

BIEVR-LIO: Robust LiDAR-Inertial Odometry through Bump-Image-Enhanced Voxel Maps

Patrick Pfreundschuh¹, Turcan Tuna², Cedric Le Gentil³, Roland Siegwart¹, Cesar Cadena², Helen Oleynikova^{1,3}

¹Autonomous Systems Lab, ETH Zürich, Switzerland, ²Robotic Systems Lab, ETH Zürich, Switzerland,

³Mobile Robotics Lab, ETH Zürich, Switzerland

Abstract—Reliable odometry is essential for mobile robots as they increasingly enter more challenging environments, which often contain little information to constrain point cloud registration, resulting in degraded LiDAR-Inertial Odometry (LIO) accuracy or even divergence. To address this, we present BIEVR-LIO, a novel approach designed specifically to exploit subtle variations in the available geometry for improved robustness. We propose a high-resolution map representation that stores surfaces as voxel-wise oriented height images. This representation can directly be used for registration without the calculation of intermediate geometric primitives while still supporting efficient updates. Since informative geometry is often sparsely distributed in the environment, we further propose a map-informed point sampling strategy to focus registration on geometrically informative regions, improving robustness in uninformative environments while reducing computational cost compared to global high-resolution sampling. Experiments across multiple sensors, platforms, and environments demonstrate state-of-the-art performance in well-constrained scenes and substantial improvements in challenging scenarios where baseline methods diverge. Additionally, we demonstrate that the fine-grained geometry captured by BIEVR-LIO can be used for downstream tasks such as elevation mapping for robot locomotion.

I. INTRODUCTION

Ego-motion estimation is a fundamental building block of many mobile robotics applications. LiDAR-Inertial Odometry (LIO) has become one of the most widely adopted solutions due to its high accuracy and robustness in both indoor and outdoor environments [1, 2]. With the advent of smaller and more affordable 3D LiDAR sensors, LIO systems are no longer confined to research platforms but are increasingly deployed in consumer products, such as robotic lawnmowers. Despite its success, the robustness of LIO degrades significantly in geometrically less informative environments, such as straight tunnels or large flat terrains, where the observed geometry does not sufficiently constrain point cloud registration. Several works have addressed this issue by detecting degeneracy [3, 4], but rely on auxiliary odometry from additional sensors, increasing system complexity. Recent methods attempt to mitigate degeneracy using LiDAR-inertial data alone [5, 6] but only address short sections of degeneracy. Another line of research leverages LiDAR intensity as an additional cue for registration [7–9]. While such intensity-augmented methods achieve promising performance, they depend on dense LiDAR data and are therefore incompatible with sensors that employ sparse or irregular scan patterns.




Fig. 1. BIEVR-LIO achieves robust odometry in a tunnel through registration against its high-resolution bump-image-enhanced voxel map. The colored bump images illustrate geometric details of wall cutouts and the train bed.

In this work, we approach robust LIO in geometrically sparse environments from a different perspective. We observe environments often considered uninformative are rarely degenerate in a strict sense, as real-world scenes typically contain subtle geometric variations that render the problem well-constrained in principle, as shown in Figure 1. However, existing LIO systems fail to exploit this information because their map representations lack the resolution required to capture such fine-grained structure under real-time constraints. We introduce a novel Bump-Image-Enhanced Voxel Representation (BIEVR) designed to explicitly encode high-resolution geometric information for efficient registration. BIEVR represents local surface geometry within each voxel as a 2D height image. Specifically, we estimate a dominant plane per voxel, which serves as an image plane, and store per-pixel deviations of observed points from this plane (“bump-image”). This representation preserves details while enabling constant-time map lookups and efficient, direct scan-to-map registration.

Although BIEVR enables high-resolution mapping, a map area can only contribute constraints to the registration if a nearby point is included in the sampled input point cloud. However, these subtle features can be missed due to down-sampling of the input scan, which is often required to reduce computational cost. Selecting relevant points is difficult because LiDARs have limited spatial resolution, particularly at longer ranges, and geometric details can be too subtle to be detectable from the incoming scan alone. As a result, these details often emerge only after integrating multiple scans into

the map. We therefore propose a map-aware sampling strategy that leverages the proposed representation. For each voxel, we compute a score derived from its height image, that reflects its potential to provide informative registration Jacobians. Given the initial pose estimate from IMU integration, we use a higher point-density in voxels with a high score. This approach focuses computation on informative areas, while using significantly fewer points than uniform high-density sampling.

We evaluate our approach extensively across multiple LiDAR sensors and robotic platforms using a single, unified set of parameters, highlighting the broad applicability of the method. The results demonstrate that our method achieves exceptional robustness in severely uninformative environments where all compared baselines fail. It also delivers state-of-the-art performance in well-constrained scenarios at competitive runtime. Finally, beyond odometry, we demonstrate how BIEVR-LIO naturally supports downstream robotics tasks, by showing real-world experiments where the map is used for footstep planning on a legged robot.

In summary, we propose two contributions to exploit the geometric details in LiDAR data: (1) a high-resolution voxel map in which each voxel stores an oriented bump image, and (2) a map-informed, dual-resolution point sampling scheme.

We open-source our implementation at <https://github.com/ethz-asl/bievr-lio>.

II. RELATED WORK

A. LiDAR-based Odometry

One of the first widely used LiDAR odometry systems was LOAM [10], which uses planar and edge features for registration. KISS-ICP [11] showed that strong performance is also achievable with a simple point-to-point ICP design. While those approaches register full scans, others treat the measurements as a point stream for continuous trajectory estimation [12–15]. IMUs enable motion compensation and provide a registration prior and are thus often fused with LiDAR. Typically, LIO systems follow a tightly-coupled fusion, using factor graphs [16–18] or filters [12, 13, 19, 20]. Different from the preintegration [21] used in these systems, RKO-LIO [22] presents a simplified IMU model that is used to regularize registration. Contrary to existing work, the IMU does not influence the registration beyond the scan-undistortion and initialization in our loosely-coupled pipeline.

A common challenge for LIO systems is geometric degeneracy, encountered in environments with little informative structure. Several degeneracy detection methods have been proposed [4, 6, 23], however, they only handle short uninformative sections or require odometry from additional sensor sources. Other work addresses degeneracy by using LiDAR intensity as an additional source of constraints [7–9]. Despite impressive results, these rely on dense LiDARs and do not support irregular or sparse scans that are often found in cheaper sensors. We instead use a high-resolution voxel-image map that increases gradient observability, combined with a map-informed point sampling strategy to prevent degeneracy.

B. Map Representations

The most direct map representation in LiDAR odometry is a point cloud map, commonly maintained using keyframes [24], submaps [17], or a globally accumulated map [19], which can also be continuously optimized [16]. These maps are typically indexed using kd-tree structures, leading to $\mathcal{O}(\log N)$ point-lookup complexity. This forces methods to downsample aggressively to stay real-time, losing detail. Faster-LIO [20] improved efficiency with $\mathcal{O}(1)$ lookups in a hashed voxel map. However, unless point-to-point registration is used [11, 22, 25], map points are not directly used in the registration objective. Instead, intermediate representations such as surface normals or covariances are estimated from local neighborhoods, which presents a computational bottleneck.

To avoid repeated neighborhood queries, voxel-based methods directly track a normal [26] or covariance [27] per voxel. Closely related are surfel maps [28–30], which represent the environment as oriented surface elements. Although efficient, such representations assume that geometry can be approximated by a plane or a Gaussian, which limits the level of detail.

Other work directly uses distance fields for registration. D-LIO [31] proposes a discrete truncated distance field updated efficiently via a binary kernel, while 2Fast-2Lamaa [32] models a continuous distance field using Gaussian Processes, enabling smooth point-to-surface distance queries.

Motivated by map compression, CURL-SLAM [33] represents the surface inside a voxel using spherical harmonics, reducing the memory footprint. Points are first projected onto an axis-aligned plane to generate a height image, which is then approximated by optimizing harmonic coefficients. Because harmonic fitting is computationally expensive, map updates are not performed at LiDAR frequency, and the resulting approximation is inherently smooth, suppressing subtle geometric cues that are valuable for registration. Registration is performed against resampled height images by minimizing the difference between the projected point height and the pixel value. Our approach builds on the same registration paradigm, but explicitly stores voxel-wise oriented height images at high resolution, which enables efficient map updates and constant-time pixel lookups without reconstruction or neighborhood queries. We do not assume an underlying geometric primitive and preserve sharp discontinuities, allowing subtle geometric gradients to be used for robust odometry.

C. Point Sampling

Modern 3D LiDARs produce more points than can be registered in real-time, necessitating downsampling. A widely used strategy is fixed-size voxel grid filtering, where the number of sampled points depends on the scale of the environment, which can result in unstable registration in narrow environments. To address this, AdaLIO [34] adapts the voxel size to the environment extent. LOAM [10] instead samples specific edge and planar features. Since informative points are usually not evenly distributed in a scan, RMS [35] presents a sampling scheme that minimizes redundancy within the point cloud.




Fig. 2. Overview of the proposed BIEVR-LIO system. Lines indicate information flow before ($\cdots$), during ($- -$) and after ($—$) the state update.

These methods sample exclusively based on the current input scan. However, degeneracy is caused by a lack of informative Jacobians in the registration objective, which also depends on the map. Consequently, other approaches [4, 36, 37] sample based on the Hessian of the optimization problem. This requires building computationally expensive correspondences and residual terms, which then necessitates coarse downsampling before score computation, and in turn can remove crucial details. Our approach avoids the calculation of residual terms by defining a metric derived from the voxel-wise height images to approximate the potential to provide informative Jacobians, and selects more points in voxels with high potential.

III. METHOD

We propose BIEVR-LIO, a system that exploits subtle geometric details while maintaining computational efficiency. At its core, it represents surfaces as voxel-wise oriented height images, enabling storage of fine-grained geometry with constant-time lookups and efficient updates. IMU integration is used for pointwise undistortion of the LiDAR scan, and provides an initial pose estimate for registration. Our system follows a loosely-coupled design in which the pose is estimated exclusively from scan-to-map registration, without any inertial residuals. We present a map-aware sampling strategy that uses a lightweight metric to increase point density in informative areas. Registration is then performed directly on the BIEVR representation by minimizing the height error between sampled input points and their corresponding voxel height images. After pose estimation, we optimize velocity, gravity, and IMU biases. Finally, we update the map using all scan points, made tractable by an efficient, parallelized process. We provide an overview of the system in Figure 2.

A. Bump-Image-Enhanced Voxel Representation

LIO systems often approximate surfaces as planes, which enables efficient registration constraints. In practice, however, choosing an appropriate spatial resolution for planar approximation is difficult, as large support regions suppress geometric detail, while small regions lead to unstable plane estimates.

Instead of assuming that a plane fully explains the surface, we explicitly capture the local deviation from a dominant plane per voxel. This plane acts as the image plane of a height image, in which pixels store the fine-grained deviation orthogonal to the plane. By tracking the deviation from a well-defined plane, this representation preserves subtle geometric cues that are lost under coarse planar models, while avoiding the instability of estimating high-resolution 3D primitives directly from points.

1) *Iterative Plane Estimation*: For the iterative plane estimation and normal update, we adopt the voxel-wise update strategy presented in iG-LIO [27]. Undistorted LiDAR points are transformed to the world frame using the pose estimate and grouped to voxels of 0.5 m length via a spatial hash function. The individual voxels are then updated in parallel. For each voxel $\mathcal{V}$, we update the sum of points $\mathbf{s}$, the sum of outer products $\mathbf{C}$, and the number of points n as

$$\mathbf{s} \leftarrow \mathbf{s} + \sum_{\mathbf{p}_i \in \mathcal{V}} \mathbf{p}_i, \quad \mathbf{C} \leftarrow \mathbf{C} + \sum_{\mathbf{p}_i \in \mathcal{V}} \mathbf{p}_i \mathbf{p}_i^\top, \quad n \leftarrow n + n_\nu, \quad (1)$$

where ν is the set of points $\mathbf{p}_i$ falling into $\mathcal{V}$, and $n_\nu = |\nu|$ is the number of points added to that voxel from the current scan. The centroid and covariance can then be extracted as:

$$\boldsymbol{\mu} = \frac{\mathbf{s}}{n}, \quad \boldsymbol{\Sigma} = \frac{\mathbf{C} - \mathbf{s} \boldsymbol{\mu}^\top}{n + 1}. \quad (2)$$

Based on eigenvalue decomposition of $\boldsymbol{\Sigma}$, we extract the voxel normal as the eigenvector with the smallest eigenvalue. We create a coordinate system with origin at $\boldsymbol{\mu}$ and set the z-axis as the normal. We define the transform¹ $\mathbf{T}_{VG} \in \text{SE}(3)$, with $\mathbf{R}_{VG} \in \text{SO}(3)$ its rotational component, that maps points from the global world frame G to the voxel coordinate frame V .

2) *Bump Image Update*: Each voxel in the BIEVR map stores an image $\mathbf{I}$ representing fine-grained surface deviations perpendicular to the local plane. We initialize a voxel once at least four points were accumulated in it. Upon initialization we determine the minimal required image size to represent the full volume of the voxel by projecting its eight corners

¹Throughout this paper, we use the shorthand ${}_{\mathbf{A}}\mathbf{p} = \mathbf{T}_{\mathbf{A}\mathbf{B}}\mathbf{B}\mathbf{p}$ to denote the homogeneous transformation $[{}_{\mathbf{A}}\mathbf{p}^\top, 1]^\top = \mathbf{T}_{\mathbf{A}\mathbf{B}}[{}_{\mathbf{B}}\mathbf{p}^\top, 1]^\top$.




Fig. 3. Map update and registration. (a) A dominant plane is iteratively updated. The image size is determined by projecting voxel corners onto this plane. Points are coloured based on plane distance. (b) Points are projected to pixels, and pixel values are updated based on height above the plane. (c) During registration, we minimize height residuals between pixel values and input points (gray). Arrows indicate the Jacobian components described in Equation (14).

onto the x - y plane of the voxel coordinate system, as shown in Figure 3. The image width and height are computed as

$$W = \frac{x_{\max} - x_{\min}}{r}, \quad H = \frac{y_{\max} - y_{\min}}{r}, \quad (3)$$

where r is the chosen pixel resolution and $x_{\min}$, $x_{\max}$, $y_{\min}$, $y_{\max}$ are the extrema of the projected corners. We define a height image coordinate system C with $\mathbf{R}_{GC} = \mathbf{R}_{GV}$ and origin at ${}_{GP}C = \mathbf{T}_{GV}[x_{\min}, y_{\min}, 0]^\top$.

To calculate the height image values (Figure 3), we project each input point in a voxel to its image coordinates:

$${}_{CP}i = \mathbf{T}_{CG} {}_{GP}i, \quad u = \frac{[{}_{CP}i]_x}{r}, \quad v = \frac{[{}_{CP}i]_y}{r}, \quad z = [{}_{CP}i]_z \quad (4)$$

where $[{}_{CP}i]_x$, $[{}_{CP}i]_y$, $[{}_{CP}i]_z$ denote the individual coordinates of ${}_{CP}i$. To accumulate information from multiple scans in the bump image $\mathbf{I}$, we update pixel values with a weighted averaging. As we use a small pixel size of 0.05 m, it is inevitable that points occasionally fall into the nearby pixels. As orientation errors result in larger Euclidean error for points further away from the LiDAR, we regard points close to the LiDAR as more reliable and assign them a higher weight:

$$\mathbf{I}(u, v) \leftarrow \frac{\mathbf{W}(u, v) \cdot \mathbf{I}(u, v) + w_i \cdot z_i}{\mathbf{W}(u, v) + w_i} \quad (5)$$

$$\mathbf{W}(u, v) \leftarrow \mathbf{W}(u, v) + w_i, \quad w_i = \min(0.5, \|\mathbf{L}_P \mathbf{p}_i\|^{-1}) \quad (6)$$

where $\mathbf{W}$ stores the pixel weights. After updating, we apply a Gaussian filter with 1 pixel radius to mitigate pixel noise.

If the voxel normal changes by more than 3° , we update the image coordinate system and recompute the image by reprojecting the previous image from the old coordinate frame C_o to the new plane frame C_n :

$${}_{C_n}P = \mathbf{T}_{C_n C_o} [u \cdot r, v \cdot r, \mathbf{I}_o(u, v)]^\top, \quad (7)$$

$$\mathbf{I}_n \left(\frac{[{}_{C_n}P]_x}{r}, \frac{[{}_{C_n}P]_y}{r} \right) = [{}_{C_n}P]_z \quad (8)$$

As this map update is computationally cheap, we are able to integrate the full undistorted input scan, which provides more information than downsampled points used in other works.

3) *Implementation*: We implement BIEVR-LIO in a hash-map, where voxels are indexed by their spatial location using Morton code [38], which achieves $\mathcal{O}(1)$ average-case lookup. Each voxel stores the transform $\mathbf{T}_{CG}$, the height image $\mathbf{I}$, pixel weights $\mathbf{W}$, the sum of points $\mathbf{s}$, the sum of outer products $\mathbf{C}$, the number of included points n and its Mean Image Distance (Section III-B). Voxels are efficiently updated in parallel.

To enforce bounded memory, we maintain a Least Recently Used cache. If the number of stored voxels exceeds a configured capacity, the least recently updated voxels are evicted.

B. Map-Informed Point Sampling

The fine-grained geometric detail captured in BIEVR can only contribute to registration if a corresponding input point is sampled in its vicinity. Uniformly sampling the input scan at high density, however, increases the registration's computational cost and introduces noise from uninformative regions. To address this, we propose a map-aware, two-stage sampling strategy that prioritizes points in informative areas.

We first quantify the potential of a voxel to provide useful gradients beyond its planar component. Conveniently, the height image itself is an indicator of non-planarity. Flat surfaces result in small deviations from the voxel plane, translating into small pixel distance, while non-planar surfaces cause large values. Thus, we use the Mean Image Distance (MID), calculated over observed pixels, as an indicator value:

$$\text{MID} = \frac{1}{|\Omega_{\text{obs}}|} \sum_{(u, v) \in \Omega_{\text{obs}}} \|\mathbf{I}(u, v)\|, \quad (9)$$

with $\Omega_{\text{obs}} = \{(u, v) : \mathbf{W}(u, v) > 0\}$. Voxels with larger MID contain richer geometric detail and are more likely to contribute informative Jacobians. This computationally efficient metric is updated during mapping, which avoids additional computation during the sampling procedure.

The sampling proceeds in two stages. First, the undistorted point cloud is downsampled at a high resolution (0.1 m) and transformed into the world frame using the IMU-based initial pose estimate. For each downsampled point, we query its corresponding map voxel and select the 300 voxels with the highest MID to retain all points within them. Second, the points in the remaining voxels are further downsampled at a coarser resolution (0.5 m). The combination of high-resolution points from informative voxels and low-resolution points elsewhere (shown in Figure 4) is then used for registration.

C. Loosely-Coupled State Estimation

BIEVR-LIO estimates the robot state defined as $\mathbf{x} \triangleq [\mathbf{R}_{\text{GI}} \ \mathbf{g}_{\text{tGI}} \ \mathbf{g}_{\text{vI}} \ \mathbf{b}_a \ \mathbf{b}_g \ \mathbf{g}_{\text{g}}] \in \text{SO}(3) \times \mathbb{R}^{12} \times S^2$ where $\mathbf{R}_{\text{GI}}$ denotes the rotation from the IMU frame I to the global frame G, $\mathbf{g}_{\text{tGI}}$ and $\mathbf{g}_{\text{vI}}$ are the IMU position and velocity expressed in the global frame, $\mathbf{b}_a$ and $\mathbf{b}_g$ represent the accelerometer and gyroscope biases, and $\mathbf{g}_{\text{g}}$ is the gravity direction in the global frame.

The majority of LIO systems follow a tightly-coupled formulation, motivated by findings [39] that IMU measurements can temporarily compensate for missing geometric constraints in short degenerate segments.

However, tight coupling introduces a weighting between modalities, requiring tuning of sensor specific noise parameters. In extended low-structure regions, we observed that inertial residuals can dominate the optimization, potentially overwhelming subtle but informative registration cues. As our work specifically aims to exploit such geometric detail, we propose a loosely-coupled strategy that preserves the IMU's benefits in motion compensation and registration initialization, while isolating the pose estimation from inertial residual influence. Specifically, $\mathbf{T}_{\text{GI}}$ is computed only from registration (Section III-C2), without any inertial residuals or covariances, and the remainder of the state is optimized separately (III-C3).

1) *IMU-Based State Prediction & Motion Compensation*: We adopt the piecewise linear continuous-time model of [40] to compute preintegrated rotation and translation increments [21] between the IMU timestamps t_j and t_k associated with two consecutive LiDAR scans, yielding $\Delta \mathbf{R}_{j,k} \in \text{SO}(3)$ and $\Delta \mathbf{p}_{j,k} \in \mathbb{R}^3$. These increments are combined with the previous state to form $\mathbf{T}_{\text{I}^k \text{I}^j}$, which is used to undistort a LiDAR points from t_i to the most recent timestamp t_k .

During this step, points are also transformed to the IMU frame for downstream processing: $\text{I}^k \mathbf{p}_i = \mathbf{T}_{\text{I}^k \text{I}^i} \mathbf{T}_{\text{I}^i \text{L}^i} \text{L}^i \mathbf{p}_i$, where $\mathbf{T}_{\text{I}^i \text{L}^i}$ denotes the calibration from the LiDAR frame to the IMU frame. The same increments are used to predict a registration prior $\mathbf{T}_{\text{GI}^k}$. Finally, they also define the inertial residuals in the optimization described in Sec. III-C.

2) *Registration*: We align the sampled point cloud directly to the map by minimizing the error between the projected heights of input points and the value of their respective pixel in the voxel height-map. This process is inspired by CURL-SLAM [33] and can be thought of as a photometric minimiza-




Fig. 4. Point sampling in Shield1. Gray points show an accumulated point map. Based on the MID metric, dense points (orange) are sampled in salient regions such as tunnel entrance corners and the track bed, while sparse points (green) are sampled in less informative regions such as the ceiling and floor.

tion in a virtual image, where the pixel values represent the height above the voxel image plane, as shown in Figure 3. Each input point $\text{I}^k \mathbf{p}_i$ is mapped to a voxel by hashing its coordinates based on the initial pose estimate. The residual is defined as:

$$r_i(\xi) = [\text{C} \mathbf{p}_i]_z(\xi) - \mathbf{I}(u_i(\xi), v_i(\xi)) \quad (10)$$

$$\text{C} \mathbf{p}_i(\xi) = \mathbf{T}_{\text{CG}} \mathbf{T}_{\text{GI}} \text{Exp}(\xi) \text{I}^k \mathbf{p}_i$$

where $[\text{C} \mathbf{p}_i]_z$ is the point height above the image plane, (u_i, v_i) are the pixel coordinates in that voxel's height image and $\text{Exp}(\xi)$ maps the pose increment $\xi \in \mathfrak{se}(3)$ to a transformation $\mathbf{T} \in \text{SE}(3)$. The image value $\mathbf{I}(u_i, v_i)$ is determined by bilinear pixel interpolation among observed pixels. If a point projects to a region where all four surrounding pixels are unobserved, it is omitted.

The Jacobian of the residual with respect to the pose perturbation ξ can be divided into two components:

$$\frac{\partial r_i}{\partial \xi} = \underbrace{\frac{\partial [\text{C} \mathbf{p}_i]_z}{\partial \xi}}_{\text{I}} - \underbrace{\frac{\partial \mathbf{I}(u_i, v_i)}{\partial \xi}}_{\text{II}} \quad (11)$$

We provide a full expansion of (14) in the supplementary material and want to highlight the geometric interpretation of the two Jacobian components. (I) can be understood as a normal-direction term, equivalent to classical point-to-plane error behavior, and describes how the residual changes as the point moves towards the image plane. Orthogonal to this, the second term (II) provides two Jacobian directions derived from the height image gradient, which describe the residual change from moving the point parallel to the plane, i.e., how the image value changes if the point projects to neighboring pixels. These two additional directions help to augment the optimization with information in directions that are unobservable under point-to-plane alignment, which increases robustness in uninformative regions, as shown in our experiments (Section IV-B).

We solve the resulting non-linear least-squares problem using Levenberg-Marquardt minimization:

$$\xi^* = \arg \min_{\xi} \sum_i \rho(\|r_i(\xi)\|), \quad (12)$$

where $\rho(\cdot)$ denotes the Huber loss function.

3) *Sliding-Window Inertial Optimization*: While poses are determined by LiDAR registration, we incorporate IMU measurements via a sliding-window optimization to enforce temporal inertial consistency using IMU residuals r_{IMU} as in [41]. Since gravity direction and IMU biases evolve slowly [42], we estimate them as window-constant variables rather than per-timestep states. All poses within the 10 s window are fixed and serve only as anchors, yielding the following optimization:

$$[\bar{\mathbf{v}}_{k-n}, \dots, \bar{\mathbf{v}}_k], \bar{\mathbf{g}}, \bar{\mathbf{b}}_a, \bar{\mathbf{b}}_g =$$

$$\arg \min \sum_{i=k-n}^{k-1} \left\| r_{\text{IMU}}(\mathbf{v}_i, \mathbf{v}_{i+1}, \mathbf{g}, \mathbf{b}_a, \mathbf{b}_g, \underbrace{\mathbf{T}_{\text{GI}^i}, \mathbf{T}_{\text{GI}^{i+1}}}_{\text{fixed}}) \right\|^2 \quad (13)$$

This preserves sensitivity to geometric details while estimating biases and avoiding covariance tuning, which improves robustness across sensor configurations and environments.

TABLE I
PARAMETER OVERVIEW

Parameter	Symbol	Value
Voxel side length	l	0.5 m
Bump-Image resolution	r	0.05 m
Downsample resolution coarse / fine	d_c / d_f	0.5 / 0.1 m
# Fine-resolution voxels	—	300
Huber Loss Delta	δ	0.1 m
IMU Optimization Window	—	10 s

IV. EXPERIMENTS

We evaluate our system against a broad set of state-of-the-art LiDAR-based odometry methods. Specifically, we compare against KISS-ICP [11], GenZ-ICP [25], Traj-LO [14], FAST-LIO2 [19], DLIO [24], iG-LIO [27], RESPLE [12] and RKO-LIO [22]. Furthermore, for datasets providing high-resolution scans (Table II III), we also include COIN-LIO [9], which additionally uses LiDAR intensity information. For CURL-SLAM [33], we only list the results (Table II) from the publication, as the code was not publicly available at the time of submission. To assess BIEVR-LIO’s generalization capabilities, we evaluate on multiple publicly available datasets spanning different platforms, LiDAR configurations, and environments. These datasets include geometrically rich scenes as well as challenging environments with limited structural information to analyze performance across a wide range of applications. We use the same parameters, listed in Table I for all experiments, datasets, and sensors without any per-scene or per-dataset tuning, which underlines the robustness of our method. For the baseline methods, we use sensor-specific parameter configurations when available and, otherwise, rely on the publicly available default parameters for each approach.

We evaluate the global trajectory accuracy using the Absolute Trajectory Error (ATE) and the relative error (RE) over trajectory segments of 10 m as an estimate of the local drift. Since the evaluated datasets only contain position ground truth, we calculate the RE using the point distance metric². Trajectory alignment and evaluation are performed using `evo`². Methods with relative error exceeding 20% are considered to have failed (×), and their ATE is not reported, as excessive drift prevents meaningful alignment to the ground truth. We conduct all experiments on an Intel i7-11800H CPU.

TABLE II
NEWER COLLEGE DATASET RESULTS

Absolute Trajectory Error RMSE (m) / Relative Error RMSE (%)				
Method	QuadHard	Cloister	Stairs	Park
KISS-ICP	0.324 / 1.9	0.297 / 2.1	× / 32.5	2.871 / 1.1
GenZ-ICP	× / 28.0	0.137 / 0.7	0.948 / 7.2	0.718 / 0.8
Traj-LO	0.133 / 1.0	0.151 / 0.6	2.043 / 9.3	0.503 / 0.6
CURL-SLAM	— / —	0.212 / —	0.169 / —	0.271 / —
FAST-LIO2	0.049 / 0.3	0.078 / 0.2	× / 3497.2	0.31 / 0.6
DLIO	0.14 / 1.0	0.162 / 0.8	0.122 / 1.0	0.335 / 0.7
iG-LIO	0.079 / 0.5	× / 88.5	× / 81.4	0.201 / 0.5
RESPLE	0.06 / 0.3	0.104 / 0.3	0.194 / 2.2	0.994 / 0.7
RKO-LIO	1.111 / 0.6	1.03 / 2.4	4.477 / 10.8	0.889 / 0.6
COIN-LIO	0.046 / 0.3	0.078 / 0.3	0.102 / 0.7	0.287 / 0.5
BIEVR-LIO	0.051 / 0.3	0.053 / 0.2	0.056 / 0.6	0.798 / 0.6

²<https://github.com/MichaelGrupp/evo>

A. Odometry Results

1) *Newer College Dataset*: We use the Newer College Dataset [43] to evaluate performance in geometrically informative environments and present the results in Table III. All sequences are captured using a handheld Ouster OS0-128 LiDAR with an integrated IMU. The *Cloister*, *Park*, and *Quad-Hard* sequences contain rich structural detail, while the *Stair* sequence is a narrow staircase with constrained geometry. Overall, LiDAR-only approaches show lower performance due to aggressive rotations and abrupt motions in the sequences. In contrast, BIEVR-LIO handles sudden motions robustly, as shown on the Quad-Hard sequence, which has particularly abrupt rotations. This demonstrates that our loosely-coupled design successfully exploits the advantages of IMU undistortion and registration priors. BIEVR-LIO achieves the lowest errors on the *Cloister* and *Stairs* sequences. While most methods struggle on the narrow *Stairs*, our approach effectively uses geometric details even in confined spaces.

2) *ENWIDE Dataset*: The ENWIDE [9] dataset is used to evaluate performance in environments with weak geometric structure. It consists of handheld outdoor sequences using an Ouster OS0-128 with an integrated IMU with both smooth and highly dynamic motions.

Due to the challenging geometry, KISS-ICP, GenZ-ICP, Traj-LO, RESPLE, and RKO-LIO diverged on all sequences and are not included in the results in Table III. Similarly, iG-LIO produced results only for the *KatzenseeD* sequence, while FAST-LIO2 performed well on *FieldS* but drifts elsewhere. DLIO achieves accurate results on the slower sequences but drifts or diverges in dynamic sequences and the extremely sparse *Runway* environment (Figure 5). BIEVR-LIO is the




Fig. 5. Resulting point cloud maps. *Top*: Inside view of the *Shield1* sequence in *Middle*. The accumulated points clearly visualize the railbed and the cutouts in the tunnel wall. *Middle*: Despite the challenging geometry, the map displays no visible drift after the 700 m trajectory. *Bottom*: After a loop around the *RunwayS*, the ground markings remain crisp, illustrating the pose accuracy.

TABLE III
ENWIDE DATASET RESULTS

Absolute Trajectory Error RMSE (m) / Relative Error RMSE (%)								
Method	IntersectionS	IntersectionD	RunwayS	RunwayD	FieldS	FieldD	KatzenseeS	KatzenseeD
FAST-LIO2	12.473 / 29.3	23.8 / 28.1	× / 53.6	× / 59.8	<u>0.163 / 0.6</u>	9.209 / 16.1	1.122 / 4.3	1.02 / 2.4
DLIO	0.177 / 0.7	3.394 / 15.6	× / 48.6	× / 41.2	0.17 / 0.6	0.617 / 3.5	0.188 / 0.5	<u>0.328 / 1.7</u>
iG-LIO	× / 53.2	× / 60.4	× / 61.6	× / 70.3	× / 46.8	× / 49.3	× / 60.8	2.13 / 14.2
RESPLE	× / 72.8	× / 79.6	× / 66.0	× / 68.7	× / 27.0	× / 64.5	× / 51.5	× / 114.9
RKO-LIO	× / 33.0	× / 38.5	× / 43.4	× / 48.0	× / 50.2	× / 36.1	× / 19.1	× / 48.7
COIN-LIO	0.466 / 1.2	<u>1.912 / 1.7</u>	<u>1.033 / 1.9</u>	2.437 / 3.0	0.232 / 0.8	<u>0.581 / 1.8</u>	0.412 / 1.0	0.592 / <u>1.6</u>
BIEVR-LIO	<u>0.231 / 0.7</u>	0.404 / 0.8	0.44 / 2.2	<u>4.35 / 10.8</u>	0.159 / 0.3	0.174 / 0.8	<u>0.194 / 0.3</u>	0.243 / 0.6

TABLE IV
GEODE DATASET RESULTS

Absolute Trajectory Error RMSE (m) / Relative Error RMSE (%)							
Method	Offroad 4	Offroad 5	Shield 1	Shield 4	Shield 5	Tunneling 1	Tunneling 2
GenZ-ICP	× / 38.6	× / 23.9	× / <u>44.2</u>	× / <u>37.5</u>	× / <u>47.9</u>	3.61 / 13.5	2.775 / 13.0
Traj-LO	0.241 / 5.7	× / 33.9	× / 524.4	× / 406.8	× / 173.2	0.339 / 1.9	2.793 / 19.3
FAST-LIO2	<u>0.138 / 1.0</u>	<u>0.142 / 0.8</u>	× / 93.0	× / 78.8	× / 88.3	0.339 / 2.0	0.576 / 4.3
DLIO	2.188 / 13.1	13.507 / 14.2	× / 68.5	× / 95.2	× / 63.9	2.542 / 9.9	× / 59.0
iG-LIO	1.648 / 4.8	0.156 / 0.8	× / 66.3	× / 66.1	× / 396.4	0.345 / <u>1.9</u>	× / 35.5
RESPLE	× / 53.0	× / 51.8	× / 393.1	× / 176.0	× / 147.2	0.45 / 4.8	0.28 / <u>2.6</u>
RKO-LIO	2.142 / 5.1	0.204 / 0.9	× / 64.8	× / 54.1	× / 63.7	0.342 / 2.1	<u>0.23 / 2.8</u>
BIEVR-LIO	0.085 / 0.7	0.116 / 0.7	0.256 / 2.1	0.275 / 2.4	0.146 / 1.3	0.34 / <u>1.9</u>	0.096 / 1.2

single geometry-only method that remains stable in all sequences and achieves the lowest error metrics on most. This demonstrates that our high-resolution mapping and informed sampling effectively leverages subtle geometric cues in geometrically uninformative environments. Notably, despite not using additional intensity information, BIEVR-LIO achieves more accurate odometry than COIN-LIO on all sequences except *RunwayD*, in which the gyroscope is saturated for multiple consecutive frames, causing a rotational jump. Even then, the odometry remains consistent (RE).

3) *GEODE Dataset*: The GEODE [44] dataset, summarized in Table IV, comprises both handheld indoor sequences (*Shield*, *Tunneling*) and UGV-based outdoor sequences (*Offroad*), all characterized by uninformative geometry captured on a Livox Avia LiDAR with an integrated IMU. KISS-ICP degenerates on all sequences and is omitted from the comparison. The *Tunneling* sequences, recorded in a mine, contain some informative structure from uneven tunnel walls and infrastructure, which most methods can exploit to achieve accurate odometry on *Tunneling1*. *Tunneling2* includes a short segment where the LiDAR faces directly towards the tunnel wall, causing several approaches to drift or diverge. BIEVR-LIO is able to use high-resolution irregularities of the stone surface and maintains accurate odometry, achieving the lowest ATE. *Offroad* sequences feature an open grass-field with limited geometric structure and pronounced sensor motion, resulting in poor performance for most methods as only FAST-LIO2 and BIEVR-LIO achieve accurate ATE in both sequences. The most challenging environment in our evaluation is *Shield*, which consists of a straight metro tunnel. Here, longitudinal constraints only arise from small-scale structures such as the train bed and cutouts in the walls. These features are too subtle for most methods to exploit. In contrast, BIEVR-LIO captures these fine-grained geometric details and maintains

robust odometry over the 700 m *Shield1* sequence (Figure 5), even under fast rotational motion (*Shield4*) or when the sensor directly faces the tunnel wall (*Shield5*).

4) *MARS-LVIG Dataset*: The MARS-LVIG dataset [45] consists of high-altitude UAV flights captured using a downward-facing Livox Avia with an integrated IMU. The sequences involve high velocities of up to 6 m/s and cover distances of up to 4.9 km. In Table V, we evaluate FAST-LIO2, iG-LIO, and BIEVR-LIO, as other methods failed during the take-off phase, where the LiDAR observes too few points for reliable state estimation. Despite its high-resolution map, BIEVR-LIO robustly handles these sequences, in which the LiDAR point clouds are sparse due to flight altitudes of approximately 80 m. It achieves performance comparable to the baselines on the *AMtown02* and *HKislandGNSS02* sequences. In the *FeaturelessGNSS02* sequence, the UAV flies at low altitude over flat terrain, causing both iG-LIO and FAST-LIO2 to diverge. Owing to its fine-grained representation, BIEVR-LIO is still able to maintain stable odometry.

5) *GrandTour Dataset*: Finally, we evaluate the GrandTour [46] dataset, which features a quadruped platform in indoor and outdoor environments. The data exhibits slow yet shaky motion from legged locomotion and is captured using a Hesai XT32 LiDAR paired with a Honeywell HG4930 IMU. As this experiment is intended solely to demonstrate generalization to legged platforms and lower resolution LiDAR,

TABLE V
MARS-LVIG DATASET RESULTS

Absolute Trajectory Error RMSE (m) / Relative Error RMSE (%)			
Method	AMtown02	HKislandGNSS02	FeaturelessGNSS02
FAST-LIO2	<u>3.471 / 1.3</u>	0.884 / 1.4	× / 32.1
iG-LIO	2.155 / 1.7	0.2 / 0.4	× / 25.6
BIEVR-LIO	4.0 / 1.2	<u>0.809 / 0.4</u>	4.93 / 12.9

TABLE VI
GRANDTOUR DATASET RESULTS

Absolute Trajectory Error RMSE (m) / Relative Error RMSE (%)				
Method	ARC-4	LEE-1	LEICA-2	TRIM-1
FAST-LIO2	0.872 / 2.6	2.949 / 7.7	0.027 / 0.3	0.019 / 0.1
DLIO	0.052 / 0.6	0.059 / 0.4	0.058 / 0.7	0.053 / 0.3
iG-LIO	$\times$ / 32.1	$\times$ / 93.8	0.051 / 0.3	0.023 / 0.2
RESPLE	0.067 / 0.5	0.107 / 0.4	0.049 / 0.4	0.038 / 0.2
RKO-LIO	0.582 / 1.1	0.048 / 0.3	0.032 / 0.2	0.031 / 0.2
BIEVR-LIO	0.051 / 0.6	0.024 / 0.1	0.015 / 0.1	0.028 / 0.2

we restrict the evaluation in Table VI to LiDAR-inertial approaches. On the *LEICA-2* and *TRIM-1* sequences, all methods achieve accurate odometry. In contrast, the confined sections present in *ARC-4* and *LEE-1* lead to drift or divergence for FAST-LIO2 and iG-LIO. BIEVR-LIO, however, remains accurate by exploiting fine-grained details, resulting in the most consistent performance across the sequences.

B. Ablation Study

In Table VIII, we present an ablation study to analyze the impact of our contributions. We perform representative experiments on the geometrically rich *Cloister*, *FieldS* with sparse structure, and the highly ambiguous *Shield1* sequence containing only subtle geometric constraints. To isolate the effect of the proposed BIEVR-Map representation, we compare three configurations: first, using point-to-plane alignment against the principal voxel planes (Plane $\times$ |HR); second, also performing point-to-plane alignment, but using each BIEVR pixel as a high-resolution plane (Bump $\times$ |HR), without the additional Jacobian direction. In this case, the registration ignores neighboring pixels, effectively treating each pixel as an infinite plane. The third configuration (Plane $\checkmark$ |HR), enables the additional Jacobian direction, allowing the registration to fully exploit the surface variations captured by our map. In the well-constrained *Cloister* environment, all configurations perform similarly, as the largely planar structure is well represented by both map formulations. In *FieldS*, the additional geometric detail provided by BIEVR leads to improved accuracy. In the most challenging *Shield1* sequence, however, the benefits of BIEVR can only be exploited when the additional Jacobian direction is used, highlighting the importance of our registration formulation. We further evaluate the proposed map-informed dual-resolution sampling (ID) and compare it against fixed high-resolution sampling (HR) as well as a random dual-resolution strategy (RD) using the same resolutions (0.1 m and 0.5 m), but selecting high-resolution voxels at random rather than based on map structure. Across all sequences, RD performs worst, as high-resolution samples are often allocated to uninformative regions. Our proposed map-informed sampling consistently outperforms both baselines by focusing computation on informative areas, effectively improving the signal-to-noise ratio, while reducing the number of points used for registration by approximately a factor of four compared to HR sampling, resulting in lower runtime. A visual comparison of the runtime can be found in the appendix.

TABLE VII
ABLATION RESULTS

Map	$\partial \mathbf{I} / \partial \xi$	Sampling	Shield1	FieldS	Cloister
Absolute Trajectory Error RMSE (m)					
Plane	$\times$	HR	$\times$	0.192	0.051
BIEVR	$\times$	HR	$\times$	0.17	0.058
BIEVR	$\checkmark$	HR	0.401	0.163	0.059
BIEVR	$\checkmark$	RD	0.604	0.218	0.113
BIEVR	$\checkmark$	ID	0.256	0.159	0.053
Runtime per Frame (ms) / Registration Points (#)					
BIEVR	$\checkmark$	HR	15.7 / 13907	17.2 / 20423	29.8 / 33320
BIEVR	$\checkmark$	ID	12.8 / 5039	12.3 / 5187	22.6 / 7574

C. Runtime Analysis

As the runtime depends on the resolution of the LiDAR, as well as the surrounding environment, we report the average runtime per frame for four representative sequences. The *Cloister* and *FieldS* sequences were recorded using an Ouster OS-0 128-beam LiDAR, *Shield1* was captured with a Livox Avia, and *LEE-1* was recorded using a Hesai XT32. The highest runtime is observed for the *Cloister* sequence, which features large geometry due to tall surrounding buildings and is recorded with a high-resolution LiDAR. Despite being recorded with the same sensor, the runtime in *FieldS* is roughly half of that, caused by very few LiDAR returns in the direction of the sky in the mostly flat environment. The lowest runtime is achieved on *LEE-1*, which uses a lower-resolution 32-beam LiDAR. As shown in Figure 6, most of the computation time across all sequences is spent in the combined state update, which consists of the inertial update and the registration step. The time spent on map updates and sampling scales with the number of observed voxels and the number of input points. Notably, the map update, and thus also the undistortion, is performed on all input points. The inertial update runtime remains approximately constant, as it depends only on the window length, which is fixed to 10 s throughout our experiments.




Fig. 6. Runtime breakdown of BIEVR-LIO.




Fig. 7. Locomotion planning based on BIEVR-LIO. The robot successfully traverses the obstacle. The extracted elevation map is partially visualized.

D. Elevation Mapping

We demonstrate the practicality of BIEVR-LIO for downstream applications in real-world locomotion experiments with a quadruped robot. The robot’s locomotion policy usually relies on an elevation map generated from depth cameras to plan its footsteps. In our experiments, we instead use a Livox Mid-360 LiDAR mounted at the front of the robot. We run BIEVR-LIO on the robot’s onboard computer and export the map in the format expected by the locomotion policy within a 4×4 m area around the robot. The generation and export of this elevation map takes approximately 1 ms. In the first experiment, visualized in Figure 7, the robot traverses an obstacle of stacked pallets. The proposed map representation provides sufficient detail for the robot to correctly place its feet on the surfaces of varying height. In the second experiment, shown in the supplementary video, the robot climbs a staircase, accurately placing its feet on each step. These results indicate that BIEVR-LIO captures sufficient geometric detail for locomotion planning and can serve as an alternative to separately constructed elevation maps, requiring only minimal additional computation beyond the odometry pipeline.

V. LIMITATIONS

While the fine-grained resolution of BIEVR-LIO enables the use of subtle geometric details for registration, it also introduces some limitations. Sufficient point density is required to populate pixels, as no infilling is applied and empty pixels cannot be used, which may reduce performance for very low-resolution LiDARs or at high sensor velocities. Additionally, a reasonably accurate initial pose is needed to compute meaningful pixel residuals. Both issues could be mitigated using multi-scale image pyramids.

Our approach also does not exploit LiDAR intensity, making it less effective in strictly degenerate environments where geometry alone is insufficient, such as the *Tunnel* scenario in the ENWIDE dataset. Intensity-based methods like COIN-LIO can still operate reliably in these cases.

While our approach reduces degeneracy, it does not explicitly detect it. As such, it could be naturally enhanced with existing degeneracy mitigation strategies as in [5].

Finally, the system does not detect or model dynamic objects, which limits performance for downstream tasks, as the elevation mapping presented in Section IV-D.

VI. CONCLUSION

In this work, we presented BIEVR-LIO, a novel LIO framework built around a high-resolution yet efficient map representation that captures subtle geometric details using voxel-based height images. The proposed pipeline further introduces a map-informed point sampling strategy that focuses computation on informative regions, improving robustness while reducing runtime compared to a global high-resolution sampling. Experiments across multiple sensors, platforms, and environments demonstrate that our approach significantly improves performance in geometrically challenging scenarios, enabling accurate registration where existing methods fail and achieving state-of-the-art results in well-structured environments. Notably, all experiments were conducted using the same parameter set, highlighting the robustness and generalizability of the proposed method across diverse conditions. In addition, we show that the fine-grained geometry encoded in BIEVR-Map is well suited for elevation mapping, while remaining computationally efficient enough to operate in real-time. In future work, we plan to incorporate LiDAR intensity information to further extend the range of environments in which BIEVR-LIO can operate reliably.

REFERENCES

- [1] Kamak Ebadi, Lukas Bernreiter, Harel Biggie, Gavin Catt, Yun Chang, Arghya Chatterjee, Christopher E Denniston, Simon-Pierre Deschênes, Kyle Harlow, Shehryar Khattak, et al. Present and future of slam in extreme environments: The darpa sub challenge. *IEEE Transactions on Robotics*, 40:936–959, 2023.
- [2] Dongjae Lee, Minwoo Jung, Woosong Yang, and Ayoun Kim. Lidar odometry survey: recent advancements and remaining challenges. *Intelligent Service Robotics*, 17(2):95–118, 2024.
- [3] Julian Nubert, Etienne Walther, Shehryar Khattak, and Marco Hutter. Learning-based localizability estimation for robust lidar localization. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 17–24. IEEE, 2022.
- [4] Turcan Tuna, Julian Nubert, Yoshua Nava, Shehryar Khattak, and Marco Hutter. X-icp: Localizability-aware lidar registration for robust localization in extreme environments. *IEEE Transactions on Robotics*, 40:452–471, 2023.
- [5] Junjie Huang, Yunzhou Zhang, Qingdong Xu, Song Wu, Jun Liu, Guiyuan Wang, and Wei Liu. La-lío: Robust localizability-aware lidar-inertial odometry for challenging scenes. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10145–10152. IEEE, 2024.
- [6] Eungchang Mason Lee, Kevin Christiansen Marsim, and Hyun Myung. Lodestar: Degeneracy-aware lidar-inertial odometry with adaptive schmidt-kalman filter and data exploitation. *IEEE Robotics and Automation Letters*, 11(1):922–929, 2026.

- [7] Nikhil Khedekar and Kostas Alexis. Pg-lío: Photometric-geometric fusion for robust lidar-inertial odometry. *arXiv preprint arXiv:2506.18583*, 2025.
- [8] Yanfeng Zhang, Yunong Tian, Wanguo Wang, Guodong Yang, Zhishuo Li, Fengshui Jing, and Min Tan. Ri-lío: Reflectivity image assisted tightly-coupled lidar-inertial odometry. *IEEE Robotics and Automation Letters*, 8(3): 1802–1809, 2023.
- [9] Patrick Pfreundschuh, Helen Oleynikova, Cesar Cadena, Roland Siegwart, and Olov Andersson. Coin-lío: Complementary intensity-augmented lidar inertial odometry. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1730–1737. IEEE, 2024.
- [10] Ji Zhang, Sanjiv Singh, et al. Loam: Lidar odometry and mapping in real-time. In *Robotics: Science and Systems*, volume 2, pages 1–9. Berkeley, CA, 2014.
- [11] Ignacio Vizzo, Tiziano Guadagnino, Benedikt Mersch, Louis Wiesmann, Jens Behley, and Cyrill 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*, 8(2):1029–1036, 2023.
- [12] Ziyu Cao, William Talbot, and Kailai Li. Resple: Recursive spline estimation for lidar-based odometry. *IEEE Robotics and Automation Letters*, 10(10):10666–10673, 2025.
- [13] Dongjiao He, Wei Xu, Nan Chen, Fanze Kong, Chongjian Yuan, and Fu Zhang. Point-lío: robust high-bandwidth light detection and ranging inertial odometry. *Advanced Intelligent Systems*, 5(7):2200459, 2023.
- [14] Xin Zheng and Jianke Zhu. Traj-lío: In defense of lidar-only odometry using an effective continuous-time trajectory. *IEEE Robotics and Automation Letters*, 9(2): 1961–1968, 2024.
- [15] Keenan Burnett, Angela P Schoellig, and Timothy D Barfoot. Continuous-time radar-inertial and lidar-inertial odometry using a gaussian process motion prior. *IEEE Transactions on Robotics*, 2024.
- [16] Kenji Koide, Masashi Yokozuka, Shuji Oishi, and Atsuhiko Banno. Glim: 3d range-inertial localization and mapping with gpu-accelerated scan matching factors. *Robotics and Autonomous Systems*, 179:104750, 2024.
- [17] Shibo Zhao, Hengrui Zhang, Peng Wang, Lucas Nogueira, and Sebastian Scherer. Super odometry: Imu-centric lidar-visual-inertial estimator for challenging environments. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8729–8736. IEEE, 2021.
- [18] Tixiao Shan, Brendan Englot, Drew Meyers, Wei Wang, Carlo Ratti, and Daniela Rus. Lío-sam: Tightly-coupled lidar inertial odometry via smoothing and mapping. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5135–5142. IEEE, 2020.
- [19] Wei Xu, Yixi Cai, Dongjiao He, Jiarong Lin, and Fu Zhang. Fast-lío2: Fast direct lidar-inertial odometry. *IEEE Transactions on Robotics*, 38(4):2053–2073, 2022.
- [20] Chungge Bai, Tao Xiao, Yajie Chen, Haoqian Wang, Fang Zhang, and Xiang Gao. Faster-lío: Lightweight tightly coupled lidar-inertial odometry using parallel sparse incremental voxels. *IEEE Robotics and Automation Letters*, 7(2):4861–4868, 2022.
- [21] Christian Forster, Luca Carlone, Frank Dellaert, and Davide Scaramuzza. On-manifold preintegration for real-time visual-inertial odometry. *IEEE Transactions on Robotics*, 33(1):1–21, 2016.
- [22] Meher V. R. Malladi, Tiziano Guadagnino, Luca Lobe-faro, and Cyrill Stachniss. Robust approach for lidar-inertial odometry without sensor-specific modeling. *IEEE Robotics and Automation Letters*, 11(6):7420–7427, 2026. doi: 10.1109/LRA.2026.3685966.
- [23] Shibo Zhao, Honghao Zhu, Yuanjun Gao, Beomsoo Kim, Yuheng Qiu, Aaron M Johnson, and Sebastian Scherer. Superloc: The key to robust lidar-inertial localization lies in predicting alignment risks superodometry. com/superloc. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 14080–14086. IEEE, 2025.
- [24] Kenny Chen, Ryan Nemiroff, and Brett T Lopez. Direct lidar-inertial odometry: Lightweight lío with continuous-time motion correction. *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3983–3989, 2023.
- [25] Daehan Lee, Hyungtae Lim, and Soohye Han. Genz-icp: Generalizable and degeneracy-robust lidar odometry using an adaptive weighting. *IEEE Robotics and Automation Letters*, 10(1):152–159, 2025.
- [26] Chongjian Yuan, Wei Xu, Xiyuan Liu, Xiaoping Hong, and Fu Zhang. Efficient and probabilistic adaptive voxel mapping for accurate online lidar odometry. *IEEE Robotics and Automation Letters*, 7(3):8518–8525, 2022.
- [27] Zijie Chen, Yong Xu, Shenghai Yuan, and Lihua Xie. ig-lío: An incremental gicp-based tightly-coupled lidar-inertial odometry. *IEEE Robotics and Automation Letters*, 9(2):1883–1890, 2024.
- [28] Jens Behley and Cyrill Stachniss. Efficient surfel-based slam using 3d laser range data in urban environments. In *Robotics: Science and Systems*, volume 2018, page 59, 2018.
- [29] Jan Quenzel and Sven Behnke. Lío-mars: Non-uniform continuous-time trajectories for real-time lidar-inertial odometry. *arXiv preprint arXiv:2511.13985*, 2025.
- [30] Thien-Minh Nguyen, Daniel Duberg, Patric Jensfelt, Shenghai Yuan, and Lihua Xie. Slic: Multi-input multi-scale surfel-based lidar-inertial continuous-time odometry and mapping. *IEEE Robotics and Automation Letters*, 8(4):2102–2109, 2023.
- [31] L. Coto-Elena, J.E. Maese, L. Merino, and F. Caballero. D-lío: 6dof direct lidar-inertial odometry based on simultaneous truncated distance field mapping. *IEEE Robotics and Automation Letters*, 11(1):169–176, 2026.
- [32] Cedric Le Gentil, Raphael Falque, Daniil Lisus, and Timothy D. Barfoot. 2fast-2lamaa: Large-scale lidar-inertial

- localization and mapping with continuous distance fields. *arXiv preprint arXiv:2410.05433*, 2025.
- [33] Kaicheng Zhang, Shida Xu, Yining Ding, Xianwen Kong, and Sen Wang. Curl-slam: Continuous and compact lidar mapping. *IEEE Transactions on Robotics*, 2025.
 - [34] Hyungtae Lim, Daebeom Kim, Beomsoo Kim, and Hyun Myung. Adalio: Robust adaptive lidar-inertial odometry in degenerate indoor environments. In *2023 20th International Conference on Ubiquitous Robots (UR)*, pages 48–53. IEEE, 2023.
 - [35] Pavel Petracek, Kostas Alexis, and Martin Saska. Rms: redundancy-minimizing point cloud sampling for real-time pose estimation. *IEEE Robotics and Automation Letters*, 9(6):5230–5237, 2024.
 - [36] Jianhao Jiao, Yilong Zhu, Haoyang Ye, Huaiyang Huang, Peng Yun, Linxin Jiang, Lujia Wang, and Ming Liu. Greedy-based feature selection for efficient lidar slam. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5222–5228. IEEE, 2021.
 - [37] Wei Li, Yu Hu, Yinhe Han, and Xiaowei Li. Kfs-lio: Key-feature selection for lightweight lidar inertial odometry. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5042–5048, 2021.
 - [38] Guy M Morton. *A computer oriented geodetic data base and a new technique in file sequencing*. International Business Machines Company, 1966.
 - [39] Haoyang Ye, Yuying Chen, and Ming Liu. Tightly coupled 3d lidar inertial odometry and mapping. In *2019 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3144–3150. IEEE, 2019.
 - [40] Cedric Le Gentil and Teresa Vidal-Calleja. Continuous latent state preintegration for inertial-aided systems. *The International Journal of Robotics Research*, 42(10):874–900, 2023.
 - [41] Tong Qin, Peiliang Li, and Shaojie Shen. Vins-mono: A robust and versatile monocular visual-inertial state estimator. *IEEE Transactions on Robotics*, 34(4):1004–1020, 2018.
 - [42] Naser El-Sheimy, Haiying Hou, and Xiaoji Niu. Analysis and modeling of inertial sensors using allan variance. *IEEE Transactions on Instrumentation and Measurement*, 57(1):140–149, 2008.
 - [43] Lintong Zhang, Marco Camurri, David Wisth, and Maurice Fallon. Multi-camera lidar inertial extension to the newer college dataset. *arXiv preprint arXiv:2112.08854*, 2021.
 - [44] Zhiqiang Chen, Yuhua Qi, Dapeng Feng, Xuebin Zhuang, Hongbo Chen, Xiangcheng Hu, Jin Wu, Kelin Peng, and Peng Lu. Heterogeneous lidar dataset for benchmarking robust localization in diverse degenerate scenarios. *The International Journal of Robotics Research*, 45(1):6–22, 2026.
 - [45] Haotian Li, Yuying Zou, Nan Chen, Jiarong Lin, Xiyuan Liu, Wei Xu, Chunran Zheng, Rundong Li, Dongjiao He, Fanze Kong, et al. Mars-lvig dataset: A multi-sensor aerial robots slam dataset for lidar-visual-inertial-gnss fusion. *The International Journal of Robotics Research*, 43(8):1114–1127, 2024.
 - [46] Jonas Frey, Turcan Tuna, Lanke Frank Tarimo Fu, Cedric Weibel, Katharine Patterson, Benjamin Krummenacher, Matthias Müller, Julian Nubert, Maurice Fallon, Cesar Cadena, and Marco Hutter. Boxi: Design Decisions in the Context of Algorithmic Performance for Robotics. In *Robotics: Science and Systems*, Los Angeles, United States, July 2025.
 - [47] Joan Sola, Jeremie Deray, and Dinesh Atchuthan. A micro lie theory for state estimation in robotics. *arXiv preprint arXiv:1812.01537*, 2018.