

LiDAR Odometry for Inland Waterways using Water-Plane Constraints

Iulian Filip¹, Jaime Godoy² and Daniel Medina¹

Abstract—LiDAR odometry is a key component of autonomous navigation, yet methods that perform well on roads and indoors often degrade in inland waterways (IWs) due to sparse and noisy water-surface returns, long-range observations, and vessel motion induced by waves and currents. We propose a lightweight LiDAR-only odometry pipeline that leverages the water surface as a domain-specific geometric prior. The approach propagates the pose with a constant-velocity motion model and fuses two complementary measurement cues within an error-state Kalman filter: (i) a water-plane observation, obtained from water-surface detection and parameterized by its normal and distance, which constrains elevation and tilt, and (ii) a scan-to-submap ICP update that refines the full 6-DoF state. We evaluate on a self-collected IW dataset spanning 14.5 km over 103 min, covering vegetation corridors, dense urban sections, port environment, narrow channels and bridge underpasses, with synchronized measurements from three LiDARs: Velodyne VLP-32, Ouster OS0-128, and Livox Avia. Compared to the state-of-the-art LiDAR-only baseline KISS-ICP, our method consistently reduces trajectory error across scenes and sensors, achieving more than 80% lower absolute position error and more than 60% lower absolute orientation error.

I. INTRODUCTION

LiDAR odometry performs well in road and indoor environments, yet its reliability in inland waterways (IWs) remains limited. This gap stems from IW-specific sensing conditions: weak water-surface constraints, limited geometric features, wave- and current-induced vessel motion, and frequent long-range returns in wide channels, all of which can degrade scan registration and increase drift. As a result, reliable vessel autonomy requires odometry methods tailored to IW operations. This need is particularly important because IWs are a key component of regional freight and sustainable logistics networks, offering cost and environmental benefits within multimodal transport systems [1], [2]. At the same time, however, IW navigation is shaped by highly variable channel and infrastructure constraints. Narrow channels complicate maneuvering of large vessels, while variable depths and bottlenecks at locks and bridges restrict traffic flow and can make planning and safe operations challenging [3]. In addition, bridge-height limits increase the risk of vessel-bridge collisions, with significant safety and recovery consequences. Taken together, these factors increase navigational complexity and strengthen the case for advanced automation and decision support. Consistent with this motivation, recent

European safety overviews identify human and organizational factors as major contributors to accidents, reinforcing the value of autonomy functions that improve situational awareness during safety-critical maneuvers [4].

A key enabling capability for such autonomy functions is reliable navigation, which depends on accurate state estimation and environment representation. While GNSS-based positioning is widely used, it can be unreliable or unavailable in operational settings due to signal blockage, multipath, and deliberate spoofing or jamming, motivating auxiliary onboard sensors that provide independent localization [5]. In practice, this is addressed through onboard perception and localization pipelines based on cameras [6], [7], LiDAR [8], [9], radar [10], [11], or their fusion, providing the pose estimates required for planning and control. Each modality offers distinct strengths and limitations. For instance, camera-based approaches exploit appearance and semantic features for visual odometry and mapping, but can be sensitive to illumination changes, low texture, and adverse weather. Conversely, radar-based methods provide robust range and velocity measurements and can operate under fog, rain, and snow, although their lower angular resolution can limit fine-grained mapping and registration. In this context, LiDAR is particularly attractive for localization and mapping because it provides fine-grained 3D returns that enable accurate geometric representation and is largely insensitive to illumination changes.

Building on these sensing advantages, LiDAR-based odometry and Simultaneous Localization and Mapping (SLAM) have become a major research focus over the past decade, with extensive validation in structured road and urban driving scenarios and in indoor environments where static geometry provides strong constraints for scan registration and consistent motion estimation. Early representative systems such as LOAM introduced feature-based LiDAR odometry by extracting edge and planar features and performing scan-to-scan and scan-to-map registration for low-drift motion estimation [12]. To further improve accuracy, particularly under fast or irregular motion, tightly coupled LiDAR–inertial approaches integrate IMU propagation with LiDAR constraints for more stable state estimation. LIO-SAM adopts a factor-graph formulation that combines IMU preintegration with LiDAR measurements to improve accuracy and long-term consistency [13]. FAST-LIO2, by contrast, follows a direct LiDAR–inertial design and uses an error-state Kalman filter (ESKF) to fuse IMU propagation with iterative updates on raw LiDAR points, enabling high-rate and accurate state estimation [14]. More recently, among LiDAR-only methods, KISS-ICP provides a simple and

*This work is supported by the German project RadarSOW.

¹Institute of Communications and Navigation, German Aerospace Center (DLR), Neustrelitz, Germany iulian.filip@dlr.de

²Department of Systems Engineering and Automation at Universidad Carlos III de Madrid, Madrid, Spain jgodoy@pa.uc3m.es

lightweight odometry pipeline based on direct scan registration, while achieving performance comparable to more complex systems in structured environments [9]. Beyond these feature-rich settings, research has also shifted toward LiDAR odometry and SLAM methods that remain reliable under degraded or weakly constrained sensing conditions, such as geometrically repetitive corridors, tunnels, and other low-structure environments where scan registration becomes ill-conditioned [15]–[17].

IWs present a comparable challenge, motivating initial efforts on perception, localization, and mapping for vessel operations in channels and near infrastructure. However, publicly available IW datasets remain limited, with only a few releases such as Pohang Canal [18] and USVInland [19]. This scarcity constrains reproducible evaluation and slows method development. Despite this constraint, recent studies demonstrate progress across complementary directions. In [20] the authors evaluated several representative LiDAR and visual odometry methods in IW scenarios and observed a consistent performance degradation relative to structured terrestrial environments. Related IW-oriented work has addressed this challenge from several angles. For localization under intermittent GNSS, a LiDAR-SLAM assisted fusion positioning method for unmanned surface vehicles switches to a LiDAR module when satellite positioning is degraded or unavailable [21]. Complementing this direction, a laser-based SLAM algorithm for IW scenarios integrates LiDAR with IMU information to support onboard localization and mapping [22]. From a mapping perspective, a recent review of high-definition mapping for IW, highlights accurate localization as a prerequisite for producing maps usable by assistance and autonomy functions [23]. At the system level, Roboat II demonstrates an autonomous vessel platform that integrates LiDAR–IMU–GPS for navigation in constrained urban waterways [24]. More recently, InlandLOAM [25] introduces an IW-tailored, feature-based LiDAR odometry and mapping framework that incorporates water-surface constraints within a factor-graph formulation to mitigate drift and generate navigation-oriented map products.

In line with InlandLOAM, we focus on exploiting the water surface as a domain-specific information to stabilize LiDAR motion estimation in IW. Motivated by the work in [20] showing degraded odometry performance under IW sensing conditions, this paper proposes a deliberately simple LiDAR-only pipeline that augments ICP-based geometric alignment with an explicit water-plane constraint, and fuses both information sources within an ESKF framework to mitigate drift. The resulting design aims to remain lightweight and practical while addressing a key source of drift in IW environments. The contributions of this paper are as follows:

- We introduce a lightweight LiDAR-only odometry pipeline for IWs that integrates a water-plane measurement update into an ESKF framework and combines it with scan-to-submap ICP.
- We provide a large-scale real-world evaluation on a self-collected IW dataset spanning 14.5 km over 103 min, covering vegetation corridors, dense urban sections, port

environments, narrow channels, and bridge underpasses, with a 6-DoF ground-truth trajectory obtained from post-processed multi-antenna GNSS.

- We demonstrate cross-sensor generalization on three LiDARs (Velodyne, Ouster, and Livox) and outperform the state-of-the-art LiDAR-only baseline KISS-ICP, achieving more than 80% lower absolute position error and more than 60% lower absolute orientation error on average.

The remainder of this paper is organized as follows. Section II presents the proposed methodology, including water-plane detection and the ESKF formulation. Section III describes the experimental setup, including the vessel platform and dataset characteristics. Section IV summarizes and discusses the experimental results. Finally, Section V concludes the paper and outlines directions for future work.

II. METHODOLOGY

This work presents a LiDAR-only odometry framework tailored to IW by introducing a water-plane constraint into an error-state Kalman filter (ESKF) pipeline. Let $\mathcal{P}^{\text{raw}} = \{\mathbf{p}_i \in \mathbb{R}^d\}$, $d \in \{3, 4, 5\}$ denote a raw LiDAR point cloud in the local frame. Each point is defined as $\mathbf{p}_i = (x_i, y_i, z_i, I_i, r_i)$, where (x_i, y_i, z_i) are the spatial coordinates, I_i is the intensity value, and r_i denotes an additional sensor-specific attribute. Depending on the LiDAR model, r_i may represent, the ring index or a tag field. We propagate a constant-velocity model and motion-compensate the scan to obtain a deskewed cloud \mathcal{P} . From \mathcal{P} , water-surface candidates are extracted using reflection cues when available (e.g., ring and intensity) or a geometry-based slicing strategy otherwise, and a plane is subsequently fitted and temporally regularized. The resulting plane parameters define a water-plane measurement update that constrains sensor height and tilt, after which a scan-to-submap ICP measurement update further refines the 3D position and attitude in the global frame.

A. Prediction Step

Let G and B denote the global and body frames. We employ an ESKF with a right-multiplicative attitude error following [26]. The nominal state is:

$$\mathbf{x}_k \triangleq [{}^G\mathbf{q}_k^\top \quad {}^G\mathbf{p}_k^\top \quad {}^B\mathbf{v}_k^\top \quad {}^B\boldsymbol{\omega}_k^\top]^\top, \quad (1)$$

where ${}^G\mathbf{q}_k$ is a unit quaternion representing the attitude of B in G , ${}^G\mathbf{p}_k \in \mathbb{R}^3$ is the position in G , and ${}^B\mathbf{v}_k, {}^B\boldsymbol{\omega}_k \in \mathbb{R}^3$ are the linear and angular velocities expressed in B . The error state is defined as:

$$\delta\mathbf{x}_k \triangleq [\delta\boldsymbol{\theta}_k^\top \quad \delta\mathbf{p}_k^\top \quad \delta\mathbf{v}_k^\top \quad \delta\boldsymbol{\omega}_k^\top]^\top, \quad (2)$$

where $\delta\boldsymbol{\theta}_k \in \mathbb{R}^3$ is the small-angle attitude perturbation associated with the right-multiplicative quaternion error. With sampling interval Δt , the predicted nominal state $\hat{\mathbf{x}}_k$ is propagated under a constant body-frame velocity model as:

$${}^G\hat{\mathbf{q}}_k = {}^G\mathbf{q}_k \otimes \text{Exp}({}^B\boldsymbol{\omega}_k \Delta t), \quad (3)$$

$${}^G\hat{\mathbf{p}}_k = {}^G\mathbf{p}_k + {}^G\hat{\mathbf{R}}_k {}^B\mathbf{v}_k \Delta t, \quad (4)$$

$${}^B\hat{\mathbf{v}}_k = {}^B\mathbf{v}_k + \mathbf{w}_{v,k}, \quad {}^B\hat{\boldsymbol{\omega}}_k = {}^B\boldsymbol{\omega}_k + \mathbf{w}_{\omega,k}, \quad (5)$$

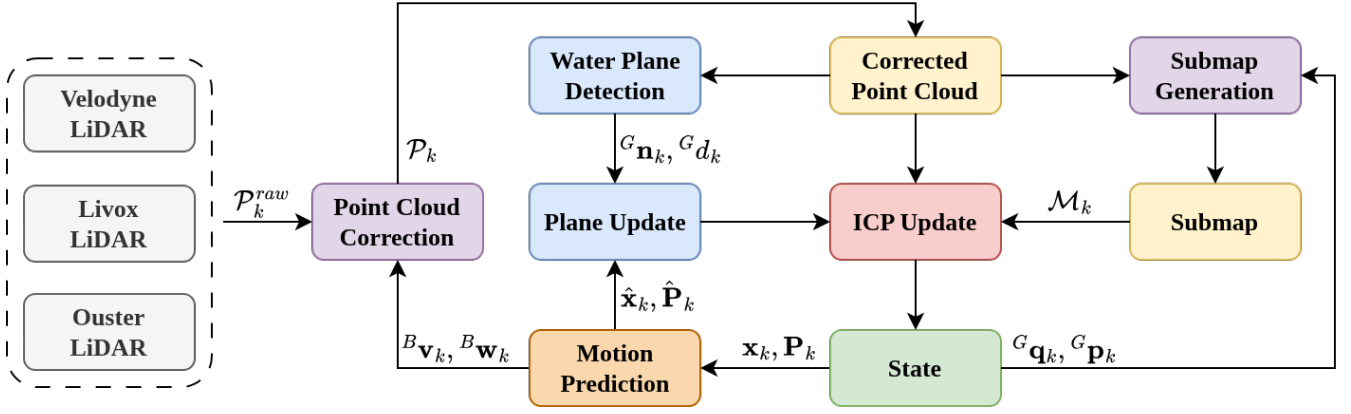


Fig. 1: System Overview: Incoming LiDAR point cloud is motion-compensated using the predicted state, followed by a water-plane detection and a scan-to-submap ICP update, producing the final state and an incrementally maintained submap.

where \otimes denotes quaternion multiplication, ${}^G\hat{\mathbf{R}}_k = \mathbf{R}({}^G\hat{\mathbf{q}}_k)$ is the rotation matrix corresponding to quaternion ${}^G\hat{\mathbf{q}}_k$, $\text{Exp}(\cdot)$ maps a rotation vector to a unit quaternion according to [26], and $\mathbf{w}_{v,k}, \mathbf{w}_{\omega,k}$ are zero-mean discrete-time Gaussian process noises. Using the predicted motion, we deskew the raw point cloud $\mathcal{P}_k^{\text{raw}}$ following [9] to obtain \mathcal{P}_k for the subsequent updates. The error covariance is propagated as:

$$\hat{\mathbf{P}}_k = \mathbf{F}_k \mathbf{P}_k \mathbf{F}_k^\top + \mathbf{Q}_k, \quad (6)$$

with first-order linearization:

$$\mathbf{F}_k = \begin{bmatrix} \mathbf{I} - [{}^B\boldsymbol{\omega}_k]_{\times} \Delta t & \mathbf{0} & \mathbf{0} & \mathbf{I} \Delta t \\ -{}^G\hat{\mathbf{R}}_k [{}^B\mathbf{v}_k]_{\times} \Delta t & \mathbf{I} & {}^G\hat{\mathbf{R}}_k \Delta t & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} \end{bmatrix}, \quad (7)$$

and process noise covariance:

$$\mathbf{Q}_k = \text{diag}(\mathbf{0}_{3 \times 3}, \mathbf{0}_{3 \times 3}, \sigma_v^2 \mathbf{I}_3, \sigma_\omega^2 \mathbf{I}_3), \quad (8)$$

where \mathbf{I} and $\mathbf{0}$ denote 3×3 identity and zero matrices, and σ_v^2 and σ_ω^2 are the random-walk variances on ${}^B\mathbf{v}_k$ and ${}^B\boldsymbol{\omega}_k$, respectively.

B. Water Plane Update

Given the deskewed LiDAR point cloud \mathcal{P}_k , we extract a subset $\mathcal{W}_k \subseteq \mathcal{P}_k$ containing water-surface candidate points prior to plane estimation. Depending on the available LiDAR attributes, \mathcal{W}_k is obtained using one of the two strategies described next.

Reflection-based filtering: if per-point intensity $i(\mathbf{p})$ and ring index $r(\mathbf{p})$ are available, we retain points from a set of valid rings $\mathcal{R}_{\text{valid}}$ and apply a low-intensity threshold:

$$\mathcal{W}_k = \{\mathbf{p} \in \mathcal{P}_k \mid r(\mathbf{p}) \in \mathcal{R}_{\text{valid}}, i(\mathbf{p}) \leq I_{\text{max}}\}. \quad (9)$$

Here, $\mathcal{R}_{\text{valid}}$ denotes the five selected downward-looking ring indices and $I_{\text{max}} = 1.0$. If $|\mathcal{W}_k| \geq N_{\text{min}}$ (we use $N_{\text{min}} = 30$), \mathcal{W}_k is passed to the plane estimator; otherwise, we fall back to the geometry-based selection.

Geometry-based selection: when reflection cues are unavailable or insufficient, the point cloud is partitioned into $\Delta x = 1$ m slices along the forward axis (body-frame x):

$$\mathcal{S}_j = \{\mathbf{p} \in \mathcal{P}_k \mid j\Delta x \leq p_x < (j+1)\Delta x\}, \quad (10)$$

where $j \in \mathbb{Z}$ indexes the slices. Each slice is split into left and right groups using p_y ,

$$\mathcal{S}_j^+ = \{\mathbf{p} \in \mathcal{S}_j \mid p_y \geq 0\}, \quad (11)$$

$$\mathcal{S}_j^- = \{\mathbf{p} \in \mathcal{S}_j \mid p_y < 0\}, \quad (12)$$

$$\mathcal{W}_k = \bigcup_{j \in \mathbb{Z}} (\text{lowest}_K(\mathcal{S}_j^+) \cup \text{lowest}_K(\mathcal{S}_j^-)), \quad (13)$$

where $\text{lowest}_K(\cdot)$ returns the K points with smallest p_z (we use $K = 5$).

Given \mathcal{W}_k , we estimate the water plane in the global frame by fitting

$${}^G\mathbf{n}_k^\top \mathbf{p} + {}^G d_k = 0, \quad \|{}^G\mathbf{n}_k\| = 1, \quad (14)$$

using RANSAC with an inlier threshold of 0.15 m and 30 iterations. To enforce temporal consistency, we align the normal direction with the previous estimate and validate the update using bounds on the tilt angle (2°) and plane offset change (0.2 m). Invalid fits are rejected.

The estimated plane parameters $({}^G\mathbf{n}_k, {}^G d_k)$ define a water-plane measurement used in the ESKF update:

$$\mathbf{z}_k^w \triangleq \begin{bmatrix} {}^G d_k \\ {}^G\mathbf{n}_k \end{bmatrix}, \quad \mathbf{h}^w(\hat{\mathbf{x}}_k) \triangleq \begin{bmatrix} {}^G\hat{p}_{k,z} \\ {}^G\hat{\mathbf{R}}_k \mathbf{e}_z \end{bmatrix}, \quad (15)$$

where $\mathbf{e}_z = [0 \ 0 \ 1]^\top$. Linearizing w.r.t. the error state $\delta\mathbf{x}_k$ yields:

$$\mathbf{H}_k^w = \begin{bmatrix} \mathbf{0}_{1 \times 3} & \mathbf{e}_z^\top & \mathbf{0}_{1 \times 3} & \mathbf{0}_{1 \times 3} \\ -{}^G\hat{\mathbf{R}}_k[\mathbf{e}_z]_{\times} & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} \end{bmatrix}. \quad (16)$$

The update is:

$$\mathbf{K}_k^w = \hat{\mathbf{P}}_k \mathbf{H}_k^{w\top} \left(\mathbf{H}_k^w \hat{\mathbf{P}}_k \mathbf{H}_k^{w\top} + \boldsymbol{\Sigma}_k^w \right)^{-1}, \quad (17)$$

$$\delta\hat{\mathbf{x}}_k = \mathbf{K}_k^w (\mathbf{z}_k^w - \mathbf{h}^w(\hat{\mathbf{x}}_k)), \quad (18)$$

$$\mathbf{P}_k = (\mathbf{I} - \mathbf{K}_k^w \mathbf{H}_k^w) \hat{\mathbf{P}}_k, \quad (19)$$

with $\Sigma_k^w = \text{diag}(\sigma_z^2, \sigma_n^2 \mathbf{I}_3)$, where σ_z^2 and σ_n^2 denote the measurement noise variances of the estimated plane offset ${}^G d_k$ and normal ${}^G \mathbf{n}_k$, respectively. The corrected nominal state is obtained by error-state injection:

$${}^G \mathbf{q}_k \leftarrow {}^G \hat{\mathbf{q}}_k \otimes \text{Exp}(\delta \hat{\boldsymbol{\theta}}_k), \quad (20)$$

$${}^G \mathbf{p}_k \leftarrow {}^G \hat{\mathbf{p}}_k + \delta \hat{\mathbf{p}}_k. \quad (21)$$

C. ICP update

After the water plane update, we refine the state by aligning the current deskewed scan against a local submap. Following [9] we adopt a scan-to-local-submap ICP formulation, which is typically more robust than scan-to-scan matching in the presence of noise and partial structure. Using the current estimate $({}^G \mathbf{q}_k, {}^G \mathbf{p}_k)$ as an initial guess, we transform the deskewed scan \mathcal{P}_k to global frame

$${}^G \mathcal{P}_k \triangleq \{ {}^G \mathbf{R}_k {}^B \mathbf{p} + {}^G \mathbf{p}_k \mid {}^B \mathbf{p} \in \mathcal{P}_k \}, \quad (22)$$

Let \mathcal{M}_k be the submap in global frame. Before registration, both ${}^G \mathcal{P}_k$ and \mathcal{M}_k are voxel-downsampled with a voxel size of 0.3m. We then run ICP between ${}^G \mathcal{P}_k$ and \mathcal{M}_k , obtaining a global-frame attitude and position measurement $({}^G \mathbf{R}_k^{\text{icp}}, {}^G \mathbf{p}_k^{\text{icp}})$. The ICP residual is formed as:

$$\mathbf{r}_k^{\text{icp}} \triangleq \begin{bmatrix} \text{Log} \left({}^G \mathbf{R}_k^{\text{icp} \top} {}^G \mathbf{R}_k^{\text{icp}} \right) \\ {}^G \mathbf{p}_k^{\text{icp}} - {}^G \mathbf{p}_k \end{bmatrix}, \quad (23)$$

and incorporated as an ESKF measurement update on $\delta \mathbf{x}_k$ with Jacobian:

$$\mathbf{H}_k^{\text{icp}} = \begin{bmatrix} \mathbf{I}_3 & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_3 & \mathbf{0} & \mathbf{0} \end{bmatrix} \in \mathbb{R}^{6 \times 12}. \quad (24)$$

We incorporate (23) with an ESKF update analogous to the water-plane update, using $\mathbf{H}_k^{\text{icp}}$ and Σ_k^{icp} , yielding the updated state $({}^G \mathbf{q}_k, {}^G \mathbf{p}_k)$. The body-frame velocities used in the next prediction step are then updated from consecutive state estimates as:

$${}^B \mathbf{v}_k = {}^G \mathbf{R}_{k-1}^{\top} \frac{{}^G \mathbf{p}_k - {}^G \mathbf{p}_{k-1}}{\Delta t}, \quad (25)$$

$${}^B \boldsymbol{\omega}_k = \frac{1}{\Delta t} \text{Log}({}^G \mathbf{R}_{k-1}^{\top} {}^G \mathbf{R}_k), \quad (26)$$

where ${}^G \mathbf{R}_k = \mathbf{R}({}^G \mathbf{q}_k)$. Lastly, the final state is then used to update the submap by re-transforming \mathcal{P}_k with (22) and fusing ${}^G \mathcal{P}_k$ into \mathcal{M}_k . To bound submap size, we remove submap points whose distance to the current sensor position exceeds 1.5 times the LiDAR's maximum range.

III. EXPERIMENTAL SETUP

To validate the proposed method, we collected data on Berlin's IW using a manually operated research vessel represented in Fig. 2. The platform carries multiple sensing modalities, but this work uses only the roof-mounted LiDAR suite, consisting of a Velodyne VLP-32, an Ouster OS0-128, and a Livox Avia. Additionally, three roof-mounted GNSS antennas record raw observations that are post-processed to generate a 6-DoF reference trajectory following [27], which is used as ground truth.

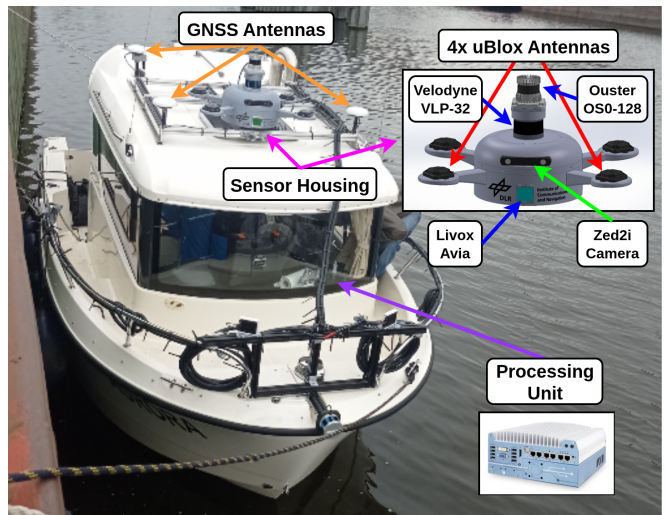


Fig. 2: Vessel sensor setup

Extrinsic calibration between the LiDARs and GNSS antennas was performed with a total station to align odometry estimates with the GNSS-derived ground truth. All sensor streams were recorded in ROS 2 on an onboard industrial Neosys Nuvo-7000 PC operating Ubuntu 22.04, which also monitored sensor status during collection. All three LiDARs operated at 10 Hz for consistency.

TABLE I: Overview of the evaluation scenes.

Scene	Duration (min)	Distance (km)	Description
scene_1	21.9	3.53	Vegetation, loop closure.
scene_2	18.6	2.65	Urban, bridge underpasses.
scene_3	14.8	2.09	Urban & vegetation, bridge underpasses.
scene_4	16.4	2.14	Port, loop closure, bridge underpasses
scene_5	19.5	2.28	Narrow channel, bridge underpasses
scene_6	12.0	1.84	Urban & vegetation, bridge underpasses.

For evaluation, we selected six representative scenes summarized in Table I. Each scene includes synchronized measurements from all three LiDARs, enabling a direct cross-sensor comparison.

IV. RESULTS

Building on Section III, we quantify odometry accuracy on six large-scale IW scenes listed in Table I, using the post-processed GNSS trajectory as ground truth. The trajectories average 2.4 km, highlighting long-term drift. We compare the proposed method against KISS-ICP, a state-of-the-art LiDAR-only odometry baseline, and report results consistently across three widely used sensors: Velodyne VLP-32, Ouster OS0-128, and Livox Avia. Table II summarizes absolute and relative trajectory errors (ATE and RTE) [28] for both position and orientation. RTE is computed using a fixed-length sliding window of 10 m along the trajectory. Fig. 3 complements these metrics with trajectories and error-time plots for two representative scenes.

As shown in Fig. 3(a,e), KISS-ICP exhibits pronounced 3D drift, consistent with its high ATE across scenes. In

TABLE II: Absolute and relative trajectory errors per scene (ATE, RTE): position (m) and orientation (deg).

Error	Scene	kiss-icp						ours					
		velodyne vlp32		ouster os0-128		livox avia		velodyne vlp32		ouster os0-128		livox avia	
		pos (m)	ori (deg)	pos (m)	ori (deg)	pos (m)	ori (deg)	pos (m)	ori (deg)	pos (m)	ori (deg)	pos (m)	ori (deg)
ATE	scene_1	336.45	9.08	75.99	7.07	92.73	2.81	19.91	0.97	11.58	0.95	10.67	1.35
	scene_2	348.10	15.15	132.96	9.43	240.52	16.09	15.60	1.71	18.56	2.57	7.73	2.02
	scene_3	227.15	13.06	44.95	4.03	115.34	3.28	20.15	2.36	3.34	2.15	15.30	2.76
	scene_4	123.96	10.75	107.70	9.06	23.45	4.25	19.73	2.57	2.74	2.16	7.02	2.43
	scene_5	614.67	25.07	72.83	6.41	154.96	4.42	18.54	2.61	9.51	4.25	17.89	3.27
	scene_6	137.22	6.43	465.62	18.54	55.26	3.97	10.14	1.99	9.21	3.99	6.01	2.71
RTE	scene_1	2.46	3.64	1.90	2.00	1.13	0.92	1.35	1.64	3.19	0.84	0.36	0.81
	scene_2	3.13	1.31	2.10	3.16	4.32	3.70	0.61	0.67	1.36	2.90	0.35	0.92
	scene_3	2.79	2.31	0.81	1.75	1.51	1.67	1.47	2.16	0.34	1.61	0.82	1.51
	scene_4	2.48	2.39	2.83	2.11	1.11	1.78	1.15	2.11	0.47	1.51	0.72	1.36
	scene_5	7.57	6.44	2.26	2.38	1.80	2.02	1.12	3.04	0.76	2.05	0.92	1.69
	scene_6	1.59	1.63	7.39	1.99	1.42	1.49	1.59	1.32	0.98	1.51	0.81	1.07

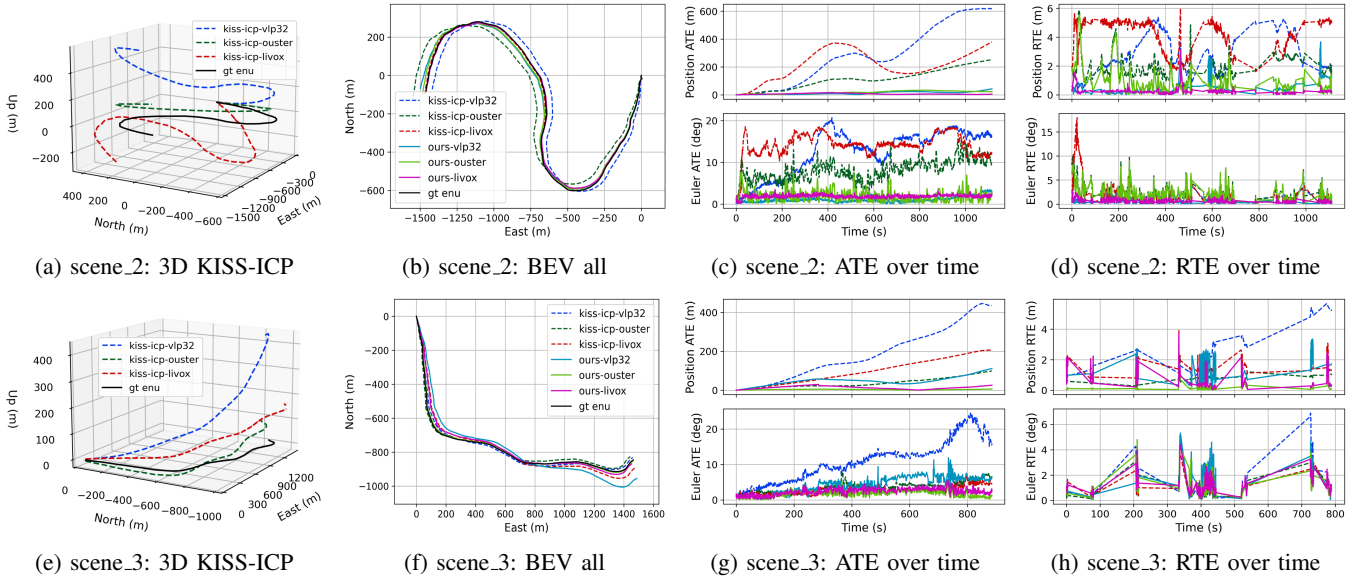


Fig. 3: Trajectory and error analysis on two IW scenes. We report 3D trajectories illustrating KISS-ICP drift, bird eye view (BEV) trajectories for all methods and sensors against ground truth, and the temporal evolution of ATE and RTE

IW conditions, water-surface returns, which play a role analogous to ground returns in road environments, are often sparse, noisy, and affected by specular reflections, providing limited constraints on height and tilt. At the same time, river banks supply most of the usable structure but are typically observed at longer ranges than roadside features, which weakens geometric leverage and increases sensitivity to ranging noise during scan alignment. Together, these factors make unmodeled vertical drift a dominant failure mode for pure scan matching in IW data. By enforcing the estimated water plane as a measurement update, our pipeline constrains elevation and tilt, reducing drift accumulation and yielding consistently lower ATE and RTE in Table II.

This behavior is also visible in the qualitative plots. The BEV projections in Fig. 3(b,f) indicate that both KISS-ICP and our method remain relatively close to the ground-truth path in the horizontal plane, with some deviation expected for LiDAR-only odometry over kilometer-scale trajectories. The 3D views in Fig. 3(a,e) show that the dominant error mode of

KISS-ICP is vertical rather than planar, with elevation drift accumulating over time, whereas the proposed water-plane constraint largely suppresses this growth. This is consistent with the temporal error curves: in Fig. 3(c,g), the ATE of KISS-ICP increases steadily, while our method exhibits a considerably slower growth rate. Similarly, Fig. 3(d,h) shows larger and more variable RTE for KISS-ICP, whereas our approach maintains lower RTE, indicating improved short-horizon consistency alongside reduced long-term drift.

Across sensors, differences are more pronounced for position than for orientation. The velodyne generally yields higher position errors than the ouster and livox, while orientation errors are of similar magnitude across all three sensors. The relative ranking between ouster and livox varies across the six scenes, indicating scenario-dependent performance driven by scene layout and the availability of nearby structure. Overall, the results suggest that higher measurement density, either through increased channel count or denser sampling within the field of view, is advantageous

in IW settings, although the best sensor choice depends on operating conditions.

These findings indicate that, although sensor characteristics affect absolute performance, incorporating a water-plane measurement mitigates a dominant IW failure mode of LiDAR-only scan matching by reducing drift that arises from weak vertical observability.

V. CONCLUSIONS

This paper presented a lightweight LiDAR-only odometry pipeline tailored to inland waterways by exploiting the water surface as a geometric prior. Plane parameters estimated from water-surface detection are fused as a measurement update in an error-state Kalman filter, and a scan-to-submap ICP update further refines the full 6-DoF pose. Experiments on a large-scale inland-waterway dataset show consistent improvements over the state-of-the-art LiDAR-only baseline KISS-ICP across all three LiDAR sensors, reducing absolute position error by more than 80% and absolute orientation error by more than 60% on average. Future work will extend the pipeline to multi-LiDAR odometry, explore LiDAR-IMU integration for improved motion modeling and observability, and incorporate robust estimation to reduce the impact of outliers and improve scan-to-map correspondences in challenging IW conditions.

ACKNOWLEDGMENT

We thank our colleagues for their contributions to data collection and platform operation.

REFERENCES

- [1] CCNR, European Commission, “Annual Report 2025: Inland Navigation in Europe Market Observation,” 2025.
- [2] C. Alias and J. Z. Felde, “Evaluating the economic performance of a decentralized waterborne container transportation service using autonomous inland vessels,” in *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, 2022, pp. 3571–3576.
- [3] J. M. Nadales, D. M. de la Peña, D. Limon, and T. Alamo, “Safe optimal vessel planning on natural inland waterways,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 9, pp. 9378–9389, 2023.
- [4] European Maritime Safety Agency (EMSA), “Annual overview of marine casualties and incidents 2025,” 2025.
- [5] D. A. Medina, C. Lass, E. P. Marcos, R. Ziebold, P. Closas, and J. García, “On gnss jamming threat from the maritime navigation perspective,” *2019 22th International Conference on Information Fusion (FUSION)*, pp. 1–7, 2019. [Online]. Available: <https://api.semanticscholar.org/CorpusID:207808717>
- [6] T. Qin, P. Li, and S. Shen, “VINS-Mono: A Robust and Versatile Monocular Visual-Inertial State Estimator,” *IEEE Transactions on Robotics*, vol. 34, no. 4, p. 1004–1020, Aug. 2018.
- [7] C. Campos, R. Elvira, J. J. G. Rodriguez, J. M. M. Montiel, and J. D. Tardos, “ORB-SLAM3: An Accurate Open-Source Library for Visual, Visual-Inertial, and Multimap SLAM,” *IEEE Transactions on Robotics*, vol. 37, no. 6, p. 1874–1890, Dec. 2021.
- [8] K. Chen, R. Nemiroff, and B. T. Lopez, “Direct LiDAR-Inertial Odometry: Lightweight LIO with Continuous-Time Motion Correction,” *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3983–3989, 2023.
- [9] I. Vizzo, T. Guadagnino, B. Mersch, L. Wiesmann, J. Behley, and C. Stachniss, “KISS-ICP: In Defense of Point-to-Point ICP – Simple, Accurate, and Robust Registration If Done the Right Way,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 8, no. 2, pp. 1029–1036, 2023.

- [10] D. C. Herraiez, M. Zeller, D. Wang, J. Behley, M. Heidingsfeld, and C. Stachniss, “Rai-slam: Radar-inertial slam for autonomous vehicles,” *IEEE Robotics and Automation Letters*, vol. 10, no. 6, pp. 5257–5264, 2025.
- [11] J. Seok, S. Kim, J. Jo, J. Lee, M. Jung, and K. Jo, “Radar4voxmap: Accurate odometry from blurred radar observations,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*, 2025, pp. 6206–6212.
- [12] J. Zhang and S. Singh, “LOAM : Lidar Odometry and Mapping in real-time,” *Robotics: Science and Systems Conference (RSS)*, pp. 109–111, 01 2014.
- [13] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, and D. Rus, “LIO-SAM: Tightly-coupled lidar inertial odometry via smoothing and mapping,” in *2020 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2020, pp. 5135–5142.
- [14] W. Xu, Y. Cai, D. He, J. Lin, and F. Zhang, “Fast-lho2: Fast direct lidar-inertial odometry,” *IEEE Transactions on Robotics*, vol. 38, no. 4, pp. 2053–2073, 2022.
- [15] J. Lee, R. Komatsu, M. Shinozaki, T. Kitajima, H. Asama, Q. An, and A. Yamashita, “Switch-slam: Switching-based lidar-inertial-visual slam for degenerate environments,” *IEEE Robotics and Automation Letters*, vol. 9, no. 8, pp. 7270–7277, 2024.
- [16] M. Nissov, N. Khedekar, and K. Alexis, “Degradation resilient lidar-radar-inertial odometry,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 8587–8594.
- [17] I. Filip, J. Pyo, M. Lee, and H. Joe, “Lidar slam with a wheel encoder in a featureless tunnel environment,” *Electronics*, vol. 12, no. 4, 2023.
- [18] D. Chung, J. Kim, C. Lee, and J. Kim, “Pohang canal dataset: A multimodal maritime dataset for autonomous navigation in restricted waters,” *The International Journal of Robotics Research*, vol. 42, no. 12, p. 1104–1114, Aug. 2023.
- [19] Y. Cheng, M. Jiang, J. Zhu, and Y. Liu, “Are We Ready for Unmanned Surface Vehicles in Inland Waterways? The USVInland Multisensor Dataset and Benchmark,” 2021.
- [20] I. Filip, C. Lass, and D. Medina, “Performance Evaluation of Visual and LiDAR SLAM for Inland Waterways,” in *2024 24th International Conference on Control, Automation and Systems (ICCAS)*, 2024, pp. 562–567.
- [21] W. Shen, Z. Yang, C. Yang, and X. Li, “A lidar slam-assisted fusion positioning method for usvs,” *Sensors*, vol. 23, no. 3, 2023.
- [22] Y. Wang, C. Liu, J. Liu, J. Wang, J. Liu, K. Zheng, and R. Zheng, “A laser-based slam algorithm of the unmanned surface vehicle for accurate localization and mapping in an inland waterway scenario,” *Journal of Marine Science and Engineering*, vol. 12, no. 12, 2024.
- [23] L. Hösch, A. Llorente, X. An, J. P. Llerena, and D. Medina, “High definition mapping for inland waterways: Techniques, challenges and prospects,” in *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*, 2023, pp. 6034–6041.
- [24] W. Wang, T. Shan, P. Leoni, D. Fernández-Gutiérrez, D. Meyers, C. Ratti, and D. Rus, “Roboat ii: A novel autonomous surface vessel for urban environments,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 1740–1747.
- [25] Z. Luo, Y. Wang, J. Swevers, P. Slaets, and H. Bruyninckx, “Inland-loam: Voxel-based structural semantic lidar odometry and mapping for inland waterway navigation,” 2025. [Online]. Available: <https://arxiv.org/abs/2508.03672>
- [26] J. Solà, “Quaternion kinematics for the error-state kalman filter,” 2017. [Online]. Available: <https://arxiv.org/abs/1711.02508>
- [27] D. Medina, J. Vilà-Valls, A. Hesselbarth, R. Ziebold, and J. García, “On the recursive joint position and attitude determination in multi-antenna GNSS platforms,” *Remote Sensing*, vol. 12, no. 12, p. 1955, 2020.
- [28] Z. Zhang and D. Scaramuzza, “A tutorial on quantitative trajectory evaluation for visual(-inertial) odometry,” in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2018.