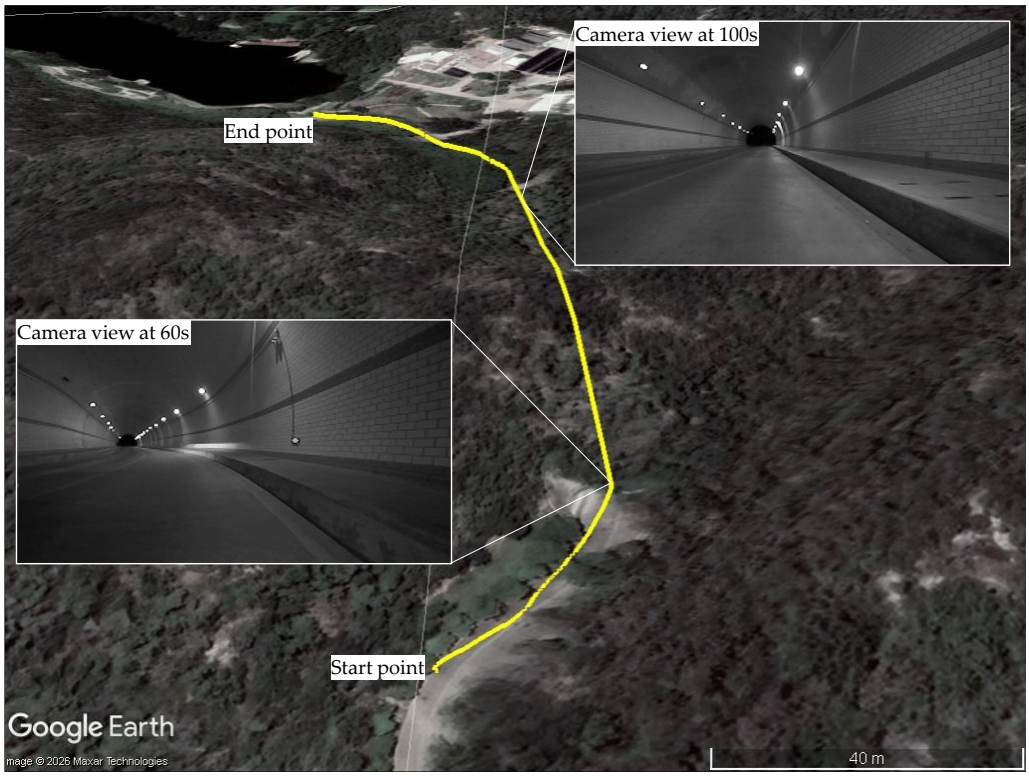


Figure 2. Tunnel dataset overview. The Google Earth background shows the mountain-road environment, and the yellow curve indicates the approximate trajectory of the robot during data collection. The overlaid onboard camera images correspond approximately to 60 s and 100 s.

Although the platform was equipped with an RTK system, reliable RTK positioning was unavailable during the experiment because mountain occlusion blocked the RTK differential signal and prevented carrier-phase differential resolution. Therefore, only single-point GNSS measurements were available outside the tunnel, with a horizontal positioning accuracy of approximately 2.5 m circular error probable (CEP). Consequently, all distance-related descriptions of the dataset are given in time terms, and the GNSS measurements are used only as a coarse positioning reference for evaluating global trajectory consistency, rather than as high-precision ground truth.

4.1.3. Compared Methods

To ensure a credible and reproducible evaluation, the compared odometry methods are selected from representative publicly available open-source algorithms, rather than re-implemented by ourselves. These methods cover the main technical categories of LiDAR-based and multi-sensor odometry. The compared methods are divided into four groups. The first group consists of representative feature-based LiDAR odometry methods, including A-LOAM [16], an advanced open-source implementation of LOAM [26], LeGO-LOAM [17], and F-LOAM [18]. The second group consists of representative direct or dense LiDAR odometry methods, including CT-ICP [19], DLO [6], KISS-ICP [7], and GenZ-ICP [20]. The third group consists of representative LiDAR-inertial and LiDAR-visual-inertial odometry methods, including LIO-SAM [21], FAST-LIO2 [5], and LVI-SAM [22]. The fourth group consists of heterogeneous odometry baselines, including KO, KAO, and VIO. For the open-source methods, their public implementations were used in the experiments, which helps avoid potential bias introduced by self-implementation and improves the reproducibility and credibility of the comparison.

Here, KO denotes encoder-based kinematic odometry. KAO denotes AHRS-aided kinematic odometry, in which the attitude information from the AHRS is used to improve the kinematic odometry estimate. VIO denotes stereo visual-inertial odometry obtained from the ZED2 camera and its embedded IMU.

Table 3 summarizes the compared methods and their sensor inputs. It should be noted that different methods use different sensor configurations. Therefore, the purpose of the comparison is not to claim superiority under identical sensor inputs, but to evaluate the practical robustness of different representative odometry systems in the same tunnel-like environment. To improve reproducibility, the exact repositories, versions (commit hashes), and principal parameters of all compared methods, together with the complete software environment, are reported in the Supplementary Materials (Tables S1–S3).

Table 3. Compared methods and sensor inputs.

Category	Method	Main Input Sensors	Note
Feature-based LO	A-LOAM	LiDAR	LOAM-style feature-based method
Feature-based LO	LeGO-LOAM	LiDAR	Ground-optimized feature-based LOAM
Feature-based LO	F-LOAM	LiDAR	Fast feature-based LiDAR odometry
Direct/dense LO	CT-ICP	LiDAR	Continuous-time ICP-based odometry
Direct/dense LO	DLO	LiDAR	Dense point-cloud-based odometry
Direct/dense LO	KISS-ICP	LiDAR	Point-to-point ICP-based odometry
Direct/dense LO	GenZ-ICP	LiDAR	Degeneracy-robust LiDAR odometry
LIO/LVIO	LIO-SAM	LiDAR + IMU	Smoothing-based LIO
LIO/LVIO	FAST-LIO2	LiDAR + IMU	Direct tightly coupled LIO
LIO/LVIO	LVI-SAM	LiDAR + IMU + camera	Smoothing-based LVIO
Heterogeneous baseline	KO	Encoder	Kinematic odometry
Heterogeneous baseline	KAO	Encoder + AHRS	AHRS-aided kinematic odometry
Heterogeneous baseline	VIO	Stereo camera + IMU	Visual-inertial odometry
Proposed	HERO-LO	LiDAR + KO	Default HERO-LO-S2S/KO-OSC

For the ablation study, HERO-LO variants are named in the form HERO-LO-[prior policy]-[correction policy]. The prior policy indicates the source of the motion prior, such as scan2scan odometry (S2S), KO, KAO, VIO, or S2S/KO. The correction policy indicates whether observable-subspace correction is enabled, denoted as OSC (correction enabled) or NoOSC (correction disabled). Detailed definitions of these variants are provided in Section 4.4.

4.1.4. Evaluation Protocol and Reproducibility

The tunnel sequence was recorded in a single physical traversal on 4 December 2020, and every algorithm was evaluated offline by replaying the identical recorded data on the same evaluation workstation. This protocol guarantees that all compared methods receive exactly the same sensor input. To verify the reproducibility of the reported results, every algorithm and every HERO-LO variant was executed three times under the same sequence and platform (fresh process launches). The three runs reproduced the same outcomes for every method: all completed runs produced bit-identical trajectory files (identical file hashes), CT-ICP failed consistently at approximately 23 s (before the tunnel entrance), and LIO-SAM diverged consistently at approximately 133 s (inside the tunnel). The HERO-LO implementation contains no random sample consensus (RANSAC), no random initialization, and no unordered parallel floating-point accumulation, so its execution is fully deterministic given identical inputs; the repeated executions of the third-party implementations likewise produced identical trajectories, which empirically confirms

reproducibility under this evaluation protocol. Because repeated executions reproduce identical trajectories, mean and standard deviation over repeated runs are not applicable, and the values reported in Section 4 correspond to a single deterministic run. We note that this addresses the repeatability of the offline algorithm execution; the repeatability of the method across multiple physical traversals is listed as future work.

4.2. Evaluation Metrics

Because high-precision ground truth is unavailable in the tunnel dataset, two types of evaluation metrics are used: endpoint horizontal error and local trajectory stability metrics.

4.2.1. Endpoint Horizontal Error

The endpoint horizontal error is used to evaluate the overall consistency of the estimated trajectory. Following the common practice in odometry evaluation [39], each estimated trajectory is first aligned with the GNSS trajectory segment before entering the tunnel. Let $\mathbf{p}_{i,xy}^{m}$ denote the horizontal position of method m at the GNSS timestamp i, and let $\mathbf{g}_{i,xy}$ denote the corresponding GNSS horizontal position. A 2D rigid transformation $(\mathbf{R}_{xy}^{m}, \mathbf{t}_{xy}^{m})$ is estimated using the pre-tunnel alignment interval $\mathcal{I}_{\text{align}}$:

$$(\mathbf{R}_{xy}^{m}, \mathbf{t}_{xy}^{m}) = \arg \min_{\mathbf{R}\in SO(2), \mathbf{t}\in\mathbb{R}^2} \sum_{i\in\mathcal{I}_{\text{align}}} \left\| \mathbf{R}\mathbf{p}_{i,xy}^{m} + \mathbf{t} - \mathbf{g}_{i,xy} \right\|^2. \tag{42}$$

The endpoint horizontal error is then computed as

$$e_{\text{end}}^{m} = \left\| \mathbf{R}_{xy}^{m}\mathbf{p}_{N,xy}^{m} + \mathbf{t}_{xy}^{m} - \mathbf{g}_{N,xy} \right\|, \tag{43}$$

where N denotes the endpoint timestamp used for evaluation. A smaller endpoint horizontal error indicates better global consistency with the GNSS reference. Since the GNSS reference is not high-precision ground truth, this metric should be interpreted as a coarse global evaluation metric.

Time and Coordinate References

All sensors were timestamped by the system clock of the onboard data-logging computer, which provides a unified time base across the LiDAR, encoders, AHRS, camera, and GNSS receiver. Raw GNSS fixes (longitude, latitude, and height on the World Geodetic System 1984 (WGS84)) were converted to a local east-north-up (ENU) Cartesian frame with the `LocalCartesian` transformation of the GeographicLib library (version 1.50.1), with the origin at the first valid fix. The LiDAR odometry world frame M is defined with its x-axis as the projection of the static initial AHRS frame's x-axis onto the horizontal plane and its z-axis aligned with the upward direction; the rotation between the initial AHRS frame A_0 and M is $\mathbf{R}_{M,A_0} = \mathbf{R}_y(\theta_{\text{pitch}})\mathbf{R}_x(\theta_{\text{roll}})$, where the roll and pitch angles are averaged over an approximately 10 s static initialization period. The world frames of the other odometry sources are aligned to M through the pre-calibrated extrinsics. The frames E and M share the same horizontal plane, and the 2D rigid transformation between them is estimated with the alignment procedure in Equation (42).

Alignment Interval and Endpoint Evaluation

The alignment interval $\mathcal{I}_{\text{align}}$ is 27–57 s, which lies after the GNSS receiver initialization (first approximately 20 s) and entirely before the tunnel entrance (approximately 59 s); it contains approximately 120 single-point GNSS fixes at the receiver's 4 Hz output rate. For each method, the estimated poses are interpolated to the GNSS timestamps (linear interpolation of translations, spherical linear interpolation of rotations) and com-

pensated by the pre-calibrated lever arm of the GNSS antenna phase centre in the LiDAR frame, $\mathbf{t}_{L,G} = [-0.6953, 0.0024, -0.2920]^\top$ m, before the 2D alignment. The endpoint horizontal error is computed at the final frame of the sequence (approximately 256 s) from the single GNSS observation nearest to that timestamp; no temporal averaging of GNSS fixes is applied. Because each single-point fix has a horizontal accuracy of approximately 2.5 m CEP, the endpoint error is a coarse global-consistency indicator: differences smaller than approximately 1 m between the best-performing methods should not be interpreted as significant.

4.2.2. Local Trajectory Stability Metrics

To evaluate local frame-to-frame stability, three metrics are defined based on the maximum unloaded speed of the mobile robot. Let $\mathbf{p}_i^m$ and $\mathbf{p}_{i-1}^m$ denote two consecutive estimated 3D positions from method m. The estimated frame-to-frame translation is

$$d_i^m = \left\| \mathbf{p}_i^m - \mathbf{p}_{i-1}^m \right\|. \tag{44}$$

The allowed maximum translation between two consecutive frames is defined as

$$d_i^{\lim} = (1+\gamma) v_{\max} \Delta t_i, \tag{45}$$

where $v_{\max} = 1.5$ m/s is the maximum unloaded speed of the robot, Δt_i is the time interval between two consecutive poses, and $\gamma = 20$ % is a tolerance factor used to absorb time synchronization errors, pose estimation noise, and small velocity fluctuations.

A frame-to-frame translation anomaly is detected when

$$d_i^m > d_i^{\lim}. \tag{46}$$

The anomaly indicator is defined as

$$a_i^m = \begin{cases} 1, & d_i^m > d_i^{\lim}, \\ 0, & \text{otherwise.} \end{cases} \tag{47}$$

Based on this definition, the frame-to-frame translation anomaly count N_a, anomaly rate R_a, and accumulated exceedance magnitude E_a are computed as

$$N_a^m = \sum_{i \in \mathcal{I}} a_i^m, \tag{48}$$

$$R_a^m = \frac{N_a^m}{|\mathcal{I}|} \times 100\%, \tag{49}$$

$$E_a^m = \sum_{i \in \mathcal{I}} a_i^m \left(d_i^m - d_i^{\lim} \right), \tag{50}$$

where $\mathcal{I}$ denotes the set of consecutive pose indices within the selected evaluation time interval.

In this paper, three local intervals of equal length are selected for evaluation: (1) the pre-degeneracy interval from 20 s to 70 s, which covers the approach segment and the first approximately 11 s inside the tunnel, during which no degeneracy is detected; (2) the degenerate interval from 85 s to 135 s, the central 50 s of the degeneracy-detected segment (approximately 83–145 s); and (3) the post-degeneracy interval from 160 s to 210 s, after the tunnel exit (approximately 144 s). Three equal-length windows are used so that the metrics are comparable across windows, and separate windows are necessary because a single overall computation would mask the contrasting behaviour of the estimators in degenerate

versus non-degenerate segments. The metrics N_a, R_a, and E_a are local trajectory stability metrics rather than ground-truth-based local accuracy metrics. Larger values indicate more frequent or more severe local motion anomalies. However, smaller values do not necessarily imply higher local accuracy.

Influence of the Output Rate on the Anomaly Metrics

The compared methods produce poses at different rates: LeGO-LOAM, A-LOAM, and LVI-SAM output at approximately 2.7 Hz, 12.6 Hz, and 5.5 Hz on average over the sequence (a consequence of their internal pipelines; no internal code was modified), VIO is deliberately capped at 10 Hz, and HERO-LO and the remaining methods output at approximately 20 Hz. The anomaly criterion is scale-invariant in the inter-frame interval: both d_i^m and $d_i^{\lim} = 1.2\, v_{\max} \Delta t_i$ scale linearly with Δt_i, so an anomaly occurs if and only if the ratio $d_i^m / \Delta t_i$ exceeds $1.2\, v_{\max}$, independent of the sampling rate. Consequently, under a uniform re-sampling of a trajectory with linear position interpolation, every sub-interval inherits the anomaly status of its parent interval, so the anomaly rate R_a and the accumulated exceedance E_a are unchanged and only the anomaly count N_a scales with the sampling density. To verify this empirically, the trajectories of the four low-rate methods were interpolated onto the 20 Hz LiDAR grid and the metrics recomputed (Tables S4 and S5 in the Supplementary Materials). For the near-uniform methods (LeGO-LOAM, LVI-SAM, VIO), R_a changed by at most 0.05 percentage points and E_a remained unchanged to within 0.001 m, while N_a scaled with the frame density (e.g., LeGO-LOAM $1 \rightarrow 8$ in the 20–70 s window, LVI-SAM $7 \rightarrow 26$ in the 85–135 s window, VIO $3 \rightarrow 6$). For A-LOAM, whose dropped frames occur in non-anomalous intervals, the anomaly events remain unchanged (N_a $56 \rightarrow 56$, E_a $0.60 \rightarrow 0.60$ m) and R_a changes from 6.4% to 5.6% only because the denominator (the number of frames) increases when the missing frames are re-inserted. The findings of this paper are, therefore, unaffected by the differences in output rate among the compared methods.

4.3. Comparison with Representative Publicly Available Open-Source Algorithms

Table 4 summarizes the overall odometry performance of all compared algorithms, including the running status and endpoint horizontal error. The status is categorized as Completed, Failed, Diverged, or Partial. "Completed" means that the method outputs a full trajectory for the entire sequence. "Failed" means that the method stops before completing the sequence. "Diverged" means that the output trajectory becomes physically unreasonable. "Partial" means that only the valid segment before failure or divergence is used for evaluation.

Table 5 reports the local trajectory stability metrics of the compared algorithms. The metrics N_a, R_a, and E_a quantify frame-to-frame translational anomalies in the pre-degeneracy, degenerate, and post-degeneracy intervals. Larger values indicate more frequent or more severe local motion anomalies. The symbol "–" indicates that the corresponding local stability metrics are not reported because the trajectory is incomplete, divergent, or dominated by large-scale drift.

Figure 3 shows the trajectory comparison of all compared algorithms. It provides an overall visual comparison of global trajectory consistency. Detailed comparisons with different categories of methods are presented in Sections 4.3.1–4.3.3.

Table 4. Overall odometry performance of all compared algorithms on the tunnel dataset.

Category	Method	Status	Endpoint Horizontal Error/m
Feature-based LO	A-LOAM	Completed	104.1
Feature-based LO	LeGO-LOAM	Completed	108.4
Feature-based LO	F-LOAM	Completed	43.8
Direct/dense LO	CT-ICP	Failed at approximately 23 s	–
Direct/dense LO	DLO	Completed	23.7
Direct/dense LO	KISS-ICP	Completed	2.8
Direct/dense LO	GenZ-ICP	Completed	3.1
LIO/LVIO	LIO-SAM	Diverged at approximately 133 s	Partial/56.4
LIO/LVIO	FAST-LIO2	Completed	2.8
LIO/LVIO	LVI-SAM	Completed	3.6
Proposed	HERO-LO	Completed	3.6

Note: single-point GNSS outside the tunnel is used as the reference, with a horizontal accuracy of approximately 2.5 m CEP. Endpoint errors of 2–4 m are of the same magnitude as this uncertainty: they show that the corresponding trajectory did not diverge, but they cannot support a fine accuracy ranking. Only endpoint errors far above this level (tens to hundreds of metres) indicate a clear deviation from the reference. The symbol "–" indicates that no endpoint error is reported (failure before the alignment interval or divergence).

Table 5. Local trajectory stability of the compared algorithms.

Category	Method	Pre-Degeneracy, 20–70 s			Degenerate, 85–135 s			Post-Degeneracy, 160–210 s		
		N_a	R_a/%	E_a/m	N_a	R_a/%	E_a/m	N_a	R_a/%	E_a/m
Feature-based LO	A-LOAM	56	6.4	0.60	0	0.0	0.00	0	0.0	0.00
Feature-based LO	LeGO-LOAM	1	0.8	0.02	0	0.0	0.00	0	0.0	0.00
Feature-based LO	F-LOAM	90	9.0	0.88	56	5.6	1.01	28	2.8	0.10
Direct/dense LO	CT-ICP	–	–	–	–	–	–	–	–	–
Direct/dense LO	DLO	157	15.7	2.48	332	33.2	25.98	161	16.1	1.75
Direct/dense LO	KISS-ICP	306	30.6	10.89	507	50.7	41.32	261	26.1	3.69
Direct/dense LO	GenZ-ICP	146	14.6	2.07	339	33.9	11.01	98	9.8	0.88
LIO/LVIO	LIO-SAM	–	–	–	–	–	–	–	–	–
LIO/LVIO	FAST-LIO2	17	1.7	0.05	38	3.8	0.27	20	2.0	0.06
LIO/LVIO	LVI-SAM	5	1.8	0.19	7	2.6	0.08	0	0.0	0.00
Proposed	HERO-LO	145	14.5	2.87	6	0.6	0.53	141	14.1	1.46

As shown in Table 4 and Figure 3, most feature-based and several direct LiDAR-only methods suffer from severe drift in the tunnel sequence, while the direct methods KISS-ICP and GenZ-ICP remain globally consistent but exhibit substantially elevated local anomaly rates inside the tunnel. HERO-LO completes the entire sequence and achieves an endpoint horizontal error in the same 2.8–3.6 m range as the best open-source LIO/LVIO and direct LiDAR odometry methods; these differences are not statistically significant given the approximately 2.5 m CEP reference uncertainty. The local stability results in Table 5 further show that HERO-LO substantially reduces frame-to-frame motion anomalies within the degenerate tunnel segment.

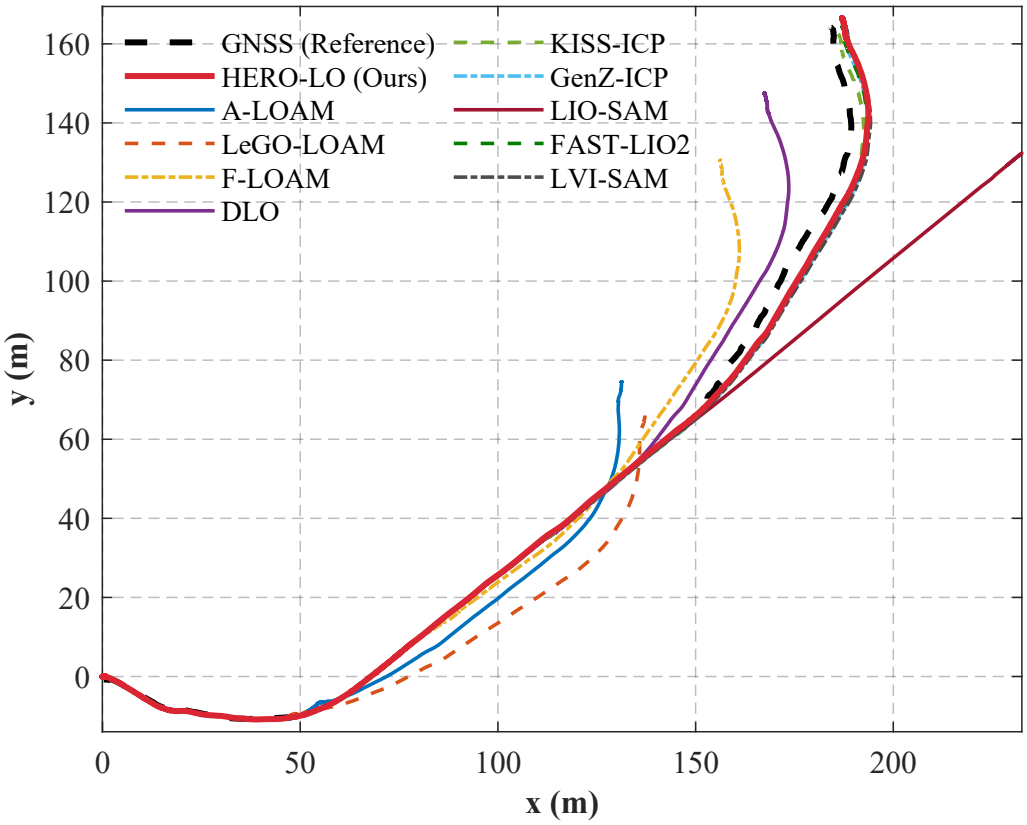


Figure 3. Trajectory comparison of all compared algorithms on the tunnel dataset.

4.3.1. Comparison with Feature-Based LiDAR Odometry

This subsection compares HERO-LO with feature-based LiDAR odometry methods, including A-LOAM, LeGO-LOAM, and F-LOAM. These methods rely on geometric edge and planar features extracted from LiDAR scans. In geometrically rich environments, such features can provide effective constraints for pose estimation. However, in tunnel-like environments, the repeated structure and weak geometric variation along the tunnel axis can make the LiDAR registration problem ill-conditioned.

As reported in Table 4, A-LOAM, LeGO-LOAM, and F-LOAM complete the sequence but produce large endpoint horizontal errors of 104.1 m, 108.4 m, and 43.8 m, respectively. Their trajectories also deviate significantly from the GNSS reference after passing through the tunnel, as shown in Figure 3. In contrast, HERO-LO maintains a continuous trajectory and obtains a much smaller endpoint error of 3.6 m.

The results indicate that conventional feature-based LiDAR odometry is vulnerable to tunnel degeneracy. Since these methods apply full LiDAR matching corrections even when some directions are weakly constrained, unreliable updates may be introduced along degenerate directions. HERO-LO addresses this issue by detecting degeneracy and applying LiDAR correction only in the observable subspace, while preserving the heterogeneous odometry prior in degenerate directions.

The endpoint errors of the three feature-based methods (104.1 m for A-LOAM, 108.4 m for LeGO-LOAM, and 43.8 m for F-LOAM) are one to two orders of magnitude larger than the uncertainty of the GNSS reference (approximately 2.5 m CEP), so they reflect systematic trajectory drift rather than reference error. The common mechanism is the application of full-space LiDAR corrections under weakly constrained axial geometry: each frame introduces a small unreliable correction along the tunnel axis, and these corrections accumulate into the large endpoint errors. The different magnitudes among the three methods reflect their different submap-management and feature-selection policies rather

than different failure mechanisms. HERO-LO avoids this accumulation mechanism: its endpoint error (3.6 m) is consistent with the reference uncertainty (note below Table 4) and, in contrast to the feature-based methods, indicates no systematic axial drift.

4.3.2. Comparison with Direct or Dense LiDAR Odometry

This subsection compares HERO-LO with direct or dense LiDAR odometry methods, including CT-ICP, DLO, KISS-ICP, and GenZ-ICP. Compared with feature-based methods, direct or dense methods generally use more point-cloud information and are often more robust when feature extraction is unreliable. Nevertheless, tunnel-like scenes remain challenging because the geometric constraints along the tunnel direction are inherently weak.

As shown in Table 4, CT-ICP fails at approximately 23 s—before the tunnel entrance—and cannot output a complete trajectory; no endpoint error is, therefore, reported. The failure occurs in the open outdoor segment and is, therefore, unrelated to tunnel degeneracy. The recorded run terminates with the terminal message "Error in ego-motion distance", i.e., the process is stopped by the ego-motion sanity check of the CT-ICP implementation: the continuous-time ego-motion model could not produce a consistent motion estimate under the geometry and motion conditions of that segment. This is a general robustness issue of the method in feature-sparse scenes and does not reflect its behaviour under degeneracy. DLO completes the sequence but has a large endpoint error of 23.7 m. KISS-ICP (2.8 m), GenZ-ICP (3.1 m), and HERO-LO (3.6 m) all achieve endpoint horizontal errors consistent with the uncertainty of the GNSS reference, indicating good global consistency; given the approximately 2.5 m CEP reference accuracy, these values are not ranked.

However, the local stability metrics in Table 5 show different behaviour inside the tunnel. KISS-ICP and GenZ-ICP exhibit significantly increased frame-to-frame anomalies in the degenerate interval. Specifically, KISS-ICP reaches $N_a = 507$, $R_a = 50.7\%$, and $E_a = 41.32$ m, while GenZ-ICP reaches $N_a = 339$, $R_a = 33.9\%$, and $E_a = 11.01$ m. In contrast, HERO-LO obtains $N_a = 6$, $R_a = 0.6\%$, and $E_a = 0.53$ m in the same interval. This indicates that HERO-LO provides more stable local motion estimation in the degenerate tunnel segment, while its endpoint error is statistically indistinguishable from those of KISS-ICP and GenZ-ICP given the reference uncertainty.

The two contrasting patterns are complementary and should be read together: KISS-ICP and GenZ-ICP remain globally consistent (endpoint errors of 2.8 m and 3.1 m) while exhibiting the highest local anomaly rates of all compared methods (50.7% and 33.9% of the frames in the degenerate interval). For KISS-ICP, 507 of the 1000 inter-frame motions in the interval exceed the physical motion limit, with median, mean, and maximum exceedances of 0.057 m, 0.082 m, and 0.386 m; for GenZ-ICP the corresponding values are 339 events and 0.025 m, 0.033 m, and 0.170 m. A maximum exceedance of 0.386 m corresponds to a single-frame translation of 0.48 m, i.e., an equivalent speed of about 9.5 m/s at the 50 ms frame interval, more than six times the maximum speed of the robot. The likely mechanism is that, under degenerate geometry, point-to-point (and metric-adaptive) registration admits many spurious local minima: the optimisation lands in a different minimum from frame to frame, producing abrupt pose jumps that partly cancel over the whole trajectory (hence the small endpoint error) but are individually physically unreasonable (hence the high anomaly rate). For applications that require smooth local motion estimates, such as inspection or control tasks, this combination of global consistency and local instability is undesirable, which is precisely the regime that the local stability metrics are designed to expose (Section 5).

Figure 4 further shows the local trajectory zoom-in comparison with direct or dense LiDAR odometry methods. The zoom-in results are consistent with the quantitative metrics and demonstrate that observable-subspace correction effectively suppresses abrupt local trajectory fluctuations in the tunnel.

4.3.3. Comparison with LIO and LVIO Methods

This subsection compares HERO-LO with representative LiDAR-inertial and LiDAR-visual-inertial odometry methods, including LIO-SAM, FAST-LIO2, and LVI-SAM. These methods use additional IMU or visual information and, therefore, have different sensor inputs from LiDAR-only methods. The comparison focuses on practical trajectory robustness under the same tunnel sequence.

As reported in Table 4, LIO-SAM diverges at approximately 133 s and only produces a partial trajectory, with a drift of 56.4 m at the divergence point. The recorded run of LIO-SAM repeatedly printed the message "Large velocity, reset IMU-preintegration!": inside the tunnel, the LiDAR odometry estimates jumped beyond the velocity gate of the implementation, forcing repeated resets of the IMU preintegration. Under sustained degeneracy the LiDAR factors carry little information along the tunnel axis, so the inertial drift is no longer corrected, and the trajectory diverges. This failure, which occurs inside the detector-flagged degenerate segment (approximately 83–145 s), is direct evidence that degeneracy can break even a tightly coupled LiDAR-inertial factor graph—the failure mechanism that HERO-LO is designed to address. FAST-LIO2 (2.8 m), LVI-SAM (3.6 m), and HERO-LO (3.6 m) complete the entire sequence with endpoint horizontal errors that are statistically indistinguishable given the reference uncertainty. Their trajectories are shown together in Figure 3.

The local stability results in Table 5 show that FAST-LIO2 and LVI-SAM maintain generally smooth local trajectories across the evaluated intervals. HERO-LO achieves the lowest anomaly count and anomaly ratio within the degenerate segment, with $N_a = 6$ and $R_a = 0.6\%$, but it shows larger local fluctuations than FAST-LIO2 and LVI-SAM in the pre-degeneracy and post-degeneracy intervals. These results suggest that HERO-LO provides competitive robustness in the degenerate tunnel segment compared with representative tightly coupled LIO/LVIO systems, even though its default configuration only uses a lightweight encoder-based kinematic odometry prior as the heterogeneous odometry source.

Two further observations on the intervals outside the degenerate segment deserve discussion. First, HERO-LO shows larger local fluctuations than FAST-LIO2 and LVI-SAM in the pre-degeneracy and post-degeneracy intervals ($R_a = 14.5\%$ in 20–70 s and 14.1% in 160–210 s, versus 1.7%/2.0% for FAST-LIO2 and 1.8%/0.0% for LVI-SAM). The per-event amplitudes of these anomalies are small (median exceedance of about 1 cm in both intervals, maximum 0.18 m), so this is a high-frequency, low-amplitude jitter. This behaviour is inherited from the feature-based scan-to-scan matcher (a LOAM-style registration), whose per-frame jitter is larger than that of tightly coupled filters; it is a property of the underlying matcher rather than of the degeneracy-handling mechanism. This is verified directly by the ablation data: HERO-LO-S2S/KO-OSC and HERO-LO-S2S-NoOSC share an identical trajectory up to the first degeneracy detection (approximately 83 s), because before that point both use the scan-to-scan prior and no projection is applied. The pre-degeneracy interval (20–70 s) lies entirely before this point, and the two variants indeed report identical metrics there ($N_a = 145$, $R_a = 14.5\%$, $E_a = 2.87$ m). The OSC mechanism is, therefore, exactly inactive when no degeneracy is detected, and the proposed modules do not alter the behaviour of normal, well-constrained scenes. Second, the post-degeneracy metrics differ among the ablation variants (e.g., 141 versus 114 events for HERO-LO-S2S/KO-OSC versus HERO-LO-S2S-NoOSC), because once the mechanism has activated, the variants carry different keyframes and submap states, accumulated during the degenerate segment, into the following interval. The contribution of this work targets the degenerate segment, where both the anomaly frequency and the accumulated exceedance are reduced (6 events, $E_a = 0.53$ m); reducing the matcher-induced jitter in well-constrained scenes could be addressed by replacing the underlying scan-to-scan matcher, which is outside the scope of this paper.

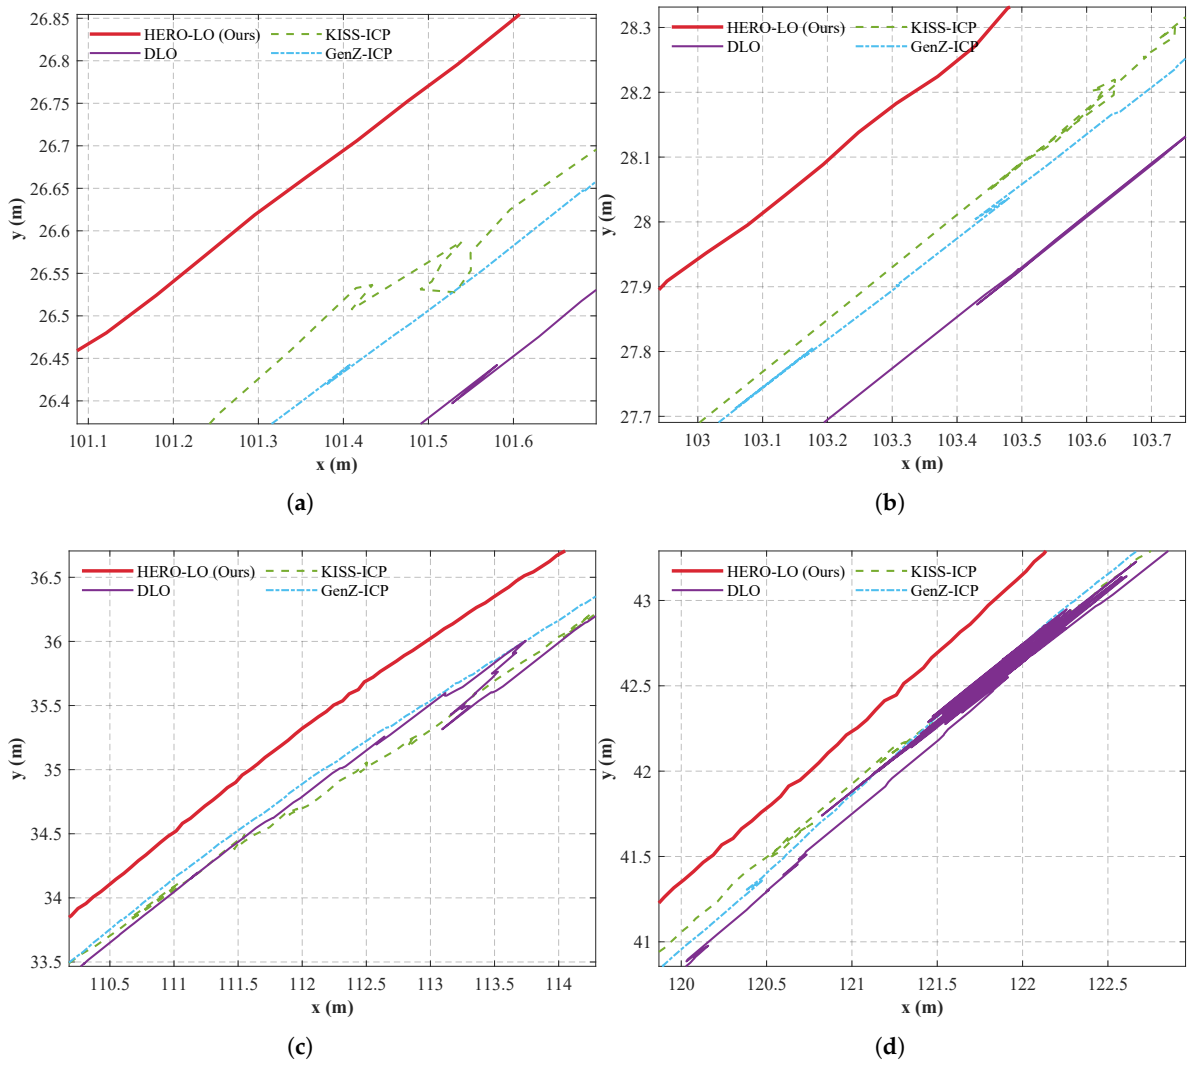


Figure 4. Local trajectory zoom-in comparison with direct or dense LiDAR odometry methods. (**a**) Zoomed-in trajectory segment around 101 s. (**b**) Zoomed-in trajectory segment around 103 s. (**c**) Zoomed-in trajectory segment around 112 s. (**d**) Zoomed-in trajectory segment around 121 s.

4.4. Ablation Study

This subsection analyzes the contributions of the observable-subspace correction (OSC) and the heterogeneous odometry prior. The HERO-LO variants evaluated in the ablation study are defined in Table 6. The overall odometry performance, local trajectory stability, and trajectories of these variants are reported in Tables 7 and 8, and Figure 5, respectively.

Table 6. Definition of HERO-LO variants used in the ablation study.

Variant	Prior Policy	OSC	Description	Purpose
HERO-LO-S2S-NoOSC	S2S	No	S2S prior with full LiDAR correction	Baseline without heterogeneous prior or OSC
HERO-LO-S2S-OSC	S2S	Yes	S2S prior with OSC	Tests OSC without a reliable heterogeneous prior
HERO-LO-KO-NoOSC	KO	No	KO prior with full LiDAR correction	Tests KO prior used mainly as initialization
HERO-LO-S2S/KO-OSC	S2S/KO	Yes	S2S in non-degenerate segments and KO in degenerate segments	Default HERO-LO; tests the combined effect of KO prior and OSC
HERO-LO-S2S/KAO-OSC	S2S/KAO	Yes	S2S in non-degenerate segments and KAO in degenerate segments	Tests the influence of KAO prior
HERO-LO-S2S/VIO-OSC	S2S/VIO	Yes	S2S in non-degenerate segments and VIO in degenerate segments	Tests the influence of VIO prior

Table 7. Overall odometry performance of HERO-LO variants.

Method	Heterogeneous Source	Status	Endpoint Horizontal Error/m
KO only	-	Completed	252.7
KAO only	-	Completed	27.6
VIO only	-	Completed	37.1
HERO-LO-S2S-NoOSC	-	Completed	77.6
HERO-LO-S2S-OSC	-	Completed	78.6
HERO-LO-KO-NoOSC	KO	Completed	78.9
HERO-LO-S2S/KO-OSC	KO	Completed	3.6
HERO-LO-S2S/KAO-OSC	KAO	Completed	2.5
HERO-LO-S2S/VIO-OSC	VIO	Completed	3.4

Note: the single-point GNSS reference (approximately 2.5 m CEP) does not support a fine ranking among the three fused variants (2.5–3.6 m).

Table 8. Local trajectory stability of HERO-LO variants.

Method	Pre-Degeneracy, 20–70 s			Degenerate, 85–135 s			Post-Degeneracy, 160–210 s		
	N_a	R_a/%	E_a/m	N_a	R_a/%	E_a/m	N_a	R_a/%	E_a/m
KO only	0	0.0	0.00	0	0.0	0.00	0	0.0	0.00
KAO only	0	0.0	0.00	0	0.0	0.00	0	0.0	0.00
VIO only	3	0.6	0.01	0	0.0	0.00	1	0.2	0.02
HERO-LO-S2S-NoOSC	145	14.5	2.87	122	12.2	4.43	114	11.4	0.78
HERO-LO-S2S-OSC	167	16.7	3.50	6	0.6	0.60	127	12.7	0.99
HERO-LO-KO-NoOSC	155	15.5	3.90	153	15.3	6.33	129	12.9	1.25
HERO-LO-S2S/KO-OSC	145	14.5	2.87	6	0.6	0.53	141	14.1	1.46
HERO-LO-S2S/KAO-OSC	155	15.5	3.90	20	2.0	1.16	146	14.6	1.47
HERO-LO-S2S/VIO-OSC	142	14.5	2.98	16	1.6	1.19	143	14.3	1.61

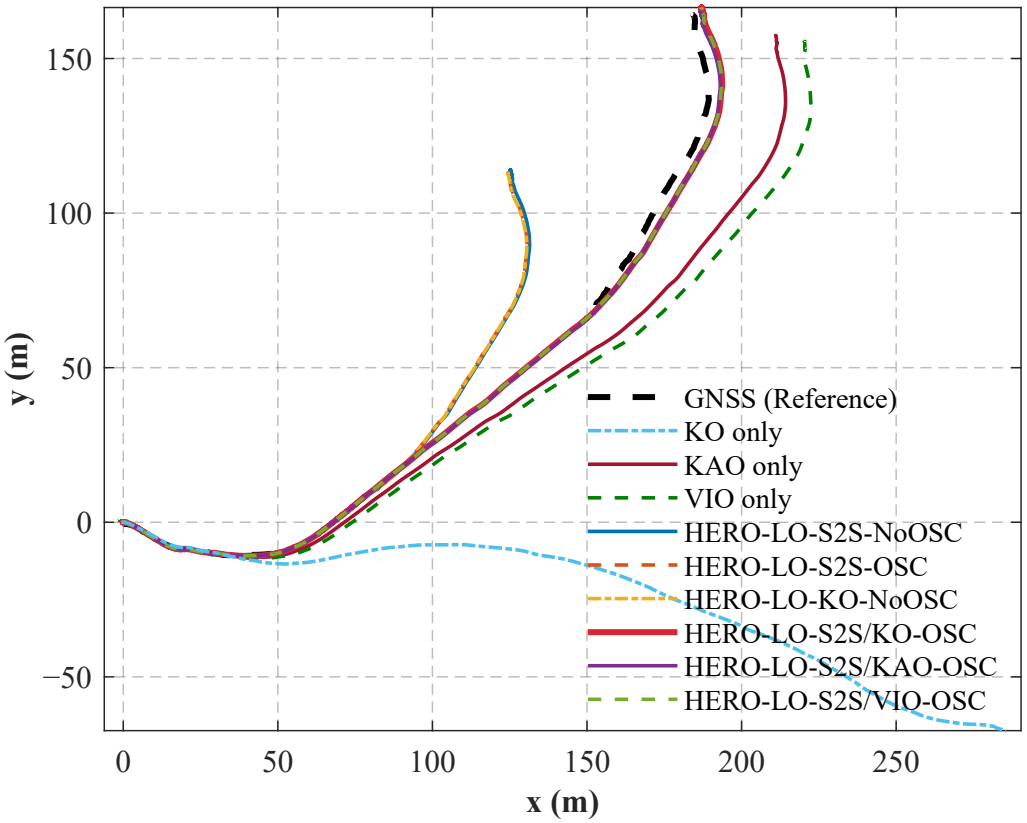


Figure 5. Trajectory comparison of HERO-LO variants with different heterogeneous odometry sources.

4.4.1. Effect of Observable-Subspace Correction Without Heterogeneous Prior

HERO-LO-S2S-NoOSC and HERO-LO-S2S-OSC are compared to evaluate the effect of OSC without a heterogeneous prior. As shown in Table 7, HERO-LO-S2S-NoOSC obtains an endpoint horizontal error of 77.6 m, while HERO-LO-S2S-OSC obtains 78.6 m. Thus, using OSC alone does not improve the overall endpoint error in this sequence.

However, Table 8 shows that OSC improves local stability in the tunnel segment. For HERO-LO-S2S-NoOSC, the degenerate-segment metrics are $N_a = 122$, $R_a = 12.2\%$, and $E_a = 4.43$ m. With OSC, HERO-LO-S2S-OSC reduces these values to $N_a = 6$, $R_a = 0.6\%$, and $E_a = 0.60$ m. These results indicate that OSC can suppress local instability in the tunnel, but it is not sufficient to reduce the global endpoint error when the retained prior is still generated by S2S LiDAR odometry.

The mechanism behind this result is that OSC does not change the source of the degenerate-direction component: when the prior is still the scan-to-scan estimate, the retained component is generated from the same weak geometry that makes the scan-to-submap registration degenerate, so it carries the same axial drift, and projection alone cannot repair it. The slightly larger endpoint error of HERO-LO-S2S-OSC (78.6 m) compared with HERO-LO-S2S-NoOSC (77.6 m) is consistent with the projection cost discussed in Section 3.6: discarding components along an approximately estimated subspace also removes part of the reliable correction. The strong local improvement in the degenerate segment itself (N_a reduced from 122 to 6) shows that, where the subspace estimate is well separated from the threshold, the projection works as intended.

4.4.2. Effect of Heterogeneous Prior Without Observable-Subspace Correction

HERO-LO-S2S-NoOSC and HERO-LO-KO-NoOSC are compared to evaluate the effect of using a heterogeneous prior without OSC. HERO-LO-KO-NoOSC uses KO as the prior but still applies the full LiDAR correction. Its endpoint horizontal error is 78.9 m, which is close to that of HERO-LO-S2S-NoOSC, 77.6 m.

The tunnel local stability metrics also do not improve without OSC. HERO-LO-KO-NoOSC has $N_a = 153$, $R_a = 15.3\%$, and $E_a = 6.33$ m in the degenerate segment, which are higher than those of HERO-LO-S2S-NoOSC. This indicates that using KO only as an initialization prior is insufficient when unreliable LiDAR corrections in degenerate directions are still fully applied.

Without OSC, the heterogeneous prior serves only as the initial guess of the registration: the full-space LiDAR correction is still applied, and in degenerate directions the optimisation can converge to a spurious minimum that overwrites the prior component that was supposed to be protected. The endpoint error (78.9 m) is, therefore, close to—and in this sequence slightly larger than—that of the scan-to-scan baseline (77.6 m), indicating that a prior used as initialisation without protection neither removes the drift nor prevents it from being amplified.

4.4.3. Combined Effect of Heterogeneous Prior and Observable-Subspace Correction

HERO-LO-S2S/KO-OSC evaluates the combined effect of the heterogeneous prior and OSC. Compared with HERO-LO-S2S-NoOSC, HERO-LO-S2S-OSC, and HERO-LO-KO-NoOSC, whose endpoint horizontal errors are 77.6 m, 78.6 m, and 78.9 m, respectively, HERO-LO-S2S/KO-OSC reduces the endpoint horizontal error to 3.6 m.

The local stability results also show improvement in the tunnel segment. HERO-LO-S2S/KO-OSC obtains $N_a = 6$, $R_a = 0.6\%$, and $E_a = 0.53$ m in the degenerate segment. Compared with HERO-LO-KO-NoOSC, E_a is reduced from 6.33 m to 0.53 m. Compared with HERO-LO-S2S-OSC, the degenerate-segment N_a and R_a remain the same, while E_a decreases from 0.60 m to 0.53 m.

These results show that the two components are complementary. The heterogeneous prior provides motion information in degenerate directions, while OSC prevents unreliable LiDAR corrections from directly overwriting these directions.

In the combined configuration, each pose component is carried by the information source that is most reliable in that direction under the tunnel geometry: the degenerate-direction component by the heterogeneous prior (hence the 0.6% anomaly rate in the degenerate segment) and the observable-direction components by the LiDAR scan-to-submap correction (hence the endpoint error consistent with the reference uncertainty). This is the per-direction error-source selection discussed in Section 3.6, now confirmed empirically.

4.4.4. Influence of Different Heterogeneous Odometry Sources

The last experiment compares different heterogeneous odometry sources, including KO, KAO, and VIO. When used alone, KO, KAO, and VIO produce endpoint horizontal errors of 252.7 m, 27.6 m, and 37.1 m, respectively. When combined with HERO-LO using the S2S/heterogeneous-prior policy and OSC, the endpoint horizontal errors are reduced to 3.6 m for HERO-LO-S2S/KO-OSC, 2.5 m for HERO-LO-S2S/KAO-OSC, and 3.4 m for HERO-LO-S2S/VIO-OSC.

Among the three fused variants, HERO-LO-S2S/KAO-OSC gives the numerically smallest endpoint horizontal error (2.5 m); however, given the approximately 2.5 m CEP reference uncertainty, the three fused variants are statistically indistinguishable by this metric. For local stability in the degenerate segment, HERO-LO-S2S/KO-OSC gives the lowest values, with $N_a = 6$, $R_a = 0.6\%$, and $E_a = 0.53$ m. The corresponding degenerate-segment metrics are $N_a = 20$, $R_a = 2.0\%$, and $E_a = 1.16$ m for HERO-LO-S2S/KAO-OSC, and $N_a = 16$, $R_a = 1.6\%$, and $E_a = 1.19$ m for HERO-LO-S2S/VIO-OSC.

Overall, the results indicate that HERO-LO can work with different heterogeneous odometry sources. The final performance depends on the selected prior source. In this sequence, KAO gives the numerically smallest endpoint error after fusion (within the reference uncertainty), while KO provides the best local stability in the degenerate segment among the fused variants.

These results also clarify the error-propagation chain of the framework. Used alone, KO accumulates 252.7 m of endpoint error over the sequence, because wheel-based kinematics drifts strongly in heading; KAO reduces this to 27.6 m, because the AHRS attitude corrects the dominant heading error; VIO reaches 37.1 m, with performance affected by the visual conditions. In the fused variants, the endpoint errors (2.5–3.6 m) are statistically indistinguishable, because the global endpoint error is governed by the LiDAR-observable part of the trajectory (common to all fused variants) together with the reference uncertainty; the differences between the priors then show up primarily in the local stability of the degenerate segment ($N_a = 6$ for KO, 16 for VIO, and 20 for KAO), which tracks the local smoothness of the prior. Practically, the framework is insensitive to the particular choice of prior source within this range—any reasonably stable 6D odometry source can be used—while the quality of the prior determines the upper bound of the degenerate-direction performance.

4.5. Runtime Analysis

The computational efficiency of HERO-LO is evaluated on a test computer equipped with an Intel Core i9-13900K CPU and 64 GB RAM. This computer is different from the onboard computer used for data collection. The runtime is measured over the degenerate tunnel segment from 85 s to 135 s, during which the robot is inside the tunnel and both the degeneracy detection and observable-subspace correction modules are active. The evaluated pipeline consists of four main stages: preprocessing, feature extraction, scan-to-scan

matching, and scan-to-submap matching. Degeneracy detection and observable-subspace correction are internal sub-modules of the scan-to-submap stage and are timed separately.

Table 9 reports the mean runtime, standard deviation, and runtime percentage of each module. Figure 6 shows the runtime distribution. The average total processing time is 9.5305 ms per frame (standard deviation 1.0685 ms), defined as the per-frame sum of the four main stages. The scan-to-submap matching stage is the main computational component, taking 7.5499 ms on average and accounting for 79.22% of the total runtime. Preprocessing takes 1.1734 ms, while feature extraction and scan-to-scan matching take 0.5155 ms and 0.2905 ms, respectively.

Table 9. Runtime analysis of HERO-LO.

Module	Mean/ms	Std./ms	Percentage/%
Preprocessing	1.1734	0.2381	12.31
Feature extraction	0.5155	0.0668	5.41
Scan-to-scan matching	0.2905	0.0970	3.05
Scan-to-submap matching	7.5499	1.0218	79.22
– Degeneracy detection *	0.0141	0.0018	0.15
– Observable-subspace correction *	0.0023	0.0022	0.02
Total	9.5305	1.0685	100.00

* Degeneracy detection and observable-subspace correction are internal sub-modules of the scan-to-submap matching stage; their runtimes are already included in the scan-to-submap matching value and are listed separately only to show that their overhead is negligible. They are not added to the total. For every frame, the total time is the per-frame sum of the four main stages; each mean and standard deviation is computed after removing the minimum and maximum over the frames of the 85–135 s degenerate segment, which is why the sum of the four stage means (9.5293 ms) differs slightly from the total mean (9.5305 ms). Percentages are the stage mean divided by the total mean.

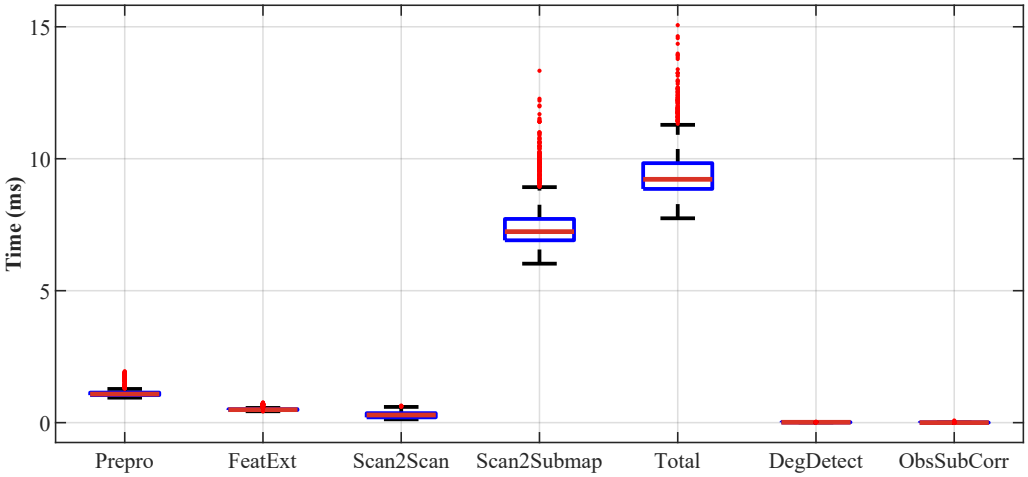


Figure 6. Runtime distribution of the main modules in HERO-LO. The blue boxes represent the interquartile range, while the black whiskers indicate the data range excluding outliers. The red lines and red dots denote the medians and outliers, respectively.

The degeneracy-related sub-modules introduce very small overhead. Degeneracy detection takes 0.0141 ms on average, and observable-subspace correction takes 0.0023 ms on average. Their runtime percentages are 0.15% and 0.02%, respectively. This is because these two sub-modules only operate on a 6×6 Hessian matrix and small tangent-space projection matrices.

Since the LiDAR input rate is 20 Hz, the corresponding frame interval is 50 ms. The average total runtime of 9.5305 ms is well below this interval, indicating that HERO-LO satisfies the real-time processing requirement under the tested computing platform. The

average total runtime of 9.5305 ms corresponds to a margin of approximately 5.2 times with respect to the 50 ms frame budget, and even the mean plus three standard deviations (approximately 12.7 ms) uses only about 25% of the budget. The two degeneracy-related sub-modules are cheap because they operate on a 6×6 Hessian and 6×6 projection matrices, operations of constant cost with respect to the number of points; the dominant per-frame cost is the construction of the Jacobian of the feature residuals.

4.6. Summary of Experimental Results

The experimental results demonstrate the robustness of HERO-LO in tunnel-like LiDAR-degenerate environments. Compared with representative publicly available open-source feature-based LiDAR odometry methods, HERO-LO avoids the severe endpoint drift observed in A-LOAM, LeGO-LOAM, and F-LOAM. Compared with direct or dense LiDAR odometry methods, HERO-LO achieves an endpoint horizontal error statistically indistinguishable from those of KISS-ICP and GenZ-ICP given the reference uncertainty, while providing substantially better local trajectory stability within the degenerate segment. In particular, although KISS-ICP and GenZ-ICP obtain endpoint errors below 4 m, their local anomaly rates in the degenerate segment reach 50.7% and 33.9%, respectively, whereas HERO-LO reduces this value to 0.6%.

Compared with representative LIO and LVIO methods, HERO-LO achieves an endpoint horizontal error in the same 2.8–3.6 m range as FAST-LIO2 and LVI-SAM, which is not a significant difference given the approximately 2.5 m CEP reference, while using only a lightweight encoder-based kinematic odometry prior in its default configuration. This suggests that explicitly preserving a stable heterogeneous prior in LiDAR-degenerate directions can be effective for robust tunnel traversal.

The ablation study further confirms the complementary roles of the heterogeneous odometry prior and observable-subspace correction (OSC). OSC alone improves local stability in the tunnel but cannot reduce the global endpoint drift when the retained prior is still generated by scan-to-scan LiDAR odometry. Conversely, using a heterogeneous prior without OSC is insufficient because unreliable full-space LiDAR corrections can still overwrite the prior in degenerate directions. Robust performance is achieved when the heterogeneous prior is preserved in LiDAR-degenerate directions and LiDAR matching corrections are applied only in the observable subspace.

The experiments with different heterogeneous odometry sources show that HERO-LO can work with KO, KAO, and VIO priors through the unified 6D pose interface. The final performance depends on the quality and stability of the selected prior source. In this sequence, the KAO-based HERO-LO variant gives the numerically smallest endpoint error (within the reference uncertainty), while the KO-based HERO-LO variant provides the best local stability in the degenerate segment.

The runtime analysis shows that HERO-LO satisfies real-time requirements on the tested platform. The average total processing time is approximately 9.5 ms per frame, which is well below the 50 ms frame interval of the 20 Hz LiDAR. The degeneracy detection and observable-subspace correction modules introduce negligible overhead, accounting for only 0.15% and 0.02% of the total runtime, respectively.

5. Discussion

5.1. Observed Failure Modes in Tunnel-like Environments

The experiments reveal four distinct failure patterns in tunnel-like environments. (i) Feature-based LiDAR odometry (A-LOAM, LeGO-LOAM, F-LOAM) completes the sequence but accumulates a systematic axial drift of 43.8–108.4 m, because full-space corrections are applied under weakly constrained axial geometry. (ii) Direct or dense

LiDAR odometry splits into two sub-patterns: CT-ICP fails early (approximately 23 s, before the tunnel entrance) when its ego-motion sanity check is triggered in the feature-sparse outdoor segment, a robustness issue unrelated to degeneracy; KISS-ICP and GenZ-ICP remain globally consistent (2.8 m and 3.1 m) but exhibit the highest local anomaly rates of all methods (50.7% and 33.9% in the degenerate interval), because the optimisation jumps between spurious local minima of the degenerate registration problem from frame to frame. (iii) The smoothing-based LIO-SAM diverges inside the tunnel (approximately 133 s, within the detector-flagged segment): the degenerate LiDAR estimates repeatedly trigger the implementation's IMU-preintegration reset, and the inertial drift is no longer corrected along the weakly constrained axis. (iv) HERO-LO completes the sequence with an endpoint error of 3.6 m (consistent with the reference uncertainty) and, among the methods that remain globally consistent, the lowest local anomaly rate in the degenerate segment (0.6%, versus 2.6–50.7% for the other four such methods: LVI-SAM, FAST-LIO2, GenZ-ICP, and KISS-ICP). The fact that the LIO-SAM divergence occurs inside the degenerate segment while HERO-LO remains locally stable in the same segment supports the interpretation that degeneracy—rather than the tunnel environment per se—is the core failure mechanism, and that a per-direction selection of the information source is an effective countermeasure.

HERO-LO addresses this problem by separating the LiDAR correction into observable and degenerate subspaces. In well-constrained directions, the LiDAR matching result is applied. In weakly constrained directions, the system preserves the heterogeneous odometry prior. This avoids directly applying unreliable LiDAR corrections along degenerate directions and explains the reduced local motion anomalies observed in the tunnel segment.

5.2. Relationship to Zhang-Type Degeneracy Handling

The observable-subspace correction in HERO-LO is closely related to the degeneracy handling method of Zhang et al. [9]. In a linearized least-squares problem, Zhang-type solution remapping analyzes the eigenstructure of the Hessian or approximate Hessian, identifies directions associated with small eigenvalues as poorly constrained, and combines the optimization update in well-conditioned directions with a prediction in degenerate directions. Under the same linearization, thresholding rule, and prediction source, the projection operation in HERO-LO is mathematically consistent with this principle.

Therefore, HERO-LO does not claim a new eigenvalue-based degeneracy criterion. Its contribution is to reinterpret and instantiate this principle for tunnel-like LiDAR odometry. First, the Hessian eigenstructure is interpreted from a state-estimation uncertainty perspective. Since the Hessian approximates the local information matrix, its inverse approximates the pose covariance. A small Hessian eigenvalue, therefore, corresponds to a long principal axis of the covariance ellipsoid, indicating large posterior uncertainty in the associated direction.

Second, HERO-LO performs degeneracy detection and correction in the tangent space of $SO(3) \times \mathbb{R}^3$. The LiDAR correction is represented as a local 6D perturbation, decomposed into observable and degenerate components, and applied through a manifold-based pose update. This makes the projection operation consistent with the pose representation used in the modern odometry system.

Third, HERO-LO specializes the prediction term in Zhang-type remapping as a deployable heterogeneous odometry prior. In non-degenerate segments, the system can use scan-to-scan LiDAR odometry as the motion prior. When tunnel-like degeneracy is detected, the prior preserved in the degenerate subspace is supplied by an auxiliary odometry source, such as KO, KAO, or VIO, through a unified 6D pose interface. From this viewpoint, HERO-LO can be regarded as a heterogeneous-prior-enhanced and manifold-based realization of Zhang-type solution remapping for tunnel-like LiDAR odometry.

The projection-based update of HERO-LO should also be distinguished from other degeneracy-aware designs. Residual re-weighting schemes (generalised ICP, and metric-adaptive ICP such as GenZ-ICP [20]) soften the influence of weak constraints by re-weighting individual residuals, but the resulting estimate still accepts part of the correction along the weakly constrained directions, and the protected directions are not stated explicitly. Tightly coupled LiDAR-inertial filters achieve robustness by adding an IMU (and camera) state and by changing the overall architecture; HERO-LO keeps a decoupled front end in which the auxiliary source is a switchable 6D pose input. A soft variant of the present update, in which a continuous gain replaces the hard observable/degenerate split, is possible; the hard selection is used here because it gives a clear protection semantics; as shown in Section 3.7, this hard projection-gain structure is also what admits a constructive common Lyapunov certificate (Proposition 3) together with the per-direction error bound of Equation (39).

5.3. Complementarity of the Prior and Observable-Subspace Correction

The ablation results show that the heterogeneous prior and OSC are complementary. When OSC is used without a reliable external prior, the local trajectory becomes smoother in the tunnel, but the endpoint error remains large because the retained prior is still affected by LiDAR degeneracy. When the KO prior is used without OSC, the system still accepts full-space LiDAR corrections, so the prior can be corrupted by unreliable updates in degenerate directions.

The combined configuration reduces both endpoint error and local anomaly metrics. This confirms that the heterogeneous prior provides motion information in weakly constrained directions, while OSC prevents unreliable LiDAR corrections from overwriting this prior.

5.4. Deployability with Heterogeneous Odometry Priors

HERO-LO uses a unified pose-level interface for heterogeneous odometry priors. This makes the framework independent of a specific auxiliary odometry source. In the experiments, KO, KAO, and VIO are all integrated into the same observable-subspace correction framework without changing the LiDAR matching pipeline.

Different priors show different error characteristics. KO is lightweight and locally smooth, but it accumulates drift when used alone. KAO improves the heading estimate and gives the numerically smallest endpoint error after fusion in this sequence (within the reference uncertainty). VIO also provides a usable prior, although its performance can be affected by visual conditions. These results indicate that HERO-LO is deployable with different odometry sources, while the final accuracy still depends on the quality and stability of the selected prior.

5.5. Interpretation of Local Stability Metrics

We note that the term "local stability" in these metrics denotes frame-to-frame kinematic consistency of the estimated trajectory; it is unrelated to the Lyapunov stability of the estimation-error dynamics analysed in Section 3.7. The local stability metrics used in this paper are not ground-truth-based local accuracy metrics. They measure physically abnormal frame-to-frame translations according to the motion limits of the robot. Therefore, small values of N_a, R_a, and E_a indicate better local consistency, but do not necessarily imply higher absolute trajectory accuracy.

For this reason, endpoint horizontal error and local stability metrics are used together. The endpoint error evaluates global trajectory consistency with respect to the reference, while the local stability metrics evaluate whether the estimated trajectory contains abrupt and physically unreasonable local fluctuations. This distinction is important because a

method may achieve a small endpoint error while still producing unstable local motion inside the tunnel.

The complementarity of the two metric families is illustrated by KISS-ICP: its endpoint error (2.8 m) is among the smallest of all methods, yet its local anomaly rate in the degenerate segment (50.7%) is the largest. A method can, therefore, appear globally successful while being locally unsuitable for smooth-motion applications; both families of metrics are needed for a complete characterisation.

5.6. Interpretation Against Application Requirements

The endpoint errors should also be interpreted against the intended applications. Endpoint errors above roughly 10 m (e.g., 23.7 m for DLO, 43.8 m for F-LOAM, 104.1 m and 108.4 m for A-LOAM and LeGO-LOAM, and 27.6–252.7 m for the standalone KO, KAO, and VIO baselines) correspond to road- or lane-level deviations and would be unacceptable for autonomous driving or high-definition map construction, which typically require sub-metre consistency. The metre-level endpoint errors of the best methods (2.8–3.6 m) are consistent with the uncertainty of the single-point GNSS reference: they show that these trajectories remained consistent with the coarse reference over the full recorded sequence without divergence. Because HERO-LO is a front-end odometry module without loop-closure detection or backend optimization, this result should not be interpreted as city-scale or long-duration global accuracy—drift would continue to accumulate over longer traversals. It is, however, directly relevant to GNSS-denied, short-duration applications that require real-time, drift-controlled local motion estimates, such as tunnel and underground-facility inspection, mining-drift mapping, and underground-parking navigation, where metre-level relative consistency over tunnel-scale traversals is typically sufficient. Centimetre-level in-tunnel positioning would require an RTK or total-station reference, which was not available in this study.

5.7. What the Present Evidence Establishes and Does Not Establish

Because the endpoint evaluation relies on a coarse single-point GNSS reference (approximately 2.5 m CEP) and the evaluation is based on one real tunnel sequence recorded in a single traversal, the present evidence should be interpreted with clear boundaries.

The experiments establish: (i) that HERO-LO remains locally stable in the degenerate tunnel segment, with an anomaly rate of 0.6% compared with up to 50.7% for the best LiDAR-only methods; (ii) that HERO-LO completes the full 256 s sequence without divergence, whereas CT-ICP fails and LIO-SAM diverges; (iii) that the endpoint horizontal error (3.6 m) is of the same magnitude as the reference uncertainty, i.e., the trajectory did not deviate from the coarse reference, without supporting a fine accuracy ranking; (iv) that the processing time (9.5 ms per frame) satisfies the real-time requirement with a large margin, with a negligible overhead of the degeneracy-handling modules; and (v) that the evaluation is deterministic and reproducible (bit-identical trajectories over three repeated executions, Section 4.1.4).

The present evidence does not establish: (i) the absolute positioning accuracy inside the tunnel, because no independent in-tunnel reference (RTK or total station) was available; (ii) that the endpoint error of HERO-LO is better than that of the other best-performing methods (2.5–3.6 m), since these differences are statistically indistinguishable given the reference uncertainty; (iii) the generalisation to other tunnel geometries (curved, sloping, intersecting, or variable-section tunnels) or to repeated physical traversals; and (iv) the drift behaviour over traversals substantially longer than the 256 s sequence. Obtaining an independent in-tunnel reference is currently not feasible, because the data-collection platform was dismantled after the original project concluded in 2021. As an intermediate

step, the complete evaluation configuration (software versions, commits, and parameters of all compared methods) is provided in the Supplementary Materials, and the recorded dataset and the implementation of HERO-LO are part of an ongoing study and are planned to be openly released in a future publication (Data Availability Statement; requests may be addressed to the corresponding author), so that such a reference can be established in future work. Tests in other tunnel geometries and a systematic sensitivity analysis of the degeneracy threshold are listed as future work in the Conclusions.

5.8. Limitations

Several limitations remain. First, the evaluation is conducted on one real, nearly straight tunnel sequence collected in a single traversal; curved, sloping, intersecting, and variable-section tunnels, corridor-like scenes, and repeated physical traversals should be tested in future work. Second, the endpoint evaluation relies on single-point GNSS outside the tunnel (approximately 2.5 m CEP) rather than centimetre-level full-trajectory ground truth; a total-station or surveyed-checkpoint reference network was beyond the scope of the original project, and the data-collection platform was dismantled after the original project concluded in 2021, so the sequence cannot currently be re-collected with better ground truth. Consequently, the absolute endpoint error is a coarse global-consistency indicator, and in-tunnel behaviour is assessed through ground-truth-free local consistency metrics only.

Third, HERO-LO requires a reasonably stable heterogeneous odometry prior in degenerate directions. If the auxiliary odometry source is severely biased or unstable, the final performance may degrade. Fourth, the current implementation depends on accurate spatial calibration and time synchronization between LiDAR and the heterogeneous odometry source. Fifth, HERO-LO focuses on odometry-level robustness and does not include loop closure or global pose-graph optimization, which could further reduce long-term drift. Finally, the observer and LMI analysis of Section 3.7 is a first-order, tangent-space interpretation in the implementation coordinates of the present system: it certifies non-amplification and per-direction error propagation of the correction step, not asymptotic convergence or uniformly bounded long-term error of the full estimation system.

6. Conclusions

This paper presented HERO-LO, a heterogeneous-odometry-enhanced robust LiDAR odometry framework for tunnel-like environments. The method builds on Zhang-type observable-subspace remapping, interprets LiDAR degeneracy from a state-estimation uncertainty perspective, and implements degeneracy detection and correction in the tangent space of $SO(3) \times \mathbb{R}^3$. The prior-preserved update is further interpreted as a projection-gain observer whose closed-loop error dynamics admit a common LMI-based Lyapunov certificate, making precise the per-direction, non-amplifying error propagation of the correction. LiDAR matching corrections are applied only in observable directions, while the prediction from a heterogeneous odometry prior, such as KO, KAO, or VIO, is preserved in degenerate directions.

Experiments on a real tunnel dataset show that HERO-LO improves the local trajectory consistency of LiDAR odometry within the degenerate tunnel segment compared with representative open-source LO, LIO, and LVIO methods. In the default KO-prior configuration, HERO-LO reduces the frame-to-frame anomaly rate inside the degenerate segment to 0.6% ($N_a = 6$, $E_a = 0.53$ m) and achieves an endpoint horizontal error of approximately 3.6 m, which is statistically indistinguishable from the best direct LiDAR odometry and LIO/LVIO methods given the approximately 2.5 m CEP uncertainty of the single-point GNSS reference.

The ablation study verifies that neither OSC alone nor the heterogeneous prior alone is sufficient. Robust tunnel performance is obtained when the heterogeneous prior and observable-subspace correction are used together. Additional experiments with KO, KAO, and VIO priors demonstrate the flexibility of the unified heterogeneous odometry interface. The runtime analysis further shows that the proposed degeneracy detection and correction modules introduce negligible computational overhead and that HERO-LO can run in real time.

Future work will evaluate HERO-LO in curved, sloping, intersecting, and variable-section tunnels, perform repeated physical traversals, incorporate higher-accuracy reference systems (RTK or total-station) for full-trajectory evaluation, investigate, in a dedicated study on degeneracy detection, the influence of the threshold value and of unit-normalisation variants on the detection and correction, and investigate adaptive prior selection when multiple heterogeneous odometry sources are available. Further improvements will also consider tighter multi-sensor integration, dynamic-object handling, and integration with loop closure or global optimization.

Supplementary Materials: The following supporting information can be downloaded at https://www.mdpi.com/article/10.3390/app162110509/s1: Video S1 shows the dynamic evolution of the estimated trajectories and point-cloud maps of all compared methods on the tunnel dataset. Table S1. Open-source implementations used for the comparison and their versions. Table S2. Principal configurations of the compared methods (all distances in metres). Table S3. Software environment used for all comparisons. Table S4. Local anomaly metrics of the low-output-rate methods at their original output rates. Table S5. Local anomaly metrics of the low-output-rate methods after interpolation to the 20 Hz LiDAR grid. Table S6. Principal parameters of the HERO-LO implementation used in the experiments. Tables S1–S3 provide the repositories, versions (commits), and dates of the compared open-source methods, their principal configurations, and the software environment used for all comparisons. Tables S4 and S5 report local anomaly metrics for the low-output-rate methods at their original rates and after interpolation to the 20 Hz LiDAR grid, respectively. Table S6 reports the principal parameters of the HERO-LO implementation used in the experiments.

Author Contributions: Conceptualization, X.C.; Methodology, D.W.; Software, D.W. and M.L.; Validation, D.W. and M.L.; Formal analysis, X.C.; Investigation, D.W., M.L. and X.C.; Resources, X.C.; Data curation, M.L.; Writing—original draft, D.W., M.L. and X.C.; Visualization, D.W. and M.L.; Supervision, X.C.; Project administration, X.C.; Funding acquisition, D.W. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Fujian Provincial Science and Technology Program (Guiding Project), grant number 2024H0016; and the Jimei University Scientific Research Startup Fund (Natural Science), grant number ZQ2024105.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The datasets and the implementation of HERO-LO presented in this article are not readily available because they are part of an ongoing study and are planned to be openly released in a future publication. The complete evaluation configuration, including the software versions, commits, and parameters of all compared methods, is provided in the Supplementary Materials. Requests for the dataset should be addressed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

6D	Six-Dimensional
AHRS	Attitude and Heading Reference System
CEP	Circular Error Probable
ENU	East-North-Up
GICP	Generalized Iterative Closest Point
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
ICP	Iterative Closest Point
IMU	Inertial Measurement Unit
KAO	AHRS-Aided Kinematic Odometry
KO	(encoder-based) Kinematic Odometry
LiDAR	Light Detection and Ranging
LIO	LiDAR-Inertial Odometry
LMI	Linear Matrix Inequality
LO	LiDAR Odometry
LVIO	LiDAR-Visual-Inertial Odometry
MAP	Maximum A Posteriori
OSC	Observable-Subspace Correction
PWL	Piecewise Linear
RANSAC	Random Sample Consensus
RTK	Real-Time Kinematic
S2S	Scan-to-Scan
SLAM	Simultaneous Localization and Mapping
VIO	Visual-Inertial Odometry
WGS84	World Geodetic System 1984

References

1. Lee, D.; Jung, M.; Yang, W.; Kim, A. LiDAR Odometry Survey: Recent Advancements and Remaining Challenges. *Intell. Serv. Robot.* **2024**, *17*, 95–118.
2. Ebadi, K.; Bernreiter, L.; Biggie, H.; Catt, G.; Chang, Y.; Chatterjee, A.; Denniston, C.E.; Deschênes, S.P.; Harlow, K.; Khattak, S.; et al. Present and Future of SLAM in Extreme Environments: The DARPA SubT Challenge. *IEEE Trans. Robot.* **2024**, *40*, 936–959.
3. Lin, J.; Zhang, F. R3LIVE: A Robust, Real-time, RGB-colored, LiDAR-Inertial-Visual tightly-coupled state Estimation and mapping package. In Proceedings of the 2022 International Conference on Robotics and Automation (ICRA), Philadelphia, PA, USA, 23–27 May 2022; pp. 10672–10678.
4. Bai, C.; Xiao, T.; Chen, Y.; Wang, H.; Zhang, F.; Gao, X. Faster-LIO: Lightweight Tightly Coupled Lidar-Inertial Odometry Using Parallel Sparse Incremental Voxels. *IEEE Robot. Autom. Lett.* **2022**, *7*, 4861–4868.
5. Xu, W.; Cai, Y.; He, D.; Lin, J.; Zhang, F. FAST-LIO2: Fast Direct LiDAR-inertial Odometry. *IEEE Trans. Robot.* **2022**, *38*, 2053–2073.
6. Chen, K.; Lopez, B.T.; Agha-mohammadi, A.a.; Mehta, A. Direct LiDAR Odometry: Fast Localization With Dense Point Clouds. *IEEE Robot. Autom. Lett.* **2022**, *7*, 2000–2007.
7. Vizzo, I.; Guadagnino, T.; Mersch, B.; Wiesmann, L.; Behley, J.; Stachniss, C. KISS-ICP: In Defense of Point-to-Point ICP—Simple, Accurate, and Robust Registration If Done the Right Way. *IEEE Robot. Autom. Lett.* **2023**, *8*, 1029–1036.
8. Yuan, C.; Xu, W.; Liu, X.; Hong, X.; Zhang, F. Efficient and Probabilistic Adaptive Voxel Mapping for Accurate Online LiDAR Odometry. *IEEE Robot. Autom. Lett.* **2022**, *7*, 8518–8525.
9. Zhang, J.; Kaess, M.; Singh, S. On Degeneracy of Optimization-Based State Estimation Problems. In *Proceedings of the 2016 IEEE International Conference on Robotics and Automation (ICRA)*; IEEE: Piscataway, NJ, USA, 2016; pp. 809–816.
10. Okawara, T.; Koide, K.; Oishi, S.; Yokozuka, M.; Banno, A.; Uno, K.; Yoshida, K. Tightly-Coupled LiDAR-IMU-Wheel Odometry With Online Calibration of a Kinematic Model for Skid-Steering Robots. *IEEE Access* **2024**, *12*, 134728–134738.
11. Zheng, C.; Xu, W.; Zou, Z.; Hua, T.; Yuan, C.; He, D.; Zhou, B.; Liu, Z.; Lin, J.; Zhu, F.; et al. FAST-LIVO2: Fast, Direct LiDAR–Inertial–Visual Odometry. *IEEE Trans. Robot.* **2025**, *41*, 326–346.
12. 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.

13. Lim, H.; Kim, D.; Kim, B.; Myung, H. AdaLIO: Robust Adaptive LiDAR-Inertial Odometry in Degenerate Indoor Environments. In Proceedings of the 2023 20th International Conference on Ubiquitous Robots (UR), Honolulu, HI, USA, 25–28 June 2023; pp. 48–53.
14. Tuna, T.; Nubert, J.; Nava, Y.; Khattak, S.; Hutter, M. X-ICP: Localizability-Aware LiDAR Registration for Robust Localization in Extreme Environments. *IEEE Trans. Robot.* **2024**, *40*, 452–471.
15. Barfoot, T.D. *State Estimation for Robotics: Second Edition*; Cambridge University Press: Cambridge, UK, 2024.
16. Qin, T.; Cao, S. A-LOAM: Advanced Implementation of LOAM. 2019. Available online: https://github.com/HKUST-Aerial-Robotics/A-LOAM (accessed on 29 April 2026).
17. 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)*; IEEE: Piscataway, NJ, USA, 2018; pp. 4758–4765.
18. Wang, H.; Wang, C.; Chen, C.L.; Xie, L. F-LOAM: Fast LiDAR Odometry and Mapping. In *Proceedings of the 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*; IEEE: Piscataway, NJ, USA, 2021; pp. 4390–4396.
19. Dellenbach, P.; Deschaud, J.E.; Jacquet, B.; Goulette, F. CT-ICP: Real-time Elastic LiDAR Odometry with Loop Closure. In *Proceedings of the 2022 International Conference on Robotics and Automation (ICRA)*; IEEE: Piscataway, NJ, USA, 2022; pp. 5580–5586.
20. Lee, D.; Lim, H.; Han, S. GenZ-ICP: Generalizable and Degeneracy-Robust LiDAR Odometry Using an Adaptive Weighting. *IEEE Robot. Autom. Lett.* **2025**, *10*, 152–159.
21. Shan, T.; Englot, B.; Meyers, D.; Wang, W.; Ratti, C.; Rus, D. LIO-SAM: Tightly-coupled Lidar Inertial Odometry via Smoothing and Mapping. In *Proceedings of the 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*; IEEE: Piscataway, NJ, USA, 2020; pp. 5135–5142.
22. Shan, T.; Englot, B.; Ratti, C.; Rus, D. LVI-SAM: Tightly-coupled Lidar-Visual-Inertial Odometry via Smoothing and Mapping. In *Proceedings of the 2021 IEEE International Conference on Robotics and Automation (ICRA)*; IEEE: Piscataway, NJ, USA, 2021; pp. 5692–5698.
23. Besl, P.J.; McKay, N.D. Method for registration of 3-D shapes. In *Proceedings of the Sensor Fusion IV: Control Paradigms and Data Structures*; SPIE: Bellingham, WA, USA, 1992; Volume 1611, pp. 586–606.
24. Rusinkiewicz, S.; Levoy, M. Efficient variants of the ICP algorithm. In Proceedings of the Proceedings Third International Conference on 3-D Digital Imaging and Modeling, Quebec City, QC, Canada, 28 May–1 June 2001; pp. 145–152.
25. Segal, A.; Haehnel, D.; Thrun, S. Generalized-ICP. In Proceedings of the Robotics: Science and Systems, Seattle, WA, USA, 28 June–1 July 2009; Volume 2, p. 435.
26. Zhang, J.; Singh, S. LOAM: Lidar Odometry and Mapping in Real-time. In Proceedings of the Robotics: Science and Systems, Berkeley, CA, USA, 12–16 July 2014; Volume 2, pp. 1–9.
27. Zhang, J.; Singh, S. Low-Drift and Real-Time Lidar Odometry and Mapping. *Auton. Robot.* **2017**, *41*, 401–416.
28. 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.
29. Wu, D.; Zhong, X.; Peng, X.; Hu, H.; Liu, Q. Multimodal Information Fusion for High-Robustness and Low-Drift State Estimation of UGVs in Diverse Scenes. *IEEE Trans. Instrum. Meas.* **2022**, *71*, 8505115.
30. Zhu, J.; Li, H.; Zhang, T. Camera, LiDAR, and IMU Based Multi-Sensor Fusion SLAM: A Survey. *Tsinghua Sci. Technol.* **2024**, *29*, 415–429.
31. Wang, Z.; Zhang, L.; Shen, Y.; Zhou, Y. D-LIOM: Tightly-Coupled Direct LiDAR-Inertial Odometry and Mapping. *IEEE Trans. Multimed.* **2023**, *25*, 3905–3920.
32. Lee, J.; Komatsu, R.; Shinozaki, M.; Kitajima, T.; Asama, H.; An, Q.; Yamashita, A. Switch-SLAM: Switching-Based LiDAR-Inertial-Visual SLAM for Degenerate Environments. *IEEE Robot. Autom. Lett.* **2024**, *9*, 7270–7277.
33. Xu, W.; Zhang, F. FAST-LIO: A Fast, Robust LiDAR-inertial Odometry Package by Tightly-Coupled Iterated Kalman Filter. *IEEE Robot. Autom. Lett.* **2021**, *6*, 3317–3324.
34. Lee, E.M.; Marsim, K.C.; Myung, H. LODESTAR: Degeneracy-Aware LiDAR-Inertial Odometry With Adaptive Schmidt-Kalman Filter and Data Exploitation. *IEEE Robot. Autom. Lett.* **2026**, *11*, 922–929.
35. Zhang, J.; Singh, S. Laser–visual–inertial odometry and mapping with high robustness and low drift. *J. Field Robot.* **2018**, *35*, 1242–1264.
36. Santamaria-Navarro, A.; How, J.P.; Carlone, L.; Agha-mohammadi, A.a. LION: Lidar-Inertial Observability-Aware Navigator for Vision-Denied Environments. In *Proceedings of the Experimental Robotics: The 17th International Symposium*; Springer Nature: Cham, Switzerland, 2021; pp. 380–390.
37. Boyd, S.; El Ghaoui, L.; Feron, E.; Balakrishnan, V. *Linear Matrix Inequalities in System and Control Theory*; SIAM: Philadelphia, PA, USA, 1994.

38. Arena, P.; Buscarino, A.; Fortuna, L.; Frasca, M. Separation and synchronization of piecewise linear chaotic systems. *Phys. Rev. E* **2006**, *74*, 026212.
39. Sturm, J.; Engelhard, N.; Endres, F.; Burgard, W.; Cremers, D. A benchmark for the evaluation of RGB-D SLAM systems. In *Proceedings of the 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*; IEEE: Piscataway, NJ, USA, 2012; pp. 573–580.

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.