

A Robust Approach for LiDAR-Inertial Odometry Without Sensor-Specific Modeling

Meher V. R. Malladi

Tiziano Guadagnino

Luca Lobefaro

Cyrill Stachniss




Fig. 1: Our robust LiDAR-inertial odometry system is directly operational in different environments, sensor configurations, and robotic platforms with distinct motion behaviours, all without any change in configuration or modeling approach. We depict the local map result of our odometry system in four distinct scenarios, shown clockwise from the top left: urban city with Ouster OS1-128 and built-in InvenSense IMU mounted on a car; mixed indoor–outdoor university buildings with Hesai QT64 and Alphasense IMU on a backpack (data from Tao et al. [25]); forest with Hesai XT32 and Xsens MTi-100 IMU mounted on a tree-harvesting machine; and parking lot with Velodyne VLP-16 and onboard IMU on a Unitree Go1 quadruped (data from Ou et al. [19]).

Abstract—Accurate odometry is a critical component in a robotic navigation stack, and subsequent modules such as planning and control often rely on an estimate of the robot’s motion. LiDAR-based odometry approaches should be robust across sensor types and deployable in different target domains, from solid-state LiDARs mounted on cars in urban-driving scenarios to spinning LiDARs on handheld packages used in unstructured natural environments. In this paper, we propose a robust LiDAR-inertial odometry system that does not rely on sensor-specific modeling. Sensor fusion techniques for LiDAR and inertial measurement unit (IMU) data typically integrate IMU data iteratively in a Kalman filter or use pre-integration in a factor graph framework, combined with LiDAR scan matching often exploiting some form of feature extraction. We propose an alternative strategy that only requires a simplified motion model for IMU integration and directly registers LiDAR scans in a scan-to-map approach. Our approach allows us to impose a novel regularization on the LiDAR registration, improving the overall odometry performance. We provide extensive experiments on different datasets covering a wide array of commonly used robotic sensors and platforms. We show that our approach works with the exact same configuration in all these scenarios, demonstrating its robustness. We have open-sourced our implementation so that the community can build further on our work and use it in their navigation stacks.

Index Terms—Mapping, SLAM

I. INTRODUCTION

ACCURATE ego-motion estimation is a key component of all autonomous mobile robot navigation stacks. A lot of robotic systems rely on range sensors such as LiDARs to some degree as they provide high-accuracy perception for robust mapping, localization and collision avoidance. Modern 3D LiDARs often also have an integrated IMU which provides complementary proprioceptive measurements that pair well with the exteroceptive sensors such as LiDARs. Sensors and especially IMUs vary significantly in quality and cost: high-end IMUs maintain accuracy over long trajectories, whereas

Manuscript received: Feb 18, 2026; Accepted: Apr 8, 2026. This paper was recommended for publication by Editor Javier Civera upon evaluation of the Associate Editor and Reviewers’ comments.

This work has partially been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy, EXC-2070 – 390732324 – PhenoRob, by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under STA 1051/5-1 within the FOR 5351 (AID4Crops), by the European Union’s Horizon Europe research and innovation programme under grant agreement No 101070405 (DigiForest), and by the German Federal Ministry of Research, Technology and Space (BMFTR) under the Robotics Institute Germany (RIG). All authors are with the University of Bonn, Center for Robotics. Cyrill Stachniss is additionally with the Lamarr Institute for Machine Learning and Artificial Intelligence.

Digital Object Identifier (DOI): see top of this page.

inexpensive units exhibit substantial noise and drift when integrated over even a few meters. Similarly, LiDAR sensors vary in design and may use standard rotating-beam scanning or atypical scanning patterns [9], [13].

In this paper, we consider the use of a 3D LiDAR-IMU pair for state estimation and the problem of developing a LiDAR-inertial odometry capable of robust performance across diverse sensor types, robotic platforms, and operational environments without relying on sensor-specific modeling. Fusing LiDAR and IMU sensor measurements for odometry or simultaneous localization and mapping (SLAM) is commonly addressed using filter-based [30] or factor graph-based methods [7]. Both approaches require detailed sensor noise modeling [5], which can demand expensive calibration procedures and may require periodic repetition. Some approaches further exploit sensor-specific characteristics such as intensity returns [20] or model the LiDAR’s scanning pattern [13]. In particular, robust modeling of IMU noise and biases is challenging due to various sources of electromechanical errors, including temperature variation, mechanical impacts, aging, and more [15], which can be difficult to characterize accurately.

The main contribution of this paper is a robust and accurate 3D LiDAR-inertial odometry pipeline that requires no sensor-specific modeling or tuning and can be deployed across diverse scenarios using the exact same parameters (see Fig. 1). In our approach, we simplify IMU modeling by taking an alternative to either filtering or pre-integration. This approach reduces the requirements of detailed noise and bias characterization. Using this model, we propose a novel adaptive regularization scheme in the ICP-based 3D LiDAR registration step, which further improves odometry accuracy. We demonstrate that our system achieves odometry results on par with, or better than, state-of-the-art methods. Finally, we open-source our implementation at https://github.com/PRBonn/rko_lio and further supplementary material is available at https://prbonn.github.io/rko_lio/suppl.html. We provide an easy-to-install package via common system package managers such as pip and apt to foster community use and further research.

In sum, we make three key claims: our approach is (i) able to perform on par or better than state-of-the-art LiDAR-inertial odometry systems on pose tracking accuracy, (ii) robust and applicable across different sensor configurations, robot platforms, and operational environments, and (iii) usable in these diverse scenarios without any changes to configuration or underlying models. These claims are backed up by the paper and our experimental evaluation.

II. RELATED WORK

The quality of IMU measurements varies widely depending on sensor hardware [11]. High-end IMUs can be accurate over long distances, while low-cost models tend to drift significantly even over short periods. IMUs are thus used either to estimate only roll and pitch [14], [15], or in sensor fusion frameworks with cameras [1], [5], LiDARs [29], or even a combination of all three [27].

Xu et al. [29] proposed FAST-LIO, a tightly coupled LiDAR-inertial odometry using an iterated extended Kalman

filter (IEKF). In the correction step of the filter, Xu et al. use planar and edge feature point based residuals [29]. FAST-LIO2 [30] improves on this system by registering raw points directly to a map using a point-to-plane ICP metric instead of extracting features. It also extends the state to include gravity [23] and the LiDAR-IMU extrinsic. LIO-EKF [28], another tightly coupled filter approach, employs a strapdown inertial navigation system for state prediction using the IMU. This more sophisticated prediction model enables it to use an EKF instead of an IEKF to achieve similar performance. In contrast to these methods, our approach does not use a filter, in order to avoid expensive sensor noise calibration and tuning requirements.

Alternatively to filter-based approaches, LIO-SAM [22] uses a factor graph [10] for fusing LiDAR and IMU data. Notably, it requires a 9-axis IMU, which our approach does not. It models LiDAR odometry as feature-point and keyframe-based factors, leveraging IMU pre-integration [5] to accumulate IMU data between keyframes into a single factor. Pre-integration efficiently summarizes numerous measurements into a single relative motion constraint, crucial to keeping factor graph approaches tractable when using high-frequency IMU data. Forster et al. [5] propose a pre-integration formulation on the $SO(3) \times \mathbb{R}^3 \times \mathbb{R}^3$ manifold that addresses non-linear uncertainty propagation in $SO(3)$. Brossard et al. [2] propose an alternative using the $SE_2(3)$ group to better represent uncertainty, showing improvement over the $SO(3)$ -based approach. Both approaches treat IMU biases independently from the navigation state. Delama et al. [4] improve on this by coupling biases and navigation state, showing further improvement. In contrast, our approach does not rely on pose graph optimization, which again requires accurate sensor noise and bias modeling. Nevertheless, our presented approach can be easily integrated into such a system to target specific platforms.

When deployed on legged robots, some odometry systems use kinematic constraints from foot contacts with the ground [19]. Leg-KILO [19] couples leg odometry, from kinematic and inertial measurements, and feature-based LiDAR odometry, building upon LIO-SAM’s factor graph approach. Similarly, VILENS [27] is a pose-graph odometry system for legged robots that uses IMU pre-integration [5], pre-integrated leg odometry, LiDAR feature tracking, local LiDAR submap-based odometry, and visual odometry factors. Their multi-sensor fusion enables improved performance, as shown in the DARPA SubT challenge.

In contrast to many prior methods relying on feature tracking for LiDAR registration, DLIO [3] proposes a geometric observer based approach for direct scan-to-map alignment. They use an adaptive keyframe-based map and register each scan using Generalized ICP. They motion-compensate the input scan by constructing a piecewise-constant linear jerk and angular acceleration trajectory from the IMU measurements to estimate a per-point deskewing transform. We show later that a much simpler motion model can still lead to improved performance in our system. Similarly, KISS-ICP [26] is a robust LiDAR-only odometry system with real-time performance across diverse robotic setups. They use an even simpler constant velocity motion model plus an adaptive association

threshold that accounts for deviations from this simple model. However, when applied to quadrupeds or backpack-mounted sensors, KISS-ICP typically struggles as its constant velocity motion model inadequately represents more aggressive motions. We build on their scan-alignment module and integrate our new approach to consider IMU measurements, making the resulting system applicable to more challenging mapping and navigation setups.

In sum, our contribution is a robust and accurate LiDAR-inertial odometry system without any sensor-specific modeling. We simplify the use of IMU measurements compared to filter-based or pre-integration based approaches, avoiding precise noise calibrations, tuning, or sensor-specific configurations that may not be available or practical for all robotic platforms. At the same time, we still maintain on par or better than state-of-the-art performance.

III. OUR APPROACH

LiDAR odometry typically relies on the alignment of successive frames, often using a form of ICP. It is well known that ICP is sensitive to the initial guess [26]. While IMUs are prone to long-term drift, they can be used well for short-term pose tracking. Thus, we frame the computation of LiDAR-inertial odometry as one of computing a good initial motion estimate from IMU data and then refining it with ICP, and incorporate an IMU-based regularization in the ICP optimization.

A. Preprocessing

We consider four relevant frames: odometry frame \mathcal{O} , LiDAR sensor frame \mathcal{L} , IMU sensor frame \mathcal{I} , and body base frame \mathcal{B} . Our goal is to estimate odometry with respect to the body frame \mathcal{B} . To this end, sensor data must first be transformed from their respective frames into \mathcal{B} , similar to Kok *et al.* [11]. For the LiDAR, this is straightforward: each point in the scan is transformed by ${}^{\mathcal{B}}\mathbf{T}_{\mathcal{L}} \in \text{SE}(3)$ from \mathcal{L} to \mathcal{B} , using the LiDAR extrinsics. Transforming the IMU data must account for the transport-rate effect on the acceleration and is given by

$$\begin{aligned} \mathbf{z}_a^{\mathcal{B}} &= {}^{\mathcal{B}}\mathbf{R}_{\mathcal{I}} \mathbf{z}_a^{\mathcal{I}} - (\dot{\mathbf{z}}_{\omega}^{\mathcal{B}} \times {}^{\mathcal{B}}\mathbf{t}_{\mathcal{I}}) - \mathbf{z}_{\omega}^{\mathcal{B}} \times (\mathbf{z}_{\omega}^{\mathcal{B}} \times {}^{\mathcal{B}}\mathbf{t}_{\mathcal{I}}), \\ \mathbf{z}_{\omega}^{\mathcal{B}} &= {}^{\mathcal{B}}\mathbf{R}_{\mathcal{I}} \mathbf{z}_{\omega}^{\mathcal{I}}, \end{aligned} \quad (1)$$

where $\mathbf{z}_a^{\mathcal{I}}$ and $\mathbf{z}_{\omega}^{\mathcal{I}}$ are the IMU acceleration and angular velocity measurements expressed in frame \mathcal{I} . ${}^{\mathcal{B}}\mathbf{R}_{\mathcal{I}} \in \text{SO}(3)$ and ${}^{\mathcal{B}}\mathbf{t}_{\mathcal{I}} \in \mathbb{R}^3$ are the rotation and translation components of the IMU-to-body extrinsic ${}^{\mathcal{B}}\mathbf{T}_{\mathcal{I}}$. The two sensor extrinsics, i.e., the position and orientation of how the sensors are mounted on the robot, are the only parameters we require the user to provide. Hereon, for notational simplicity, we drop the right superscript on vectors as they will all be expressed in the body frame unless otherwise specified.

B. Motion Model

Given a linear acceleration \mathbf{a} and angular velocity $\boldsymbol{\omega}$, the continuous-time rigid body kinematics is given by

$$\begin{aligned} \dot{\mathbf{p}} &= \mathbf{v}, \\ \dot{\mathbf{v}} &= \mathbf{a}, \\ \dot{\mathbf{R}} &= \mathbf{R} [\boldsymbol{\omega}]_{\times}, \end{aligned} \quad (2)$$

where \mathbf{p} is the position, \mathbf{v} is the velocity, $\mathbf{R} := {}^{\mathcal{O}}\mathbf{R}_{\mathcal{B}}$ is the rotation of the body frame with respect to the odometry frame, and $[\cdot]_{\times}$ denotes the skew-symmetric matrix corresponding to a vector [6]. We use the motion model to obtain a relative pose estimate between two successive LiDAR frames to serve as the initial guess for ICP registration. For a discrete time interval $[t, t + \Delta t]$, Euler integration yields the relative changes in position and rotation as:

$$\begin{aligned} \Delta \mathbf{p} &= \mathbf{v}(t) \Delta t + \frac{1}{2} \mathbf{a} \Delta t^2, \\ \Delta \mathbf{R} &= \exp([\boldsymbol{\omega} \Delta t]_{\times}), \end{aligned} \quad (3)$$

where $\exp(\cdot)$ is the exponential map on $\text{SO}(3)$. If \mathbf{a} and $\boldsymbol{\omega}$ are constant over Δt , the Euler integration is exact.

Here, we make our key assumption: we assume a constant linear acceleration and angular velocity motion between successive LiDAR frames. Typical LiDARs, even solid-state ones [13], operate at or above 10 Hz [9]. Deviations from our model due to linear jerk \mathbf{j} or angular acceleration $\boldsymbol{\alpha}$ introduce errors that scale cubically or quadratically with time, respectively. Over such a short horizon as 0.1 s, even with conservative upper bounds for the magnitude of \mathbf{j} and $\boldsymbol{\alpha}$ as 5 m/s³ and 180 deg/s², the resulting errors in estimating relative position and rotation (angle) changes are ~ 0.83 mm and ~ 0.9 deg respectively. Thus, even with such extreme limits for \mathbf{j} and $\boldsymbol{\alpha}$, for typical robots the position error especially is neglectable and the two errors can still be corrected, as we refine the initial guess by using LiDAR scan-to-map alignment (see Sec. III-C).

The IMU gives a measurement directly related to the controls of our motion model: the accelerometer measures specific force on the sensor and the gyroscope measures angular velocity of the sensor with respect to a fixed inertial frame. We can neglect Earth's curvature and rotation effects, as they are insignificant compared to typical sensor measurement noise, and consider the inertial frame to be coincident with \mathcal{O} [11]. In the time interval $[t_{l-1}, t_l]$ between successive LiDAR scans, we denote the set of IMU measurements as $\mathcal{S} = \{(\mathbf{z}_{a,1}, \mathbf{z}_{\omega,1}), \dots, (\mathbf{z}_{a,n}, \mathbf{z}_{\omega,n})\}$. For each of these, the IMU measurement model [5] can be expressed as

$$\begin{aligned} \mathbf{z}_a &= \mathbf{a} - \mathbf{R}^{\top} \mathbf{g}^{\mathcal{O}} + \mathbf{b}_a + \boldsymbol{\eta}_a, \\ \mathbf{z}_{\omega} &= \boldsymbol{\omega} + \mathbf{b}_{\omega} + \boldsymbol{\eta}_{\omega}, \end{aligned} \quad (4)$$

where $\mathbf{g}^{\mathcal{O}}$ is the gravity vector in the odometry frame. We consider two sources of error: \mathbf{b}_a , \mathbf{b}_{ω} are accelerometer and gyroscope biases, typically modeled as being constant [15] or slowly varying [29], and $\boldsymbol{\eta}_a$, $\boldsymbol{\eta}_{\omega}$ are the Gaussian white noise of the accelerometer and gyroscope measurements.

First, we assume a constant bias affecting sensor measurements and estimate it through the following initialization procedure. We use the first measurement set \mathcal{S} assuming no motion between the first two LiDAR frames, by which we can estimate the biases, and we additionally estimate an initial orientation by aligning accelerometer measurements to $\mathbf{g}^{\mathcal{O}}$. As the initialization is short and in the order of 0.1 s, we can produce pose estimates from the second frame onwards with no delay. We show in Sec. IV-B that even on trajectories

longer than 50 km, we obtain state-of-the-art results compared to approaches that perform online bias estimation.

Second, to estimate the relative motion using Eq. (3) with $\Delta t = t_l - t_{l-1}$, we estimate the control inputs $\bar{\mathbf{a}}$ and $\bar{\boldsymbol{\omega}}$ by computing the mean of the measurements in \mathcal{S} while accounting for the bias and gravity terms in Eq. (4). This also reduces the effect of sensor measurement noise $\boldsymbol{\eta}$ by a factor proportional to the number of IMU measurements. This follows observations made by Bloesch et al. [1], who reported no loss in performance in their visual-inertial system by averaging IMU measurements between camera frames. We also show in Sec. IV-C that averaging the set \mathcal{S} leads to no loss in performance, and in some cases is better than integrating measurements independently.

C. LiDAR Scan Registration

Following KISS-ICP [26], we employ a scan-to-map ICP alignment scheme to register each LiDAR scan. Our local map \mathcal{M} is a voxel grid implemented using the VDB data structure [17] with voxel size v in meters, where each voxel stores a fixed number of points. Each LiDAR scan provides a time-stamped point cloud given by $\mathcal{P} = \{(\mathbf{p}_i, t_i) \mid \mathbf{p}_i \in \mathbb{R}^3, t_i \in \mathbb{R}\}$, where \mathbf{p}_i is a point measurement at time t_i during the scan interval $[t_{l-1}, t_l]$. We account for scan distortion caused by platform motion during $[t_{l-1}, t_l]$, by deskewing each measured point \mathbf{p}_i with the transform $\mathbf{T}(t_l, t_i) = \mathbf{T}(t_{l-1}, t_i)^{-1} \mathbf{T}(t_{l-1}, t_i)$, where $\mathbf{T}(t_1, t_2)$ maps the base frame from its pose at t_2 to t_1 . The intermediate transformations are estimated using $\bar{\mathbf{a}}$ and $\bar{\boldsymbol{\omega}}$ with appropriate Δt in Eq. (3). We denote the deskewed scan \mathcal{P}^* and downsample it to $\mathcal{P}_{\text{merge}}^*$ with voxel size $0.5v$. This intermediate cloud is used to update the local map \mathcal{M} later, once the odometry pose has been estimated. We then downsample $\mathcal{P}_{\text{merge}}^*$ further, with voxel size $1.5v$, to obtain points \mathcal{Q} for registration against \mathcal{M} . We refer to this step as double-downsampling and if known that the LiDAR scanner produces only sparse scans, the user may optionally disable this step (see Sec. IV-C). For ICP using the points \mathcal{Q} for registration against the map, we estimate an initial pose $\hat{\mathbf{T}}$ as

$$\hat{\mathbf{T}} = \mathbf{T}_{l-1} \Delta \mathbf{T}, \quad (5)$$

where \mathbf{T}_{l-1} is the previous estimate of the pose, and $\Delta \mathbf{T}$ is computed using $\bar{\mathbf{a}}$ and $\bar{\boldsymbol{\omega}}$ in Eq. (3) with $\Delta t = t_l - t_{l-1}$. In each iteration, we transform \mathcal{Q} into the odometry frame with the current estimate \mathbf{T} and compute the correspondence set $\mathcal{C} = \{(\mathbf{q}, \mathbf{m}) \mid \mathbf{q} \in \mathcal{Q}, \mathbf{m} \in \mathcal{M}\}$ via a fast nearest-neighbor search with a fixed association threshold, exploiting the VDB data structure [17]. The residual for a map point \mathbf{m} and a source point \mathbf{q} is

$$\mathbf{r}(\mathbf{T}) = \mathbf{T} \mathbf{q} - \mathbf{m}. \quad (6)$$

We define the ICP cost as the mean squared residual over \mathcal{C} :

$$\chi_{\text{icp}}(\mathbf{T}) = \frac{1}{|\mathcal{C}|} \sum_{(\mathbf{q}, \mathbf{m}) \in \mathcal{C}} \|\mathbf{r}(\mathbf{T})\|_2^2, \quad (7)$$

where $|\mathcal{C}|$ denotes the cardinality of the correspondence set.

D. Adaptive ICP Regularization

Orientation filters [12], [14] estimate a system's roll and pitch by integrating gyroscope measurements and using accelerometer measurements to correct for drift (yaw is unobservable). In our case, we can similarly exploit the accelerometer measurements further given our motion model. We propose adding a novel regularization to the optimization using the following orientation cost:

$$\chi_{\text{ori}}(\mathbf{T}) = \|\mathbf{R}(\bar{\mathbf{z}}_a - \mathbf{a}_b) + \mathbf{g}^{\mathcal{O}}\|_2^2, \quad (8)$$

where $\bar{\mathbf{z}}_a$ is the average of the accelerometer measurements from \mathcal{S} and \mathbf{a}_b is the true body acceleration. A typical approach in orientation filters to estimate \mathbf{a}_b is to simply assume static motion or to use a low-pass filter [12]. Instead, we use a Kalman filter to estimate \mathbf{a}_b , while aiming to keep parameter tuning to a minimum. We again assume constant acceleration for the prediction step, but model a jerk (magnitude) uniformly distributed over $[-j, j]$ for the process noise in the covariance propagation. The process noise covariance is then given by $(j^2 \Delta t^2 / 3) \mathbf{I}_3$, where $\Delta t = t_l - t_{l-1}$ and \mathbf{I}_3 is the identity matrix. For the correction step, we use $\bar{\mathbf{a}}$ as the measurement. We also use the set \mathcal{S} to compute σ_a , the standard deviation of the acceleration magnitude, and set the measurement covariance to $(\sigma_a^2 / 3) \mathbf{I}_3$, similar to Lee et al. [12]. Through this approach, only one parameter needs to be set: the maximum expected jerk j , a reflection of the expected motion. In our experiments, we set j to 3 m/s^3 as a default starting point applicable to most platforms.

We add χ_{ori} to the optimization for a combined cost:

$$\chi(\mathbf{T}) = \chi_{\text{icp}}(\mathbf{T}) + \frac{1}{\beta} \chi_{\text{ori}}(\mathbf{T}), \quad (9)$$

where β is inversely proportional to the amount of regularization we want to impose [8]. Exploiting σ_a , we set

$$\beta = \beta_0 (1 + \sigma_a^2), \quad (10)$$

where β_0 is a constant and the minimum amount of regularization imposed. Eq. (10) adaptively regularizes the optimization to prefer the LiDAR measurement more when σ_a is high, indicating that either the IMU measurements are noisy or the constant acceleration assumption does not hold for that interval. We then minimize the total cost $\chi(\mathbf{T})$ in an iterative least squares fashion by solving

$$\begin{aligned} \delta \mathbf{x}^* &= \arg \min_{\delta \mathbf{x}} \chi(\mathbf{T} \boxplus \delta \mathbf{x}), \\ \mathbf{T} &\leftarrow \mathbf{T} \boxplus \delta \mathbf{x}^*, \end{aligned} \quad (11)$$

where $\delta \mathbf{x} \in \mathbb{R}^6$ is a correction vector and \boxplus denotes the operator applying this vector to the pose estimate [24]. After convergence, we update the map with $\mathcal{P}_{\text{merge}}^*$ transformed by \mathbf{T} and estimate the linear velocity as $\mathbf{v} = \log_{\mathbf{t}}(\mathbf{T}_{l-1}^{-1} \mathbf{T}) / \Delta t$, where $\log(\cdot)$ is the logarithm map on $\text{SE}(3)$ and $\log_{\mathbf{t}}$ denotes the translational component.

IV. EXPERIMENTAL EVALUATION

We present our experiments to show the capabilities of our method for robust and accurate LiDAR-inertial odometry that

operates reliably across diverse environments and robotic platforms, without relying on sensor-specific configuration. The results of our experiments also support our key claims, which are that our system is: (i) able to perform on par or better than state-of-the-art LiDAR-inertial odometry systems on pose tracking accuracy, (ii) robust and applicable to different sensor configurations, robot platforms, and operational environments, and (iii) usable in diverse scenarios without any change in configuration or modeling.

A. Experimental Setup

We use various publicly available datasets and use established methodologies to evaluate our system’s performance. The Oxford Spires dataset [25] uses a backpack-mounted Hesai QT64 LiDAR paired with a cellphone-grade IMU operating at 400 Hz. It was recorded in a university campus and includes both outdoor and outdoor–indoor sequences. Ground truth trajectories are computed by registering each undistorted scan to an accurate terrestrial laser-scanning map.

Ou *et al.* [19] open-source four sequences recorded using a Unitree Go1 quadruped equipped with a Velodyne VLP-16 and an onboard IMU operating at 500 Hz. We refer to this as the Leg-KILO dataset. Most sequences are indoors, with one in a parking lot and another featuring running motion. Ground truth for the “Indoor” sequence is estimated using a prior map; the other sequences use an offline optimization process exploiting loop closures.

HeLiPR [9] is an urban driving dataset featuring four LiDAR sensors with different ranging technologies and an Xsens MTi-300 IMU operating at 100 Hz. We focus on the Aeva Aeries II and Livox Avia scanners due to their sparsity and unique scan patterns. Ground truth trajectories for each sensor are estimated using an RTK-GPS INS based system.

DigiForests [16] is a multi-session forestry dataset recorded with a backpack sensor rig. Different sessions provide LiDAR and IMU data recorded with different sensor setups, including data from Hesai XT32, QT32, and QT64. Reference trajectories are provided and have been estimated using VILENS [27], incorporating GNSS and loop closures [18] in an offline pose estimation pipeline.

DRZ Living Lab [21] is a drone dataset recorded using an Ouster OS-0 mounted on a DJI M210 v2 platform. The dataset provides nine sequences with ground truth trajectories acquired using a motion capture system.

We also recorded own challenging car datasets, using an OS1-128 LiDAR, the built-in InvenSense IMU operating at 100 Hz, and an SBG Ellipse-D GNSS-INS. We recorded five sequences of varying length; notably, the “Forest” sequence is 20 km long, while “Rural” is a particularly challenging 52 km sequence. This allows us to evaluate odometry approaches on significantly long driving scenarios. We generated the reference poses through offline LiDAR bundle adjustment incorporating RTK-GPS data.

To demonstrate the robustness of our approach and validate our claims, we run our system with a single configuration on all datasets: we set the VDB voxel size v to 1.0 m, data association threshold to 0.5 m, maximum expected jerk j to

3 m/s^3 , and regularization term β_0 to 200. To account for any occlusions around the LiDAR sensor, we clip the scan up to a 1 m radius around the sensor origin. Only in the ablation studies in Sec. IV-C do we change our configuration to analyze our design choices. Additionally, in the Oxford Spires dataset, as the sequences start in motion, we disable initialization routines for all approaches.

For evaluating localization accuracy, we use two widely adopted metrics: absolute trajectory error (ATE) and relative pose error (RPE) from the KITTI Odometry benchmark [26]. ATE provides a measure of global drift in the estimated trajectories, and RPE measures the average translational error between estimated and reference trajectories over various segment lengths, reported as a percentage. This metric is commonly used in autonomous driving, and we use standard segment lengths for evaluating approaches in urban driving scenarios. We also use datasets featuring trajectories of much shorter length, e.g., the shortest sequence in Oxford Spires is 283 m. Hence, for reporting results on the Oxford Spires, Leg-KILO, and DigiForests datasets, we use smaller interval lengths of 1, 2, 5, 10, 20, 50, and 100 m for RPE [8].

B. Odometry Performance

In this first experiment, we compare our approach against various baselines to compare the pose tracking performance. For each baseline, we use the default configuration, including any dataset- or sensor-specific settings. For FAST-LIO2, we disable online extrinsic calibration, and for all approaches, we use the sensor extrinsics provided by the datasets to enable a fair comparison of the odometry performance.

Tab. I reports the results on the Oxford Spires dataset, comparing our system against KISS-ICP, DLIO, FAST-LIO2, and VILENS-SLAM. Tao *et al.* [25] provide results for the closed-source VILENS-SLAM system, which in their experiments was the best performing online SLAM approach. We report the same results here only as a reference for the performance achievable by SLAM systems. All other reported approaches are odometry-only and do not include loop closures. KISS-ICP, a LiDAR-only odometry, generally underperforms compared to other methods, which are otherwise all LiDAR-inertial. It uses a constant velocity motion assumption which can fail to sufficiently model the motion profile of a backpack-mounted sensor. DLIO uses continuous-time deskewing along with a geometric observer for fusing IMU and LiDAR odometry. However, it has worse performance than FAST-LIO2 and our approach in terms of RPE on most sequences, showing local odometry inconsistency. FAST-LIO2 is the next-best performing odometry in our evaluation. It uses a tightly-coupled IEKF with an 18 dimensional state, including online bias and gravity direction estimation. It achieves the best result on the “Bodleian” sequence, better even than the SLAM approach. Notably, however, it has worse RPE results on the challenging “Christ” sequence, which is 2.2 km long across 4 recordings including stairs accessing different indoor levels. Our approach shows consistent performance on the entire dataset. It is generally the best performing odometry on both metrics, and is mainly on par with the SLAM approach.

Method	Blenheim	Bodleian	Christ	Keble	Radcliffe
KISS-ICP [26]	1.16/46.99	2.15/19.10	0.96/11.01	9.48/38.48	<u>0.46/3.58</u>
DLIO [3]	11.35/46.87	<u>0.40/1.11</u>	0.17/4.17	<u>0.36/4.16</u>	0.92/4.28
FAST-LIO2 [30]	0.94/1.11	0.26/0.26	0.72/23.64	6.06/5.42	0.46/1.40
Ours	0.20/0.95	<u>0.86/0.74</u>	<u>0.25/0.89</u>	0.08/0.96	0.15/0.70
VILENS-SLAM [25]	0.56/1.13	1.11/1.68	0.11/0.38	0.12/0.39	0.07/0.39

TABLE I: Quantitative results on the Oxford Spires dataset. We report the ATE in meters and the RPE in percent as [m] / [%]. VILENS-SLAM [25] is listed separately from the other odometry approaches as it is a SLAM system. The best odometry results are in bold, the second best are underlined. We report average values of the metrics over all sequences of each scene.

Method	Urban	Hill	Residential	Forest	Rural
KISS-ICP [26]	8.17/ <u>0.19</u>	17.69/0.78	36.02/ <u>0.11</u>	178.49/0.64	1460.35/0.25
DLIO [3]	4.94/0.54	12.56/0.50	<u>15.87/0.34</u>	79.90/1.10	-/-
FAST-LIO2 [30]	<u>3.53/0.16</u>	<u>6.16/0.07</u>	15.46/0.09	<u>71.69/0.07</u>	<u>1086.51/0.08</u>
Ours	3.52/0.16	6.05/0.06	<u>24.00/0.11</u>	50.11/0.06	714.82/0.07
no-AVG, no-AR	4.16/0.16	8.49/0.08	28.67/0.12	58.02/0.06	833.48/0.07
no-AR	<u>5.00/0.15</u>	<u>5.42/0.06</u>	<u>23.54/0.11</u>	<u>50.55/0.06</u>	<u>732.79/0.07</u>

TABLE II: Quantitative results on our self-recorded car dataset. We report the ATE in meters and the RPE in percent as [m] / [%]. The best results are in bold, the second best are underlined. A dash indicates failure to run on the sequence. Below “Ours”, the ablation rows (“no-AVG, no-AR” and “no-AR”) test disabling of components. Italicized values here indicate improvement over the full method.

Odometry	Corridor	Indoor	Parking	Running
KISS-ICP [26]	10.59/175.08	0.94/n.a.	18.80/348.80	4.29/n.a.
DLIO [3]	2.70/2.83	<u>0.07/n.a.</u>	0.90/5.56	0.24/-
FAST-LIO2 [30]	0.28/1.94	0.21/n.a.	0.24/1.50	0.09/n.a.
Leg-KILO [19]	0.17/0.64	0.05/n.a.	0.16/0.43	0.05/n.a.
Ours	<u>0.22/1.51</u>	0.05/n.a.	<u>0.23/1.80</u>	0.10/n.a.

TABLE III: Quantitative results on the Leg-KILO dataset. We report the ATE in meters and the RPE in percent as [m] / [%]. The best results are in bold, the second best are underlined. For the Indoor and Running sequences, the RPE metric is not applicable (n.a.) due to the short trajectory length.

On our self-recorded car dataset, the odometry results reported in Tab. II retain the general trend from Oxford Spires. KISS-ICP underperforms compared to LiDAR-inertial methods, which can be expected as the other methods can exploit IMU measurements for a better motion model. DLIO fails on the 52 km long “Rural” sequence, which includes underpass sections and frequent elevation changes over 100 m. FAST-LIO2 outperforms our approach on the “Residential” sequence and is very close to our results on the “Urban” and “Hill” sequences. On most sequences, our approach achieves the best results, especially on the “Forest” and “Rural” sequences, showing the robustness of our system in challenging environments. We discuss the ablation results from Tab. II in the next section.

For the Leg-KILO dataset (see Tab. III), the “Indoor” and “Running” sequences are both shorter than 100 m so the RPE metric, even with 1 - 100 m intervals, is not applicable. Here, the ATE metric alone can give insight into both global drift and local consistency. KISS-ICP shows poor results, reflecting that the constant velocity model is not well applicable here. Leg-KILO [19] gives the best results on this dataset. This approach is based on LIO-SAM [22] and incorporates quadruped leg kinematics and foot contact height measurements. These additional sensor requirements lead to a strong performance on legged robots but also prevent the method from being applied to the other datasets. Our approach is purely LiDAR-inertial and we show the next-best performance after Leg-KILO.

We report results on the DigiForests dataset in Tab. IV.

Odometry	2023-03	2023-10	2024-07
KISS-ICP [26]	-/-	-/-	-/-
DLIO [3]	<u>0.27/1.55</u>	<u>0.12/1.45</u>	<u>0.09/0.72</u>
FAST-LIO2 [30]	-/-	-/-	0.08/0.51
Ours	0.18/0.79	0.11/1.14	0.10/0.87

TABLE IV: Quantitative results on the DigiForests dataset. We report ATE in meters and RPE in percent as [m] / [%]. The best results are in bold, the second best are underlined. A dash indicates failure on at least one of the sequences from the recording period.

The forest environment here is challenging and cluttered, with dense vegetation and thick tree canopies leading to noisy range measurements. The dataset provides multiple sequences recorded across three different seasons, and we report average results for each season in Tab. IV. Especially, the first season provides an additionally challenging situation as the LiDAR is mounted at a 45 degree inclination on the sensor setup. KISS-ICP fails on almost all sequences in this dataset. FAST-LIO2 as well fails on three of seven sequences in the first season, and on all sequences in the second season. It achieves the best results on the third season, but it is also close to our approach and DLIO. Our approach and DLIO provide consistent results across the full dataset, with our approach outperforming DLIO on two of three seasons. We note that the reference trajectories are obtained with the VILENS-SLAM system [25] and all approaches reported here are purely odometry. A qualitative result of our odometry in a forest environment is shown in Fig. 1, where additionally the robot platform is again different from that in the DigiForests dataset.

In Tab. V, we present results on the HeLiPR dataset using the Aeva Aeries II and Livox Avia sensors, which feature challenging scanning patterns, with especially the Aeva being a low field-of-view solid-state LiDAR. No single odometry performs clearly better than others on this dataset. DLIO struggles on the Bridge sequences but performs best on the Town Aeva sequences. FAST-LIO2 achieves the best results on Bridge sequences, while KISS-ICP and our approach are among the best on the Roundabout and Town sequences. Importantly, ours is the only approach that performs consistently and without failures across all datasets presented thus far. Note

Method	Bridge		Roundabout		Town	
	Aeva	Avia	Aeva	Avia	Aeva	Avia
KISS-ICP [26]	144.92/0.94	152.97/0.29	34.90/0.65	9.73/0.26	30.92/0.99	21.62/0.68
DLIO [3]	-/-	-/-	40.50/1.17	13.82/0.27	25.95/0.79	54.88/1.06
FAST-LIO2 [30]	115.17/0.62	51.41/0.07	39.69/0.69	8.39/ 0.12	39.77/0.84	12.11/0.35
Ours	<u>132.77/0.97</u>	133.89/0.14	38.54/0.70	8.38/0.13	<u>45.22/1.02</u>	<u>13.05/0.30</u>
Ours, no-DD	<u>141.02/0.65</u>	<u>118.53/0.13</u>	<u>36.93/0.58</u>	<u>8.79/0.13</u>	<u>29.64/0.86</u>	11.11/0.30

TABLE V: Quantitative results on the Bridge, Roundabout, and Town sequences of the HeLiPR dataset using the Aeva and Avia LiDAR sensors. We report the ATE in meters and the RPE in percent as [m] / [%]. The best results are in bold, the second best are underlined. We report the average values of the metrics over each scene as three different runs are performed per scene. A dash indicates failure on at least one run in the scene. Additionally, we report an ablation of our method where “no-DD” disables the double-downsampling.

Method	Blenheim	Bodleian	Christ	Keble	Radcliffe
no-AVG, no-AR	0.48/1.06	0.81/0.80	0.29/1.10	0.08/1.05	0.11/0.72
no-AR	<u>0.22/0.94</u>	0.81/0.77	0.24/0.79	0.08/1.00	0.19/0.71
Ours	0.20/0.95	<u>0.86/0.74</u>	<u>0.25/0.89</u>	0.08/0.96	0.15/0.70

TABLE VI: Ablation study on the Oxford Spires dataset. We report the ATE in meters and the RPE in percent as [m] / [%]. The best results are in bold, the second best are underlined. “no-AVG” denotes disabling the IMU averaging and “no-AR” disables the adaptive regularization.

that “Ours, no-DD” in Tab. V is an ablation variant discussed in the next section.

To further show the robustness of our approach, we report results of our system on the DRZ Living Lab drone dataset [21]. We achieve an average APE of 0.05 m and average RPE of 2.64% over the entire dataset (9 sequences), which features rapid and dynamic maneuvers by the drone. Our results are competitive with those reported by Quenzel et al. [21] in their work. The experiments presented in this section support our first claim that our method achieves pose tracking accuracy on par or better than state-of-the-art LiDAR-inertial odometry systems.

C. Ablation Study

In this last set of experiments, we systematically remove or modify components of our system and analyze how these changes affect performance. Other than the below-mentioned changes, the configuration is identical to that used in Sec. IV-B. In the following, “no-AVG, no-AR” denotes integrating individual IMU measurements instead of averaging the set \mathcal{S} (Sec. III-B). Disabling averaging prevents computing $\bar{\mathbf{a}}$, $\bar{\boldsymbol{\omega}}$, and σ_a (Sec. III-D), so we effectively disable our adaptive regularization as well. “no-AR” disables only the adaptive regularization while retaining the averaging. As detailed in Sec. III-B, averaging the measurements relates to using a constant acceleration and constant angular velocity motion model, which can reduce noise in the effective control inputs $\bar{\mathbf{a}}$ and $\bar{\boldsymbol{\omega}}$. “Ours” reports the results of our complete system, including both averaging and adaptive regularization.

Tab. II, Tab. VI and Tab. VII report ablation study results on our self-recorded car dataset, Oxford Spires and Leg-KILO datasets, respectively. Comparing “no-AVG, no-AR” to “no-AR”, we see improved or on par ATE performance in 12 of 14 sequences across all three tables and improved RPE performance in all sequences with “no-AR” over “no-AVG, no-AR”. Our motion model is enabled in “no-AR”, and these results support our argument that a constant acceleration and angular velocity motion model can improve

Method	Corridor	Indoor	Parking	Running
no-AVG, no-AR	0.22/1.80	<u>0.05/n.a.</u>	0.23/2.34	<u>0.09/n.a.</u>
no-AR	0.22/1.59	<u>0.05/n.a.</u>	0.22/1.98	<u>0.09/n.a.</u>
Ours	0.22/1.51	<u>0.05/n.a.</u>	0.23/ 1.80	0.10/n.a.
no-DD	0.22/1.35	0.04/n.a.	0.19/2.32	0.08/n.a.

TABLE VII: Ablation study on the Leg-KILO dataset. We report the ATE in meters and the RPE in percent as [m] / [%]. The best results are in bold, the second best are underlined. For the Indoor and Running sequences, the RPE metric is not applicable (n.a.) due to the short trajectory length. “no-AVG” denotes disabling the IMU averaging, “no-AR” disables the adaptive regularization, and “no-DD” disables the double-downsampling.

performance, as averaging the measurements can reduce signal noise, making the system more robust.

By exploiting this motion model, we can additionally enable our adaptive regularization in “Ours”. Comparing “no-AR” to “Ours”, we see improved or on par ATE on 8 of 14 sequences and improved or on par RPE on 9 of 12 sequences with “Ours” over “no-AR” (and on par or outperforming “no-AVG, no-AR” on 11 of 14 sequences for ATE and all sequences for RPE). The regularization leads to reduced drift by prioritizing the IMU measurements when their variance is low, and on the significantly longer car dataset sequences this improvement is reflected in ATE (e.g. Urban and Rural sequences), whereas on the much shorter Leg-KILO sequences this is reflected in RPE.

Additionally, since the Leg-KILO dataset uses a relatively sparse Velodyne VLP-16 LiDAR, we evaluate the effect of disabling only double-downsampling for sparse LiDARs. We evaluate the same also for the HeLiPR dataset using the Aeva and Avia sensors. This is reported as “no-DD” in Tab. VII and Tab. V, and we see that, in general, the results improve further. Our double-downsampling is intended to speed up processing for dense LiDARs and disabling it for sparse LiDARs can improve accuracy with no notable loss in runtime performance, given that the scan is sparse anyhow. We note that tuning our system in this manner, i.e., disabling double-downsampling, incorporates knowledge of the sensor in use. However, we argue this is trivial for the user to determine and provide the option for the user as a boolean flag to optionally test. By default, no assumption on sensor resolution is made and double-downsampling is enabled.

The ablation studies presented in this section validate our design choices and, combined with the results from Sec. IV-B, support our second and third claims. These experiments show that our approach is robust and usable with diverse sensor setups and operational scenarios without any configuration

changes. Moreover, our approach operates faster than the sensor frame rate on all presented datasets. Thus, we supported all our claims with this experimental evaluation.

V. CONCLUSION

In this paper, we presented a novel approach to LiDAR-inertial odometry, designed for general applicability of LiDAR-IMU combinations. Our system is robust, accurate, and requires only knowledge about the extrinsics of the LiDAR and IMU. By using a simplified motion model, we average high-frequency IMU measurements, reducing the sensor noise without a notable loss in performance. Based on this model, we propose an adaptive regularization on the ICP-based LiDAR scan registration, leading to improved odometry performance. We evaluated our approach on several public datasets featuring a wide range of environments, sensor configurations, and motion profiles. We provide extensive comparisons to other existing techniques and supported all claims made in this paper. The experiments suggest that our approach is robust, adaptable to different setups with a default initial configuration, and able to achieve strong performance out of the box. Finally, we have open-sourced our implementation and provide an easy-to-install package available via common system package managers to encourage community development and further research.

ACKNOWLEDGMENTS

We thank Ignacio Vizzo, Benedikt Mersch, Louis Wiesmann, Lucas Nunes, and Jens Behley for developing and maintaining the sensor platform used for recording the car data. We thank Fang Nan for providing data from the forestry harvesting machine.

REFERENCES

- [1] M. Bloesch, M. Burri, S. Omari, M. Hutter, and R. Siegwart. Iterated extended kalman filter based visual-inertial odometry using direct photometric feedback. *Intl. Journal of Robotics Research (IJRR)*, 36(10):1053–1072, 2017.
- [2] M. Brossard, A. Barrau, P. Chauchat, and S. Bonnabel. Associating uncertainty to extended poses for on lie group imu preintegration with rotating earth. *IEEE Trans. on Robotics (TRO)*, 38(2):998–1015, 2022.
- [3] K. Chen, R. Nemiroff, and B.T. Lopez. Direct lidar-inertial odometry: Lightweight lio with continuous-time motion correction. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2023.
- [4] G. Delama, A. Fornasier, R. Mahony, and S. Weiss. Equivariant imu preintegration with biases: A galilean group approach. *IEEE Robotics and Automation Letters (RA-L)*, 10(1):724–731, 2025.
- [5] C. Forster, L. Carlone, F. Dellaert, and D. Scaramuzza. On-manifold preintegration for real-time visual-inertial odometry. *IEEE Trans. on Robotics (TRO)*, 33(1):1–21, 2017.
- [6] G. Grisetti, T. Guadagnino, I. Aloise, M. Colosi, B. Della Corte, and D. Schlegel. Least Squares Optimization: From Theory to Practice. *Robotics*, 9(3), 2020.
- [7] T. Guadagnino, B. Mersch, S. Gupta, I. Vizzo, G. Grisetti, and C. Stachniss. KISS-SLAM: A Simple, Robust, and Accurate 3D LiDAR SLAM System With Enhanced Generalization Capabilities. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2025.
- [8] T. Guadagnino, B. Mersch, I. Vizzo, S. Gupta, M. Malladi, L. Lofebaro, G. Doisy, and C. Stachniss. Kinematic-ICP: Enhancing LiDAR Odometry with Kinematic Constraints for Wheeled Mobile Robots Moving on Planar Surfaces. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2025.
- [9] M. Jung, W. Yang, D. Lee, H. Gil, G. Kim, and A. Kim. Helipr: Heterogeneous lidar dataset for inter-lidar place recognition under spatiotemporal variations. *Intl. Journal of Robotics Research (IJRR)*, 43(12):1867–1883, 2024.
- [10] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. Leonard, and F. Dellaert. isam2: Incremental smoothing and mapping with fluid relinearization and incremental variable reordering. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2011.
- [11] M. Kok, J.D. Hol, and T.B. Schön. *Using Inertial Sensors for Position and Orientation Estimation*. Now Publishers Inc., 2017.
- [12] J.K. Lee, E.J. Park, and S.N. Robinovitch. Estimation of attitude and external acceleration using inertial sensor measurement during various dynamic conditions. *IEEE Trans. on Instrumentation and Measurement*, 61(8):2262–2273, 2012.
- [13] K. Li, M. Li, and U.D. Hanebeck. Towards high-performance solid-state-lidar-inertial odometry and mapping. *IEEE Robotics and Automation Letters (RA-L)*, 6(3):5167–5174, 2021.
- [14] H.J. Luinge and P.H. Veltink. Measuring orientation of human body segments using miniature gyroscopes and accelerometers. *Medical and Biological Engineering and Computing*, 43(2):273–282, 2005.
- [15] R. Mahony, T. Hamel, and J.M. Pfifflin. Nonlinear complementary filters on the special orthogonal group. *IEEE Trans. on Automatic Control*, 53(5):1203–1218, 2008.
- [16] M.V. Malladi, N. Chebrolu, I. Scacchetti, L. Lofebaro, T. Guadagnino, B. Casseau, H. Oh, L. Freißmuth, M. Karpainen, J. Schweier, S. Leutenegger, J. Behley, C. Stachniss, and M. Fallon. DigiForests: A Longitudinal LiDAR Dataset for Forestry Robotics. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2025.
- [17] K. Museth, J. Lait, J. Johanson, J. Budsberg, R. Henderson, M. Alden, P. Cucka, D. Hill, and A. Pearce. OpenVDB: An Open-source Data Structure and Toolkit for High-resolution Volumes. In *Proc. of the ACM SIGGRAPH Courses*, 2013.
- [18] H. Oh, N. Chebrolu, M. Mattamala, L. Freißmuth, and M. Fallon. Evaluation and Deployment of LiDAR-based Place Recognition in Dense Forests. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2024.
- [19] G. Ou, D. Li, and H. Li. Leg-kilo: Robust kinematic-inertial-lidar odometry for dynamic legged robots. *IEEE Robotics and Automation Letters (RA-L)*, 9(10):8194–8201, 2024.
- [20] P. Pfreundschuh, H. Oleynikova, C. Cadena, R. Siegwart, and O. Andersson. Coin-lio: Complementary intensity-augmented lidar inertial odometry. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2024.
- [21] J. Quenzel and S. Behnke. Real-time Multi-Adaptive-Resolution-Surfel 6D LiDAR Odometry using Continuous-time Trajectory Optimization. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2021.
- [22] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, and D. Rus. LIO-SAM: Tightly-coupled Lidar Inertial Odometry via Smoothing and Mapping. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2020.
- [23] J. Solà. Quaternion kinematics for the error-state kalman filter. *arXiv preprint*, arXiv:1711.02508, 2017.
- [24] J. Solà, J. Deray, and D. Atchuthan. A micro lie theory for state estimation in robotics. *arXiv preprint*, arXiv:1812.01537, 2021.
- [25] Y. Tao, M.Á. Muñoz-Bañón, L. Zhang, J. Wang, L.F.T. Fu, and M. Fallon. The oxford spires dataset: Benchmarking large-scale lidar-visual localisation, reconstruction and radiance field methods. *Intl. Journal of Robotics Research (IJRR)*, 2025.
- [26] 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 (RA-L)*, 8(2):1029–1036, 2023.
- [27] D. Wisth, M. Camurri, and M. Fallon. VILENS: Visual, Inertial, Lidar, and Leg Odometry for All-Terrain Legged Robots. *IEEE Trans. on Robotics (TRO)*, 39(1), 2023.
- [28] Y. Wu, T. Guadagnino, L. Wiesmann, L. Klingbeil, C. Stachniss, and H. Kuhlmann. LIO-EKF: High frequency LiDAR-inertial odometry using extended Kalman filters. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2024.
- [29] 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 (RA-L)*, 6(2):3317–3324, 2021.
- [30] W. Xu, Y. Cai, D. He, J. Lin, and F. Zhang. Fast-lio2: Fast direct lidar-inertial odometry. *IEEE Trans. on Robotics (TRO)*, 38:2053–2073, 2021.