

# IMU-Centric to LiDAR-Centric Decoupled Multi-Stage Extrinsic Calibration of Repetitive Scanning LiDARs in Mobile Mapping Systems

Ying Jiang, Han Hu\*, Jiahao Zhou, Qiang Wang, Pengxin Gu, Zhaoxuan Zhou, Qing Zhu

Faculty of Geosciences and Engineering, Southwest Jiaotong University, Chengdu 611756, China

---

## Abstract

Extrinsic calibration of commercial-grade repetitive scanning LiDARs remains challenging because their fixed and sparse scan patterns provide weak geometric constraints for direct registration in multi-LiDAR mobile mapping systems. More informative geometric structures can be obtained by accumulating scans along a moving trajectory. However, scan accumulation itself depends on accurate extrinsics; poor initialization can introduce motion distortion and structural blurring, creating a circular dependency between accumulation quality and calibration accuracy. This paper presents a multi-stage calibration method that resolves this dependency by combining inertial trajectory constraints with accumulated geometric evidence. The proposed method first estimates a coarse calibration from the rigid-body relationship between the LiDAR-integrated IMUs and a reference IMU, providing reliable initial values for motion-compensated multi-frame aggregation without relying on sparse LiDAR geometry. A factor-graph averaged multiway-registration prior is then constructed to initialize the subsequent geometric refinement. The extrinsics are further refined using inter-LiDAR consistency constraints and correspondences to a global reference point cloud. Experiments on a mobile mapping system equipped with four RoboSense Airy LiDARs demonstrate substantial improvements over existing calibration baselines. The proposed method reduces the overall inter-LiDAR consistency error to 0.048 m, compared with 0.318 m for LI-Calib and 0.262 m for iKalibr, and reduces the global reference error to 0.099 m, compared with 0.385 m and 0.427 m for the two baselines. The calibrated extrinsics also remain stable across additional scenes without recalibration.

**Keywords:** Multi-LiDAR Calibration, Repetitive Scanning LiDAR, Rigid Body Extrinsic Initialization, Global Reference Point Cloud, Mobile Mapping System, Continuous-Time Optimization

---

## 1. Introduction

Repetitive scanning LiDARs have been increasingly used in localization, SLAM, mobile mapping, and three-dimensional reconstruction because their compact sensing units, stable scanning patterns, and wide field-of-view configurations are suitable for mobile platforms operating in complex environments (Elhashash et al., 2022; Khoramshahi et al., 2025; Zhou et al., 2021). In practical mobile mapping systems (MMS), a single repetitive scanning LiDAR is often insufficient to observe complex road scenes, building facades, vegetation, and roadside objects from all directions, especially under occlusions and limited instantaneous geometric coverage (Elhashash et al., 2022; Pan et al., 2023). Deploying multiple LiDARs with complementary viewing directions can enlarge the effective sensing coverage, increase point cloud density, improve geometric completeness, and support more reliable multi-sensor fusion for mapping and perception (Zhu et al., 2023; Khoramshahi et al., 2025). These advantages, however, rely on accurate extrinsic calibration among all LiDARs and the platform reference frame, since small installation errors can directly degrade point cloud fusion and downstream reconstruction quality (Wu et al., 2023; Zhu et al., 2023).


Extrinsic calibration of repetitive scanning LiDARs differs from that of non-repetitive scanning sensors because the former repeatedly samples fixed scan lines during static acquisition, whereas the latter can gradually cover richer surface structures over time (Aijazi and Checchin, 2024; Zhou et al., 2021). As a result, a static point cloud from a repetitive scanning LiDAR often contains only sparse and line-like observations, providing weak three-dimensional constraints for inter-sensor registration. In a multi-LiDAR setup, this sparsity makes the overlapping regions between different scanners ambiguous, so direct static calibration based on point cloud matching becomes unreliable and inaccurate (Wu et al., 2023; Kim et al., 2024). As illustrated in Fig. 1, the limited structural information in repetitive scanning observations leads to unstable correspondences and even wrong registration results. This limitation indicates that repetitive multi-LiDAR calibration cannot rely solely on static scan matching and instead requires additional motion or trajectory constraints (Pöppel et al., 2024).

To compensate for the sparse static observations, repetitive scanning LiDARs usually require scan accumulation along the moving trajectory to form denser and more geometrically informative point clouds. This moving accumulation process, however, depends on both the platform trajectory and the initial LiDAR extrinsics, so installation errors are directly propagated into the reconstructed local map. When the initial extrinsics are inaccurate, the accumulated point cloud can suffer from


---

\*Corresponding author at Faculty of Geosciences and Engineering, Southwest Jiaotong University, Chengdu, China


Email address: han.hu@swjtu.edu.cn (Han Hu)



(a) Non-repetitive scan pattern.




(b) Repetitive scan pattern.




(c) Failed registration with repetitive scans.


**Fig. 1.** Illustration of the geometric limitations of repetitive scanning LiDARs. Compared with non-repetitive scanning, repetitive scanning produces sparse and fixed scan patterns, which provide weaker local geometric constraints and may lead to unreliable inter-LiDAR registration.



(a) Duplicated structures.



(b) Blurred boundaries.



(c) Inconsistent local geometry.

**Fig. 2.** Accumulation–calibration interdependency in repetitive-scan multi-LiDAR calibration. Multi-frame accumulation provides richer geometry for correspondence construction but depends on accurate initial extrinsics; poor initialization introduces duplicated structures, blurred boundaries, and unreliable global-reference matching.

progressive frame misalignment, ghosting artifacts, and blurred structural boundaries, as illustrated in Fig. 2 (Lee et al., 2024). This creates a severe circular dependency: reliable accumulation requires accurate extrinsics, whereas accurate extrinsic calibration for repetitive scanning LiDARs requires sufficiently dense accumulated geometry.

To address this circular dependency, we organize the calibration pipeline into two stages: IMU-centric initialization and LiDAR-centric refinement. *The IMU-centric initialization does not rely on LiDAR point clouds, and therefore breaks the circular dependency between point-cloud integration and extrinsic calibration.* In the IMU-centric stage, the trajectories of the LiDAR-integrated IMUs are aligned with the reference IMU in a continuous-time form under rigid-body constraints, thereby estimating temporal offsets and coarse rotational and translational extrinsics before any LiDAR point-cloud integration. This decouples the initial estimate from the quality of accumulated LiDAR geometry. In the LiDAR-centric stage, the initialized extrinsics are used for motion compensation and multi-frame LiDAR integration (accumulation), which builds denser local maps for correspondence construction. Graph-based averaging of multiway-registration results first provides pose-

level extrinsic priors, and the final optimization then uses original point-level inter-LiDAR and global-reference matches in a common world frame. This final optimization keeps the reliable translations fixed and refines only the rotational components, preventing physically implausible translational compensation under weak geometric observability (Aijazi and Checchin, 2024).

The main contributions of this paper are as follows.

- (1) We formulate the extrinsic calibration of repetitive scanning multi-LiDAR MMS as an IMU-centric to LiDAR-centric framework that breaks the circular dependency between accurate LiDAR-to-body extrinsics and reliable motion-compensated scan accumulation.
- (2) We introduce a continuous-time IMU-centric initialization strategy that aligns the LiDAR-integrated IMUs with the reference IMU to estimate temporal offsets and physically plausible coarse LiDAR-to-body extrinsics without relying on sparse LiDAR geometry. These estimates support robust deskewing and multi-frame accumulation before geometric calibration.
- (3) We design a LiDAR-centric refinement strategy that combines factor-graph averaged multiway-registration priors

with point-level inter-LiDAR and global-reference optimization to improve rotational alignment while preserving the rigid baseline geometry of the multi-LiDAR rig.

## 2. Related work

### 2.1. Extrinsic calibration of multiple LiDAR sensors

Multi-LiDAR systems are widely adopted on mobile mapping and robotic platforms to enlarge the sensing coverage, increase point cloud density, and reduce blind spots. Since the LiDARs are usually mounted at different positions and oriented toward different viewing directions, their relative poses must be accurately calibrated before multi-sensor point clouds can be fused into a consistent platform-level representation (Lyu et al., 2024). Inaccurate inter-LiDAR extrinsics directly lead to duplicated structures, blurred surfaces, and systematic offsets in the fused map, which further degrade downstream mapping, localization, and measurement tasks.

Existing multi-LiDAR calibration methods mainly construct constraints from the spatial consistency among point clouds observed by different sensors (Abdullah, 2023). A common strategy is to register overlapping point clouds or local submaps using ICP, GICP, NDT, or related variants (Koide et al., 2021; Xu et al., 2023; Li et al., 2022). Other methods introduce geometric primitives such as planes, edges, corners, or poles, and formulate calibration as feature correspondence matching followed by nonlinear optimization (Lee and Chung, 2022; Jiao et al., 2021). Recent systems further improve automation and scalability. For example, Multi-LiCa combines feature-based coarse alignment with GICP-based refinement for automatic targetless LiDAR-to-LiDAR calibration (Kulmer et al., 2024), while unified frameworks such as iKalibr extend spatiotemporal calibration to heterogeneous sensor suites with multiple LiDARs, IMUs, cameras, and radars (Chen et al., 2025).

Despite this progress, most multi-LiDAR calibration methods still rely on sufficient common-view geometry, reliable local features, or informative motion constraints (Yang and Holst, 2025). In vehicle-mounted configurations, LiDARs are often installed with divergent viewing directions to maximize coverage, which naturally reduces the overlap between neighboring sensors. Under small-overlap or large-viewpoint-difference conditions, direct point cloud registration becomes sensitive to initialization, noise, and outliers, and the extracted features may not provide enough independent constraints for full six-degree-of-freedom calibration (Ling and Qin, 2022; Li et al., 2023b).

These limitations become more pronounced for vehicle-mounted systems equipped with repetitive-scan LiDARs. Unlike conventional spinning LiDARs that provide relatively uniform angular coverage, repetitive-scan sensors observe the environment with fixed and sparse scanning patterns. Consequently, the effective overlap between adjacent sensors can be extremely sparse, and the feature correspondences required by registration-based, primitive-based, or SLAM-based calibration methods are often weak, unevenly distributed, or absent. This makes translational parameter estimation particularly ill-conditioned, because small rotational errors and sparse overlap-

ping structures can be compensated by physically implausible translations during optimization.

Therefore, robust extrinsic calibration for vehicle-mounted repetitive-scan multi-LiDAR systems cannot rely solely on local inter-LiDAR overlap or intra-frame geometric features. Additional constraints are required to compensate for insufficient feature richness and weak common-view geometry. This motivates the use of inertial initialization to provide physically plausible coarse extrinsics, followed by motion-compensated temporal accumulation and correspondence construction with an external global reference point cloud. The global reference provides an absolute geometric anchor for the LiDAR rig, while inter-LiDAR correspondences preserve the relative consistency among onboard LiDARs.

### 2.2. Extrinsic calibration between multi-modal sensors

LiDAR-IMU extrinsic calibration is a fundamental prerequisite for mobile mapping and robotic perception systems. The IMU provides high-frequency inertial measurements, while the LiDAR provides geometric observations of the surrounding environment. Accurate spatial and temporal calibration between the two sensors is essential for motion compensation, LiDAR-inertial odometry, and the construction of a geometrically consistent point cloud map. Errors in the LiDAR-IMU extrinsics can be directly propagated into scan undistortion, state estimation, and multi-sensor map fusion.

Existing LiDAR-IMU calibration methods commonly exploit motion consistency between inertial measurements and LiDAR-based geometric observations. Some approaches estimate LiDAR motion through scan registration or LiDAR odometry and align it with the inertial motion using hand-eye calibration or motion correlation. Their accuracy depends on reliable LiDAR odometry and sufficient motion excitation; otherwise, some extrinsic components become weakly observable (Lv et al., 2022; Zhu et al., 2022).

To handle asynchronous measurements and improve estimation accuracy, continuous-time trajectory representations have been widely adopted. LI-Calib formulates LiDAR-IMU calibration as a continuous-time batch estimation problem by jointly minimizing IMU residuals and LiDAR point-to-surfel distances (Lv et al., 2020; Wang et al., 2025). Observability-aware calibration further selects informative data segments or updates only identifiable directions to cope with degenerate motions (Lv et al., 2022). Robust initialization methods such as LI-Init estimate temporal offsets, LiDAR-IMU extrinsics, gravity, and IMU biases by aligning LiDAR-based motion with IMU measurements (Zhu et al., 2022). Recent work in mobile mapping and multi-sensor adjustment also highlights the importance of tightly coupled GNSS/IMU/LiDAR trajectories and scalable hybrid adjustment for accurate georeferencing (Zhou et al., 2021; Wu et al., 2023; Jonassen et al., 2023). More recently, unified frameworks such as iKalibr extend continuous-time spatiotemporal calibration to heterogeneous sensor suites involving IMUs, LiDARs, cameras, and radars (Chen et al., 2025).

Nevertheless, most LiDAR-IMU calibration methods still require LiDAR measurements to provide sufficiently stable

geometric constraints. This assumption is usually reasonable for conventional spinning LiDARs in structured environments, where planar surfaces, edges, and local surfels can be repeatedly observed. However, it becomes problematic for repetitive-scan LiDARs. Their fixed and sparse scanning patterns lead to limited intra-frame geometric diversity and weaken the LiDAR residuals used in calibration (Aijazi and Checchin, 2024; Kim et al., 2024). In addition, vehicle-mounted platforms are often dominated by forward motion with limited roll, pitch, and vertical excitation, which further reduces the observability of translational extrinsic parameters (Lv et al., 2022; Wu et al., 2023). Therefore, for repetitive-scan multi-LiDAR MMS, inertial constraints are needed to obtain physically plausible coarse translations, while external global-reference correspondences are needed to refine the geometric alignment without allowing unconstrained translational compensation.

### 2.3. Motion compensation

Motion compensation is a fundamental step in LiDAR odometry, mapping, and multi-frame point cloud integration. Because LiDAR points are acquired sequentially within a scan, treating all points as if they were captured at a single pose introduces motion distortion when the platform is moving. Existing LiDAR and LiDAR-inertial systems therefore estimate the platform motion from odometry, IMU measurements, or their fused trajectory, and use it to deskew each point before scan registration or map accumulation (Zhang and Singh, 2014; Xu et al., 2022, 2023). For repetitive scanning LiDARs, this step is especially important because a single static frame is sparse, and useful geometric structures are usually obtained only after motion-compensated temporal accumulation.

*Motivation for accumulated error in motion compensation.* However, motion compensation implicitly assumes that the extrinsic transformation between the body frame in which the motion trajectory is defined and the sequentially sampled LiDAR frame is already accurate. The body frame is generally the reference IMU frame  $B$ , whose trajectory is represented by  $\mathbf{T}_B^W(t)$ , and the  $i$ -th LiDAR frame is  $L_i$ . For a point  $\mathbf{p}^{L_i}(t_p)$  sampled by the  $i$ -th LiDAR with synchronized timestamp  $\tau_p$ , the ideal world-frame position used for accumulation is

$$\mathbf{p}^W(\tau_p) = \mathbf{R}_B^W(\tau_p) \left( \mathbf{R}_{L_i}^B \mathbf{p}^{L_i}(t_p) + \mathbf{t}_{L_i}^B \right) + \mathbf{t}_B^W(\tau_p), \quad (1)$$

where  $\mathbf{T}_{L_i}^B = \{\mathbf{R}_{L_i}^B, \mathbf{t}_{L_i}^B\}$  is the LiDAR extrinsic transformation. Suppose that the extrinsic used for deskewing and accumulation contains a small rotation perturbation  $\delta\theta$  and a translation perturbation  $\delta\mathbf{t}$ . Using the first-order approximation  $\exp([\delta\theta]_{\times}) \approx \mathbf{I} + [\delta\theta]_{\times}$ , the induced world-frame point error can be written as

$$\mathbf{e}_p^W \approx \mathbf{R}_B^W(\tau_p) \left( [\delta\theta]_{\times} \mathbf{q}_p + \delta\mathbf{t} \right), \quad \mathbf{q}_p = \mathbf{R}_{L_i}^B \mathbf{p}^{L_i}(t_p). \quad (2)$$

For temporal accumulation, the critical effect is not the absolute error of a single point, but the relative inconsistency between points acquired along the platform trajectory. For two points acquired at synchronized timestamps  $\tau_p$  and  $\tau_r$  within an accumulation interval  $\Delta t_{p,r} = |\tau_p - \tau_r|$ , this inconsistency is

$$\mathbf{e}_{p,r}^W = \mathbf{R}_B^W(\tau_p) \mathbf{u}_p - \mathbf{R}_B^W(\tau_r) \mathbf{u}_r, \quad \mathbf{u}_p = [\delta\theta]_{\times} \mathbf{q}_p + \delta\mathbf{t}. \quad (3)$$

Equation (3) indicates that the distortion induced by installation-parameter errors is coupled with the platform motion during accumulation. As the platform moves from  $\tau_r$  to  $\tau_p$ , the same local extrinsic perturbation is propagated through different poses along the trajectory. The resulting distortion therefore depends on both the traveled distance and the attitude variation within the accumulation window. For a vehicle-mounted MMS, this trajectory segment cannot be treated as negligible, because multi-frame accumulation usually spans multiple scans and covers a finite driving distance. Consequently, inaccurate initial extrinsics can lead to duplicated structural edges, surface thickening, and blurred boundaries in the accumulated point cloud, as illustrated in Fig. 2. This dependency motivates reliable coarse initialization before LiDAR deskewing and subsequent geometric refinement.

## 3. Methodology

### 3.1. Overview

As shown in Fig. 3, the proposed method estimates the extrinsic parameters between each LiDAR and the platform reference frame in an MMS equipped with repetitive scanning LiDARs. For such sensors, static or single-frame observations are too sparse for reliable direct inter-LiDAR registration, so useful geometric constraints must be obtained through motion-compensated multi-frame integration. This integration, however, depends on the initial extrinsic parameters and can be distorted by poor initialization. The proposed framework therefore first performs IMU-centric initialization using the LiDAR-integrated IMUs and the reference IMU, decoupling the coarse extrinsic estimate from LiDAR point-cloud integration and thereby breaking the circular dependency. These extrinsics are then used in the LiDAR-centric stage to accumulate denser local maps, construct inter-LiDAR and global-reference correspondences, and perform fine calibration.

Because dynamic point cloud accumulation requires motion compensation, inaccurate installation parameters can distort the aggregated local map before reliable geometric correspondences are established. The framework therefore avoids relying on sparse LiDAR geometry for the initial estimate and instead recovers coarse extrinsics through continuous-time trajectory alignment between the LiDAR-integrated IMUs and the reference IMU. The POS trajectory then provides the platform motion for scan compensation, while a global reference point cloud in the same world frame provides the external geometric source for final matching. The framework outputs the temporal offsets and the LiDAR extrinsics with respect to the platform reference frame.

The proposed framework consists of three stages: First, continuous-time multi-IMU initialization estimates temporal offsets and coarse rotational and translational extrinsics under rigid-body constraints; Second, with these initial parameters and the POS trajectory, LiDAR scans are synchronized, deskewed, and accumulated into geometrically informative local maps, from which a multiway-registration extrinsic prior, inter-LiDAR consistency constraints, and global-reference correspondences are constructed; Finally, the correspondences are

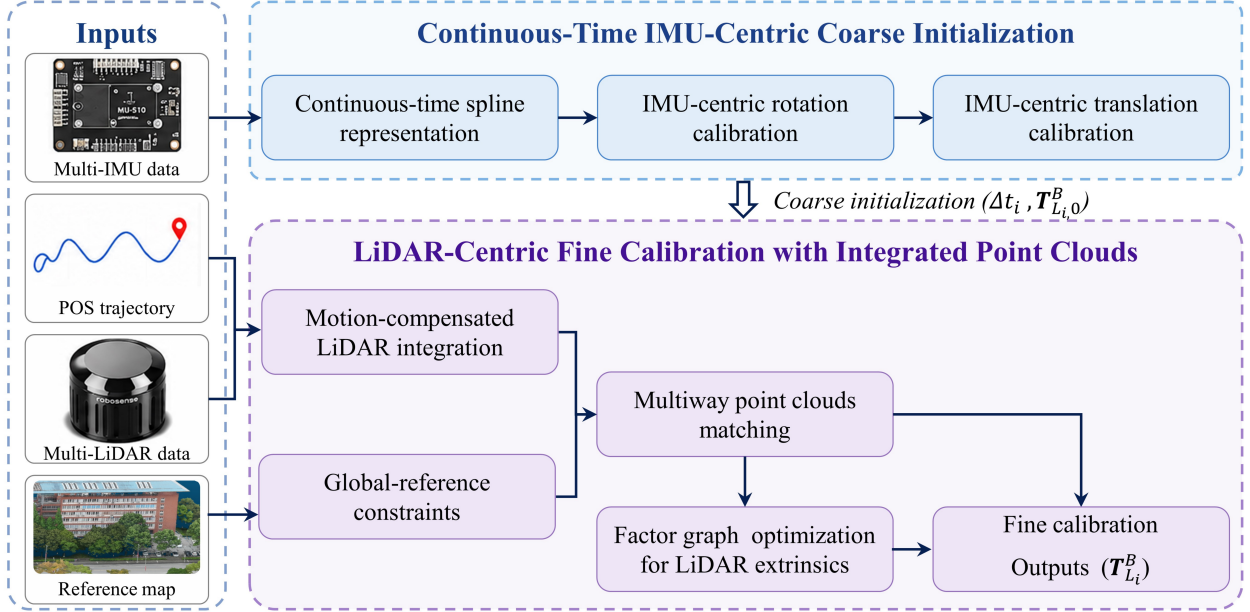


Fig. 3. Overview of the proposed calibration framework.

Table 1: Main hardware components of the vehicle-mounted MMS.

Sensor	Model
GNSS board	Unicore UM982
Reference IMU	Forsense FSS-IMU618-C
LiDAR	RoboSense Airy
LiDAR-integrated IMU	Built-in IMU of RoboSense Airy

incorporated into a rotation refinement with fixed multiway-registration translations, where only rotational extrinsics are optimized.

### 3.2. Hardware setup

The MMS combines a POS unit with four repetitive-scanning multi-beam LiDARs mounted on a rigid frame, as shown in Fig. 4. The POS unit integrates a GNSS receiver and a reference IMU, and its trajectory is post-processed using base-station corrections obtained through network RTK. The resulting trajectory provides attitude accuracy on the order of 1 arcmin and is used as the platform motion reference for scan deskewing and map accumulation. Each LiDAR also contains a built-in IMU, a common configuration in compact LiDAR units, and these inertial measurements are used for continuous-time coarse initialization relative to the reference IMU. The main hardware components are summarized in Table 1.

Hardware time synchronization is implemented using the GNSS time reference and the acquisition carrier board shown in Fig. 4. The GNSS board and the reference IMU are synchronized through the PPS (Pulse Per Seconds) signal and GN-RMC messages, after which the reference IMU internally outputs synchronized inertial measurements at 200 Hz. The LiDAR units are also synchronized to GNSS time through PPS and GNRMC signals. However, because the low-cost timing circuits inside compact LiDAR units may introduce residual

systematic delays, hardware synchronization alone does not completely remove the temporal offset between the LiDAR-integrated IMUs and the reference IMU. This residual offset is therefore explicitly estimated and compensated during the continuous-time IMU-centric initialization.

The nomenclature of coordinate frames and main mathematical symbols used in this work is summarized in Table 2. Throughout this paper,  $\mathbf{T}_A^B = \{\mathbf{R}_A^B, \mathbf{t}_A^B\} \in SE(3)$  denotes the rigid transformation that maps a point from frame  $A$  to frame  $B$ , i.e.,

$$\mathbf{p}^B = \mathbf{R}_A^B \mathbf{p}^A + \mathbf{t}_A^B. \quad (4)$$

The purpose of this work is to estimate the LiDAR-to-body extrinsic transformation and temporal offset for each LiDAR unit,

$$\Theta_i = \{\mathbf{T}_{L_i}^B, \Delta t_i\}, \quad \mathbf{T}_{L_i}^B = \{\mathbf{R}_{L_i}^B, \mathbf{t}_{L_i}^B\}. \quad (5)$$

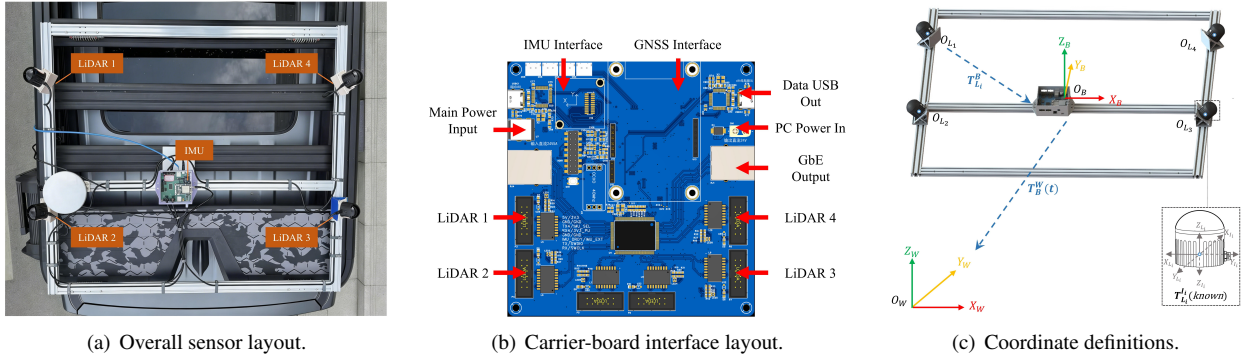
Here  $\Delta t_i$  is the temporal offset between the  $i$ -th LiDAR unit and the reference IMU. Since each LiDAR and its built-in IMU are integrated as a single unit, no additional LiDAR-IMU time offset is assumed. A LiDAR point sampled at its raw timestamp  $t_p$  is therefore evaluated at the synchronized timestamp

$$\tau_p = t_p + \Delta t_i. \quad (6)$$

In the coarse initialization stage, the extrinsic from  $I_i$  to  $B$ ,  $\mathbf{T}_{L_i}^B$ , is first estimated from inertial measurements. The rigid transformation between each LiDAR and its paired LiDAR-integrated IMU, denoted by  $\mathbf{T}_{L_i}^{I_i}$ , is treated as a fixed known quantity provided by the sensor manufacturer. The LiDAR-to-body extrinsic is then obtained as

$$\mathbf{T}_{L_i}^B = \mathbf{T}_{L_i}^{I_i} \mathbf{T}_{L_i}^{I_i}, \quad (7)$$

for  $i = 1, \dots, 4$ . These coarse LiDAR extrinsics are subsequently used in LiDAR deskewing, global-reference matching, and rotation refinement.



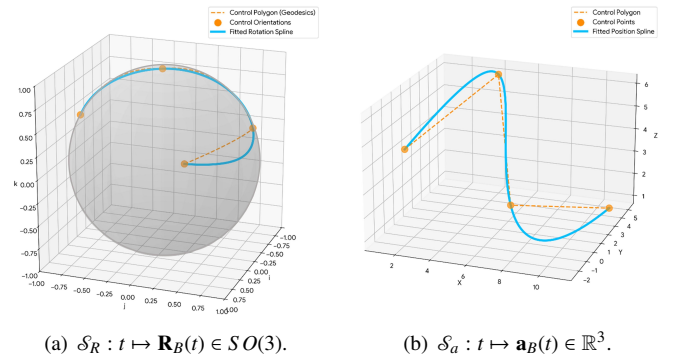
**Fig. 4.** Vehicle-mounted MMS hardware setup and coordinate definitions. (a) Overall sensor layout on the rigid frame, where four LiDARs and the reference IMU are mounted to provide multi-view coverage. (b) Interface layout of the custom carrier board, which supports LiDAR, IMU, and GNSS integration for synchronized data acquisition. (c) Coordinate frames used for the LiDAR-integrated IMUs, LiDARs, reference IMU, and world frame.

Table 2: Nomenclature of coordinate frames and mathematical symbols. The frame indices of transformations are omitted for brevity.

Symbol	Description
$\{W\}$	World frame.
$\{B\}$	Platform reference frame.
$\{I_i\}$	$i$ -th LiDAR-integrated IMU frame.
$\{L_i\}$	$i$ -th LiDAR frame.
$i, j$	Sensor indices; e.g., the $i$ -th or $j$ -th LiDAR/IMU.
$m$	Retained segment or integration-window index.
$p, r$	Point indices; e.g., the $p$ -th and $r$ -th LiDAR points or correspondences.
$\mathbf{T}$	Rigid transformation.
$\mathbf{R}$	Rotation matrix.
$\mathbf{t}$	Translation vector.
$\mathbf{p}$	3D point position.
$\mathbf{R}(t)$	Continuous-time rotation trajectory.
$\mathbf{a}(t)$	Continuous-time acceleration trajectory.
$\boldsymbol{\omega}(t)$	Angular velocity.
$\boldsymbol{\alpha}(t)$	Angular acceleration.
$\mathbf{v}$	Linear velocity.
$\mathbf{g}$	Gravity vector.
$\mathbf{b}_g$	Gyroscope bias.
$\mathbf{b}_a$	Accelerometer bias.
$\Delta t$	Temporal offset.
$\boldsymbol{\Theta}$	Calibration variables.
$\delta\boldsymbol{\theta}$	Local rotation perturbation.

### 3.3. Continuous-time IMU-centric coarse initialization

To obtain a reliable coarse initialization before LiDAR deskewing and global-reference matching, we perform a continuous-time, IMU-centric multi-IMU initialization (Li et al., 2024; Chen et al., 2025). The reference IMU defines the body frame  $B$  and is not assigned extrinsic variables. Its inertial measurements are used to construct the body-frame continuous-time motion spline, while the LiDAR-integrated IMUs are aligned to this reference motion under rigid-body constraints. This stage estimates temporal offsets  $\Delta t_i$  and the coarse rotational/translational extrinsics  $\mathbf{R}_{L_i}^B, \mathbf{t}_{L_i}^B$  without relying on LiDAR geometry, which improves robustness under repetitive-scan constraints and provides a reliable initial value for the sub-



**Fig. 5.** Illustration of continuous-time spline-based motion representation. The rotation component is modeled on the  $SO(3)$  manifold, while the acceleration component is modeled in Euclidean space, corresponding to the two splines in Eq. (8).

sequent LiDAR multiway registration.

#### 3.3.1. Continuous-time spline representation

To align asynchronous IMU streams and estimate temporal offsets during coarse initialization, we represent the platform motion using continuous-time B-spline trajectories (Lv et al., 2020; Li et al., 2023a). In implementation, the spline residuals and their analytical Jacobians are generated using existing symbolic and continuous-time optimization libraries, including Hyperion and SymForce (Hug et al., 2024; Martiros et al., 2022). The B-spline representation also supports analytical differentiation and integration, so motion quantities such as angular velocity, angular acceleration, and linear acceleration can be obtained consistently at arbitrary offset-corrected timestamps. These properties make it convenient to construct gyroscope and accelerometer residuals for asynchronous inertial alignment and temporal-offset estimation.

In this work, the platform motion is not modeled as a single continuous-time spline on  $SE(3)$ . Instead, it is decomposed into a rotational spline on  $SO(3)$ , which models the continuous rotation trajectory of  $B$  with respect to  $W$ , and a Euclidean spline in  $\mathbb{R}^3$ , which models the continuous linear acceleration of  $B$ .

Specifically, the two continuous splines are written as

$$\mathcal{S}_R : t \mapsto \mathbf{R}_B(t) \in SO(3), \quad \mathcal{S}_a : t \mapsto \mathbf{a}_B(t) \in \mathbb{R}^3. \quad (8)$$

For brevity, the superscript  $W$  is omitted in  $\mathbf{R}_B(t)$  and  $\mathbf{a}_B(t)$  in this subsection; they denote the body-frame motion with respect to the world frame, i.e.,  $\mathbf{R}_B(t) \equiv \mathbf{R}_B^W(t)$  and  $\mathbf{a}_B(t) \equiv \mathbf{a}_B^W(t)$  when the target frame is written explicitly. The two spline components are illustrated in Fig. 5. This parameterization is adopted to support both the IMU-centric initialization and the subsequent LiDAR point-cloud calibration in a unified continuous-time form. For the IMU-centric stage, the angular velocity  $\omega_B(t)$  and angular acceleration  $\alpha_B(t)$  are obtained by differentiating  $\mathbf{R}_B(t)$ , while  $\mathbf{a}_B(t)$  is evaluated directly from the acceleration spline for the accelerometer residuals. For the LiDAR stage, the same continuous motion representation provides offset-corrected motion quantities for deskewing, accumulation, and geometric refinement.

Therefore, the continuous-time model serves as a shared motion reference for temporal-offset estimation, rotational and translational extrinsic initialization, and subsequent LiDAR processing. In practice, well-conditioned spline fitting and extrinsic estimation require sufficiently excited motion across multiple degrees of freedom, which is achieved by collecting the calibration sequence along a route with diverse dynamic maneuvers.

### 3.3.2. IMU-centric calibration of $\mathbf{R}_I^B$

The rotational initialization estimates the rotational extrinsics  $\mathbf{R}_I^B$  and temporal offsets  $\Delta t_i$  by enforcing gyroscope consistency between the LiDAR-integrated IMUs and the reference IMU (Li et al., 2024; Chen et al., 2025). The reference IMU remains fixed and defines the body motion spline. After compensating for the relative rotation and time offset, the angular velocity derived from the reference motion should agree with the gyroscope measurements in  $\{I_i\}$ .

Given the rotational spline  $\mathbf{R}_B(t)$ , the body-frame angular velocity of  $B$  is obtained from its analytical derivative:

$$[\omega_B(t)]_\times = \mathbf{R}_B^\top(t) \dot{\mathbf{R}}_B(t), \quad (9)$$

where  $[\cdot]_\times$  denotes the skew-symmetric matrix operator. For a gyroscope sample from the  $i$ -th LiDAR-integrated IMU, the predicted measurement and residual are

$$\begin{aligned} \hat{\omega}_i(t) &= (\mathbf{R}_{I_i}^B)^\top \omega_B(t + \Delta t_i) + \mathbf{b}_{g,i}, \\ \mathbf{r}_g^{(i)}(t) &= \hat{\omega}_i(t) - \omega_i^m(t). \end{aligned} \quad (10)$$

where  $\omega_i^m(t)$  is the measured angular velocity and  $\mathbf{b}_{g,i}$  is the gyroscope bias.

The rotational initialization is formulated as a nonlinear least-squares problem:

$$\begin{aligned} \min_{\mathcal{X}_R} & \sum_{i=1}^{N_L} \sum_{k \in \mathcal{K}_i} \|\mathbf{r}_g^{(i)}(t_{i,k})\|^2, \\ \mathcal{X}_R &= \{\mathcal{S}_R, \mathbf{R}_I^B, \Delta t_i, \mathbf{b}_{g,i}\}_{i=1}^{N_L}. \end{aligned} \quad (11)$$

where  $\mathcal{S}_R$  denotes the rotational spline control variables,  $N_L$  is the number of LiDAR-integrated IMUs, and  $\mathcal{K}_i$  is the gyroscope sample set of  $I_i$ . The frame  $B$  is fixed to remove the rotational gauge freedom.

### 3.3.3. IMU-centric calibration of $\mathbf{t}_I^B$

Unlike the rotational extrinsics, the translational extrinsics  $\mathbf{t}_I^B$  cannot be recovered from gyroscope measurements alone (Lv et al., 2022). They become observable through the lever-arm effect. Specifically, when the origin of  $I_i$  is spatially offset from the reference body frame  $B$ , the angular motion of the platform induces tangential and centripetal acceleration terms that depend on  $\mathbf{t}_I^B$ . Let

$$\mathbf{D}_B(t) = [\alpha_B(t)]_\times + [\omega_B(t)]_\times^2, \quad (12)$$

where  $\omega_B(t)$  and  $\alpha_B(t)$  are obtained analytically from the rotational spline  $\mathcal{S}_R$  estimated above. The ideal accelerometer measurement of the  $i$ -th IMU is predicted as

$$\hat{\mathbf{a}}_i(t) = (\mathbf{R}_{I_i}^B)^\top \left[ \mathbf{R}_B^\top(t) (\mathbf{a}_B(t) - \mathbf{g}) + \mathbf{D}_B(t) \mathbf{t}_I^B \right] + \mathbf{b}_{a,i}, \quad (13)$$

where  $\mathbf{a}_B(t)$  is directly evaluated from the acceleration spline  $\mathcal{S}_a$  and denotes the linear acceleration of the origin of  $B$  expressed in  $W$ ,  $\mathbf{g}$  is the gravity vector expressed in  $W$ , and  $\mathbf{b}_{a,i}$  is the accelerometer bias.

Although Eq. (13) gives the point-wise accelerometer model, the translational initialization does not directly optimize a free acceleration spline together with free endpoint velocities, following the treatment in iKalibr (Chen et al., 2025). Instead, the explicit acceleration spline is eliminated by integrating the accelerometer model over a time interval  $\mathcal{T}_n = [t_n, t_{n+1}]$  in the world frame, where the interval-wise velocity increments are introduced as common unknowns. For the  $i$ -th IMU, define

$$\begin{aligned} \mathbf{C}_n^{(i)}(\mathbf{b}_{a,i}) &= \int_{t_n}^{t_{n+1}} \mathbf{R}_B(\tau) \mathbf{R}_{I_i}^B (\mathbf{a}_i^m(\tau - \Delta t_i) - \mathbf{b}_{a,i}) d\tau, \\ \mathbf{A}_n &= \int_{t_n}^{t_{n+1}} \mathbf{R}_B(\tau) \mathbf{D}_B(\tau) d\tau. \end{aligned} \quad (14)$$

Then the world-frame velocity variation of the reference body satisfies

$$\mathbf{C}_n^{(i)}(\mathbf{b}_{a,i}) - \mathbf{A}_n \mathbf{t}_I^B = \Delta \mathbf{v}_n - \mathbf{g} \Delta t_n, \quad (15)$$

where  $\Delta \mathbf{v}_n = \mathbf{v}_B(t_{n+1}) - \mathbf{v}_B(t_n)$  and  $\Delta t_n = t_{n+1} - t_n$ . Since  $\Delta \mathbf{v}_n$  is common to all rigidly mounted IMUs over the same interval, stacking Eq. (15) over multiple IMUs provides constraints on their relative translational extrinsics.

The translational initialization is therefore formulated as

$$\begin{aligned} \min_{\mathcal{X}_T} & \sum_{i=1}^{N_L} \sum_{n \in \mathcal{J}} \|\mathbf{C}_n^{(i)}(\mathbf{b}_{a,i}) - \mathbf{A}_n \mathbf{t}_I^B - \Delta \mathbf{v}_n + \mathbf{g} \Delta t_n\|_{\Sigma_{a,i}}^2, \\ \mathcal{X}_T &= \{\mathbf{t}_I^B, \mathbf{b}_{a,i}, \Delta \mathbf{v}_n, \mathbf{g}\}. \end{aligned} \quad (16)$$

To remove the translational gauge freedom, the body-frame origin is fixed by definition at the reference IMU, i.e., the reference IMU translation is set to zero. In addition, the gravity magnitude is constrained as  $\|\mathbf{g}\| = g_0$ . The translational extrinsics

are observable only under sufficient angular excitation, since the rank of the accumulated lever-arm matrix  $\mathbf{A}_n$  determines whether different components of  $\mathbf{t}_i^B$  can be separated from the common body acceleration and accelerometer biases.

### 3.4. LiDAR-centric fine calibration with integrated point clouds

After the continuous-time IMU-centric initialization, the coarse extrinsics and temporal offsets are used together with the PPK-IMU trajectory to compensate LiDAR motion distortion and integrate sequential scans. The resulting accumulated point clouds provide denser local structures than individual repetitive scans and are therefore used to establish refined correspondences for fine calibration. In this stage, global-reference correspondences anchor the mobile LiDAR observations to the world-frame reference point cloud, while inter-LiDAR correspondences enforce relative consistency among the onboard LiDARs. These two types of correspondences are finally combined in a translation-constrained multiway optimization (Aras-tounia, 2026).

The calibration target is therefore twofold. On the one hand, the point clouds acquired by different LiDARs should be mutually consistent in their overlapping regions, so that the fused multi-LiDAR point cloud does not contain relative misalignment. On the other hand, the complete LiDAR rig must be correctly aligned with the POS-defined body frame, which requires an external absolute geometric reference rather than inter-LiDAR consistency alone. In this work, the POS trajectory is obtained from tightly coupled GNSS/IMU post-processing with base-station corrections, and a global reference point cloud represented in the same world frame is used as the external benchmark after motion-compensated LiDAR accumulation. The calibration data are collected under normal MMS driving conditions to maintain reliable GNSS/IMU coupling in the post-processed trajectory.

#### 3.4.1. Motion-compensated LiDAR integration for richer local geometry

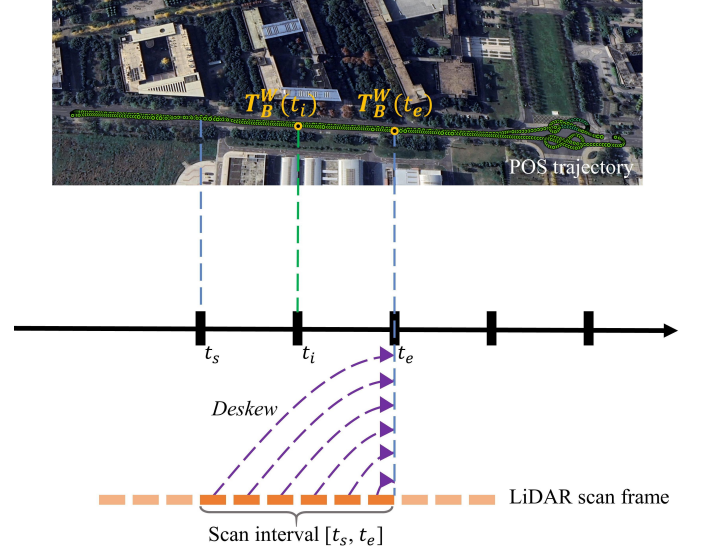
To solve the sparsity issue of repetitive LiDAR scans, the 10 Hz raw scans are accumulated over 1 s integration windows to form local point clouds. For the  $m$ -th integration window  $[t_m - 1, t_m]$ , all LiDAR points are projected to the end time  $t_m$  of the window using the coarse extrinsics, temporal offsets, and POS trajectory.

As illustrated in Fig. 6, each LiDAR point has its own sampling time  $t_p$ . Let  $\mathbf{T}_B^W(t)$  denote the pose of the body frame  $B$  in the world frame, and let  $\mathbf{T}_{L_i,0}^B$  denote the coarse extrinsic transformation from LiDAR  $L_i$  to  $B$ . With the estimated temporal offset  $\Delta t_i$ , the corrected point timestamp is  $\tau_p = t_p + \Delta t_i$ . The POS pose at any corrected timestamp is obtained by interpolation, where translation is linearly interpolated and rotation is interpolated using quaternion SLERP. To remove motion distortion, all points in the integration window are reprojected to the LiDAR frame at the window end time. For the  $m$ -th window, the reference time is  $\tau_m = t_m + \Delta t_i$ . The deskewed point is

therefore computed as

$$\tilde{\mathbf{p}}_{p,m}^{L_i} = (\mathbf{T}_{L_i,0}^B)^{-1} (\mathbf{T}_B^W(\tau_m))^{-1} \mathbf{T}_B^W(\tau_p) \mathbf{T}_{L_i,0}^B \mathbf{p}^{L_i}(t_p). \quad (17)$$

Equation (17) compensates motion distortion and places all points in the LiDAR frame corresponding to the window end time.



**Fig. 6.** Illustration of LiDAR deskewing. A LiDAR scan or integration window is acquired over a finite time interval, where each point has its own sampling timestamp  $t_p$ . Using the corrected timestamp  $\tau_p$  and the interpolated platform poses  $\mathbf{T}_B^W(\tau_p)$  and  $\mathbf{T}_B^W(\tau_m)$ , all points are reprojected to the common window end time  $t_m$  to remove motion distortion.

The resulting local point cloud is then expressed in the world frame at the window end time as

$$\mathcal{P}_{i,m}^W = \{ \mathbf{T}_B^W(\tau_m) \mathbf{T}_{L_i,0}^B \tilde{\mathbf{p}}_{p,m}^{L_i} \}. \quad (18)$$

Thus, ten consecutive 10 Hz scans are converted into  $\mathcal{P}_{i,m}^W$  for each 1 s integration window, providing richer local structures for subsequent correspondence construction (Sampath et al., 2025). The qualitative effect of this integration is shown in Fig. 7.

The effect of deskewing is qualitatively demonstrated in Fig. 7. Without motion compensation, the accumulated point cloud exhibits geometric artifacts and blurred structural boundaries, especially along curved trajectory sections. After deskewing, motion-induced distortions are effectively reduced, resulting in a geometrically consistent local map suitable for subsequent correspondence construction.

#### 3.4.2. Multiway point clouds matching

Each  $\mathcal{P}_{i,m}^W$  is generated using the coarse extrinsics and may still be affected by residual extrinsic errors. To reduce this influence, only nearly straight driving segments are retained by requiring the POS heading variation within each segment to be smaller than  $\tau_h$ , which is set to  $3^\circ$  in our experiments. Meanwhile, because the reference point cloud matching accuracy is about 3 cm and the POS attitude accuracy is about 1 arcmin, reliable calibration requires observations from multiple heading

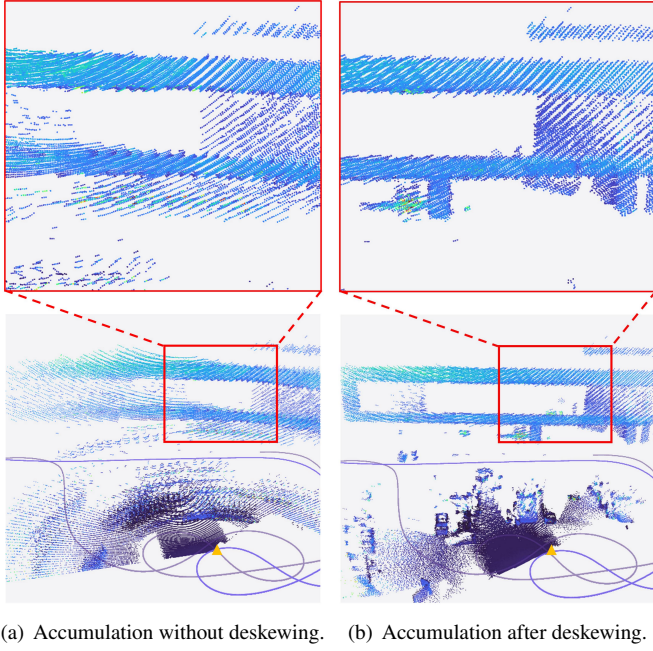


Fig. 7. Qualitative comparison of accumulated LiDAR point clouds before and after deskewing. Motion compensation reduces geometric artifacts and produces a sharper  $\mathcal{P}_{i,m}^W$ .

directions. This diversity is controlled during data acquisition by driving the MMS platform along different directions.

For each retained window  $m$ , multiway matching is performed using the retained  $\mathcal{P}_{i,m}^W$  from all LiDARs. The matching contains two types of constraints: LiDAR-reference matching between LiDAR  $L_i$  and the global reference point cloud, denoted by  $\mathcal{S}_{i0}$  with index 0 representing the reference, and inter-LiDAR matching between overlapping point clouds from LiDARs  $L_i$  and  $L_j$ , denoted by  $\mathcal{S}_{ij}$ . Both types of matches are obtained using GICP (Koide et al., 2021). Besides point correspondences, GICP also estimates registration transformations. For LiDAR-reference matching, the transformation gives the pose of the accumulated local LiDAR frame with respect to the world reference frame. For inter-LiDAR matching, the transformation gives the relative pose between the accumulated local frames of LiDARs  $L_i$  and  $L_j$ .

For the  $m$ -th retained integration window, the complete matching records are written compactly as

$$\begin{aligned} \mathcal{S}_{i0}^m &= \left\{ \hat{\mathbf{T}}_{i,m}^W, \left\{ \left( \mathbf{p}_p^{L_i}, \mathbf{q}_p^W, t_p \right) \right\} \right\}, \\ \mathcal{S}_{ij}^m &= \left\{ \hat{\mathbf{T}}_{ij,m}, \left\{ \left( \mathbf{p}_p^{L_i}, \mathbf{p}_p^{L_j}, t_{i,p}, t_{j,p} \right) \right\} \right\}. \end{aligned} \quad (19)$$

Here,  $\hat{\mathbf{T}}_{i,m}^W$  denotes the GICP-estimated pose of the accumulated LiDAR frame with respect to the world frame, and  $\hat{\mathbf{T}}_{ij,m}$  denotes the GICP-estimated relative transformation between the accumulated frames of LiDARs  $L_i$  and  $L_j$ . The points  $\mathbf{p}_p^{L_i}$  and  $\mathbf{p}_p^{L_j}$  are original LiDAR-frame points,  $\mathbf{q}_p^W$  is the matched reference point, and  $t_p$ ,  $t_{i,p}$ , and  $t_{j,p}$  are point-level LiDAR timestamps corrected by the temporal offset in Eq. (6). These timestamps are inherited from the original LiDAR measurements

rather than from the end time of the 1 s integration window, which is used only to construct  $\mathcal{P}_{i,m}^W$  for matching. In the following,  $\mathcal{S}_{i0}$  and  $\mathcal{S}_{ij}$  denote the retained matching sets over all valid windows, and the calibration residuals directly use their point-correspondence entries. To keep the final point-wise adjustment tractable, the correspondences in each retained window are randomly sampled with at most  $n = 800$  points. Segments with failed GICP convergence, poor fitness, or insufficient correspondences are discarded.

### 3.4.3. Factor graph averaging for LiDAR extrinsics

The IMU-centric initialization mainly estimates the installation parameters of the LiDAR-integrated IMUs relative to the reference IMU. Although the manufacturer-provided LiDAR-to-IMU transformation gives a practical LiDAR extrinsic prior, small internal offsets between the LiDAR optical frame and its built-in IMU may remain. To improve the compatibility of the subsequent point-level optimization, we perform multiway registration independently for each retained window using Open3D (Zhou et al., 2018). For the  $m$ -th window, the retained point clouds  $\mathcal{P}_{i,m}^W$  from all LiDARs are jointly registered, and Open3D returns the optimized LiDAR poses  $\hat{\mathbf{T}}_{i,m}^W$  together with the registration fitness scores  $f_{i,m}$ . These pose estimates provide a direct per-window estimate of the LiDAR-to-body extrinsic.

Specifically, the extrinsic estimate of LiDAR  $L_i$  from window  $m$  is computed as

$$\hat{\mathbf{T}}_{L_i,m}^B = \left( \mathbf{T}_B^W(\tau_{i,m}) \right)^{-1} \hat{\mathbf{T}}_{i,m}^W, \quad \tau_{i,m} = t_m + \Delta t_i. \quad (20)$$

The final pose-level extrinsic prior  $\bar{\mathbf{T}}_{L_i}^B = \{ \bar{\mathbf{R}}_{L_i}^B, \bar{\mathbf{t}}_{L_i}^B \}$  is obtained by averaging the valid per-window estimates. Windows with failed registration or poor fitness are discarded, and in implementation the average is weighted by the Open3D fitness score  $f_{i,m}$ . The rotation is averaged on  $SO(3)$  using a Lie-algebra perturbation mean, while the translation is averaged as a weighted Euclidean mean. The averaged extrinsic is then used as the initial value for the posterior point-correspondence refinement.

### 3.4.4. Fine calibration with refined point correspondences

The final refinement uses the point correspondences in  $\mathcal{S}_{i0}$  and  $\mathcal{S}_{ij}$  to improve the rotational alignment of the LiDARs. These point-level matches provide the most direct and fundamental geometric observations for calibration. The translations are fixed to the averaged multiway-registration estimates  $\bar{\mathbf{t}}_{L_i}^B$ , since allowing translation to vary under weak geometric constraints can absorb rotational errors and break the physical rigidity of the sensor rig. The optimized extrinsic is parameterized as

$$\mathbf{T}_{L_i}^B = \left\{ \exp(\delta\theta_i^\wedge) \bar{\mathbf{R}}_{L_i}^B, \bar{\mathbf{t}}_{L_i}^B \right\}, \quad (21)$$


where  $\delta\theta_i$  is the rotation increment.

For a global-reference correspondence  $(\mathbf{p}^{L_i}, \mathbf{q}^W, t_p) \in \mathcal{S}_{i0}$ , the residual is

$$\mathbf{r}^{i0} = \mathbf{T}_B^W(t_p) \mathbf{T}_{L_i}^B \mathbf{p}^{L_i} - \mathbf{q}^W. \quad (22)$$

For an inter-LiDAR correspondence  $(\mathbf{p}^{L_i}, \mathbf{p}^{L_j}, t_{i,p}, t_{j,p}) \in \mathcal{S}_{ij}$ ,  $i \neq j$ , the residual is

$$\mathbf{r}^{ij} = \mathbf{T}_B^W(t_{i,p}) \mathbf{T}_{L_i}^B \mathbf{p}^{L_i} - \mathbf{T}_B^W(t_{j,p}) \mathbf{T}_{L_j}^B \mathbf{p}^{L_j}. \quad (23)$$



**Fig. 8.** Survey platform and reconstructed 3D point cloud used as the global reference point cloud in the experiments. Green markers indicate camera positions, while colored points represent the reconstructed 3D point cloud.

For compact notation in the objective function,  $\mathbf{r}_m^{ref}$  denotes the  $m$ -th global-reference residual defined by Eq. (22), and  $\mathbf{r}_n^{ll}$  denotes the  $n$ -th inter-LiDAR residual defined by Eq. (23). The rotation-only refinement is then formulated as

$$\min_{\{\delta\theta_i\}} \sum_i \sum_{\mathcal{S}_{i0}} \rho(\|\mathbf{r}_m^{ref}\|^2) + \sum_{i<j} \sum_{\mathcal{S}_{ij}} \rho(\|\mathbf{r}_n^{ll}\|^2) \quad (24)$$

where  $\rho(\cdot)$  is the Huber loss. After convergence, the optimized rotations and fixed translations form the final LiDAR extrinsics.

## 4. Experimental evaluations and analyses

### 4.1. Dataset description

The experimental setup comprises a vehicle-mounted MMS and an external reference-map survey, as shown in Fig. 4 and Fig. 8, respectively. The vehicle-mounted platform integrates four RoboSense Airy LiDARs and a reference IMU mounted on a common rigid frame, with a custom multi-sensor fusion carrier board enabling synchronized multi-sensor data acquisition. Each Airy sensor provides a 360° horizontal and 90° vertical field of view (FOV) at 10 Hz, recording range, intensity, and timestamp information. Hardware-level synchronization is achieved using the GNSS time reference: the GNSS board and the reference IMU are synchronized through PPS and GNRMC messages, and the LiDAR units are also synchronized to GNSS time through PPS and GNRMC signals before residual temporal offsets are estimated in the proposed initialization. In addition to the reference IMU, each LiDAR unit includes a LiDAR-integrated IMU for high-frequency motion measurement. The vehicle dataset was collected along the main roads of the Xipu

campus of Southwest Jiaotong University, covering approximately 3 km at a driving speed of about 20–30 km/h.

For the global reference point cloud, a DJI M4E drone equipped with its built-in RGB camera was used to acquire oblique imagery, from which a dense 3D point cloud was reconstructed as the external reference for global-reference matching and M3C2-based deviation analysis. According to the reconstruction quality report, the photogrammetric block contains one connected component, with 2881 images in the connected region, 7,207,984 tie-point projections, and 969,365 tie points. The reported reprojection RMS is 1.057 pixels, and the georeferencing RMSE is 0.029 m. The reconstructed dense point cloud is therefore used as the global geometric reference rather than as a set of manually selected control points. The external-reference consistency reported below measures agreement with this photogrammetric reference.

More interactive point-cloud visualizations and the experimental dataset are available online<sup>1</sup>.

### 4.2. Comparison with other methods

To evaluate the reliability of the proposed calibration method on the Airy multi-LiDAR platform, we compare it with two representative LiDAR-IMU calibration frameworks: LI-Calib (Lv et al., 2020) and iKalibr (Chen et al., 2025). LI-Calib represents the classical single-LiDAR-IMU calibration setting, where each LiDAR is calibrated independently with respect to the IMU. iKalibr represents a multi-sensor joint optimization framework that simultaneously optimizes multiple LiDARs with a shared IMU. In contrast, the proposed method initializes the extrinsics using IMU-centric rigid-body constraints and further refines them by exploiting multi-LiDAR geometric consistency and a dense global reference point cloud reconstructed from oblique imagery.

Unlike conventional spinning LiDARs, the Airy LiDARs used in our platform generate sparse and repetitive scanning patterns. Such measurements provide weaker and more ambiguous geometric constraints, especially for translation estimation. As a result, a low local LiDAR-IMU residual does not necessarily indicate a physically reliable extrinsic calibration. We therefore evaluate the calibration results from three complementary perspectives: (i) the estimated extrinsic parameters compared against the mechanical design reference as a physical plausibility check; (ii) the internal consistency of the fused multi-LiDAR point cloud; and (iii) the deviation with respect to the dense global reference point cloud. All extrinsics are represented as LiDAR-to-IMU transformations. The mechanical design reference is used only as a physical plausibility check, while the internal consistency and global reference error are used as the primary quantitative evaluation metrics. Euler angles are reported only for readability.

#### 4.2.1. Quantitative comparison

As shown in Table 3, the estimated extrinsic parameters reveal clear differences among the three calibration methods. LI-Calib calibrates each LiDAR independently with respect to the

<sup>1</sup><https://vrlab.org.cn/~hanhu/projects/mmscal/>.

Table 3: LiDAR-to-IMU extrinsic parameters estimated by different calibration methods.

LiDAR	Method	$t_x$ (m)	$t_y$ (m)	$t_z$ (m)	yaw ( $^\circ$ )	pitch ( $^\circ$ )	roll ( $^\circ$ )
lidar1	Design	-0.520	0.375	0.030	150.000	45.000	0.000
	LI-Calib	-0.336	0.244	-0.045	149.790	44.570	1.497
	iKalibr	-0.487	-0.033	-0.135	152.192	47.032	0.514
	Ours	-0.531	0.374	0.028	149.714	45.630	0.236
lidar2	Design	-0.520	-0.110	0.030	180.000	45.000	0.000
	LI-Calib	-0.464	-0.105	-0.082	179.130	43.083	1.602
	iKalibr	-0.417	-0.061	-0.129	-177.207	45.498	0.635
	Ours	-0.523	-0.112	0.030	-179.644	45.636	0.303
lidar3	Design	0.600	-0.110	0.030	0.000	45.000	0.000
	LI-Calib	0.421	-0.081	0.083	-1.501	43.289	-1.943
	iKalibr	-0.073	-0.072	-0.073	0.058	45.497	0.516
	Ours	0.591	-0.113	0.032	-1.765	45.106	-1.455
lidar4	Design	0.600	0.375	0.030	30.000	45.000	0.000
	LI-Calib	0.265	0.237	-0.085	27.380	46.681	0.761
	iKalibr	-0.030	0.125	-0.076	30.689	45.796	-1.618
	Ours	0.601	0.374	0.027	29.350	45.480	-0.829

IMU and therefore does not explicitly enforce geometric consistency among different LiDARs. Although the local LiDAR-IMU residuals can be minimized, the estimated translations show noticeable deviations from the mechanical design reference. The most evident case occurs for lidar4, where the translation differs by about 0.380 m from the nominal installation. This indicates that the optimizer may compensate for weak geometric constraints by introducing physically implausible translational offsets.

iKalibr performs joint multi-sensor optimization and produces rotations that are relatively close to the design reference. However, its translation estimates are still unstable on the Airy multi-LiDAR platform. For example, lidar3 and lidar4 are estimated at positions that deviate significantly from the expected vehicle installation geometry. This result suggests that joint optimization alone cannot fully resolve the weak translation observability caused by sparse and repetitive LiDAR scanning patterns.

In contrast, the proposed method keeps the estimated extrinsics physically consistent with the mechanical design reference across all four LiDARs. The translations remain close to the nominal sensor layout, and the rotations are also consistent with the expected mounting configuration. Since the mechanical design reference is only used as a plausibility check, we further evaluate whether these extrinsics lead to improved internal consistency in the fused multi-LiDAR point cloud.

Table 4 and Fig. 9 provide quantitative evidence for the internal consistency improvement. The reported M3C2 value is the standard deviation of the distance distribution, which reflects the dispersion of local surface deviations. LI-Calib achieves an overall consistency of 0.318 m, indicating that independently calibrated LiDARs cannot be reliably fused into a coherent multi-LiDAR map. iKalibr reduces the overall consistency error to 0.262 m, showing that joint optimization improves rela-

Table 4: Quantitative comparison of different calibration methods. Overall consistency and global reference error are computed from M3C2 distance statistics using the CloudCompare M3C2 plugin.

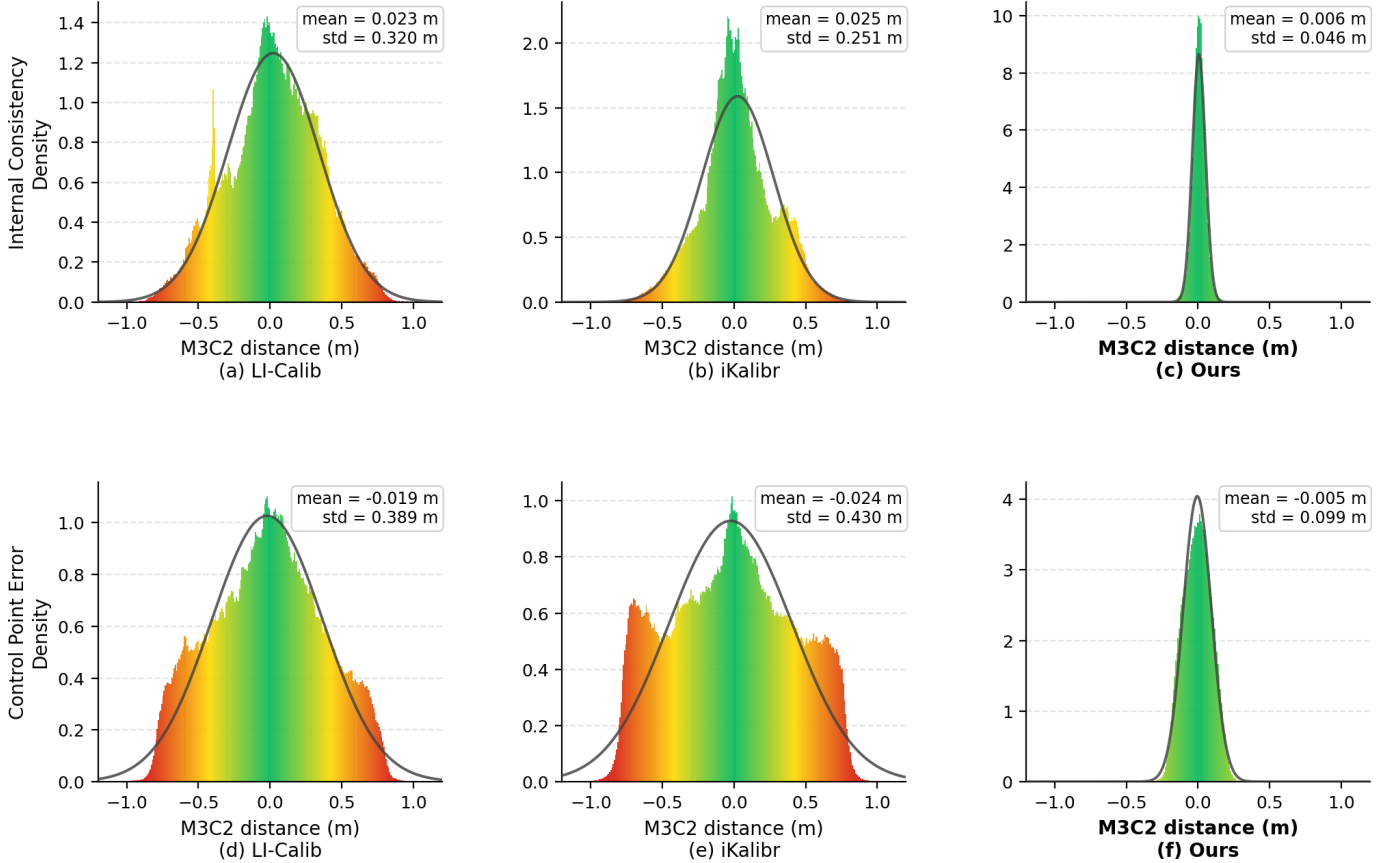
Method	Overall consistency (m)	Global reference error (m)
LI-Calib	0.318	0.385
iKalibr	0.262	0.427
Ours	<b>0.048</b>	<b>0.099</b>

tive alignment to some extent. However, its M3C2 distribution remains broad and irregular, which suggests that the estimated extrinsics are still not globally consistent across the full platform.

The proposed method achieves an overall consistency of 0.048 m, about 5.5–6.6 times lower than the two classical baselines. The corresponding M3C2 distribution is narrow and well-centered, with most internal consistency errors concentrated around zero. This demonstrates that the proposed calibration strategy effectively suppresses cross-LiDAR misalignment and produces a geometrically coherent fused point cloud. The global reference error is also reduced from 0.385 m and 0.427 m for the two baselines to 0.099 m for the proposed method, further indicating that the improved internal consistency does not come at the cost of degraded absolute accuracy.

#### 4.2.2. Qualitative comparison

Figure 10 further verifies the internal consistency from a qualitative perspective. Different colors denote point clouds from different LiDARs. For LI-Calib and iKalibr, clear color discontinuities and structural misalignments can be observed at the boundaries between adjacent LiDAR coverage regions. Typical artifacts include duplicated edges, blurred facades, and split pole structures. These visual artifacts are consistent with



**Fig. 9.** M3C2 distance distributions for internal consistency and global reference error under different calibration methods. Top row: M3C2 distances computed between overlapping LiDAR pairs (lidar1-lidar2, lidar1-lidar4, lidar3-lidar4), merged per method. Bottom row: M3C2 distances computed between each LiDAR point cloud and the global reference point cloud, merged across all four LiDARs per method. The Gaussian fit is overlaid for reference. Classical methods produce broad and irregular distributions, while the proposed method achieves narrow, well-centered distributions in both metrics.

the large translational deviations in Table 3 and the broad M3C2 (Lague et al., 2013) distributions in Fig. 9.

By contrast, the proposed method produces a structurally coherent fused map with significantly reduced boundary artifacts. The overlapping regions from different LiDARs are better aligned, and the reconstructed facades, poles, and vegetation boundaries appear more compact and continuous. This qualitative result confirms that the proposed calibration method improves the internal consistency of the fused multi-LiDAR point cloud, which is critical for high-quality mapping with the Airy multi-LiDAR platform.

### 4.3. Analysis of multi-stage calibration

#### 4.3.1. Internal precision analysis of adjustment stages

To further evaluate the internal precision of the estimated parameters, we analyze the covariance matrices obtained from the converged least-squares adjustments. For a linearized adjustment model with residual vector  $\mathbf{v}$ , Jacobian matrix  $\mathbf{A}$ , and weight matrix  $\mathbf{P}$ , the unit weight standard deviation is computed as


$$\hat{\sigma}_0 = \sqrt{\frac{\mathbf{v}^T \mathbf{P} \mathbf{v}}{n - u}}, \quad (25)$$

where  $n$  is the residual dimension and  $u$  is the number of effective unknowns. The covariance matrix of the estimated parameters is then expressed as

$$\Sigma_{\hat{\mathbf{x}}} = \hat{\sigma}_0^2 (\mathbf{A}^T \mathbf{P} \mathbf{A})^{-1}. \quad (26)$$

The square roots of the diagonal elements of  $\Sigma_{\hat{\mathbf{x}}}$  are reported as the internal standard deviations of the corresponding extrinsic parameters.

Table 5 summarizes the internal standard deviations of the IMU-centric initialization and the posterior refinement. For completeness, the unit weight standard deviations computed from Eq. (25) are 7.6276 for the IMU-centric initialization and 0.0256 for the posterior refinement. These two values are not directly comparable because the two stages use different observation models, residual definitions, and weighting schemes. They are therefore used only to describe the internal residual consistency of each adjustment stage. In the IMU-centric initialization, the reference IMU is fixed and the LiDAR-related extrinsics are estimated with both rotation and translation. The resulting translational standard deviations are at the sub-millimeter level, while the rotational standard deviations are approximately  $0.011^\circ$ – $0.020^\circ$ . This supports the use of the IMU-centric extrinsics as reliable initial values for the subsequent LiDAR



**Fig. 10.** Qualitative comparison of internal consistency in fused multi-LiDAR maps under different calibration methods. Different colors denote point clouds from different LiDARs. The red box highlights a representative LiDAR boundary region for evaluating cross-LiDAR structural consistency. For visual clarity, ground points are omitted where applicable because the nearly planar ground surface provides weak visual separation between different calibration results. LI-Calib and iKalibr show visible color discontinuities and structural misalignments, while the proposed method produces a more coherent fused map with reduced boundary artifacts.

multiway registration.

In the posterior refinement, the translations are fixed to the averaged multiway-registration estimates and only the three rotational components of each LiDAR are optimized, resulting in 12 effective unknowns. Accordingly, the translational parameters are excluded from the covariance analysis of this stage. The posterior refinement substantially reduces the uncertainty of the  $r_z$  component and maintains comparable precision for the  $r_x$  and  $r_y$  components. After refinement, all rotational standard deviations remain below  $0.013^\circ$ , indicating that the constructed inter-LiDAR and global-reference correspondences provide sufficient constraints for the rotation-only refinement.

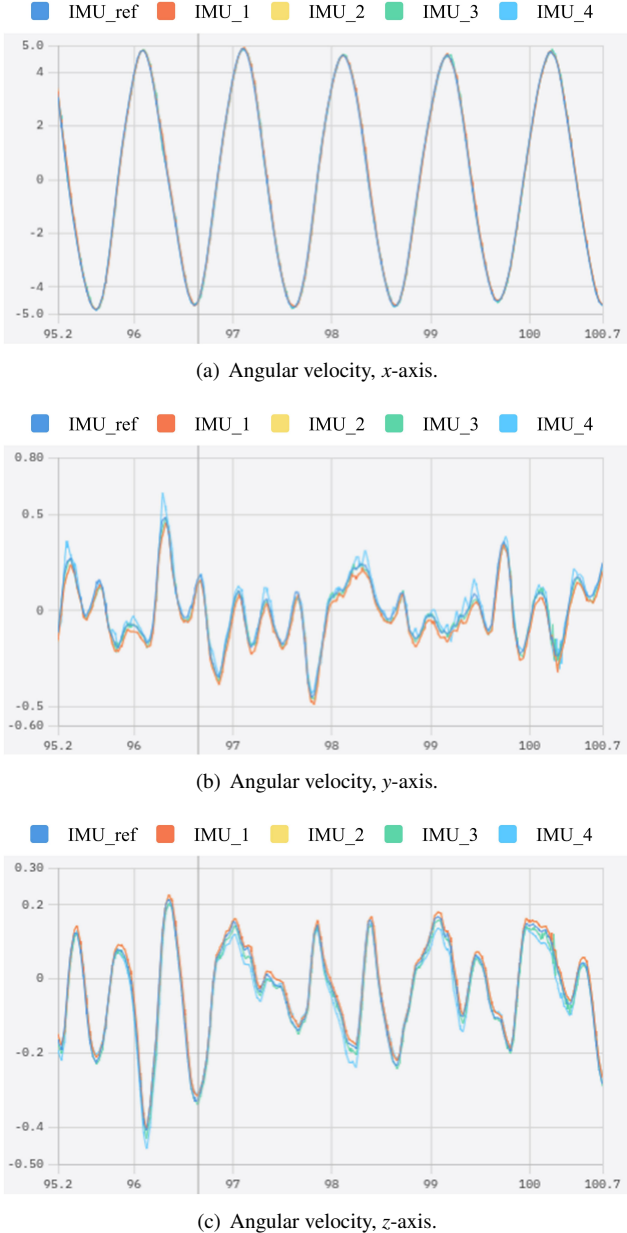
To further validate the IMU-centric initialization, Fig. 11 visualizes the angular velocity measurements after transforming the LiDAR-integrated IMUs into the reference IMU frame. These measurements correspond to the gyroscope consistency constraint used for rotational initialization. Since all IMUs are rigidly mounted on the same platform, their angular velocity measurements should be consistent after temporal-offset and rotational calibration. The close agreement of the angular velocity curves along all three axes therefore directly supports the estimated temporal offsets and rotational extrinsics.

Table 5: Internal standard deviations of the estimated extrinsic parameters. Rotations are in degrees and translations are in millimeters. Values are rounded to meaningful precision.

Rotation std. (deg)				
Stage	LiDAR	$\sigma_{r_x}$	$\sigma_{r_y}$	$\sigma_{r_z}$
IMU-centric init.	lidar1	0.012	0.012	0.020
	lidar2	0.012	0.012	0.020
	lidar3	0.012	0.012	0.019
	lidar4	0.011	0.012	0.018
Posterior refine	lidar1	0.011	0.012	0.005
	lidar2	0.011	0.012	0.004
	lidar3	0.012	0.012	0.010
	lidar4	0.012	0.012	0.009
Translation std. of IMU-centric init. (mm)				
Stage	LiDAR	$\sigma_{t_x}$	$\sigma_{t_y}$	$\sigma_{t_z}$
IMU-centric init.	lidar1	0.3	0.3	0.3
	lidar2	0.3	0.3	0.3
	lidar3	0.3	0.3	0.2
	lidar4	0.3	0.3	0.3

#### 4.3.2. Effect of multi-frame accumulation under reliable initialization

To validate the role of multi-frame accumulation when a reliable initial estimate is available, we conduct a controlled ablation study using Multi-LiCa (Kulmer et al., 2024), an open-source inter-LiDAR consistency calibration framework that op-



**Fig. 11.** Zoomed-in angular velocity consistency after IMU-centric initialization. The LiDAR-integrated IMUs are transformed into the reference IMU frame using the estimated temporal offsets and rotational extrinsics. The three axes are shown over the same time window, and the close agreement of the angular velocity curves validates the temporal and rotational alignment.

erates without IMU participation. By fixing the calibration backend and using the same motion-compensated input and initialization strategy, this experiment isolates whether a single corrected frame provides enough geometric structure for reliable inter-LiDAR registration, or whether multi-frame accumulation is required to supply sufficient constraints.

Two input settings are compared. In Setting 1, motion compensation is applied, and the nominal design values are used as initialization, while only a single corrected frame is used. In Setting 2, multi-frame accumulation is further incorporated on top of the motion-compensated input with the same initializa-

Table 6: Evaluation of single-frame and multi-frame inputs in Multi-LiCa under the same motion compensation and initialization conditions.

Setting	M.C.	Init.	M.F.	Avg. RMSE	Status
S1	✓	✓	×	0.381	Local optimum
S2	✓	✓	✓	0.213	Stable

tion. Table 6 summarizes the quantitative results, where M.C., Init., and M.F. denote motion compensation, coarse initialization, and multi-frame accumulation, respectively. Fig. 12 provides the corresponding qualitative comparisons.


As shown in Fig. 12 and Table 6, with motion compensation and design initialization in Setting 1, the fused point clouds show reasonable consistency before refinement, indicating that the initial conditions are not severely degraded. However, noticeable local ghosting and structural splitting can still be observed in the calibrated result, implying that the geometric constraints provided by a single corrected frame remain insufficient. This is further confirmed by the estimated extrinsics of the representative challenging pair  $L_4 \rightarrow L_1$ : the estimated translation drifts to  $(-17.60, -17.89, -4.38)$  m with rotations of  $(30.76^\circ, 14.51^\circ, -12.81^\circ)$ , which is clearly physically implausible, indicating convergence to a local optimum under single-frame constraints.

By further introducing multi-frame accumulation in Setting 2, the overlap among different LiDARs becomes significantly tighter, especially around tree trunks, canopy boundaries, and other vertical structures. Quantitatively, the average inlier RMSE decreases to 0.213, substantially lower than Setting 1. The same challenging pair  $L_4 \rightarrow L_1$  converges to  $(-0.314, -0.044, 0.128)$  m and  $(-1.00^\circ, -0.96^\circ, 0.37^\circ)$ , returning to a physically reasonable range. These results validate the necessity of multi-frame accumulation under reliable initialization, showing that the accumulated geometry is the critical factor that resolves the geometric ambiguity caused by insufficient single-frame constraints from repetitive scanning LiDARs.

#### 4.3.3. Effect of fixing multiway-registration translations during refinement

The Airy LiDARs used in our platform generate sparse and repetitive scanning patterns, which result in weak and ambiguous geometric constraints, particularly for the translational components of the extrinsic parameters. When the translation is freely optimized together with rotation during refinement, the optimizer may reduce the local alignment residual by adjusting the relative positions of the LiDARs rather than correcting the physically meaningful rotational error. Although this can produce visually plausible local overlaps, the resulting translation estimates may deviate from the true rigid installation geometry and deform the inter-LiDAR baseline structure.

To isolate this effect, we compare two refinement strategies initialized from the same multiway-registration extrinsic prior: (i) unconstrained 6-DoF refinement, where both rotation and translation are freely optimized; and (ii) the proposed fixed-translation refinement, where the multiway-registration trans-



**Fig. 12.** Qualitative comparison of two progressive settings in Multi-LiCa. (a) Motion-compensated single-frame input with design initialization, and (b) motion-compensated multi-frame accumulation with the same initialization. The upper row shows the point clouds before calibration, and the lower row shows the corresponding results after Multi-LiCa calibration. Ground points are omitted where applicable to make the non-ground alignment differences more visible.

lation is fixed and only the rotational parameters are refined. Both settings use the same refinement objective and the same initialization.

Table 7 reports the translational deviation from the mechanical design reference and the inter-LiDAR baseline change for each setting. The inter-LiDAR baseline length between LiDAR  $i$  and LiDAR  $j$  is defined as:

$$b_{ij} = \|\hat{\mathbf{t}}_{L_i}^B - \hat{\mathbf{t}}_{L_j}^B\|_2, \quad (27)$$


and the baseline change is computed as  $|\hat{b}_{ij} - b_{ij}^0|$ , where  $b_{ij}^0$  is

the design reference baseline. This quantity serves as a physical consistency indicator for the rigid multi-LiDAR rig: any significant change in inter-LiDAR baseline length indicates that the optimizer has deformed the assumed rigid structure, which is physically impossible for a fixed mechanical installation.


In the unconstrained 6-DoF setting, the maximum per-LiDAR translational deviation reaches 0.158 m (lidar2), and the maximum inter-LiDAR baseline change reaches 0.157 m (lidar2-lidar4 pair, design baseline: 1.221 m). These values indicate that the optimizer has effectively deformed the rigid

Table 7: Effect of freely optimizing or fixing translations during posterior refinement. Translation deviation and baseline change are measured with respect to the mechanical design reference. “Reg.” denotes the multiway-registration extrinsic prior initialized from the IMU-centric coarse extrinsics. All values are in meters.

Setting	Optimized parameters	Max trans. dev. (m)	Max baseline change (m)
Reg. + 6-DoF refine	$R + t$	0.158	0.157
Ours: Reg. $t$ fixed + rot. refine	$R$ only	<b>0.011</b>	<b>0.012</b>



(a) Unconstrained refinement distorts physical sensor geometry.



(b) Fixed-translation refinement retains the IMU-centric sensor layout.

**Fig. 13.** Qualitative comparison of unconstrained and fixed-translation refinement strategies, overlaid on the aerial photogrammetric reference (true-color dense point cloud). LiDAR 1 points are shown in yellow and LiDAR 2 points in cyan. Red rectangles indicate representative reference structures. The unconstrained refinement produces physically implausible translational drift and distorts the relative sensor layout, while fixing the multiway-registration translations retains the initialized sensor geometry and achieves better alignment with the reference structures, consistent with the quantitative comparison in Table 7.

multi-LiDAR geometry during refinement, absorbing residual errors through physically implausible translational shifts. Such deformation is inconsistent with the fixed mechanical installation of the rig and suggests that the free translation is compensating for errors that should be attributed to other sources, such as residual rotational misalignment or unmodeled sensor noise.

In contrast, when the multiway-registration translations are fixed, the maximum translational deviation remains 0.011 m and the maximum baseline change remains 0.012 m, while the rotation is still refined using geometric and reference constraints. Figure 13 provides a qualitative comparison of the two refinement strategies overlaid on the aerial photogrammetric reference. LiDAR 1 (yellow) and LiDAR 2 (cyan) are displayed alongside the dense point cloud reconstructed from the aerial oblique photogrammetry (true-color background). The red rectangles highlight representative reference structures used for visual alignment assessment.


As shown in Fig. 13, both settings produce visually similar overall structures when viewed at the scene level. However, close inspection of the highlighted boundary regions reveals a clear difference: in the unconstrained 6-DoF result, the yellow (LiDAR 1) and cyan (LiDAR 2) point clouds exhibit residual

color separation at the edges of the highlighted structures, indicating that the two LiDARs remain locally misaligned despite the optimizer converging to a low residual. In the proposed method, the two LiDAR point clouds align more consistently with the reference edges, and the color transition at the boundaries is more coherent.

These results demonstrate that freely optimizing the translation during refinement, even from a good IMU-centric initialization, can lead to physically implausible translational drift that deforms the rigid rig geometry without producing a correspondingly better geometric alignment. The posterior stage therefore keeps the multiway-registration translations fixed and refines only the rotational parameters, ensuring that the optimization does not absorb residual errors through weakly observable translational updates.

#### 4.4. Evaluation on different scenes of the calibrated extrinsics

To evaluate the stability of the calibrated extrinsics across different mapping scenes, we directly apply three levels of extrinsic parameters to three representative scenes beyond the calibration area, covering different roadside structures including building facades, tree-lined roads, and open street intersections.



**Fig. 14.** Qualitative evaluation of three extrinsic parameter sets in different scenes. From left to right: manually designed parameters, IMU-centric initialization, and the final refined calibration. From top to bottom: three representative scenes with different roadside structures beyond the calibration area. No scene-specific recalibration or optimization is applied. Ground points are omitted where applicable to emphasize non-ground structural differences.

The compared parameter sets are: (1) the manually designed extrinsics, (2) the IMU-centric initialized extrinsics, and (3) the final refined extrinsics obtained by the proposed method. No additional recalibration or scene-specific optimization is performed in this experiment. For visual clarity, only the most representative local regions from each scene are shown. Ground points are omitted where applicable because the ground surface is difficult to visually distinguish across calibration results, whereas non-ground structures more clearly reveal color separation, duplicated edges, and local misalignment.


As shown in Fig. 14, the manually designed parameters already provide a coarse but usable alignment in all three scenes. However, noticeable color separation, edge thickening, and local ghosting can be observed around structural boundaries such as roof edges, facade contours, repeated vertical structures, and roadside poles, indicating that the design values are sufficient only for rough alignment and remain inadequate for accurate multi-LiDAR fusion in different scenes.

After replacing the designed parameters with the IMU-centric initialized extrinsics, the overlap among different Li-

DARs becomes consistently tighter across all three scenes. The major structural outlines become better aligned, and the obvious misalignment visible in column (a) is substantially reduced, demonstrating that the IMU-centric initialization markedly improves coarse extrinsic quality without requiring any scene-specific data. Nevertheless, residual misalignment can still be observed in fine structures and sharp geometric boundaries, where the fused points remain slightly thickened or split.

With the final refined calibration, the fused point clouds show the best consistency across all three scenes. Compared with the first two columns, the refined results exhibit more compact overlap and sharper alignment along repeated vertical structures, roof boundaries, facade edges, and other high-frequency geometric details, confirming that the calibrated extrinsics remain stable in different roadside scenes.

These results show a consistent improvement across the three calibration stages, which agrees with the quantitative findings in Section 4.2. To further evaluate the calibrated extrinsics in additional scenes, Fig. 15 provides a direct before-and-after comparison in three additional scenes with different roadside struc-



**Fig. 15.** Before-and-after qualitative comparison in three additional scenes beyond the calibration area. The upper row shows fused multi-LiDAR maps using the manually designed parameters, and the lower row shows the corresponding results after applying the calibrated extrinsics from the proposed pipeline. The calibrated results exhibit reduced color separation, thinner structural edges, and more coherent overlaps across different LiDARs. For visual clarity, ground points are omitted where applicable so that non-ground structural differences are more evident.

tures, where “before” refers to the manually designed parameters and “after” refers to the final calibrated extrinsics from the proposed pipeline.

As shown in Fig. 15, the calibrated extrinsics consistently reduce duplicated edges, color separation, and structural ghosting across all three additional scenes. Together with the stage-wise comparison in Fig. 14, these results demonstrate that the calibrated extrinsics remain stable across different roadside scenes without scene-specific recalibration.

## 5. Conclusion

This paper presents a hierarchical extrinsic calibration method for vehicle-mounted multi-LiDAR MMS equipped with repetitive scanning LiDARs. The method is motivated by two key challenges: sparse single-frame observations provide insufficient geometric constraints for stable inter-sensor calibration, while temporal multi-frame accumulation is highly sensitive to inaccurate initial extrinsics and may produce ghosting artifacts and blurred structural boundaries.

To address this circular dependency, continuous-time IMU-centric initialization is first used to estimate temporal offsets and obtain physically plausible coarse rotational and translational extrinsics by exploiting rigid-body constraints between  $I_i$  and  $B$ . Based on this initialization, motion-compensated multi-frame accumulation is performed to enhance local geometric

structure, support factor-graph averaging of multiway registration results, and construct reliable inter-LiDAR and global-reference correspondences. The posterior refinement fixes the multiway-registration translations and optimizes only the rotational components, preventing compensatory translational drift under weak geometric observability.

Experiments on a four-LiDAR RoboSense Airy platform show that the proposed method improves both internal point cloud consistency and absolute geometric accuracy compared with existing calibration methods. The overall consistency error is reduced from 0.318 m and 0.262 m for LI-Calib and iKalibr to 0.048 m, and the global reference error is reduced from 0.385 m and 0.427 m to 0.099 m, respectively. Additional ablation studies confirm the necessity of multi-frame accumulation under reliable initialization, and evaluations on different scenes indicate that the calibrated extrinsics remain stable across different environments. The current method relies on an external global reference point cloud for posterior refinement, which may limit its applicability when such a reference is unavailable. Future work will explore online recalibration, reference-light refinement strategies, and additional sensing modalities to further strengthen geometric constraints.

## Acknowledgments

This study was supported in part by the National Natural Science Foundation of China (Project No. U25A20772, 42230102)

and the Natural Science Foundation of Sichuan Province under Grant 2026NSFSCZY0054.

## References

- Abdullah, Q., 2023. The asprs positional accuracy standards, edition 2: The geospatial mapping industry guide to best practices. Photogrammetric Engineering & Remote Sensing 89.
- Aijazi, A.K., Checchin, P., 2024. Non-repetitive scanning lidar sensor for robust 3d point cloud registration in localization and mapping applications. Sensors 24, 378.
- Arastounia, M., 2026. Accuracy evaluation of airborne and terrestrial point cloud registration. Photogrammetric Engineering & Remote Sensing 92, 155–166.
- Chen, S., Li, X., Li, S., Zhou, Y., Yang, X., 2025. ikalibr: Unified targetless spatiotemporal calibration for resilient integrated inertial systems. IEEE Transactions on Robotics .
- Elhashash, M., Albanwan, H., Qin, R., 2022. A review of mobile mapping systems: From sensors to applications. Sensors 22, 4262.
- Hug, D., Alzugaray, I., Chli, M., 2024. Hyperion – A fast, versatile symbolic Gaussian Belief Propagation framework for Continuous-Time SLAM, in: Computer Vision – ECCV 2024, Springer Nature Switzerland, Cham. pp. 215–231.
- Jiao, J., Ye, H., Zhu, Y., Liu, M., 2021. Robust odometry and mapping for multi-lidar systems with online extrinsic calibration. IEEE Transactions on Robotics 38, 351–371.
- Jonassen, V.O., Kjörsvik, N.S., Gjevestad, J.G.O., 2023. Scalable hybrid adjustment of images and lidar point clouds. ISPRS Journal of Photogrammetry and Remote Sensing 202, 652–662.
- Khoramshahi, E., Nezami, S., Pellikka, P., Honkavaara, E., Chen, Y., Habib, A., 2025. A taxonomy of sensors, calibration and computational methods, and applications of mobile mapping systems: A comprehensive review. Remote Sensing 17, 1502.
- Kim, T., Pak, G., Kim, E., 2024. Gril-calib: Targetless ground robot imu-lidar extrinsic calibration method using ground plane motion constraints. IEEE Robotics and Automation Letters 9, 5409–5416.
- Koide, K., Yokozuka, M., Oishi, S., Banno, A., 2021. Voxelized gicp for fast and accurate 3d point cloud registration, in: 2021 IEEE international conference on robotics and automation (ICRA), IEEE. pp. 11054–11059.
- Kulmer, D., Tahiraj, I., Chumak, A., Lienkamp, M., 2024. Multi-lica: A motion- and targetless multi-lidar-to-lidar calibration framework, in: 2024 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI), IEEE. pp. 1–7.
- Lague, D., Brodu, N., Leroux, J., 2013. Accurate 3d comparison of complex topography with terrestrial laser scanner: Application to the rangitikei canyon (nz). ISPRS journal of photogrammetry and remote sensing 82, 10–26.
- Lee, E., Park, S., Jang, H., Choi, W., Sohn, H.G., 2024. Enhancement of low-cost uav-based photogrammetric point cloud using mms point cloud and oblique images for 3d urban reconstruction. Measurement 226, 114158.
- Lee, H., Chung, W., 2022. Extrinsic calibration of multiple 3d lidar sensors by the use of planar objects. Sensors 22, 7234.
- Li, L., Yang, F., Zhu, H., Li, D., Li, Y., Tang, L., 2022. Point cloud registration and localization based on voxel plane features. ISPRS Journal of Photogrammetry and Remote Sensing 188, 363–374.
- Li, S., Li, X., Chen, S., Zhou, Y., Wang, S., 2023a. Two-step lidar/camera/imu spatial and temporal calibration based on continuous-time trajectory estimation. IEEE Transactions on Industrial Electronics 71, 3182–3191.
- Li, S., Li, X., Chen, S., Zhou, Y., Wang, S., 2024. Targetless spatiotemporal calibration of multi-lidar multi-imu system based on continuous-time optimization. IEEE Transactions on Industrial Informatics 20, 7565–7575.
- Li, Z., Zhang, X., Wang, Y., Li, J., Li, X., Yang, B., Stilla, U., 2023b. MuCo-Graph: Multi-scale constraints for graph-based global registration of terrestrial point clouds. ISPRS Journal of Photogrammetry and Remote Sensing 202, 628–644.
- Ling, X., Qin, R., 2022. A graph-matching approach for cross-view registration of over-view and street-view based point clouds. ISPRS Journal of Photogrammetry and Remote Sensing 185, 2–15.
- Lv, J., Xu, J., Hu, K., Liu, Y., Zuo, X., 2020. Targetless calibration of lidar-imu system based on continuous-time batch estimation, in: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), IEEE. pp. 9968–9975.
- Lv, J., Zuo, X., Hu, K., Xu, J., Huang, G., Liu, Y., 2022. Observability-aware intrinsic and extrinsic calibration of lidar-imu systems. IEEE Transactions on Robotics 38, 3734–3753.
- Lyu, M., Yang, J., Qi, Z., Xu, R., Liu, J., 2024. Rigid pairwise 3d point cloud registration: A survey. Pattern Recognition 151, 110408.
- Martiros, H., Miller, A., Bucki, N., Solliday, B., Kennedy, R., Zhu, D., Dang, T., Pattison, D., Cheng, J., Strathearn, S., Teyssier, M., Manchester, Z., Carlone, L., Sverdllov, A., 2022. SymForce: Symbolic computation and code generation for robotics, in: Proceedings of Robotics: Science and Systems, p. 041.
- Pan, D., Shao, J., Zhang, S., Zhang, S., Chang, B., Xiong, H., Zhang, W., 2023. Slam-based forest plot mapping by integrating imu and self-calibrated dual 3-d laser scanners. IEEE Transactions on Geoscience and Remote Sensing 61, 1–13.
- Pöpl, F., Ullrich, A., Mandlbürger, G., Pfeifer, N., 2024. A flexible trajectory estimation methodology for kinematic laser scanning. ISPRS journal of photogrammetry and remote sensing 215, 62–79.
- Sampath, A., Irwin, J.R., Stoker, J.M., 2025. Evaluating three-dimensional elevation program lidar consistency and accuracy at scale using cloud-native, open-source methods. Photogrammetric Engineering & Remote Sensing 91, 777–785.
- Wang, Z., Huang, S., Butt, J.A., Cai, Y., Varga, M., Wieser, A., 2025. Cross-modal feature fusion for robust point cloud registration with ambiguous geometry. ISPRS Journal of Photogrammetry and Remote Sensing 227, 31–47.
- Wu, W., Li, J., Chen, C., Yang, B., Zou, X., Yang, Y., Xu, Y., Zhong, R., Chen, R., 2023. Afli-calib: Robust lidar-imu extrinsic self-calibration based on adaptive frame length lidar odometry. ISPRS Journal of Photogrammetry and Remote Sensing 199, 157–181.
- Xu, N., Qin, R., Song, S., 2023. Point cloud registration for lidar and photogrammetric data: A critical synthesis and performance analysis on classic and deep learning algorithms. ISPRS open journal of photogrammetry and remote sensing 8, 100032.
- Xu, W., Cai, Y., He, D., Lin, J., Zhang, F., 2022. Fast-lid2: Fast direct lidar-inertial odometry. IEEE Transactions on Robotics 38, 2053–2073.
- Yang, Y., Holst, C., 2025. Piecewise-icp: Efficient and robust registration for 4d point clouds in permanent laser scanning. ISPRS Journal of Photogrammetry and Remote Sensing 227, 481–500.
- Zhang, J., Singh, S., 2014. LOAM: Lidar odometry and mapping in real-time, in: Robotics: Science and Systems, Berkeley, CA. pp. 1–9.
- Zhou, Q.Y., Park, J., Koltun, V., 2018. Open3D: A modern library for 3D data processing. arXiv preprint arXiv:1801.09847 .
- Zhou, T., Hasheminasab, S.M., Habib, A., 2021. Tightly-coupled camera/lidar integration for point cloud generation from gnss/ins-assisted uav mapping systems. ISPRS Journal of Photogrammetry and Remote Sensing 180, 336–356.
- Zhu, F., Ren, Y., Zhang, F., 2022. Robust real-time lidar-inertial initialization, in: 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), IEEE. pp. 3948–3955.
- Zhu, J., Li, H., Zhang, T., 2023. Camera, lidar, and imu based multi-sensor fusion slam: A survey. Tsinghua Science and Technology 29, 415–429.