Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Reliable LiDAR–inertial Odometry-based 3D Mapping in Long Straight Degenerated Mining Tunnels with Kinematic Regularization

Yikun Cheng
China University of Mining and Technology - Beijing
Shaocong Zhang
China University of Mining and Technology - Beijing
Xingyuan Miao
China University of Mining and Technology - Beijing
Fanxin Wang
Fanxin.Wang@xjtlu.edu.cn

Xi'an Jiaotong-Liverpool University

Research Article

Keywords:

Posted Date: November 24th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11764029/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Reliable LiDAR–inertial Odometry-based 3D Mapping in Long Straight Degenerated Mining Tunnels with Kinematic Regularization

Yikun Cheng[1], Shaocong Zhang[1], Xingyuan Miao[1], Fanxin Wang[2*]

[1]School of Mechanical and Electrical Engineering, China University of Mining & Technology-Beijing, Ding No.11 Xueyuan Road, Beijing, 100083, P. R. China.

[2]Academy of Artificial Intelligence and Advanced Technology, Xi'an Jiaotong-Liverpool University, No. 111, Ren'ai Road, Suzhou, 215123, Jiangsu Province, P. R. China.

*Corresponding author(s). E-mail(s): Fanxin.Wang@xjtlu.edu.cn;
Contributing authors: yk.cheng@cumtb.edu.cn; zhshc0106@163.com; miaoxingyuan96@126.com;

Abstract

Autonomous mining robots depend on reliable localization and mapping to operate in underground environments. Mining tunnels support the transportation of personnel, equipment, and materials, provide ventilation, and connect working areas. These networks are often dominated by long, straight sections with limited geometric variation along their axes, providing weak constraints on axial motion and challenging the accuracy of LiDAR–inertial odometry. To address this problem, we propose a LiDAR–inertial odometry framework that incorporates degeneracy-aware kinematic regularization. The framework characterizes translational geometric information through a regularized Schur complement constructed from the anisotropic components of correspondence precision matrices, accounting for coupling between rotation and translation. The normalized spectrum of this information defines a continuous deficiency operator that assigns greater weight to relatively weak directions without requiring a hard degeneracy threshold. The operator from the preceding scan weights a body-frame velocity-smoothness prior referenced to the previous posterior estimate within an iterated error-state Kalman filter. We compare the proposed method with four baseline methods on eleven sequences from the public GEODE γ tunnel benchmark and

assess mapping performance on two self-collected tunnel datasets. Among the reported results for the six shield-tunnel sequences, the proposed method achieves the lowest absolute trajectory error (ATE) on five sequences and the lowest translational relative pose error (RPE) on all six. The framework uses only LiDAR and inertial measurements, with a mean processing time of **74.3** ms per scan on the evaluation platform.

Keywords: Autonomous mining systems, Underground mining, LiDAR–inertial odometry, Geometric degeneracy, Robotic perception

1 Introduction

Robotic systems have been developed for shaft inspection [1], autonomous haulage [2, 3], and search-and-rescue support [4] in underground coal mines. For autonomous mobile platforms, reliable onboard localization supports navigation, while consistent maps provide a spatial reference for planning and monitoring. Simultaneous localization and mapping (SLAM) [5, 6] addresses these needs by jointly estimating a robot's pose and a map of its environment, providing a basis for navigation and control. A quantitative comparison of SLAM systems in an underground test field shows that map accuracy and reliability vary substantially across algorithms, affecting their suitability for geomonitoring [7]. Underground environments present additional challenges because satellite positioning is unavailable, illumination is often poor, and airborne dust and fog can degrade the performance of visible-light cameras [8, 9]. These conditions motivate the use of light detection and ranging (LiDAR) for localization and mapping, while practical deployments must operate within fixed computational and memory budgets [10]. LiDAR-only odometry estimates motion through scan registration [11–13]. LiDAR–inertial odometry (LIO) combines LiDAR measurements with data from an inertial measurement unit (IMU) to predict motion between scans, compensate for scan distortion, and provide a prior for registration [14–16]. Despite these advances, reliable localization in mining tunnels and other feature-poor environments remains a challenge.

LIO estimates each pose by aligning a new scan with a map built from earlier scans [17, 18]. In a roadway with a near-constant cross-section, the ribs, roof and floor face inward, so their surface normals are perpendicular to the roadway axis (Fig. 1). Shifting a scan across the roadway pushes points into those surfaces. Sliding it along the axis leaves every point on the same surface at the same distance [19–21]. The registration residual therefore changes little under axial displacement, and a small residual does not imply an accurate axial position. Degeneracy of this kind arises from insufficient constraint along specific directions, not from a lack of points [22–24]. The axial estimate falls back on inertial integration, drift along the roadway accumulates, and the failure is silent because the scan still fits the walls. Three major approaches address this issue from different aspects of SLAM.

The first approach identifies weak directions and controls the update along them. Zhang et al. and other researchers [22, 25, 26] analyze the constraint directions of

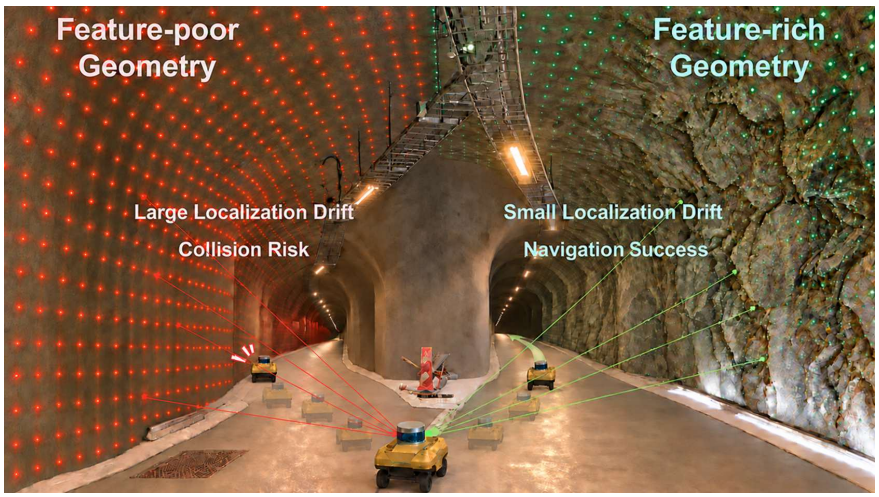


Fig. 1: Geometric degeneracy in underground tunnels. Conventional LIO methods are susceptible to localization drift in geometrically feature-poor environments. DARK-LIO is designed to mitigate this drift while maintaining computational efficiency in both feature-poor and feature-rich environments.

the optimization and retain the solution update mainly in well-constrained directions. Along the restricted directions the estimate then depends on the initial value or prior, which does not recover missing observations. DCReg [27] characterizes degeneracy with a Schur-complement analysis that accounts for rotation–translation coupling and improves the numerical conditioning of the solver with a structured preconditioner. However, it cannot restore geometric constraint absent from the scan, and the objective is left unchanged. These methods require a decision about which directions count as degenerate, and they act on the solver rather than on the information available along the weak direction.

The second approach relies on extracting more from the measurements already available. LIO-SAM [14] keeps a window of past keyframes in a factor graph with preintegrated IMU factors, so that earlier scans continue to constrain the current pose. GLIM [16] does the same through fixed-lag smoothing with GPU-accelerated registration factors, so that several scans contribute jointly to the local estimate. In a straight roadway the retained scans see the same parallel surfaces as the current one, so the window strengthens transverse constraint without adding axial constraint. AdaLIO [28] adapts voxelization and normal-estimation parameters when a narrow environment yields too few correspondences. The added correspondences lie on the same walls and therefore strengthen the same directions. GenZ-ICP [29] blends point-to-plane and point-to-point residuals according to local geometry, which improves the conditioning of the registration for a fixed correspondence set but does not by itself establish correct axial correspondences. COIN-LIO [30] uses the intensity channel of the same scan and selects intensity patches along geometrically weak directions. This requires a LiDAR dense enough to form intensity images and surfaces with trackable

intensity structure. These methods improve how the already available information is used, but a direction left unobserved for a long time is not thereby recovered[31].

The third approach focuses on modeling motion within and between scans. FAST-LIO2 [15] propagates the IMU measurements to predict and deskew each scan and uses the propagated state as the prior of an iterated Kalman update. KISS-ICP [11] does the same with a constant-velocity model and no IMU at all. In both, the motion model enters with a fixed weight in every direction. DLIO [32] combines continuous-time coarse-to-fine motion correction with direct scan-to-map generalized iterative closest point (GICP). More accurate deskewing resolves the timing model but does not provide an independent observation along the axis of a uniform tunnel. Point-LIO [33] updates the state point by point and treats IMU readings as measurements of a motion state, enabling it to handle highly dynamic motion and IMU saturation. However, handling rapid motion does not by itself resolve the weak geometric constraints along the axis of a long, straight tunnel. A learning-based method named ALIVE-LIO [34] learns body-frame velocity from inertial data to assist an error-state Kalman filter along degenerate directions. The learned velocity shares the IMU and state compensation with the LIO, and its validity depends on the training conditions. What these leave open is how a motion prior should be allocated across directions by the geometry.

In this paper, we present Degeneracy-Aware Regularized Kinematics for LiDAR–Inertial Odometry (DARK-LIO), which combines the identification of weak geometric directions with a motion prior using only onboard LiDAR and IMU measurements. The prior is referenced to the previous posterior estimate and represents a motion-model constraint rather than an independent observation. It penalizes changes in body-frame velocity between consecutive scans, with directional weights determined by the geometry and overall strength controlled by a single scalar. The contributions are threefold. First, we derive a continuous directional deficiency operator from a registration information matrix assembled from the anisotropic parts of the correspondence precisions. A regularized Schur complement accounts for rotation–translation coupling, and the normalized translational spectrum assigns greater weight to relatively weaker directions without requiring a degeneracy threshold. The operator is invariant to uniform positive scaling of the translational information and transforms consistently under rotations of the global frame. Second, we introduce an additive kinematic regularizer that uses the operator computed from the preceding scan to weight a body-frame velocity-smoothness prior within an iterated error-state Kalman filter (IESKF). The regularizer contributes positive semidefinite curvature to the normal equations while retaining the LiDAR information term at each fixed linearization state and correspondence set. Third, we evaluate the method on shield and mined tunnel sequences, supported by ablation and sensitivity analyses and mapping experiments on self-collected data. The directional error analysis shows small cross-axis errors despite substantially larger overall trajectory errors, which is consistent with predominantly axial drift. Reporting the two tunnel families separately allows us to examine how the regularizer affects performance under different levels of axial geometric constraint. Removing the prior increases errors on the shield sequences while producing substantially smaller changes on the mined sequences.

The remainder of this paper is organized as follows. Section 2 presents the DARK-LIO formulation and overall algorithm. Section 3 describes the experimental setup and evaluates localization accuracy, degeneracy handling, computational cost, and map consistency through benchmark tests, ablation and sensitivity analyses, and real-world experiments. The final section summarizes the findings and discusses the assumptions underlying the method and directions for future work.

2 Methodology

2.1 Overview of the Proposed Framework

DARK-LIO augments LiDAR–inertial odometry with a directionally weighted body-frame velocity-smoothness prior to mitigate drift in long, straight tunnels. The front end combines an IESKF with GICP scan-to-map registration. To characterize directional geometric constraints, we construct a separate information matrix from the anisotropic component of each correspondence precision matrix, excluding its direction-independent baseline. A regularized Schur complement of this matrix accounts for rotation–translation coupling and provides a local measure of translational geometric information. The normalized spectrum of this measure defines a continuous deficiency operator that assigns larger weights to relatively weaker directions. The operator computed from the preceding scan weights a body-frame velocity-smoothness regularizer referenced to the previous posterior state. This regularizer adds positive semidefinite (PSD) curvature and a residual-dependent correction to the normal equations, while leaving the LiDAR term unchanged at each linearization point. The formulation uses only LiDAR and IMU data and introduces a motion-model constraint to mitigate drift without requiring the recovery of geometric information absent from the scan.

2.2 State Representation and IMU Propagation

Let G, I, and L denote the global, IMU, and LiDAR frames, respectively. The fixed extrinsic transform ${}^{I}T_{L} = ({}^{I}R_{L}, {}^{I}\mathbf{p}_{L})$ maps a LiDAR point into the IMU frame as $\mathbf{q}^{I} = {}^{I}R_{L}\mathbf{q}^{L} + {}^{I}\mathbf{p}_{L}$, where ${}^{I}\mathbf{p}_{L}$ is the position of the LiDAR origin expressed in I. The global frame is initialized with its positive z-axis pointing in the direction of gravity. The state is defined as

$$\mathbf{x} = (R,\ \mathbf{p},\ \mathbf{v},\ \mathbf{b}_g,\ \mathbf{b}_a,\ \mathbf{g}) \in SO(3) \times \mathbb{R}^{15}, \tag{1}$$

where R maps vectors from I to G, and $\mathbf{p}, \mathbf{v} \in \mathbb{R}^3$ are the position and velocity of the IMU origin expressed in G. The gyroscope and accelerometer biases, $\mathbf{b}_g, \mathbf{b}_a \in \mathbb{R}^3$, are expressed in I, and $\mathbf{g} \in \mathbb{R}^3$ is the gravity vector expressed in G. Gravity alignment sets the initial orientation of the global frame. The corresponding local error state is

$$\delta\mathbf{x} = \left[\delta\boldsymbol{\theta}^{\top}\ \delta\mathbf{p}^{\top}\ \delta\mathbf{v}^{\top}\ \delta\mathbf{b}_g^{\top}\ \delta\mathbf{b}_a^{\top}\ \delta\mathbf{g}^{\top}\right]^{\top} \in \mathbb{R}^{18}. \tag{2}$$

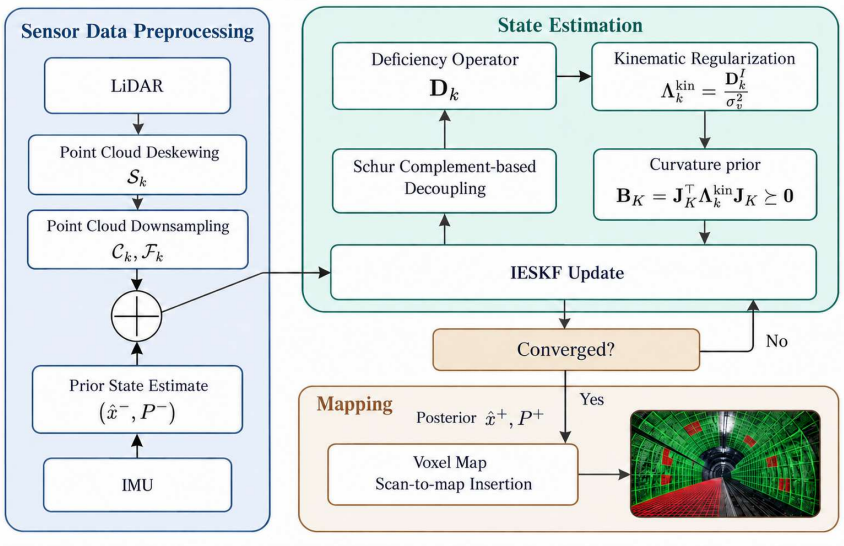


Fig. 2: System overview of DARK-LIO.

We apply a right-multiplicative rotation perturbation, $R \leftarrow R\,\mathrm{Exp}(\delta\boldsymbol{\theta})$, and additive perturbations to the remaining nominal state variables. $\delta\boldsymbol{\theta}$ is the local rotation error vector expressed in the IMU frame, and $\mathrm{Exp}(\mathbf{a}) = \exp(\mathbf{a}_\times)$. The skew-symmetric matrix $\mathbf{a}_\times$ satisfies $\mathbf{a}_\times \mathbf{b} = \mathbf{a} \times \mathbf{b}$ for $\mathbf{a}, \mathbf{b} \in \mathbb{R}^3$. We denote the state update operator by $\boxplus$ and the corresponding local difference operator by $\boxminus$ [35]. The error-state covariance is $\mathbf{P} \in \mathbb{R}^{18\times18}$. Hats denote estimates, and superscripts $-$ and $+$ denote prior and posterior quantities. For a symmetric positive semidefinite matrix $\mathbf{M}$, we use $\|\mathbf{r}\|_{\mathbf{M}}^2 = \mathbf{r}^\top \mathbf{M}\mathbf{r}$. Let $\tilde{\boldsymbol{\omega}}$ and $\tilde{\mathbf{a}}$ denote the measured angular velocity and calibrated specific force, respectively, both expressed in I. Between consecutive IMU samples separated by Δt, the nominal state is propagated as

$$\hat{R}^- \leftarrow \hat{R}\,\mathrm{Exp}\Big((\tilde{\boldsymbol{\omega}} - \hat{\mathbf{b}}_g)\Delta t\Big), \tag{3}$$

$$\hat{\mathbf{p}}^- \leftarrow \hat{\mathbf{p}} + \hat{\mathbf{v}}\Delta t + \tfrac{1}{2}\left[\hat{R}(\tilde{\mathbf{a}} - \hat{\mathbf{b}}_a) + \hat{\mathbf{g}}\right]\Delta t^2, \tag{4}$$

$$\hat{\mathbf{v}}^- \leftarrow \hat{\mathbf{v}} + \left[\hat{R}(\tilde{\mathbf{a}} - \hat{\mathbf{b}}_a) + \hat{\mathbf{g}}\right]\Delta t. \tag{5}$$

These equations [36, 37]s assume constant bias-corrected angular velocity and approximately constant acceleration in the global frame over the integration interval. The biases are modeled as random walks, with their nominal estimates held constant during each propagation step. The gravity vector is also held constant during propagation. Linearizing the complete propagation model, including the bias and gravity dynamics, gives $\delta\mathbf{x}^- = \mathbf{F}\delta\mathbf{x} + \mathbf{n}$, where $\mathbf{F} \in \mathbb{R}^{18\times18}$ is the discrete-time error-state transition matrix and $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$ is the process noise accumulated over the integration interval, with covariance $\mathbf{Q} \in \mathbb{R}^{18\times18}$. The error-state covariance is propagated

as $\mathbf{P}^- = \mathbf{F}\mathbf{P}\mathbf{F}^\top + \mathbf{Q}$. The index k denotes a LiDAR scan. For each scan k, the propagated IMU trajectory and fixed extrinsic transform are used to deskew the LiDAR returns to a common scan reference time Δt_{scan}. In the following derivation, $\mathcal{S}_k$ denotes the deskewed point set expressed in the IMU frame at that reference time. The deskewed point coordinates in the scan-reference IMU frame remain fixed throughout registration.

2.3 Scan-to-Map Registration and the IESKF Update

The registration target is an octree-based voxel map following the mapping setup in [38, 39]. Each accepted source point $\mathbf{q}_i \in \mathcal{S}_k$ is associated with a map voxel whose mean $\boldsymbol{\mu}_i \in \mathbb{R}^3$ and covariance $\boldsymbol{\Sigma}_{M,i} \in \mathbb{R}^{3\times3}$ are expressed in G. Each source point also has a covariance $\boldsymbol{\Sigma}_{S,i} \in \mathbb{R}^{3\times3}$ expressed in the IMU frame at the scan reference time. Following the GICP covariance model [20, 40], we define the target-minus-source residual and its precision matrix as

$$\mathbf{e}_i = \boldsymbol{\mu}_i - (R\mathbf{q}_i + \mathbf{p}), \tag{6}$$

$$\boldsymbol{\Omega}_i = \left(\boldsymbol{\Sigma}_{M,i} + R\boldsymbol{\Sigma}_{S,i}R^\top + \epsilon\mathbf{I}_3\right)^{-1}, \tag{7}$$

where $\mathbf{I}_n$ is the $n \times n$ identity matrix. The positive scalar $\epsilon > 0$ provides isotropic covariance regularization, ensuring that the covariance sum is positive definite and invertible. At each iteration, correspondences and precision matrices are evaluated at the current state estimate and held fixed while solving the linearized system. With the right-multiplicative rotation perturbation of (2), the residual Jacobian with respect to the pose perturbation $\boldsymbol{\xi} = [\delta\boldsymbol{\theta}^\top \ \delta\mathbf{p}^\top]^\top \in \mathbb{R}^6$ is

$$\mathbf{J}_i = \begin{bmatrix} R(\mathbf{q}_i)_\times & -\mathbf{I}_3 \end{bmatrix} \in \mathbb{R}^{3\times6}. \tag{8}$$

Each scan is incorporated into the state estimate through an IESKF update [41]. This update combines the LiDAR residuals with the prior propagated from the IMU. The resulting linear system is solved repeatedly at successive linearization points. We stack the accepted residuals into $\mathbf{r}_L \in \mathbb{R}^{3|\mathcal{C}|}$ and lift the pose Jacobians to the full error state as $\mathbf{J}_L = [\,\mathbf{J}_i \ \ \mathbf{0}_{3\times12}\,]$. We denote the block-diagonal stacked precision matrix by $\boldsymbol{\Omega}$. The LiDAR term contributes the information matrix $\mathbf{H}_L = \mathbf{J}_L^\top\boldsymbol{\Omega}\mathbf{J}_L$ and the negative-gradient vector $\mathbf{b}_L = -\mathbf{J}_L^\top\boldsymbol{\Omega}\mathbf{r}_L$. At iteration ℓ of scan k, the increment is obtained by solving

$$\mathcal{A}_0\,\delta\mathbf{x} = \mathbf{b}_L, \qquad \mathcal{A}_0 = (\mathbf{P}_k^-)^{-1} + \mathbf{H}_L, \qquad \mathbf{P}_k^+ = \mathcal{A}_0^{-1}. \tag{9}$$

The state estimate is updated as $\hat{\mathbf{x}}_k^{(\ell+1)} = \hat{\mathbf{x}}_k^{(\ell)} \boxplus \delta\mathbf{x}$. The residuals and Jacobians are recomputed at each iteration, while the propagated prior covariance $\mathbf{P}_k^-$ remains fixed throughout the iterations. The matrix $\mathcal{A}_0$ is the filter information matrix. Its inverse, $\mathcal{A}_0^{-1}$, is evaluated only at the final iterate and provides a local covariance approximation in the coordinates of $\hat{\mathbf{x}}_k^+$.

2.4 Identification of Translational Deficiency

The update in (9) condenses every correspondence into a single information matrix, in which directions the scan constrains well and directions it constrains weakly are summed together. Placing the kinematic prior where the scan is silent therefore requires a measure that resolves the available translational constraint direction by direction. This measure is the anisotropic part of the correspondence precision. Let $\mathbf{n}_i$ be the surface normal, that is the eigenvector of the largest eigenvalue $\kappa_{n,i}$ of $\mathbf{\Omega}_i$, let $\mathbf{m}_i$ be the eigenvector of its second eigenvalue $\kappa_{m,i}$, and let $\kappa_{t,i} = \lambda_{\min}(\mathbf{\Omega}_i)$ be its smallest eigenvalue. The precision decomposes as

$$\mathbf{\Omega}_i = \kappa_{t,i}\mathbf{I}_3 + (\kappa_{n,i} - \kappa_{t,i})\mathbf{n}_i\mathbf{n}_i^\top + (\kappa_{m,i} - \kappa_{t,i})\mathbf{m}_i\mathbf{m}_i^\top, \tag{10}$$

where $\kappa_{t,i} = \lambda_{\min}(\mathbf{\Omega}_i)$ is the isotropic term that the regularized covariance of (7) leaves under every correspondence, regardless of its orientation, and the last two terms carry the directional shape of the correspondence. The anisotropy is rank two in general and reduces to rank one only when the two in-plane scales coincide, which is the idealised plane model of [40] rather than what a voxel-fitted target covariance paired with a per-point source covariance generally gives.

Rotation and translation are coupled in the pose problem, and ignoring that coupling can overestimate the translational information available when rotation is also uncertain. A plane is unbounded, so $\kappa_{t,i}\mathbf{I}_3$ describes no position along the surface: it is the cost of modelling that surface as a finite patch, and $\kappa_{t,i} = 1/\lambda_{\max}(\mathbf{\Sigma}_{M,i} + R\mathbf{\Sigma}_{S,i}R^\top + \epsilon\mathbf{I}_3)$ is set by the in-plane extent of the patch rather than by ϵ. In the translational block of the information matrix it contributes $\left(\sum_{i\in\mathcal{C}} \kappa_{t,i}\right)\mathbf{I}_3$ regardless of the surface normals, so that $\mathbf{H}_{tt}^{full} = \left(\sum_{i\in\mathcal{C}} \kappa_{t,i}\right)\mathbf{I}_3 + \tilde{\mathbf{H}}_{tt}$ exactly. The reduction does not separate these two parts additively, but the complement it returns obeys $\mathbf{S}_t^{full} \preceq \mathbf{H}_{tt}^{full}$, with equality only when the coupling block of $\mathbf{H}^{full}$ vanishes. The isotropic term is therefore carried into the complement, where it grows with the number of accepted points while the shape of the tunnel says nothing along its axis; forming the complement from the anisotropic remainder instead leaves only the geometric terms. Applying the Schur-complement approach in [27] to the anisotropic remainder $\tilde{\mathbf{\Omega}}_i = \mathbf{\Omega}_i - \kappa_{t,i}\mathbf{I}_3$ eliminates the rotational block and gives the regularized translational complement

$$\begin{aligned} \tilde{\mathbf{H}} &= \sum_{i\in\mathcal{C}} \mathbf{J}_i^\top \tilde{\mathbf{\Omega}}_i \mathbf{J}_i, \qquad \tilde{\mathbf{S}}_t = \tilde{\mathbf{H}}_{tt} - \tilde{\mathbf{H}}_{Rt}^\top (\tilde{\mathbf{H}}_{RR} + \epsilon_R \mathbf{I}_3)^{-1} \tilde{\mathbf{H}}_{Rt}, \\ \epsilon_R &= 10^{-6}\ \mathrm{tr}(\tilde{\mathbf{H}}_{RR})/3, \end{aligned} \tag{11}$$

where $\tilde{\mathbf{\Omega}}_i$ is positive semidefinite. The damping factor ϵ_R is set to 10^{-6} times the mean curvature of the rotational block. The matrix $\tilde{\mathbf{H}}$ is partitioned into the 3×3 blocks $\tilde{\mathbf{H}}_{RR}$, $\tilde{\mathbf{H}}_{Rt}$, $\tilde{\mathbf{H}}_{Rt}^\top$ and $\tilde{\mathbf{H}}_{tt}$, which carry the curvature with respect to rotation and to translation with the other variable held fixed, and the coupling between them. The damping makes the rotational block invertible even when it is rank deficient. When $\tilde{\mathbf{H}}_{RR} = \mathbf{0}$, positive semidefiniteness of $\tilde{\mathbf{H}}$ implies $\tilde{\mathbf{H}}_{Rt} = \mathbf{0}$ and the complement reduces to $\tilde{\mathbf{S}}_t = \tilde{\mathbf{H}}_{tt}$ without an inversion [42]. The result is positive semidefinite

and describes translational curvature in the global frame. For a unit vector $\mathbf{u}$ along the tunnel axis, the along-axis entry of the anisotropic information is $\mathbf{u}^\top \tilde{\mathbf{H}}_{tt}\mathbf{u} = \sum_{i\in\mathcal{C}} \left[(\kappa_{n,i} - \kappa_{t,i})(\mathbf{u}^\top \mathbf{n}_i)^2 + (\kappa_{m,i} - \kappa_{t,i})(\mathbf{u}^\top \mathbf{m}_i)^2\right]$. When every surface normal is perpendicular to $\mathbf{u}$ the first summand vanishes and what remains is the in-plane anisotropy $\kappa_{m,i} - \kappa_{t,i}$, which is zero only under the idealised plane model. The weak direction is identified by the geometry, and its magnitude is read from the spectrum of $\tilde{\mathbf{S}}_t$.

For each direction, the deficiency operator $\mathbf{D}_k$ assigns a dimensionless weight that increases as the scan constrains that direction more weakly. Because $\tilde{\mathbf{S}}_t$ is positive semidefinite, it has eigenvalues $0 \le \tilde{\lambda}_1 \le \tilde{\lambda}_2 \le \tilde{\lambda}_3$ with orthonormal eigenvectors. The eigenvector associated with the smallest eigenvalue identifies the least-constrained translational direction, which is approximately aligned with the tunnel axis in a long, straight tunnel. Normalizing by the largest eigenvalue gives the dimensionless relative eigenvalues $\tilde{\rho}_i = \tilde{\lambda}_i/\tilde{\lambda}_3 \in [0, 1]$. The operator applied during scan k is

$$\mathbf{D}_k = \mathbf{I}_3 - \frac{\tilde{\mathbf{S}}_{t,k-1}}{\tilde{\lambda}_{3,k-1}}, \qquad \tilde{\lambda}_{3,k-1} = \lambda_{\max}(\tilde{\mathbf{S}}_{t,k-1}) > 0, \tag{12}$$

whose eigenvalues $d_{i,k} = 1 - \tilde{\rho}_{i,k-1}$ are the directional weights. The best-determined direction receives exactly zero weight, no direction receives a negative one, and a direction with $\tilde{\lambda}_i = 0$ receives unit weight, so no weight is clipped and no tuning threshold is introduced [22]. A numerical test that $\tilde{\lambda}_3 > 0$ is still required, and when $\tilde{\mathbf{S}}_t = \mathbf{0}$ the expression (12) is undefined; we then set $\mathbf{D}_k = \mathbf{I}_3$, which is a convention rather than the limit of that expression. Because $\mathbf{D}_k \succeq \mathbf{0}$, it remains positive semidefinite under positive scaling and under a change of basis.

The complement in (12) is evaluated at the final pose of the preceding scan, so the operator is available before the current scan is solved. This relies on the assumption that the weak directions and their relative strengths change slowly between scans. Under a change of global orientation $U \in SO(3)$, the operator transforms as $\mathbf{D}'_k = U\mathbf{D}_k U^\top$, so its eigenvectors need not be aligned with the coordinate axes. The operator remains uniquely defined when repeated eigenvalues make individual eigenvectors nonunique. Normalizing by $\tilde{\lambda}_3$ also makes the operator invariant to uniform scaling of $\tilde{\mathbf{S}}_t$, so it carries relative directional weight rather than an absolute information level.

2.5 Directionally Weighted Kinematic Regularization

The scan update is regularized by penalizing changes in velocity resolved in the body frame. The velocity components in that frame are assumed to change slowly between consecutive scans. The penalty is anisotropic, and the operator $\mathbf{D}_k$ in (12) carries its directional factor, with a weight of $d_{i,k} = 0$ along the best-determined direction and a weight approaching one along a direction the scan leaves unconstrained. The prior therefore acts where the scan is silent, while the measurements dominate elsewhere. We take the reference $\mathbf{v}_{\mathrm{ref},k} = (\hat{R}^+_{k-1})^\top \hat{\mathbf{v}}^+_{k-1}$ from the preceding posterior and hold it

fixed throughout the update. This gives the residual

$$\mathbf{h}_{\mathrm{kin},k}(\mathbf{x}) = R^\top \mathbf{v} - \mathbf{v}_{\mathrm{ref},k}, \tag{13}$$

which represents a velocity difference expressed in $\mathrm{m\,s^{-1}}$.

The weight matrix of the prior is then formed by expressing that directional factor in the predicted IMU frame, using the propagated orientation $\hat{R}_k^-$ at the scan reference time, and scaling it by a single constant,

$$\mathbf{\Lambda}_k^{\mathrm{kin}} = \frac{1}{\sigma_v^2} (\hat{R}_k^-)^\top \mathbf{D}_k \hat{R}_k^-, \tag{14}$$

whose middle factor is denoted $\mathbf{D}_k^I$. Multiplication on both sides by the same matrix preserves positive semidefiniteness. Holding $\mathbf{D}_k^I$ fixed throughout the iterations keeps the weight independent of the current attitude. The constant $\sigma_v > 0$, expressed in $\mathrm{m\,s^{-1}}$, sets only the magnitude of the penalty. The relative directional weights are carried entirely by $\mathbf{D}_k^I$. Therefore, changing σ_v rescales the penalty without changing which directions it favours. It is a regularization parameter of the estimator rather than a noise standard deviation. We report its value with the experimental configuration.

For the state ordering in (2), the residual Jacobian is

$$\mathbf{J}_K = \begin{bmatrix} (R^\top \mathbf{v})_\times & \mathbf{0} & R^\top & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix} \in \mathbb{R}^{3\times 18}, \tag{15}$$

so only the rotation and velocity blocks are nonzero. The effects on position, biases and gravity arise through coupling in the full system.

The scan update then solves

$$\delta\mathbf{x}^\star = \arg\min_{\delta\mathbf{x}} \left\{ \|\mathbf{r}_L + \mathbf{J}_L \delta\mathbf{x}\|_{\mathbf{\Omega}}^2 + \|\mathbf{h}_{\mathrm{kin},k} + \mathbf{J}_K \delta\mathbf{x}\|_{\mathbf{\Lambda}_k^{\mathrm{kin}}}^2 + \|\delta\mathbf{x}\|_{(\mathbf{P}_k^-)^{-1}}^2 \right\}, \tag{16}$$

with the objective evaluated at $\hat{\mathbf{x}}_k^{(\ell)}$. The first term is scan-to-map registration, the second is the kinematic prior, and the third damps the increment with the propagated covariance. Because that third term penalizes the increment rather than the offset from the propagated state, the iteration is a damped Gauss–Newton step: it starts from $\hat{\mathbf{x}}_k^-$, allows at most four iterations, stops early when the largest component of the increment falls below 10^{-3}, and its fixed point is where the scan and kinematic gradients balance. Setting the derivative with respect to $\delta\mathbf{x}$ to zero gives

$$\begin{aligned} \mathcal{A}_k^{(\ell)} \delta\mathbf{x} &= \mathbf{b}_L - \mathbf{J}_K^\top \mathbf{\Lambda}_k^{\mathrm{kin}} \mathbf{h}_{\mathrm{kin},k}, \\ \mathcal{A}_k^{(\ell)} &= \mathcal{A}_0 + \mathbf{J}_K^\top \mathbf{\Lambda}_k^{\mathrm{kin}} \mathbf{J}_K, \end{aligned} \tag{17}$$

so the regularizer changes the update in (9) in exactly two places: the information matrix gains the PSD block $\mathbf{B}_K = \mathbf{J}_K^\top \mathbf{\Lambda}_k^{\mathrm{kin}} \mathbf{J}_K \succeq \mathbf{0}$, and the gradient gains $-\mathbf{J}_K^\top \mathbf{\Lambda}_k^{\mathrm{kin}} \mathbf{h}_{\mathrm{kin},k}$. The LiDAR information $\mathbf{H}_L$ remains unchanged for a fixed linearization state and correspondence set. The state is updated as before, and the curvature

of the regularized solve is $\mathcal{A}_k^{(\ell)}$, whose inverse at the final iterate serves as the local covariance approximation about $\hat{\mathbf{x}}_k^+$.

The regularizer therefore adds to the filter instead of replacing part of it. Because $\mathbf{B}_K$ is PSD, the regularized system remains positive definite and uniquely solvable even though $\mathbf{\Lambda}_k^{\text{kin}}$ is singular by construction, with its null space containing the best-determined direction. The regularization changes the curvature of the solve. Whether it improves the estimate depends on the velocity-smoothness assumption, which the experiments test.

2.6 Main Algorithm

Algorithm 1 summarizes the DARK-LIO scan update after LiDAR–IMU initialization and map initialization. For each scan k, the update takes the deskewed point set $\mathcal{S}_k$, the propagated prior $(\hat{\mathbf{x}}_k^-, \mathbf{P}_k^-)$, the preceding posterior state $\hat{\mathbf{x}}_{k-1}^+$, the deficiency operator $\mathbf{D}_k$ obtained from the preceding scan, and the current voxel map as inputs. The reference body-frame velocity $\mathbf{v}_{\text{ref},k}$ is first computed from the preceding posterior. The deficiency operator is then expressed in the predicted IMU frame using $\hat{R}_k^-$ and scaled to construct the kinematic weight matrix $\mathbf{\Lambda}_k^{\text{kin}}$. The reference velocity, deficiency operator, and weight matrix remain fixed throughout the current scan update. Starting from $\hat{\mathbf{x}}_k^{(0)} = \hat{\mathbf{x}}_k^-$, each iteration forms accepted scan-to-map correspondences and accumulates the LiDAR information matrix $\mathbf{H}_L$ and negative-gradient vector $\mathbf{b}_L$. The kinematic residual $\mathbf{h}_{\text{kin},k}$ and its Jacobian $\mathbf{J}_K$ are evaluated at the current state estimate. These terms are combined with the propagated prior in the regularized normal equations, which are solved for the increment $\delta\mathbf{x}^\star$. The state estimate is updated through the $\boxplus$ operator, and the process repeats until convergence or the iteration limit is reached. The final iterate becomes the posterior state $\hat{\mathbf{x}}_k^+$, at which the system is relinearized to obtain the local covariance approximation $\mathbf{P}_k^+$. The LiDAR-only translational complement $\tilde{\mathbf{S}}_{t,k}$ is then constructed at the final pose and normalized to obtain $\mathbf{D}_{k+1}$ for the next scan, with the initialization and fallback rules specified in Algorithm 1 applied when necessary. This diagnostic is computed before map insertion, so it uses correspondences between the current scan and the existing map. Finally, the accepted scan samples are transformed using $\hat{\mathbf{x}}_k^+$ and inserted into the voxel map.

Algorithm 1 DARK-LIO scan update after initialization

Require: deskewed scan $\mathcal{S}_k$; propagated prior $\hat{\mathbf{x}}_k^-, \mathbf{P}_k^-$; previous posterior $\hat{\mathbf{x}}_{k-1}^+$; the operator $\mathbf{D}_k$ from the preceding scan; voxel map
Ensure: posterior $\hat{\mathbf{x}}_k^+, \mathbf{P}_k^+$; next operator $\mathbf{D}_{k+1}$; updated map
Construct $\mathbf{v}_{\mathrm{ref},k}$ from $\hat{\mathbf{x}}_{k-1}^+$ using (13)
Set $\mathbf{\Lambda}_k^{\mathrm{kin}}$ using (14), rotating $\mathbf{D}_k$ into the predicted IMU frame; hold it fixed for this scan
$\hat{\mathbf{x}}_k^{(0)} \leftarrow \hat{\mathbf{x}}_k^-$; $\ell \leftarrow 0$
repeat
    Form accepted correspondences; accumulate $\mathbf{H}_L, \mathbf{b}_L$ from them
    Evaluate $\mathbf{h}_{\mathrm{kin},k}, \mathbf{J}_K$ using (13) and (15)
    Solve (17) for $\delta\mathbf{x}^\star$
    $\hat{\mathbf{x}}_k^{(\ell+1)} \leftarrow \hat{\mathbf{x}}_k^{(\ell)} \boxplus \delta\mathbf{x}^\star$; $\ell \leftarrow \ell + 1$
until convergence or the iteration limit is reached
$\hat{\mathbf{x}}_k^+ \leftarrow \hat{\mathbf{x}}_k^{(\ell)}$
Relinearize at $\hat{\mathbf{x}}_k^+$ and form the local covariance approximation $\mathbf{P}_k^+$
Form $\tilde{\mathbf{S}}_{t,k}$ using (11), with $\tilde{\mathbf{S}}_{t,k} = \tilde{\mathbf{H}}_{tt}$ when $\tilde{\mathbf{H}}_{RR} = \mathbf{0}$
Form $\mathbf{D}_{k+1}$ using (12) ▷ no valid operator yet: $\mathbf{0}$; $\tilde{\lambda}_{3,k} = 0$: $\mathbf{I}_3$, i.e. isotropic at full strength; otherwise retain the last valid operator
Insert accepted scan samples into the map using $\hat{\mathbf{x}}_k^+$

3 Experiments

3.1 Experimental setup

3.1.1 Hardware Platform and datasets

DARK-LIO is evaluated in tunnel environments, with particular emphasis on long, geometrically repetitive sections that provide strong translational constraints across the tunnel but limited constraints along its axis. The evaluation combines quantitative comparisons on the public GEODE γ tunnel benchmark with qualitative mapping results from two self-collected tunnel datasets. The public benchmark comprises eleven tunnel sequences from the GEODE γ subset [43]. Six sequences were recorded in shield tunnels lined with precast concrete segments, resulting in relatively uniform geometry along the tunnel axis. The remaining five sequences were recorded in tunnels constructed using the mining method, referred to here as mined tunnels. These tunnels have sprayed-concrete support, rougher surfaces and less uniform cross-sections. The bag recording and survey use unsynchronized clocks, so a time offset is required for the GEODE data. The offset is a property of the recording and not of the estimator. We therefore fix *one* value per sequence and use it for all methods. Results are reported separately for shield and mined tunnels to assess each method's performance under different levels of axial geometric constraint. Sequence durations range from 232 to 557 s for shield tunnels and from 153 to 219 s for mined tunnels, providing approximately 50 min of tunnel data in total.

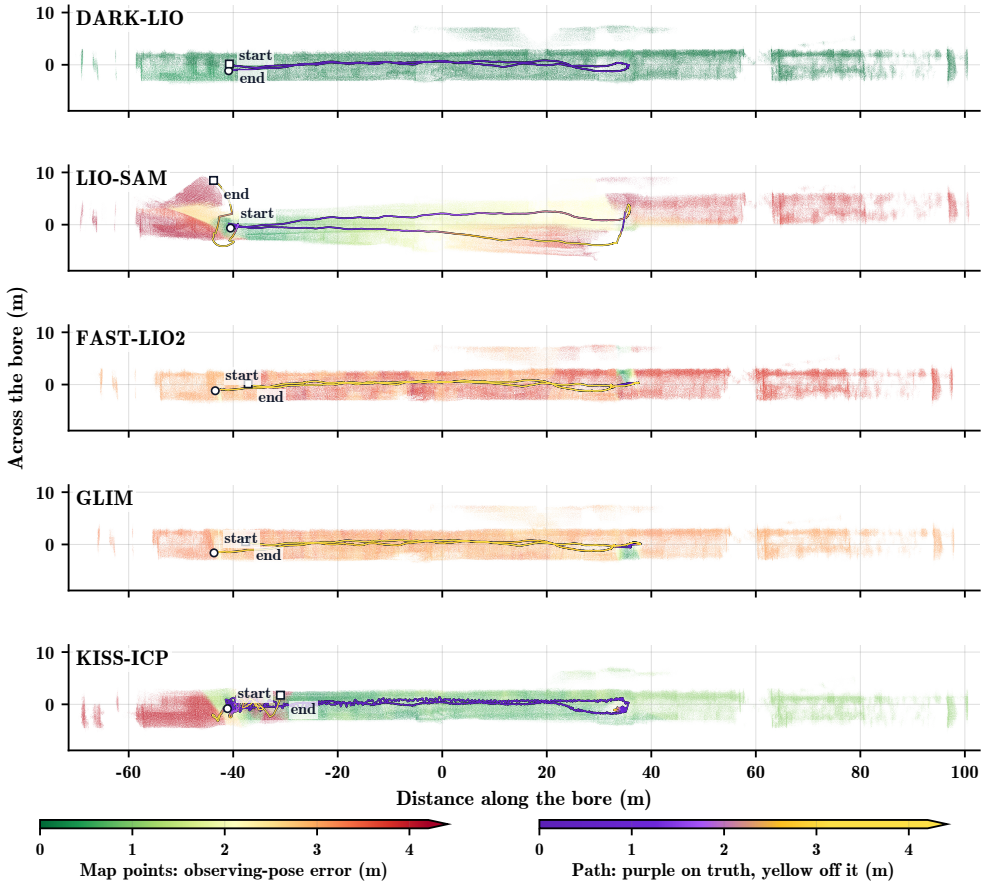
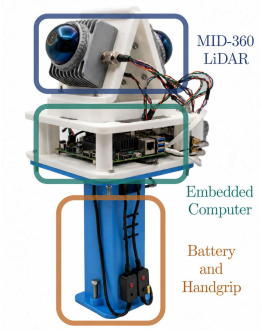


Fig. 3: Plan-view comparison of mapping results on mined tunnel 2. All five methods are shown within the same spatial window. Point colors indicate the position error of the estimated pose associated with each mapped point, using a common linear scale across all panels. The estimated trajectory is overlaid using a separate color gradient.

Ground-truth positions are obtained from total-station measurements. Accordingly, the quantitative trajectory evaluation considers translational accuracy only, and orientation accuracy is not evaluated. These measurements cover only part of each recording, so the corresponding portions of the estimated trajectories must be identified for evaluation. The evaluation metrics and alignment convention are described in Section 3.1.2.

To complement the public benchmark, we collected data in two additional environments using a handheld platform equipped with two Livox Mid-360 LiDAR sensors (Fig. 4a). The test environments comprised a shield tunnel under construction (Fig. 4b) and a mining-tunnel test facility (Fig. 4c). These datasets are used to qualitatively assess mapping performance beyond the public benchmark. The layout of the mining-tunnel test facility and the corresponding map reconstructed by DARK-LIO are shown in Figs. 4d and 4e, respectively.



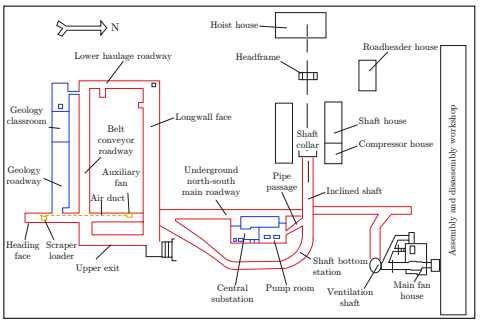
(a) Handheld LIO test platform.



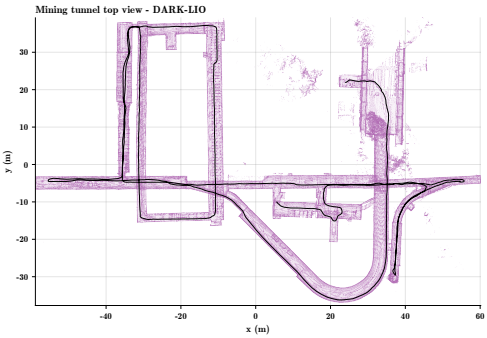
(b) Shield tunnel.



(c) Mine tunnel.



(d) Mine roadway layout.



(e) Map reconstructed by DARK-LIO.

Fig. 4: Handheld data acquisition platform, self-collected test environments, and mapping results: (a) platform equipped with two Livox Mid-360 LiDAR sensors, (b) shield tunnel under construction, (c) unlit mining-tunnel test facility, (d) layout of the test facility, and (e) map reconstructed by DARK-LIO. Purple points represent the reconstructed map, and black lines indicate the sensor trajectory.

3.1.2 Metrics and evaluation convention

We compare DARK-LIO with four cutting-edge SLAM methods. Three criteria are quantitative — absolute trajectory error, relative pose error and per-scan cost — and a fourth checks the self-consistency of the reconstructed maps, which we assess on the map views of Figure X rather than on a scalar score.

Absolute trajectory error (ATE). We report the position root mean square error (RMSE) after rigid SE(3) alignment without scale adjustment (Table 1). Any scale error is therefore retained in the evaluation.

Relative pose error (RPE). We report the translation error over 10 m reference segments as a percentage, using the same shared time offset (Table 2). RPE is insensitive to global alignment and therefore assesses drift separately from global accuracy.

Computational cost. We measure the central processing unit (CPU) time per scan in milliseconds on an Intel i9 processor running at 2.60 GHz. For each method, measurements are taken at every output, serially and without instrumentation (Table 5). The same timing procedure is used for all methods. Reported values are not taken from the methods' own internal timers.

Map self-consistency. A tunnel has a fixed cross-sectional shape, so a self-consistent map forms a single thin shell. The same wall observed during the outbound and return passes must map to the same surface.

3.1.3 Baselines

We compare DARK-LIO with four state-of-the-art LiDAR–inertial or LiDAR-only odometry systems, including LIO-SAM [14], FAST-LIO2 [15], GLIM [16] and KISS-ICP [11]. These methods were selected to represent filtering and optimization approaches to LiDAR–inertial estimation, together with LiDAR-only odometry, allowing comparison of localization accuracy, map consistency, and computational cost across different estimation frameworks. All four provide publicly available official source code, facilitating reproducible comparisons using the same recorded data and a common evaluation protocol. Furthermore, each method is evaluated by replaying the same bag at the same rate, with no method-specific tuning of the evaluation. KISS-ICP cannot process a 10 Hz Livox stream in real time on four shield sequences. Its measured throughput on Shield tunnel 1 decreases from 10.02 Hz to 2.43 Hz. These runs were therefore replayed at a rate of 0.2. Because the per-scan cost depends on the scans rather than their arrival rate and each recording is sample-complete, the reported cost values remain valid. KISS-ICP is the only method among those evaluated that does not run in real time on these sequences.

3.2 Localization accuracy in tunnels

Table 1 reports ATE RMSE on all eleven GEODE γ tunnel sequences. All sequences exhibit degeneracy. The six shield sequences, with either repetitive shield patterns or walls with few features, exhibit the most severe degeneracy. Compared with the four baselines, DARK-LIO ranks first on five of the six shield sequences and second on the remaining sequence.

The results differ on the five mined sequences. DARK-LIO ranks first on Mined tunnels 1 and 2 and second on 3, 4 and 5, where FAST-LIO2 leads by margins of 0.006–0.048 m. These margins are an order of magnitude below the offset band on those sequences and are *not* used to establish a ranking. These results therefore show that DARK-LIO did not degrade performance in tunnels with favorable geometric features and was not specifically tuned for featureless environments.Figure 5 presents the ATE RMSE values from Table 1 on a common scale. The off-scale FAST-LIO2 entry for Shield tunnel 5 is annotated rather than plotted.

Figure 3 shows Mined tunnel 2 in plan view within the same spatial window for all five methods. Color encodes the position error of the pose associated with each mapped point, using a common linear scale across all panels. The trajectory is shown using a separate color gradient. On this sequence, DARK-LIO's error is concentrated

Table 1: ATE RMSE in meters. The value in parentheses is the surveyed distance traveled by the platform on that sequence. When no run is available for a method on a sequence, its entry is shown as – rather than omitted or set to zero. **Bold type** indicates the best value in each row, and underlining indicates the second best.

sequence (traversed, m)	DARK-LIO	LIO-SAM	FAST-LIO2	GLIM	KISS-ICP
Shield tunnel 1 (691.8)	**2.243**	384.221	138.082	90.718	102.216
Shield tunnel 2 (673.8)	**25.768**	–	174.468	85.017	94.946
Shield tunnel 3 (506.8)	**1.651**	599.896	203.068	65.444	71.489
Shield tunnel 4 (501.4)	**10.486**	72.363	111.135	64.005	125.979
Shield tunnel 5 (201.3)	30.103	–	55692.188	41.609	**27.525**
Shield tunnel 6 (205.2)	**2.882**	–	33.678	30.571	30.844
Mined tunnel 1 (188.3)	**0.365**	2.010	0.425	0.405	37.787
Mined tunnel 2 (155.5)	**0.133**	2.270	3.346	2.989	0.787
Mined tunnel 3 (182.5)	0.445	–	**0.397**	56.803	50.827
Mined tunnel 4 (207.0)	0.190	–	**0.184**	0.279	0.446
Mined tunnel 5 (150.7)	0.158	–	**0.149**	4.689	31.201

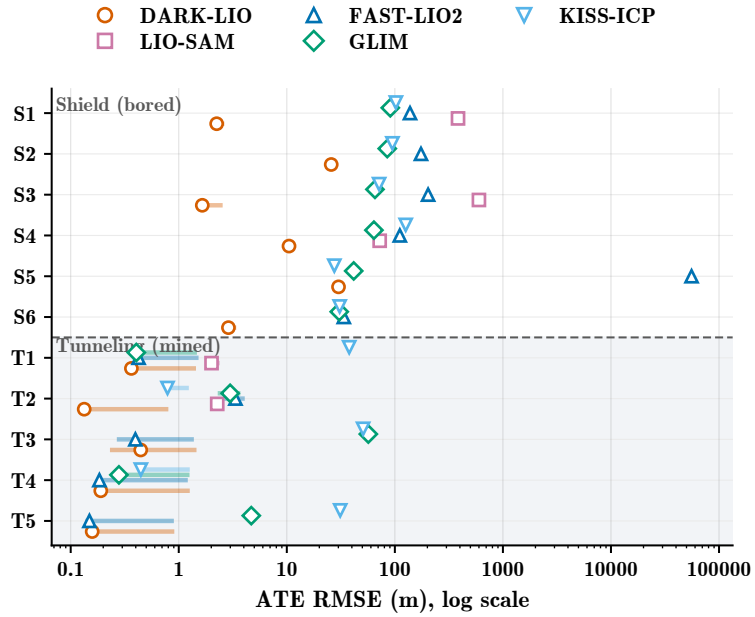


Fig. 5: ATE accuracy on a logarithmic scale for DARK-LIO and the four published baselines, with one group per sequence and the same shared offset.

along the bore and remains small throughout the run, while the baselines accumulate along-axis error that grows monotonically with distance travelled.

RPE (Table 2) further distinguishes the methods. On the shield sequences, DARK-LIO's median RPE is 9.491 %, compared with 99.463 (GLIM), 211.359 (KISS-ICP) and 302.524 (FAST-LIO2). On the mined sequences, it is 2.393 %, compared with

3.308 (FAST-LIO2) and 9.295 (GLIM). The two error measures therefore show the same direction of improvement on every shield sequence. An ATE improvement that was not reflected in RPE would indicate that the gain came from global alignment rather than from tracking.

Table 2: Relative pose error over 10 m reference segments, expressed as a percentage, at the same shared offset. **Bold type** indicates the best value in each row, and underlining indicates the second best.

sequence	DARK-LIO	LIO-SAM	FAST-LIO2	GLIM	KISS-ICP
Shield tunnel 1	**6.437**	1152.264	328.850	97.334	222.716
Shield tunnel 2	**16.249**	–	277.416	89.674	363.376
Shield tunnel 3	**5.219**	1447.082	327.632	94.689	200.003
Shield tunnel 4	**11.608**	596.999	106.366	101.591	421.614
Shield tunnel 5	**84.808**	–	82324.706	173.324	97.107
Shield tunnel 6	**7.375**	–	153.970	105.700	106.249
Mined tunnel 1	3.019	5.675	3.308	**3.004**	87.001
Mined tunnel 2	**1.663**	14.641	10.955	9.295	6.593
Mined tunnel 3	**2.887**	–	3.325	87.982	78.865
Mined tunnel 4	**2.393**	–	2.490	2.492	3.890
Mined tunnel 5	2.280	–	**2.248**	13.734	72.292

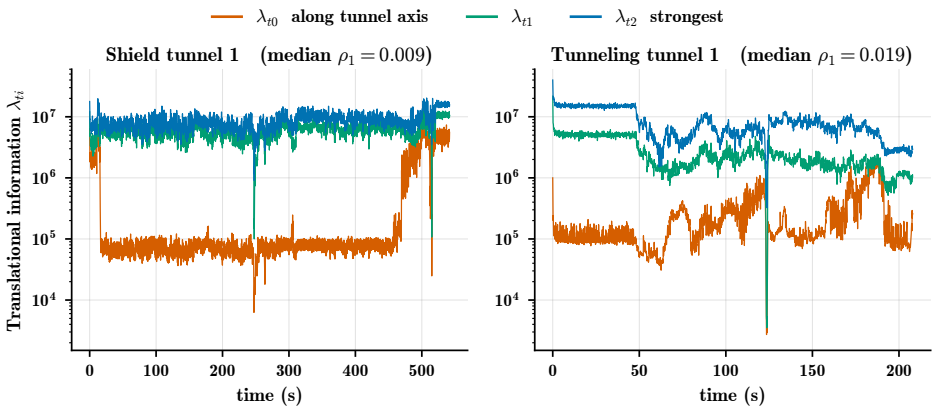
3.3 Degeneracy analysis based on residual error

Table 1 alone cannot explain *why* a run failed. This requires distinguishing a loss of geometric tracking across the tunnel from drift along the bore, because the two types of error require different remedies. The former is a front-end or calibration problem, whereas the latter is a degeneracy problem. Figure 6 shows the geometric origin of the latter: the translational information carried by the LiDAR measurements collapses along the tunnel axis.

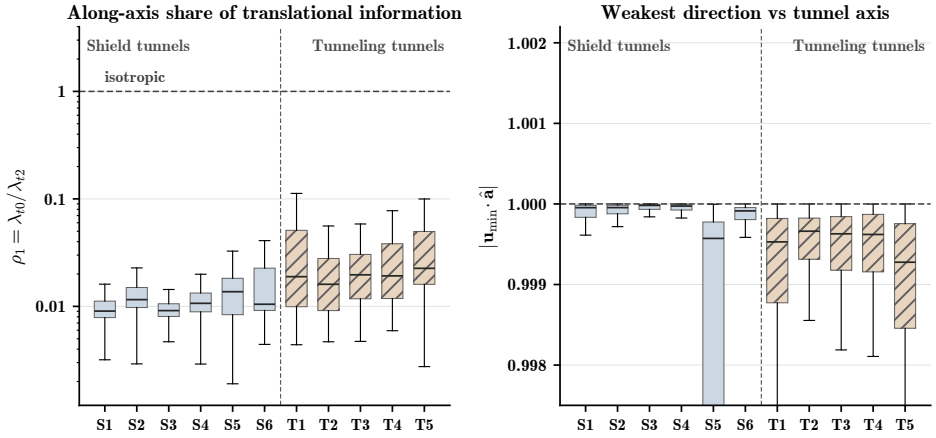
DARK-LIO's cross-axis error is small and nearly constant on every shield sequence: 0.396, 0.438, 0.335, 0.319, 0.318 and 0.320 m on Shield tunnels 1–6, respectively. The error thus falls within 0.32–0.44 m across all six sequences, including Shield tunnel 5, where its ATE reaches 30.103 m. The entire error is therefore along the bore: the cross-section is tracked to better than half a meter while the position along the axis drifts by tens of meters. This is the empirical content of the degeneracy argument, and it has a practical consequence: a front end that matches the tunnel walls ever more accurately cannot recover a direction that those walls do not constrain. Improvement here requires acting on the unobserved direction, which is what the kinematic regularization does.

3.4 Ablation and sensitivity analysis

We ablate the directionally weighted kinematic prior, which is the component this method adds to the estimator, and then sweep the single scalar that sets its strength. The two analyses are reported in separate tables because they represent different types of comparison. The ablation is a discrete comparison with the component present or absent, whereas the sweep is a curve over a continuous parameter. Readers should



(a) Translational spectrum over time for one shield tunnel and one mined tunnel.



(b) The along-axis deficit across all eleven sequences.

Fig. 6: Degeneracy trends: translational information collapses along the tunnel axis.

not interpret "off" as another point on that curve. Every configuration was run on all eleven sequences using a single binary, so only one configuration key differs between any two configurations.

Removing the component. Table 3 compares the adopted configuration with a configuration in which the prior is disabled. Disabling the prior increases the estimation error on all six shield sequences by +19.5 % to +4558 % and leaves the errors on the five mined sequences statistically unchanged (−9.4 % to +5.3 %, with a median of −0.5 %). This is the behavior expected from an ablation of this component. The component acts where the geometry is deficient and is neutral where the geometry already supplies the constraint. The largest effect is on Shield tunnel 5, where removing the prior increases the error from 30.103 m to 1402.043 m. Without the prior, the estimate on that sequence does not merely degrade but approaches divergence.

Table 3: Ablation of the kinematic prior. ATE RMSE in meters with the component present and absent.

Sequence	Prior OFF	Prior ON (adopted)	Diff. Percentage
Shield tunnel 1	2.831	**2.243**	+26.2 %
Shield tunnel 2	30.789	**25.768**	+19.5 %
Shield tunnel 3	25.134	**1.651**	+1422.4 %
Shield tunnel 4	24.567	**10.486**	+134.3 %
Shield tunnel 5	1402.043	**30.103**	+4557.6 %
Shield tunnel 6	3.895	**2.882**	+35.1 %
Mined tunnel 1	**0.363**	0.365	−0.5 %
Mined tunnel 2	0.140	**0.133**	+5.3 %
Mined tunnel 3	**0.403**	0.445	−9.4 %
Mined tunnel 4	**0.189**	0.190	−0.5 %
Mined tunnel 5	**0.156**	0.158	−1.3 %

Sensitivity to σ_v. Table 4 reports the results of sweeping the scale. On the shield sequences, the median ATE is 6.684 m at the adopted $\sigma_v = 0.05$, 9.266 m at 0.10 and 17.467 m at 0.20. Weakening the regularizer degrades accuracy over the range tested, and the adopted value is the best of those tested. The results on the mined sequences remain flat over the same range (0.188–0.190 m), consistent with the prior being inactive where the geometry already supplies the constraint. A doubling of σ_v costs 38.6 %, and a quadrupling costs 161.3 %. We state the sensitivity rather than claim insensitivity because the parameter matters. The sweep is one-sided and covers only $\sigma_v \geq 0.05\,\mathrm{m\,s^{-1}}$. Results at the stiff end are not tabulated.

Table 4: Sensitivity to the regularizer's scale σ_v ($\mathrm{m\,s^{-1}}$). ATE RMSE in meters.

sequence	σ_v=0.05 **(adopted)**	0.10	0.20
Shield tunnel 1	2.243	**2.230**	2.410
Shield tunnel 2	**25.768**	30.483	30.624
Shield tunnel 3	**1.651**	2.326	10.837
Shield tunnel 4	**10.486**	14.707	24.097
Shield tunnel 5	**30.103**	48.974	636.120
Shield tunnel 6	**2.882**	3.826	4.038
Mined tunnel 1	0.365	**0.363**	0.365
Mined tunnel 2	**0.133**	0.141	0.169
Mined tunnel 3	0.445	0.409	**0.402**
Mined tunnel 4	0.190	**0.188**	**0.188**
Mined tunnel 5	0.158	0.157	**0.156**

3.5 Computational cost and real-time operation

Table 5 reports the per-scan CPU cost. DARK-LIO has a mean cost of 74.3 ms, a median of 68.0 ms, and a 95th percentile of 100.0 ms. The mean corresponds to 13.5 Hz, indicating real-time operation with a 10 Hz sensor. Its mean cost is lower than those of GLIM (78.9 ms) and LIO-SAM (206.1 ms) and higher than that of FAST-LIO2 (46.5 ms). KISS-ICP's mean of 589.4 ms on four sequences played at a reduced rate

is not directly comparable and is reported for completeness only. Figure 7 shows the variation across sequences rather than reporting only a mean. DARK-LIO's mean per-scan cost ranges from 26.1 ms on Shield tunnel 5 to 100.9 ms on Shield tunnel 1, a factor of 3.9. Its processing rate falls below 10 Hz on two of the eleven sequences. The method therefore operates within the real-time budget on most sequences, but not all.

Table 5: Per-scan CPU cost in milliseconds. *rate* denotes the playback rate of the sequence, *p95* denotes the 95th percentile, *max* denotes the maximum per-scan cost, and *n* denotes the number of sequences timed. *source* indicates whether the value was measured in this work. **Bold type** indicates a mean cost below the period of the 10 Hz sensor.

method	rate	mean	median	p95	max	Hz	n	source
DARK-LIO	1	**74.3**	68.0	100.0	260	13.5	11	timed
LIO-SAM	1	206.1	208.3	750.0	4360	4.9	4	timed
FAST-LIO2	1	**46.5**	45.0	59.1	140	21.5	11	timed
GLIM	1	**78.9**	84.3	111.8	350	12.7	11	timed
KISS-ICP	0.2,1	589.4	169.0	1648.2	83640	1.7	11	timed

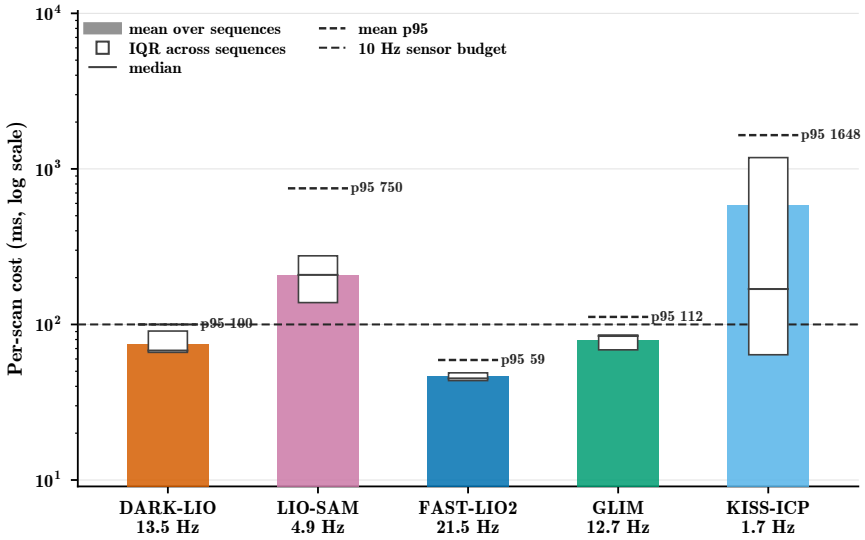
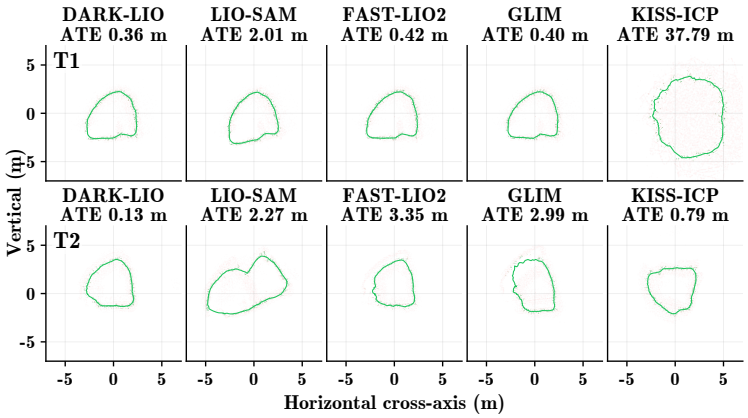


Fig. 7: Per-scan cost showing the mean for each method over the sequences on which it was run and the interquartile range across those sequences. The period of the 10 Hz sensor is shown as a dashed line.

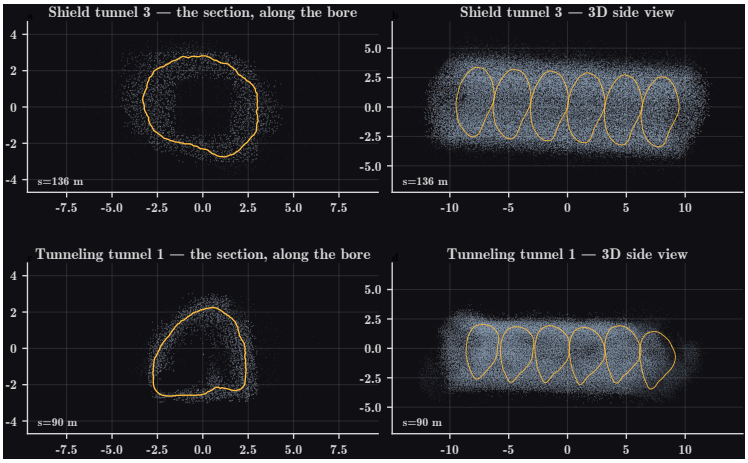
3.6 Map consistency

ATE does not capture the type of failure addressed in this work. A trajectory can pass through the surveyed points while the surrounding walls are reconstructed twice. These

two outcomes are independent, as an estimate that maintains positional accuracy can still reconstruct the same wall from two passes with an offset caused by drift.



(a) Map self-consistency on the mined tunnel sequences.



(b) Tunnel section views and three-dimensional side views produced by DARK-LIO.

Fig. 8: Comparison of map consistency and map representations for the GEODE tunnel sequences.

Each panel of Figure 8a shows a map section in the tunnel frame, with the along-axis direction shown horizontally and the cross-section vertically. On the mined sequences, the region produced by DARK-LIO is a thin shell that closes on itself. LIO-SAM, FAST-LIO2, and GLIM largely retain a similar shape. On Mined tunnel 1, their ATE values are 2.01 m, 0.42 m, and 0.40 m, respectively. In contrast, the map produced by KISS-ICP on that sequence encloses approximately twice the area of the maps produced by the other methods, with an ATE of 37.79 m. Figure 8b presents

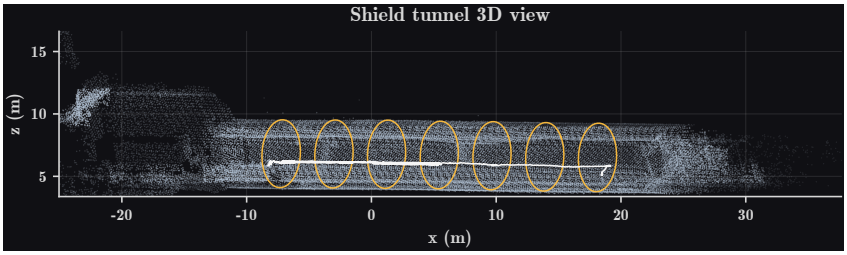


Fig. 9: Shield tunnel quality illustrated by yellow ellipses along the tunnel (partially shown).

the point cloud from which the values above are computed. The left column shows a single cross-section of the map, cut along the bore in the same frame and following the same rules as Figure 8a. The right column shows a long stretch of the same tunnel viewed from the side and slightly above, so the walls recede and the map appears as a bore rather than a ring. The paired views show that the wall forms a thin shell rather than a band. In the side view, the surface forms a close-fitting skin over the bore, with the cross-sections appearing as regular rings on the surface. The section closes on itself with no interior structure or second layer.

3.7 Real-world experiments

We now present the details of our real-world experiments on two self-collected tunnel sequences. All self-collected data were recorded with the handheld unit shown in Figure 4a, which carries **two Livox Mid-360 LiDARs**. One Mid-360 covers $360° \times 59°$, so a single unit leaves a cone in the near field unobserved, including the floor and roof directly below and above the sensor. The second unit is mounted to cover this region, and the two units are calibrated as a single rigid body. Both the shield-tunnel and mine-tunnel recordings use the same unit, so the two scenes differ in their environments rather than in their sensors. Each recording is replayed using one selected LiDAR and its own IMU. This IMU is the Mid-360's internal unit, whose accelerations are normalized by g and rescaled accordingly.

The first self-collected scene is a shield tunnel under construction (Figure 4b). It is bored and lined with precast concrete segments and is therefore geometrically regular along the bore. It belongs to the same family as the six shield sequences of the benchmark and represents the environment in which along-axis information is scarcest. The second scene is an unlit mining tunnel test field (Figure 4c). It is excavated and supported by sprayed concrete, so its surface is rougher and its cross-section is less uniform. It belongs to the mined family of the benchmark and represents the environment targeted by this work.

Shield tunnel quality analysis. Fitting a free ellipse to each of 14 stations along the traversed stretch, partially shown in Figure 9, gives semi-axes of 2.75 m and 2.73 m and an aspect ratio of 1.01, meaning the section is round to within one percent. The aspect ratio is not constrained to take this value because both axes and the orientation are freely fitted. The section's radius remains at 2.73–2.75 m over the entire stretch, a spread of under one percent, so the bore is reconstructed at a constant scale along the traversed portion. **This is evidence that the map's scale is consistent along**

the run, but it is not evidence that the map is accurate in absolute terms. A reconstruction with a uniformly incorrect scale would also be round and maintain its radius.

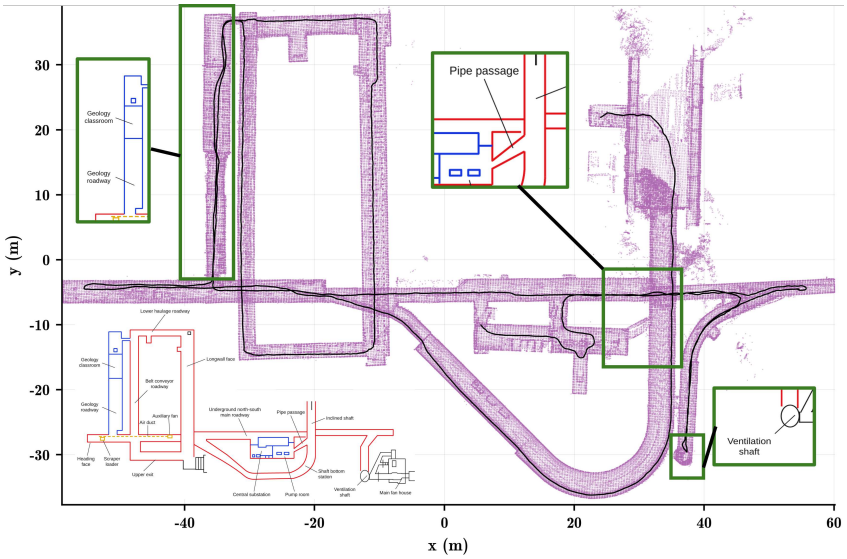


Fig. 10: Comparison of the mined tunnel layout with enlarged views of three geometric features.

Mine tunnel quality analysis. The mining tunnel has a known roadway layout, reproduced in Figure 4d from the mine's own plan. Figure 10 shows the map reconstructed by this method for comparison. **The reconstructed map reproduces the layout's topology**, including the same junctions, turns, and smaller geometric features such as the pipe passage, ventilation shaft, and geology roadway, as shown in the enlarged regions. We include both scenes to validate the target environments, as they lie outside the benchmark's two original families. Moreover, a map is a pipeline output that trajectory error alone cannot evaluate. A run can track the ground truth along a corridor while duplicating its walls. The mine roadway is the more challenging of the two scenes. In this scene, the reconstructed map is compared against an independent description of the same place.

4 Conclusions and Discussion

We presented DARK-LIO, a degeneracy-aware regularized kinematic LiDAR–inertial odometry framework for long, straight underground tunnels. In such environments, surrounding surfaces strongly constrain motion across the tunnel but provide little information along its axis, making conventional LIO prone to silent axial drift. DARK-LIO mitigates this failure mode without additional sensors, prior maps, or learned modules.

The core idea is to use directional geometric information to weight a motion prior. We derived a continuous directional deficiency operator from the anisotropic part

of the registration precision through a regularized Schur complement. The operator accounts for rotation–translation coupling, is invariant to uniform positive scaling of the translational information, and transforms consistently under rotations of the global frame. It assigns greater weight to relatively weaker directions without requiring a degeneracy threshold. Computed from the preceding scan, the operator weights a body-frame velocity-smoothness prior

inside an iterated error-state Kalman filter. The prior is referenced to the previous posterior estimate and contributes positive semidefinite curvature to the normal equations. At each linearization point, the LiDAR information term is retained for the same state and correspondence set. The prior penalizes changes in body-frame velocity between consecutive scans, with greater weight assigned to relatively weak directions. A single scalar controls its overall strength.

We evaluated DARK-LIO on the public GEODE γ tunnel benchmark and on two self-collected tunnel datasets. Across the six shield sequences, DARK-LIO achieves the lowest ATE on five sequences and the lowest RPE on all six among the evaluated methods. Relative to GLIM, it reduces median ATE by 89.7 % and median RPE by 90.5 %. Removing the prior increases shield-tunnel ATE by approximately 19.5 % to 4558 %, while producing substantially smaller changes on the mined-tunnel sequences. The degeneracy analysis shows cross-axis errors of approximately 0.32–0.44 m despite overall trajectory errors reaching tens of meters, which is consistent with predominantly axial drift. On the real-time experimental datasets, the shield-tunnel map exhibits consistent cross-sectional dimensions, while the mine map reproduces major structural features shown in an independent mine plan. These results support map consistency and qualitative agreement with the mine layout, but do not establish absolute mapping accuracy.

The approach assumes that body-frame velocity varies slowly between consecutive scans and that the preceding scan remains informative about the current weak directions. Future work will extend the framework in three directions. First, we will investigate incorporating the regularizer into a sliding-window or factor-graph backend to examine its effects on loop closure and global consistency. Second, we will combine the kinematic regularizer with intensity or semantic cues available in nonuniform tunnels. Third, we will evaluate DARK-LIO on broader underground datasets, including mines with varying cross-sections, fault zones, and dynamic machinery.

Author Contributions. Y.C. and F.W. contributed to the writing of the manuscript. S.Z. collected experimental data. Y.C. and F.W conducted simulations and experimental investigations. All authors reviewed the manuscript. Y.C. and F.W. performed the final proofreading and approved the manuscript prior to submission.

Funding. This project is supported by the National Natural Science National Natural Science Foundation of China (5260040159).

Data Availability. The public GEODE γ tunnel benchmark dataset used in this study is available from https://github.com/PengYu-Team/GEODE_dataset. The two self-collected tunnel datasets generated and/or analyzed during the current study are not publicly available due to [ethical restrictions], but are available from the corresponding author on reasonable request.

Declarations

Competing interests. The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] Tang, C., Gao, E., Li, Y., Li, M., Bai, D., Tang, H., Zhou, G.: Inspection robot and wall surface detection method for coal mine wind shaft. Applied Sciences **13**(9), 5662 (2023) https://doi.org/10.3390/app13095662

[2] Androulakis, V., Sottile, J., Schafrik, S., Agioutantis, Z.: Navigation system for a semi-autonomous shuttle car in room and pillar coal mines based on 2D LiDAR scanners. Tunnelling and Underground Space Technology **117**, 104149 (2021) https://doi.org/10.1016/j.tust.2021.104149

[3] Sottile, J., Rose, S., Rajvanshi, A., Schafrik, S., Agioutantis, Z., Sizintsev, M., Chiu, H.-P.: Concepts for autonomous navigation of underground mine face haulage equipment using depth cameras. International Journal of Coal Science & Technology **13**(1), 67 (2026) https://doi.org/10.1007/s40789-026-00903-8

[4] Wang, W., Dong, W., Su, Y., Wu, D., Du, Z.: Development of search-and-rescue robots for underground coal mine applications. Journal of Field Robotics **31**(3), 386–407 (2014) https://doi.org/10.1002/rob.21501

[5] Dissanayake, M.W.M.G., Newman, P., Clark, S., Durrant-Whyte, H.F., Csorba, M.: A solution to the simultaneous localization and map building (SLAM) problem. IEEE Transactions on Robotics and Automation **17**(3), 229–241 (2001) https://doi.org/10.1109/70.938381

[6] Jiang, H., Cheng, Y., Dai, W., Wan, W., Liu, Q., Wang, F.: GPS-denied LiDAR-based SLAM: A survey. IET Cyber-Systems and Robotics **7**(1), 70031 (2025) https://doi.org/10.1049/csy2.70031

[7] Li, J., Benndorf, J., Trybała, P.: Quantitative analysis of different SLAM algorithms for geo-monitoring in an underground test field. International Journal of Coal Science & Technology **12**(1), 7 (2025) https://doi.org/10.1007/s40789-025-00745-w

[8] Jacobson, A., Zeng, F., Smith, D., Boswell, N., Peynot, T., Milford, M.: What localizes beneath: A metric multisensor localization and mapping system for autonomous underground mining vehicles. Journal of Field Robotics **38**(1), 5–27 (2021) https://doi.org/10.1002/rob.21978

[9] Khattak, S., Papachristos, C., Alexis, K.: Visual-thermal landmarks and inertial fusion for navigation in degraded visual environments. In: IEEE Aerospace

Conference, pp. 1–9 (2019). https://doi.org/10.1109/AERO.2019.8741787

[10] Reinke, A., Palieri, M., Morrell, B., Chang, Y., Ebadi, K., Carlone, L., Aghamohammadi, A.-a.: LOCUS 2.0: Robust and computationally efficient lidar odometry for real-time 3D mapping. IEEE Robotics and Automation Letters **7**(4), 9043–9050 (2022) https://doi.org/10.1109/LRA.2022.3181357

[11] Vizzo, I., Guadagnino, T., Mersch, B., Wiesmann, L., Behley, J., Stachniss, C.: KISS-ICP: In defense of point-to-point ICP – simple, accurate, and robust registration if done the right way. IEEE Robotics and Automation Letters **8**(2), 1029–1036 (2023) https://doi.org/10.1109/LRA.2023.3236571

[12] Arnold, E., Mozaffari, S., Dianati, M.: Fast and robust registration of partially overlapping point clouds. IEEE Robotics and Automation Letters **7**(2), 1502–1509 (2022) https://doi.org/10.1109/LRA.2021.3137888

[13] Yang, H., Shi, J., Carlone, L.: TEASER: Fast and certifiable point cloud registration. IEEE Transactions on Robotics **37**(2), 314–333 (2021) https://doi.org/10.1109/TRO.2020.3033695

[14] Shan, T., Englot, B., Meyers, D., Wang, W., Ratti, C., Rus, D.: LIO-SAM: Tightly-coupled lidar inertial odometry via smoothing and mapping. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 5135–5142 (2020). https://doi.org/10.1109/IROS45743.2020.9341176

[15] Xu, W., Cai, Y., He, D., Lin, J., Zhang, F.: FAST-LIO2: Fast direct LiDAR-inertial odometry. IEEE Transactions on Robotics **38**(4), 2053–2073 (2022) https://doi.org/10.1109/TRO.2022.3141876

[16] Koide, K., Yokozuka, M., Oishi, S., Banno, A.: GLIM: 3D range-inertial localization and mapping with GPU-accelerated scan matching factors. Robotics and Autonomous Systems **179**, 104750 (2024) https://doi.org/10.1016/j.robot.2024.104750

[17] Shan, T., Englot, B., Ratti, C., Rus, D.: LVI-SAM: Tightly-coupled lidar-visual-inertial odometry via smoothing and mapping. In: IEEE International Conference on Robotics and Automation (ICRA), pp. 5692–5698 (2021). https://doi.org/10.1109/ICRA48506.2021.9561996

[18] Lin, J., Zhang, F.: R^3LIVE: A robust, real-time, RGB-colored, LiDAR-inertial-visual tightly-coupled state estimation and mapping package. In: IEEE International Conference on Robotics and Automation (ICRA), pp. 10672–10678 (2022). https://doi.org/10.1109/ICRA46639.2022.9811935

[19] Li, K., Li, M., Hanebeck, U.D.: Towards high-performance solid-state-LiDAR-inertial odometry and mapping. IEEE Robotics and Automation Letters **6**(3), 5167–5174 (2021) https://doi.org/10.1109/LRA.2021.3070251

[20] Chen, Z., Xu, Y., Yuan, S., Xie, L.: iG-LIO: An incremental GICP-based tightly-coupled LiDAR-inertial odometry. IEEE Robotics and Automation Letters **9**(2), 1883–1890 (2024) https://doi.org/10.1109/LRA.2024.3349915

[21] Gao, Y., Zhao, L., Zhang, B.: ININ-LIO: Informer-based neural inertial network-aided LiDAR-inertial odometry. IEEE Transactions on Instrumentation and Measurement **74**, 1–10 (2025) https://doi.org/10.1109/TIM.2025.3546399

[22] Zhang, J., Kaess, M., Singh, S.: On degeneracy of optimization-based state estimation problems. In: IEEE International Conference on Robotics and Automation (ICRA), pp. 809–816 (2016). https://doi.org/10.1109/ICRA.2016.7487211

[23] Lee, E.M., Marsim, K.C., Myung, H.: LODESTAR: Degeneracy-aware LiDAR-inertial odometry with adaptive Schmidt-Kalman filter and data exploitation. IEEE Robotics and Automation Letters **11**(1), 922–929 (2026) https://doi.org/10.1109/LRA.2025.3634906

[24] Ji, M., Shi, W., Cui, Y., Liu, C., Chen, Q.: Adaptive denoising-enhanced LiDAR odometry for degeneration resilience in diverse terrains. IEEE Transactions on Instrumentation and Measurement **73**, 1–15 (2024) https://doi.org/10.1109/TIM.2024.3365168

[25] Lee, J., Komatsu, R., Shinozaki, M., Kitajima, T., Asama, H., An, Q., Yamashita, A.: Switch-SLAM: Switching-based LiDAR-inertial-visual SLAM for degenerate environments. IEEE Robotics and Automation Letters **9**(8), 7270–7277 (2024) https://doi.org/10.1109/LRA.2024.3421792

[26] Ebadi, K., Palieri, M., Wood, S., Padgett, C., Agha-mohammadi, A.-a.: DARE-SLAM: Degeneracy-aware and resilient loop closing in perceptually-degraded environments. Journal of Intelligent & Robotic Systems **102**(1), 2 (2021) https://doi.org/10.1007/s10846-021-01362-w

[27] Hu, X., Chen, X., Jia, M., Wu, J., Tan, P., Waslander, S.L.: DCReg: Decoupled Characterization for Efficient Degenerate LiDAR Registration. arXiv:2509.06285 (2025)

[28] Lim, H., Kim, D., Kim, B., Myung, H.: AdaLIO: Robust adaptive LiDAR-inertial odometry in degenerate indoor environments. In: 20th International Conference on Ubiquitous Robots (UR), pp. 48–53 (2023). https://doi.org/10.1109/UR57808.2023.10202252

[29] Lee, D., Lim, H., Han, S.: GenZ-ICP: Generalizable and degeneracy-robust LiDAR odometry using an adaptive weighting. IEEE Robotics and Automation Letters **10**(1), 152–159 (2025) https://doi.org/10.1109/LRA.2024.3498779

[30] Pfreundschuh, P., Oleynikova, H., Cadena, C., Siegwart, R., Andersson, O.: COIN-LIO: Complementary intensity-augmented LiDAR inertial odometry. In:

IEEE International Conference on Robotics and Automation (ICRA), pp. 1730–1737 (2024). https://doi.org/10.1109/ICRA57147.2024.10610938

[31] Pan, Y., Zhong, X., Wiesmann, L., Posewsky, T., Behley, J., Stachniss, C.: PIN-SLAM: LiDAR SLAM using a point-based implicit neural representation for achieving global map consistency. IEEE Transactions on Robotics **40**, 4045–4064 (2024) https://doi.org/10.1109/TRO.2024.3422055

[32] Chen, K., Nemiroff, R., Lopez, B.T.: Direct LiDAR-inertial odometry: Lightweight LIO with continuous-time motion correction. In: IEEE International Conference on Robotics and Automation (ICRA), pp. 3983–3989 (2023). https://doi.org/10.1109/ICRA48891.2023.10160508

[33] He, D., Xu, W., Chen, N., Kong, F., Yuan, C., Zhang, F.: Point-LIO: Robust high-bandwidth light detection and ranging inertial odometry. Advanced Intelligent Systems **5**(7), 2200459 (2023) https://doi.org/10.1002/aisy.202200459

[34] Kim, S., Lee, D., Hong, J., Park, S., Jo, H., Han, S.: ALIVE-LIO: Degeneracy-Aware Learning of Inertial Velocity for Enhancing ESKF-Based LiDAR-Inertial Odometry. arXiv:2604.02706 (2026)

[35] Hertzberg, C., Wagner, R., Frese, U., Schröder, L.: Integrating generic sensor fusion algorithms with sound state representations through encapsulation of manifolds. Information Fusion **14**(1), 57–77 (2013) https://doi.org/10.1016/j.inffus.2011.08.003

[36] Xu, W., Zhang, F.: FAST-LIO: A fast, robust LiDAR-inertial odometry package by tightly-coupled iterated kalman filter. IEEE Robotics and Automation Letters **6**(2), 3317–3324 (2021) https://doi.org/10.1109/LRA.2021.3064227

[37] Forster, C., Carlone, L., Dellaert, F., Scaramuzza, D.: On-manifold preintegration for real-time visual-inertial odometry. IEEE Transactions on Robotics **33**(1), 1–21 (2017) https://doi.org/10.1109/TRO.2016.2597321

[38] Wang, L., Zhang, X., Li, C., He, D., Pan, Y., Yi, J.: Super-lio: A robust and efficient lidar-inertial odometry system with a compact mapping strategy. IEEE Robotics and Automation Letters **11**(3), 2666–2673 (2026) https://doi.org/10.1109/LRA.2026.3653372

[39] Yuan, C., Xu, W., Liu, X., Hong, X., Zhang, F.: Efficient and probabilistic adaptive voxel mapping for accurate online LiDAR odometry. IEEE Robotics and Automation Letters **7**(3), 8518–8525 (2022) https://doi.org/10.1109/LRA.2022.3187250

[40] Segal, A., Haehnel, D., Thrun, S.: Generalized-ICP. In: Robotics: Science and Systems (RSS) (2009). https://doi.org/10.15607/RSS.2009.V.021

[41] Bell, B.M., Cathey, F.W.: The iterated kalman filter update as a gauss-newton method. IEEE Transactions on Automatic Control **38**(2), 294–297 (1993) https://doi.org/10.1109/9.250476

[42] Boyd, S., Vandenberghe, L.: Convex Optimization. Cambridge University Press, Cambridge (2004). https://doi.org/10.1017/CBO9780511804441

[43] Chen, Z., Qi, Y., Feng, D., Zhuang, X., Chen, H., Hu, X., Wu, J., Peng, K., Lu, P.: Heterogeneous LiDAR dataset for benchmarking robust localization in diverse degenerate scenarios. The International Journal of Robotics Research **45**(1), 6–22 (2026) https://doi.org/10.1177/02783649251344967