

Extending Ground-Constraint LiDAR-IMU Calibration to Tilted Surfaces in a Continuous-Time Framework

Vassili Korotkine¹, Pierre Chamoun¹, Mohammed Ayman Shalaby², and James Richard Forbes¹

Abstract—This paper presents a novel method that extends targetless LiDAR-IMU calibration for ground vehicles to non-flat environments. Calibration typically necessitates full excitation of the sensor rig, a requirement that is not fulfilled by ground vehicles in normal operation. To address the degenerate planar motion, state-of-the-art methods propose residuals that assume the colinearity of the gravity and physical surface normal vectors, restricting usage to cases where the ground is assumed flat. This paper proposes ground-plane residuals that do not require this assumption, and are applicable for planar motion on a tilted surface. Results are demonstrated on a dataset collected from a Husky ground vehicle, on the M2DGR dataset, as well as on an offroad vehicle dataset. Repeatability is shown to be improved both in tilted and flat-ground scenarios, with strong improvement demonstrated for the tilted case. The implementation and experiments are open-sourced at https://anonymous.4open.science/r/licalib_tilted_ground-9885.

I. INTRODUCTION

Localization from sensor data is a central task for robot operation. Multiple sensors are commonly used to provide redundancy, as well as to complement each other. In particular, Light Detection And Ranging (LiDAR) and Inertial Measurement Unit (IMU) sensors are commonly used together, where the IMU is used for motion prediction as well as to account for LiDAR point cloud distortion [1–3]. However, a prerequisite for fusing information from the two sensors is accurate *sensor calibration*, which provides the ability to resolve data from the different sensors in the same reference frame and relative to the same clock. In particular, the extrinsic pose transformation (rotation and translation) as well as time offset between the two sensors must be known. Inaccurate sensor calibration can cause downstream issues in a robotics pipeline, making reliable and accurate calibration important for robot operation [4].

General calibration methods may be split into target-based and targetless approaches. Target-based approaches, while accurate, necessitate custom targets that may not be readily available during deployment, and can be expensive to manufacture. On the other hand, targetless approaches [5–7] do not require any prior setup for operation, but may be

This work was funded via the Natural Sciences and Engineering Research Council of Canada (NSERC) Alliance Grant in collaboration with Denso Corporation, as well as Rheinmetall Canada.

¹The authors are with the Department of Mechanical Engineering, McGill University, Montreal, Canada (e-mails: vassili.korotkine@mail.mcgill.ca, pierre.chamoun@mail.mcgill.ca, james.richard.forbes@mcgill.ca).

²The author is with Rheinmetall Provectus, Ottawa, ON, Canada (e-mail: mohammed.shalaby@mail.mcgill.ca).

Table I

COMPARISON OF TARGETLESS LiDAR-IMU CALIBRATION METHODS.

Method	Continuous-Time	Ground Constraint	Open Source
LI-Init [6]	No	No	Yes
OA-Calib [7]	Yes	No	Yes
GRIL-Calib [11]	No	Flat only	Yes
Slope-Based [9]	No	No	No
GCT-LiCalib [8]	Yes	Flat only	No
Proposed	Yes	Tilted surfaces	Yes

less accurate due to the lack of geometric information provided by a known target. Recent LiDAR-IMU calibration methods simultaneously solve for the IMU trajectory and calibration parameters using a nonlinear-least-squares solver. The IMU trajectory may be parametrized as a B-Spline [5, 7–9], a Gaussian process [10], or as a sequence of discrete poses [6, 11]. Together, the B-Spline and Gaussian process parametrizations are termed continuous-time methods, while the discrete sequence parametrization is termed a discrete-time method. Existing discrete-time methods [6, 11] introduce discrete robot states at the LiDAR timestamps, and linearly interpolate the IMU data to match the lower-frequency LiDAR stamps. The continuous-time methods are advantageous in that they naturally allow for use of all IMU measurements.

Generally, calibration methods require full motion excitation of the sensor rig to ensure observability of the extrinsic calibration parameters [12]. This is an issue for ground vehicle sensor calibration, since ground vehicles commonly undergo the single-axis rotation motion profile, which causes the position calibration extrinsic to become unobservable along the axis of rotation [9, 13, 14]. The ground vehicle motion degeneracy has prompted research into ground constraints that can be added into the problem to ensure observability.

Existing methods for ground-vehicle LiDAR-IMU calibration, GRIL-Calib [11] and GCT-Calib [8], extract the ground plane from the LiDAR point cloud, yielding a normal vector and distance from the ground. Colinearity of the gravity and surface normal vector is then assumed to create a constraint on the transformation between the LiDAR and the IMU. However, in the case of a tilted surface, the gravity is no longer colinear with the physical surface normal and this assumption is no longer valid, limiting their applicability to flat ground. Furthermore, the ground-constraint residuals of [8, 11] overparametrize the state by introducing ground reference frames that are unique only up to a yaw transformation. While this works in practice, it adds

unnecessary interpretation and implementation complexity. The slope-based method of [9] proposes an observability analysis that shows that the unobservabilities disappear with pitch excitation of the vehicle, and proposes pitch excitation as a solution to restore calibration parameter observability. However, the requirement for pitch excitation does not address the case of strictly planar motion. Table I presents an overview of LiDAR-IMU calibration systems and their relevant characteristics.

In summary, existing targetless LiDAR-IMU calibration systems either require full motion excitation or rely on a flat-ground assumption for introduction of ground constraints. To address these limitations, this work proposes ground-constraint residuals that are independent of the flat ground assumption. A distance residual is formulated by comparing the LiDAR-IMU height difference to the position extrinsic projection onto the surface normal, while an inclination residual is formed as the dot product between the surface normal and the gravity vector. The distance residual directly constrains the position extrinsic along the surface normal, addressing the unobservability caused by planar motion. The inclination residual extends the orientation ground-constraint residual of [11] to the tilted surface case.

To the best of the author's knowledge, this is the first work that addresses the planar motion observability loss while remaining valid for tilted terrain. Furthermore, while the proposed work is not the first continuous-time ground-constraint system, it contributes the first open-source implementation of a continuous-time ground-constraint system for the community to utilize and validate against.

The contributions of this paper thus consist of

- 1) a ground-plane constraint formulation, consisting of a distance and inclination residual, that allows for LiDAR-IMU spatial-temporal calibration during single-axis degenerate motion *without assuming that the ground is flat*,
- 2) a simplified formulation of the flat-ground orientation residual that does not require introduction of extraneous reference frames, and
- 3) an open-source implementation, built within the framework of [7], available at https://anonymous.4open.science/r/licalib_tilted_ground-9885.

The proposed method is compared to GRIL-Calib [11], an open-source ground-constraint-based LiDAR-IMU calibration system. The proposed method is evaluated on self-collected flat-ground and inclined-ground datasets from two different vehicle platforms, as well as on the openly available M2DGR dataset [15]. The self-collected datasets are openly available at [16]. The inclined-ground sequences are used to demonstrate degraded performance for the baseline while the proposed method retains accurate calibration. Furthermore, due to the continuous-time nature of the proposed method compared to the discrete-time baseline, improved performance is also observed on flat ground.

The rest of this paper is organized as follows. Section II reviews the continuous-time LiDAR-IMU calibration method

used as the base system for the proposed method. Section III describes the proposed method including the novel residuals as well as the measured quantities that are required. Section IV presents results of the proposed approach, including an ablation and misinitialization studies. Section V concludes and presents directions for future work.

II. CONTINUOUS-TIME LiDAR-IMU CALIBRATION

This section presents a very review of the continuous-time LiDAR-IMU calibration system OA-Calib [7], used as the base for the proposed method. The choice of OA-Calib as the base method contrasts with GRIL-Calib, which builds upon the discrete-time system of LI-Init [6]. OA-Calib is a mature continuous-time LiDAR-IMU calibration framework that naturally uses all measurements without downsampling or finite-differencing, whose improved repeatability on self-collected datasets is shown in [7]. Since the contribution of the proposed work is in the ground-constraint residuals, the rest of the proposed calibration system is inherited unchanged from [7].

OA-Calib utilizes a B-Spline with N control points used to parametrize the IMU trajectory. Each control point is parametrized by the pose transformation [17]

$$\mathbf{T}_{ab_i} = \begin{bmatrix} \mathbf{C}_{ab_i} & \mathbf{r}_a^{b_i a} \\ \mathbf{0} & 1 \end{bmatrix}, \quad (1)$$

where $\mathcal{F}_{b_i}$ denotes the reference frame of the i 'th control point, and b_i denotes its physical position. The world frame $\mathcal{F}_a$ and the world reference position a are set to be the IMU reference frame and reference position, respectively, at the time that the first IMU measurement is received. Formally, the world reference frame $\mathcal{F}_a = \mathcal{F}_{b_0}$, and the world physical reference point is $a = b_0$. The control points of the B-Spline are thus given by $\mathcal{X}_{\text{spline}} = (\mathbf{C}_{ab_0}, \mathbf{r}_a^{b_0 a}, \dots, \mathbf{C}_{ab_i}, \mathbf{r}_a^{b_i a}, \dots, \mathbf{C}_{ab_N}, \mathbf{r}_a^{b_N a}; t_i)$, where the spline time knots t_i are *fixed* throughout the optimization. The full optimization state is given by

$$\mathcal{X}_{\text{full}} = (\mathcal{X}_{\text{spline}}, \mathcal{X}_{\text{intr}}, \mathbf{b}_{\text{IMU}}, \mathbf{g}_a, \mathbf{C}_{bl}, \mathbf{r}_b^{lb}, \tau_{bl}), \quad (2)$$

where $\mathcal{X}_{\text{intr}}$ denotes intrinsic IMU and LiDAR parameters, $\mathbf{b}_{\text{IMU}}$ denotes the IMU accelerometer and gyroscope biases, assumed constant over the calibration time interval, $\mathbf{g}_a$ is the gravity vector resolved in the world frame, and $\mathbf{C}_{bl}, \mathbf{r}_b^{lb}, \tau_{bl}$ are the spatiotemporal LiDAR-IMU extrinsics. The states and constant-bias assumptions are inherited as-is from the OA-Calib system used as the base for the proposed approach.

Angular velocity and linear acceleration residuals are obtained by comparing IMU measurements to differentiated quantities obtained from the spline. Furthermore, a LiDAR residual is constructed using a point-to-plane error. The planes are extracted from the map constructed using the LiDAR. The points are computed from each LiDAR point measurement that is transformed into the LiDAR map frame using the LiDAR-IMU extrinsics and the B-Spline parametrizing the IMU trajectory. The reader is referred to [7] for details on the continuous-time calibration aspect, since it is left unchanged in the proposed ground vehicle LiDAR-IMU calibration approach.

Table II
SUMMARY OF RESIDUALS

Residual	Type	Tilted Ground	Uses IMU Inclination
$\mathbf{e}_{\text{flat}}$	Orientation	No	No
$\mathbf{e}_{\text{dist}}$	Distance	Yes	No
$\mathbf{e}_{\text{incl}}$	Orientation	Yes	Yes

III. PROPOSED METHOD

The proposed method consists of extracting ground plane information from the LiDAR point cloud and forming ground plane residuals that involve calibration parameters. These ground plane residuals are added into the OA-Calib pipeline. The ground residuals are computed assuming a given time instant. Therefore, the time subscripts are omitted for clarity in this section. Furthermore, the pose quantities $\mathbf{C}_{ab}, \mathbf{r}_a^{ba}$ at a given time instant are computed as a function of the spline parameters in the state (2).

Three residuals are described and summarized in Table II. An orientation residual $\mathbf{e}_{\text{flat}}$ assuming a flat surface is first proposed, which assumes colinearity of gravity and surface normal, in the manner of GRIL-Calib. Then two residuals are proposed that do not depend on the flat ground assumption. A distance residual $\mathbf{e}_{\text{dist}}$ is proposed to constrain the extrinsic translation along the axis of rotation, and an orientation residual $\mathbf{e}_{\text{incl}}$ is proposed that uses the ground inclination angle. It should be noted that the distance residual $\mathbf{e}_{\text{dist}}$ by itself resolves the unobservability of the position extrinsic caused by the ground vehicle planar motion. The inclination residual $\mathbf{e}_{\text{incl}}$ provides an additional orientation constraint, and is included as it is a direct extension to the tilted case of the orientation ground constraint in previous work [8, 11]. Ablation studies are presented in Section IV that disambiguate the relative contribution of each residual. Both $\mathbf{e}_{\text{dist}}$ and $\mathbf{e}_{\text{incl}}$ are novel residuals that work on tilted ground. The flat ground orientation residual $\mathbf{e}_{\text{flat}}$ carries the same information as the orientation ground constraint as [8, 11], but avoids the creation of extraneous ground frames.

Each constraint in the following sections may be rewritten as a nonlinear error to be used in a nonlinear-least-squares solver. The nonlinear errors are written in the form $\mathbf{e}(\mathcal{X}_{\text{dep}}; \mathcal{Y}_{\text{meas}})$, where $\mathcal{X}_{\text{dep}}$ denotes the state subset that this error depends on, and $\mathcal{Y}_{\text{meas}}$ denotes the measurements used to construct this error. Throughout this section, $\tilde{\mathbf{g}} = -\mathbf{g} / \|\mathbf{g}\|$ denotes the unit vector pointing opposite to gravity. The normal vector $\underline{\mathbf{n}}$ is used to denote the unit normal vector directed upward from the ground.

A. Required Quantities and Assumptions

In the same manner as previous ground-constraint methods [8, 11], knowledge of the IMU height d_b is required, corresponding to the shortest distance from the IMU to the local ground plane. This quantity is measured physically beforehand, as it cannot be refined on strictly flat ground. On strictly planar surfaces, the accuracy of d_b directly affects the accuracy of the position extrinsic along the axis of rotation,

and has to be known as accurately as possible. On the LiDAR side, the ground plane is extracted from the LiDAR point cloud using Patchwork++ [18] at each time instant yielding the surface normal resolved in the LiDAR frame $\mathbf{n}_l$ as well as the distance of the LiDAR above the ground d_l . Therefore, the ground must be reasonably planar for the ground plane extraction to function. Robustness is provided through the use of robust loss functions in the calibration optimization to reject outliers.

Furthermore, the proposed orientation residual $\mathbf{e}_{\text{incl}}$ on tilted ground requires knowledge of the ground inclination $\cos(\phi) = \tilde{\mathbf{g}} \cdot \underline{\mathbf{n}}$. The ground inclination may be obtained from a separate sensor or estimator that provides pitch and roll. A method for doing so using knowledge of the IMU-resolved gravity on flat ground is provided.

B. Orientation Residual Assuming Flat Surface

This subsection describes a flat-ground residual that carries the same information as the orientation residual of [11]. However, it avoids the introduction of extraneous ground reference frames that are used in [11], simplifying practical interpretation and implementation.

If the surface being driven on is flat, the surface normal is colinear with the gravity vector,

$$\underline{\mathbf{n}} = \tilde{\mathbf{g}}, \quad (3)$$

Resolving (3) in $\mathcal{F}_l$ yields

$$\mathbf{n}_l = \mathbf{C}_{bl}^T \tilde{\mathbf{g}}_b, \quad (4)$$

where $\tilde{\mathbf{g}}_b = \mathbf{C}_{ab} \tilde{\mathbf{g}}_a$ depends on the IMU rotation computed from the B-Spline $\mathbf{C}_{ab}$ and the world frame gravity vector $\mathbf{g}_a$. The corresponding error vector is thus

$$\mathbf{e}_{\text{flat}}(\mathbf{C}_{ab}, \mathbf{g}_a, \mathbf{C}_{bl}; \mathbf{n}_l) = \mathbf{n}_l - \mathbf{C}_{bl}^T \tilde{\mathbf{g}}_b \quad (5)$$

$$= \mathbf{n}_l + \mathbf{C}_{bl}^T \frac{\mathbf{g}_b}{\|\mathbf{g}_b\|_2}. \quad (6)$$

C. Distance Residual

The distance residual is constructed by considering the ground plane expressed using each sensor. It does not require the ground to be flat, and constrains the position extrinsic along the axis of rotation of the vehicle. Given the IMU height d_b and the LiDAR height d_l , their difference is equal to the position extrinsic projected onto the surface normal,

$$\underline{\mathbf{n}} \cdot \underline{\mathbf{r}}^{lb} = d_l - d_b, \quad (7)$$

and resolving both sides of (7) in $\mathcal{F}_l$ yields

$$\mathbf{r}_l^{lbT} \mathbf{n}_l = d_l - d_b, \quad (8)$$

which constrains the position extrinsic along the direction of the normal to the surface. This constraint is fully independent of the flat ground assumption. The choice to resolve in $\mathcal{F}_l$ makes this residual independent of the flat ground assumption, since the surface normal can be extracted from ground segmentation of the LiDAR point cloud, in the LiDAR frame regardless of inclination. The corresponding error is written

$$\mathbf{e}_{\text{dist}}(\mathbf{C}_{bl}, \mathbf{r}_b^{lb}; d_b, d_l, \mathbf{n}_l) = d_l - d_b - (\mathbf{C}_{bl}^T \mathbf{r}_b^{lb})^T \mathbf{n}_l. \quad (9)$$

The distance residual constrains the position extrinsic along the surface normal. For typical operating conditions, the surface normal corresponds to the axis of rotation of the vehicle. Under this assumption, this directly corresponds to the part of the position extrinsic that is not observable in the single-axis rotation motion typically experienced by ground vehicles, resolving the position extrinsic unobservability.

D. Inclination Residual

This section proposes an extension of (4) to the case where the surface is tilted and the gravity vector is no longer colinear with the surface normal. For a given sequence, the inclination residual *or* the flat-ground residual of Section III-B should be used, depending on whether the ground is assumed flat. For flat ground, both residuals are valid. However, on tilted ground, the inclination residual remains valid while the flat-ground residual of Section III-B does not, and they will disagree with each other.

The proposed inclination residual requires knowledge of the surface inclination, $\cos(\phi)$. This quantity has to be provided through a separate sensor or estimation method. In the present work, it is computed by assumption of knowledge of the gravity vector resolved in the IMU frame on flat ground, denoted $\mathbf{g}_f$, together with its normalized version $\tilde{\mathbf{g}}_f = -\mathbf{g}_f / \|\mathbf{g}_f\|_2$. The knowledge of $\mathbf{g}_f$ is a strong assumption. In practice, this method of computing the inclination requires a single initial flat ground calibration sequence. The resulting $\mathbf{g}_f$ can then be reused for subsequent calibration sequences collected on tilted ground, under the assumption that the IMU is rigidly mounted on the vehicle and has not shifted. In the proposed experiments, it is computed by first running the algorithm on a known flat-ground sequence, and extracting the optimized $\mathbf{g}_a$. With knowledge of $\mathbf{g}_f$, the inclination is computed as

$$\cos(\phi) = \tilde{\mathbf{g}}_b^T \tilde{\mathbf{g}}_f. \quad (10)$$

With $\cos(\phi)$ is provided, the inclination residual can be written as

$$\cos(\phi) = \tilde{\mathbf{g}}_{\rightarrow} \cdot \mathbf{n}_{\rightarrow}. \quad (11)$$

Resolving in $\mathcal{F}_l$ yields

$$\cos(\phi) = \tilde{\mathbf{g}}_b^T \mathbf{C}_{lb}^T \mathbf{n}_l. \quad (12)$$

The corresponding error term is

$$\mathbf{e}_{\text{ori}} = \cos(\phi) + \frac{\mathbf{g}_a^T \mathbf{C}_{ab}^T \mathbf{C}_{bl} \mathbf{n}_l}{\|\mathbf{g}_a\|_2}. \quad (13)$$

The inclination residual is written by substituting the inclination angle (10) into the orientation residual (13), which yields

$$\mathbf{e}_{\text{incl}}(\mathbf{C}_{ab}, \mathbf{g}_a, \mathbf{C}_{bl}; \mathbf{g}_f, \mathbf{n}_l) = \frac{\mathbf{g}_a^T \mathbf{C}_{ab}^T}{\|\mathbf{g}_a\|_2} (\tilde{\mathbf{g}}_f - \mathbf{C}_{bl} \mathbf{n}_l). \quad (14)$$




Fig. 1. Problem setup: LiDAR-IMU calibration on a tilted planar surface.

IV. RESULTS

The proposed algorithm was compared to GRIL-Calib [11] on a number of datasets. While the closest work is the continuous-time GCT-Calib [8], their algorithm source code is not openly available for benchmarking, and thus cannot be reasonably reproduced. The proposed distance residual requires IMU height, in the same manner as GRIL-Calib. The inclination residual replaces the flat-ground assumption with the requirement of a surface inclination measurement. The IMU height was measured physically. The IMU-resolved flat-ground gravity vector was extracted by taking a known flat-ground sequence, running the calibration algorithm with the flat-ground assumption, and using the resulting gravity vector. Furthermore, both the proposed method and GRIL-Calib require the distance to the ground from the LiDAR, as well as the normal to the ground resolved in the LiDAR frame. The Patchwork++ ground segmentation algorithm [18] was used to segment the ground and to extract the distance and surface normal, in the same manner as GRIL-Calib.

The measured extrinsics, used as an initial guess for the calibration algorithms, are presented in Table III. This is a key difference from the result evaluation in [11], that used heavily perturbed measured extrinsics as an initial guess. This is proposed as a more realistic reflection of real-world operation, where rough sensor placement is known from engineering drawings, and automated procedures are used for refinement.

A. Metrics

Physical measurements of extrinsic calibration parameters are used for initializing the calibration procedures, but are of insufficient accuracy to be used as ground truth to evaluate against. Instead, repeatability is used as a metric for calibration quality. Multiple sequences are collected from the same sensor rig, and the spread of results across these sequences is used to assess quality of the calibration algorithm. This spread should nevertheless be interpreted with caution, as the amount of sequences used is relatively low, ranging between three and seven for each dataset. Nevertheless, it is proposed that this metric is closest to what would be used by an engineer “in the wild” to evaluate the quality of the calibration. The measured extrinsic parameters are provided in Table III. For the rotation extrinsic, the mean and covariance are computed on the $SO(3)$ group, and converted to Euler angles and degrees for readability in the results tables. Formally, given a sequence of results

Table III
ESTIMATED EXTRINSICS ACROSS ALL DATASETS

Method	$\mathbf{r}_b^{lb}$ (cm)	$\mathbf{C}_{bl}$ (Euler, deg)
Husky UGV, Flat Ground		
Measured	[24.77, -6.99, -46.36]	[180.00, 0.00, -90.00]
GRIL-Calib	[25.17, -3.38, -43.55]	[179.98, -0.21, -90.11]
Proposed	[26.00, -7.20, -43.53]	[179.98, -0.17, -87.48]
Husky UGV, Incline (Note: GRIL-Calib Diverges.)		
Measured	[24.77, -6.99, -46.36]	[180.00, 0.00, -90.00]
GRIL-Calib	[31.29, 17.95, -46.43]	[182.17, -1.54, 199.22]
Proposed	[25.76, -8.26, -43.19]	[179.93, -0.15, -87.07]
M2DGR		
Measured	[27.26, 0.00, 17.95]	[0.00, 0.00, 0.00]
GRIL-Calib	[25.61, 6.71, 17.81]	[-0.37, -0.08, 0.54]
Proposed	[25.00, 0.52, 17.36]	[-0.47, 0.98, -0.77]
Offroad		
Measured	[-63.40, 132.60, 34.05]	[0.06, 0.00, -0.26]
GRIL-Calib	[-61.14, 133.61, 34.83]	[0.15, 0.43, -0.68]
Proposed	[-59.82, 132.72, 33.78]	[0.18, 0.70, -1.25]

$(\mathbf{C}_1, \dots, \mathbf{C}_i, \dots, \mathbf{C}_{n_{\text{res}}})$, the mean rotation in the Riemannian sense is [19]

$$\bar{\mathbf{C}} = \arg \min_{\mathbf{C} \in SO(3)} \sum_{i=1}^{n_{\text{res}}} \|\log(\mathbf{C}^T \mathbf{C}_i)\|_F^2. \quad (15)$$

Denoting the differences in the Lie algebra as $\delta \xi_i = \text{Log}(\bar{\mathbf{C}}^T \mathbf{C}_i)$, it can be shown [19] that $\bar{\mathbf{C}}$ must satisfy

$$\sum_{i=1}^{n_{\text{res}}} \delta \xi_i = 0. \quad (16)$$

As such, the mean rotation $\bar{\mathbf{C}}$ is the rotation matrix that ensures the mean difference in the Lie algebra, $\delta \bar{\xi} = \frac{1}{n_{\text{res}}} \sum_{i=1}^{n_{\text{res}}} \delta \xi_i$, is zero. An algorithm for computing $\bar{\mathbf{C}}$ is provided in [20]. The empirical covariance matrix is computed as

$$\Sigma = \frac{1}{n_{\text{res}} - 1} \sum_{i=1}^{n_{\text{res}}} \delta \xi_i \delta \xi_i^T, \quad (17)$$

with the per-component standard deviation given by

$$\sigma = [\sqrt{\Sigma_{11}} \quad \sqrt{\Sigma_{22}} \quad \sqrt{\Sigma_{33}}]. \quad (18)$$

In subsequent tables, the mean rotation entry corresponds to the 1 – 2 – 3 Euler angle decomposition of $\bar{\mathbf{C}}$ in (15), while the standard deviations used to quantify result spread correspond to σ in (18). Both the mean rotation Euler angles and the standard deviations are normalized to units of degrees. A lower standard deviation demonstrates improved repeatability across calibration sequences, which is taken to be the performance metric of the tested calibration methods. While the standard deviations are considered the metric for performance, the mean calibration is also reported as a sanity check to compare to the previously measured extrinsics for each dataset.

The mean calibration results for each dataset and proposed method are consolidated in Table III and discussed below.

B. Husky Clearpath Ground Vehicle

A Clearpath Husky unmanned ground vehicle equipped with a Velodyne VLP-16 LiDAR and a 3DM-GX5-AHRS attitude-heading-reference-system for IMU data, pictured in




Fig. 2. Husky UGV used for algorithm validation.




Fig. 3. Per-sequence position extrinsic results for the flat-ground Husky UGV dataset.

Figure 2, was driven on flat and inclined ground to validate the proposed algorithm. Three sequences are collected indoors on flat ground in a lab environment to demonstrate nominal flat-ground performance for both methods. Seven other sequences are collected on inclined ground in an outdoors urban environment on a university campus to demonstrate performance improvement on inclined ground. The previously measured extrinsics are obtained from a visual inspection for the rotation extrinsic, and a tape measure for the position extrinsic.

The calibration results for the flat-ground sequences are presented in Figure 3. The extrinsic spread is plotted in Figure 4. The computed mean extrinsics are presented in Table III. Both GRIL-Calib and the proposed method converge to values reasonably close to the measured extrinsics from Table III. The proposed method achieves increased repeatability across all extrinsic parameters. Since the sequences are collected on flat ground, the performance improvement is attributed to the continuous-time nature of the overall system. The performance increase is especially pronounced in the yaw orientation and the $x - y$ position extrinsics. This is




Fig. 4. Extrinsic spread for the flat-ground Husky UGV dataset.

Proposed		
	$\mathbf{r}_b^{tb}$ (cm)	$\mathbf{C}_{bl}$ (Euler, deg)
Outdoors1	[25.600, -7.700, -46.600]	[179.965, -0.179, -88.440]
Outdoors2	[25.400, -8.200, -41.000]	[179.969, -0.029, -87.740]
Outdoors3	[29.300, -5.400, -46.700]	[179.638, -0.306, -87.790]
Outdoors4	[23.700, -14.100, -42.400]	[-179.983, -0.103, -82.450]
Outdoors5	[26.600, -8.100, -43.600]	[180.000, -0.150, -86.010]
Outdoors6	[24.500, -8.300, -41.800]	[179.947, -0.179, -89.120]
Outdoors7	[24.700, -6.100, -40.200]	[-179.996, -0.160, -87.700]
GRIL-Calib		
	$\mathbf{r}_b^{tb}$ (cm)	$\mathbf{C}_{bl}$ (Euler, deg)
Outdoors1	[49.304, -9.349, -46.908]	[-179.746, 3.644, 77.519]
Outdoors2	[-6.948, 16.054, -42.245]	[179.561, -3.068, -69.199]
Outdoors3	[29.370, 51.323, -51.922]	[-174.973, -0.909, 155.730]
Outdoors4	[-13.562, 50.277, -47.859]	[-175.170, 0.754, 136.414]
Outdoors5	[99.302, 8.624, -47.950]	[177.023, -1.507, -173.926]
Outdoors6	[34.369, 2.119, -43.992]	[-179.190, -0.707, -89.989]
Outdoors7	[27.174, 6.580, -44.160]	[178.368, -1.326, -81.869]

Table IV
RESULTS FOR EACH SEQUENCE OF INCLINED HUSKY UGV DATASET




Fig. 5. Per-sequence position extrinsic results for the M2DGR dataset.

likely due to the planar motion excitation in the vehicle. The sensors are mounted flat on the vehicle, with yaw and $x-y$ directions corresponding to the vehicle motion directions.

The results for the inclined ground sequences are presented in Table IV, and are the primary evidence for the performance improvement of the proposed method on inclined ground. Direct comparison to the measured extrinsic of Table III shows that the GRIL-Calib method yields an incorrect answer, while the proposed method still converges to values that are reasonably close both to the measured extrinsics of Table III and the mean flat-ground calibration results of Figure III.

C. M2DGR

The M2DGR dataset [15], equipped with a Velodyne VLP-32C and Handsfree A9 IMU, collected on flat ground, is used to demonstrate the advantages of the continuous-time nature of the proposed method. The sequences used and their respective time segments are chosen to directly correspond to the ones used in the analysis of GRIL-Calib [11]. The previously measured extrinsics are extracted from the calibration provided by the dataset paper [15].

The per-sequence position extrinsics, the spread in the extrinsics, and the mean computed extrinsics are presented in Figure 5, Figure 6, and Table III, respectively.

The mean extrinsics of both methods in Table III are comparable to the measured extrinsics of Table III. The $x-y$ position extrinsic repeatability for the proposed is significantly improved compared to GRIL-Calib in Figure 6,




Fig. 6. Extrinsic spread for the M2DGR dataset.




Fig. 7. Per-sequence position extrinsic results for the Offroad dataset.

mirroring the results obtained on the Husky UGV flat-ground sequences. The spread change is mostly driven by the street_08 and rotation_02 sequences, for the x and y position extrinsic respectively as seen from Figure 5. The yaw extrinsic repeatability is comparable; however it is slightly degraded compared to the GRIL-Calib baseline. Both methods yield about one degree of standard deviation in the yaw calibration extrinsic. This differs from the Husky UGV sequences, where the proposed method had a standard deviation of a tenth of a degree. It is also worth noting that the GRIL-Calib spread in Figure 6 is comparable in magnitude compared to the RMSE values presented in [11].

D. Offroad Vehicle Dataset

An offroad vehicle equipped with an Ouster OS1-32 LiDAR and a XSSENS MTI-200-2A8G4 IMU was driven in various environments, both indoors and outdoors. The mean extrinsics of the two methods, presented in Table III are very close to each other. The spread in the obtained extrinsics is improved by the proposed method across the board, as shown in Figure 8, except for the y -component of the orientation error that is slightly worsened. Per-sequence position extrinsic results are presented in Figure 7, where the dashed line shows the measured extrinsic used to initialize the calibration. The insideGarage sequence shows deviation with respect to the other sequences. It is also worth noting that the z -component of the position extrinsic in Figure 7 is systematically higher for GRIL-Calib than for the proposed method. Both methods use the same Patchwork++ code [18] with the same configuration parameters, which means that the bias is due to some other convergence properties of the method. A top-down visualization of the offroad trajectories is provided in Figure 9, where the trajectory starting position is normalized to zero for all sequences.




Fig. 8. Extrinsic spread for the Offroad dataset.




Fig. 9. Offroad dataset trajectory visualization.

E. Ablation Study on Offroad Dataset

An ablation study was run on the offroad dataset in order to analyze the relative contribution of each residual to the calibration solution. The results are presented in Figure 10 and Figure 11 for the extrinsic parameter spread and per-sequence position extrinsic values, respectively. The compared methods correspond to combinations of the flat-ground (6) and inclined orientation residuals (13), and the distance residual (9). The flat-ground residuals yield significantly higher spread in the position extrinsic, which seems to be significantly driven by the `fields` sequence of Figure 11. The baseline of OA-Calib has the lowest spread among the proposed methods. This is conjectured to be caused by the initialization of the method, where the method is initialized with the pre-existing measured extrinsics, and does not drift much from that measured extrinsic. This also shows the limitations of the repeatability metric used in this work. OA-Calib cannot by itself refine the axis-of-rotation position extrinsic, as it is not observable without ground constraints. However, the spread in Figure 10 shows that any spurious updates present in the nonlinear-least-squares solver are quite low in magnitude. The z -extrinsic remains essentially unchanged from the original value of 34.05 cm for all sequences.

F. Offroad With Misinitialized Extrinsics

The results of Section IV-E show that the system without any additional constraints performs better in terms of repeatability, which shows the limits of repeatability as a metric. The updates to an unobservable state direction in calibration are due to spurious updates from the solver, and Section IV-E showed that the z -position extrinsic, which corresponds to an unobservable direction for this calibration problem,




Fig. 10. Extrinsic spread for the Offroad (ablation) dataset.




Fig. 11. Per-sequence position extrinsic results for the Offroad (ablation) dataset.

remains relatively close to its initial value. This section shows that for misinitialized extrinsics, the calibration algorithm without a ground constraint does not update the unobservable direction. To study the effect of misinitialization, the offroad dataset was run with a misinitialized position extrinsic, set to $\mathbf{r}_b^{lb} = [-0.6 \ 1.0 \ 0.6]$ meters. Figure 12 shows the spread in the resultant calibration extrinsics. The orientation extrinsic remains unchanged. However, the position extrinsic demonstrates a far larger spread for cases where the distance residual is used, reflecting the convergence behaviour of an algorithm initialized far from the true value. While the z -position extrinsic has the smallest spread for the “Neither” and “Only Inclination” cases, where the distance residual is not used, Figure 13 shows that this value is heavily biased. Figure 13 plots the position errors with respect to the measured values in Table III. The error of around 25 centimeters in the z -extrinsic corresponds to the z -extrinsic not being refined from its initial value in the case of residuals that do not use the distance residual. Furthermore, the `insideGarage` sequence shows a large error even for the case of the distance residual. This is caused by the system setup. The distance residual is wrapped inside of a robust Cauchy loss [17]. Depending on how far the system is initialized and what shape parameter is used for the robust loss, the robust loss may cause entire rejection of the distance residual update. The exact parameter choice depends on the problem setup. A lenient robust loss may cause updates from spurious ground measurements, while a harsh robust loss may cause the system to favour misinitialized values.




Fig. 12. Extrinsic spread for ablation study on offroad dataset with misinitialized extrinsics.




Fig. 13. Per-sequence position error for ablation study on offroad dataset with misinitialized extrinsics.

V. CONCLUSION

This paper presents a continuous-time LiDAR-IMU targetless calibration method for planar motion on non-flat ground. Compared to the state-of-the-art flat ground calibration system, the proposed method demonstrates robustness to tilted ground. Improved performance is also demonstrated on flat ground due to the continuous-time nature of the proposed method. However, the method has limitations. The IMU height still has to be known for the distance residual, and the IMU-resolved gravity vector on flat ground has to be known to compute the inclination of the vehicle. Future work includes treating the IMU height as a state parameter instead of a measured, fixed quantity as well as different methods to compute inclination. Furthermore, the lack of precise ground truth, the difficulty of measuring impact of calibration on estimation accuracy, as well as sensitivity of the calibration algorithm to initialization makes evaluating the quality of calibration algorithms, and the resulting calibration, a non-trivial task. While this work advocates for repeatability as a primary metric, the misinitialization study also shows that previously measured values are necessary as a sanity check. This shows that robust metrics of calibration quality, whether based on error, repeatability, or downstream estimator performance, are an important direction of future research.

VI. ACKNOWLEDGEMENTS

The authors would like to acknowledge Nicholas Dahdah for help with the dataset collection and Mitchell Cohen for valuable discussions.

REFERENCES

- [1] W. Xu and F. Zhang, "Fast-lio: A fast, robust lidar-inertial odometry package by tightly-coupled iterated kalman filter," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3317–3324, 2021.
- [2] W. Xu, Y. Cai, D. He, J. Lin, and F. Zhang, "Fast-lio2: Fast direct lidar-inertial odometry," *IEEE Transactions on Robotics*, vol. 38, no. 4, pp. 2053–2073, 2022.
- [3] I. Vizzo, T. Guadagnino, B. Mersch, L. Wiesmann, J. Behley, and C. Stachniss, "Kiss-icp: In defense of point-to-point icp – simple, accurate, and robust registration if done the right way," *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 1029–1036, 2023.
- [4] S. Li, X. Li, Y. Zhou, and C. Xia, "Targetless extrinsic calibration of lidar-imu system using raw gnss observations for vehicle applications," *IEEE Transactions on Instrumentation and Measurement*, vol. 72, pp. 1–11, 2023.
- [5] J. Lv, J. Xu, K. Hu, Y. Liu, and X. Zuo, "Targetless calibration of lidar-imu system based on continuous-time batch estimation," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 9968–9975.
- [6] F. Zhu, Y. Ren, and F. Zhang, "Robust real-time lidar-inertial initialization," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 3948–3955.
- [7] J. Lv, X. Zuo, K. Hu, J. Xu, G. Huang, and Y. Liu, "Observability-aware intrinsic and extrinsic calibration of lidar-imu systems," *IEEE Transactions on Robotics*, vol. 38, no. 6, pp. 3734–3753, 2022.
- [8] T. He, P. Wang, D. Das, and B. T. Maharaj, "Gct-licolib: Targetless online lidar-imu calibration via ground plane motion constraints and continuous-time optimization," in *2025 5th International Conference on Communication Technology and Information Technology (ICC-TIT)*, 2025, pp. 238–243.
- [9] J. Xiao, J. Guo, C. Ouyang, J. Shi, Y. Xu, and M. Wu, "A slope-based targetless extrinsic calibration method for lidar-imu systems on ground vehicles," *IEEE Transactions on Instrumentation and Measurement*, vol. 74, pp. 1–14, 2025.
- [10] S. Li, L. Wang, J. Li, B. Tian, L. Chen, and G. Li, "3d lidar/imu calibration based on continuous-time trajectory estimation in structured environments," *IEEE Access*, vol. 9, pp. 138 803–138 816, 2021.
- [11] T. Kim, G. Pak, and E. Kim, "Gril-calib: Targetless ground robot imu-lidar extrinsic calibration method using ground plane motion constraints," *IEEE Robotics and Automation Letters*, vol. 9, no. 6, pp. 5409–5416, 2024.
- [12] S. Chen, X. Li, S. Li, Y. Zhou, and X. Yang, "Ikalibr: Unified targetless spatiotemporal calibration for resilient integrated inertial systems," *IEEE Transactions on Robotics*, vol. 41, pp. 1618–1638, 2025.
- [13] W. Lee, C. Chen, and G. Huang, "Degenerate motions of multisensor fusion-based navigation," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 8113–8119.
- [14] Y. Yang, P. Geneva, X. Zuo, and G. Huang, "Online self-calibration for visual-inertial navigation: Models, analysis, and degeneracy," *IEEE Transactions on Robotics*, vol. 39, no. 5, pp. 3479–3498, 2023.
- [15] J. Yin, A. Li, T. Li, W. Yu, and D. Zou, "M2dgr: A multi-sensor and multi-scenario slam dataset for ground robots," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2266–2273, 2022.
- [16] A. Anonymous, *LiDAR-IMU Data for Sensor Calibration on Tilted Ground*, 2026. [Online]. Available: <https://doi.org/10.5281/zenodo.22098085>
- [17] T. D. Barfoot, *State estimation for robotics*. Cambridge University Press, 2024.
- [18] S. Lee, H. Lim, and H. Myung, "Patchwork++: Fast and robust ground segmentation solving partial under-segmentation using 3d point cloud," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 13 276–13 283.
- [19] M. Moakher, "Means and averaging in the group of rotations," *SIAM Journal on Matrix Analysis and Applications*, vol. 24, no. 1, pp. 1–16, 2002.
- [20] J. Manton, "A globally convergent numerical algorithm for computing the centre of mass on compact lie groups," in *ICARCV 2004 8th Control, Automation, Robotics and Vision Conference, 2004.*, vol. 3, 2004, 2211–2216 Vol. 3.