

FAR-AVIO: Fast and Robust Schur-Complement Based Acoustic-Visual-Inertial Fusion Odometry With Sensor Calibration

Hao Wei¹, Graduate Student Member, IEEE, Peiji Wang, Qianhao Wang²,
Yang Gu¹, Graduate Student Member, IEEE, Tong Qin³, Senior Member, IEEE, Fei Gao², Member, IEEE,
and Yulin Si¹, Member, IEEE

I. INTRODUCTION

Abstract—Underwater environments impose severe challenges to visual-inertial odometry systems, as strong light attenuation, marine snow and turbidity, together with weakly exciting motions, degrade inertial observability and cause frequent tracking failures over long-term operation. While tightly coupled acoustic-visual-inertial fusion, typically implemented through an acoustic Doppler Velocity Log (DVL) integrated with visual-inertial measurements, can provide accurate state estimation, the associated graph-based optimization is often computationally prohibitive for real-time deployment on resource-constrained platforms. Here we present FAR-AVIO, a Schur-Complement based, tightly coupled acoustic-visual-inertial odometry framework tailored for underwater robots. FAR-AVIO embeds a Schur complement formulation into an Extended Kalman Filter(EKF), enabling joint pose-landmark optimization for accuracy while maintaining constant-time updates by efficiently marginalizing landmark states. On top of this backbone, we introduce Adaptive Weight Adjustment and Reliability Evaluation(AWARE), an online sensor health module that continuously assesses the reliability of visual, inertial and DVL measurements and adaptively regulates their sigma weights, and we develop an efficient online calibration scheme that jointly estimates DVL-IMU extrinsics, without dedicated calibration manoeuvres. Numerical simulations and real-world underwater experiments consistently show that FAR-AVIO outperforms state-of-the-art underwater SLAM baselines in both localization accuracy and computational efficiency, enabling robust operation on low-power embedded platforms. Our implementation has been released as open source software.

Index Terms—Acoustic-Visual-Inertial fusion, error state Kalman filter, underwater robotics.

Received 24 December 2025; accepted 9 April 2026. Date of publication 27 April 2026; date of current version 15 May 2026. This article was recommended for publication by Associate Editor M. J. Islam and Editor A. Kim upon evaluation of the reviewers' comments. This work was supported in part by the National Key R&D Program of China under Grant 2023YFC2809300, in part by the National Natural Science Foundation of China under Grant 52571313, and in part by Zhejiang Provincial Natural Science Foundation under Grant LZ24E090001. (Corresponding author: Yulin Si.)

Hao Wei, Peiji Wang, Yang Gu, and Yulin Si are with the State Key Laboratory of Ocean Sensing, and Ocean College, Zhejiang University, Zhoushan 316021, China (e-mail: isweihao@zju.edu.cn; peijiwang@zju.edu.cn; yanggu@zju.edu.cn; yulinsi@zju.edu.cn).

Qianhao Wang and Fei Gao are with the State Key Laboratory of Industrial Control Technology, and the Institute of Cyber-Systems and Control, Zhejiang University, Hangzhou 310027, China (e-mail: qhwangaa@zju.edu.cn; fgaoaa@zju.edu.cn).

Tong Qin is with the Global Institute of Future Technology, Shanghai Jiao Tong University, Shanghai 200240, China (e-mail: qingtong@sjtu.edu.cn).

Digital Object Identifier 10.1109/LRA.2026.3688065

MARINE robotics, such as Autonomous Underwater Vehicles (AUVs) and Remotely Operated Vehicles (ROVs), have become indispensable platforms for subsea infrastructure inspection, offshore energy maintenance, and ocean exploration [1], [2]. For these missions, accurate and drift-bounded localization is essential, yet satellite navigation is unavailable in underwater environments and external positioning infrastructure is typically absent or severely constrained in practice. Cameras, as compact and low-cost payloads, are therefore increasingly adopted on underwater robots, and visual Simultaneous Localization and Mapping (SLAM) has become a natural candidate for precise underwater localization.

State-of-the-art visual and visual-inertial SLAM systems, such as ORB-SLAM3 [3], VINS-Mono [4], DM-VIO [5], OV2-SLAM [6], OKVIS2-X [7] have demonstrated impressive performance in ground and aerial scenarios with abundant visual texture and well-excited inertial motion. However, underwater environments present substantial challenges to these methods. Rapid light attenuation, poor illumination, and dense “marine snow” often cause visual degradation that persists for tens of seconds to minutes, exceeding the temporal horizon over which pure visual tracking remains reliable [2]. Moreover, underwater robots frequently operate with gentle, quasi-static motions, leading to low accelerometer signal-to-noise ratios and difficult IMU initialization, as a result, conventional visual-inertial pipelines struggle to maintain consistent accuracy during prolonged visual degradation [8], [9].

To mitigate these limitations, numerous underwater SLAM and navigation frameworks have incorporated additional proprioceptive and exteroceptive sensors, such as DVL, depth/pressure sensors and sonar. Recent tightly coupled visual-inertial-DVL [10] and sonar-based [11] systems have demonstrated remarkable robustness and accuracy in challenging underwater scenarios. However, most of these approaches rely on large-scale graph-based optimization or factor-graph SLAM backends, which incur substantial computational overhead and are difficult to deploy in real time on resource-constrained embedded platforms. In contrast, existing filter-based multi-sensor fusion methods typically simplify visual and DVL information into low-dimensional pose or velocity measurements [12], [13],




Fig. 1. Real-world deployment of the proposed system in underwater environments. (a) illustrates the experimental setup, where an ROV inspects the target underwater structure. (b), (c), and (d) presents typical visual challenges encountered underwater, including motion blur, longtime textureless regions, and marine snow. (e) shows the dense reconstruction result and estimated trajectory produced by FAR-AVIO.

without fully exploiting point-level constraints or performing online calibration of multi-sensor extrinsics, which limits their ultimate accuracy and long-term consistency.

These observations highlight a persistent gap between accuracy and efficiency in current underwater localization systems. To mitigate this gap, this letter presents FAR-AVIO, a **Fast and Robust Acoustic-Visual-Inertial Odometry with online calibration and an AWARE module**. FAR-AVIO is inspired by Schur-Complement based sliding-window filters [14], [15], and employs the Schur complement within an EKF framework, enabling joint pose-landmark optimization while leveraging landmark independence for efficient, constant-time updates. This design bridges the gap between optimization-based accuracy and filter-based efficiency. The main contributions of this work are as follows:

- We propose the first Schur-Complement based tightly coupled Acoustic-Visual-Inertial odometry framework that rigorously models DVL measurements from doppler-shift principles and embeds them in filter based backend optimization, while jointly performing online DVL-IMU extrinsic calibration, thereby achieving high localization accuracy and real-time performance on resource-constrained underwater platforms without dedicated calibration procedures.
- FAR-AVIO introduces an AWARE module, performing an online health score mechanism that dynamically adjusts sensor fusion sigma scale based on real-time reliability assessment, enabling robust operation under sensor degradation and failure conditions.
- FAR-AVIO has been validated through extensive experiments including numerical simulations, and real-world underwater experiments, demonstrating superior accuracy and computational efficiency compared to state-of-the-art methods (see running result example in Fig. 1). To benefit the research community, we release the complete implementation as open-source software.

The proposed system architecture is illustrated in Fig. 2, and the rest of the paper is organized as follows: Section II reviews related literature. Section III describes the proposed FAR-AVIO framework, including the Schur-Complement based visual update, DVL measurement update, online sensor calibration, and




Fig. 2. Proposed system architecture of FAR-AVIO. Input is from an underwater robot equipped with a stereo camera, an IMU, and a DVL.

the AWARE module. Section IV presents experimental evaluation on real-world underwater datasets, Section V concludes the paper and discusses future work.

II. RELATED WORK

We briefly discuss the related works in this section. Visual and visual-inertial SLAM have been extensively studied in terrestrial and aerial domains [3], [4], [5], [6], [16], [17], [18], [19]. These approaches provide a strong baseline when visual texture is abundant and inertial excitation is sufficient, but their performance degrades severely in underwater environments as discussed in the Section I. In the following, we focus on underwater multi-sensor fusion methods, with particular emphasis on DVL-based localization.

Early underwater SLAM frameworks combined a DVL, stereo vision and gyroscopic measurements in a loosely integrated pipeline [13], [20], where DVL outputs served only as external velocity priors to the visual estimator. Although this configuration can deliver reasonable pose estimates in difficult underwater scenes, it neither exploits accelerometer data nor models gyroscope bias, which leads to drifting roll and pitch and ultimately limits the long-term consistency of the solution. To improve this, several tightly coupled visual-DVL and visual-inertial-DVL frameworks have been proposed. A visual-DVL fusion method in [21] directly injects DVL velocities into a factor-graph backend to jointly optimize camera poses and DVL measurements, but does not incorporate IMU data, restricting robustness under rapid motion and severe visual degradation. A filter-based tightly coupled visual-inertial-DVL odometry method [12] jointly fuses all three modalities, yet assumes known, fixed extrinsics, making it sensitive to calibration drift or hardware changes.

More recent work extends DVL fusion to richer multi-sensor SLAM systems. A graph-based LiDAR-VI-DVL framework for autonomous surface vehicles [22] achieves accurate localization on water-surface trajectories, but is tailored to 2-D surface motion with strong LiDAR constraints and is not directly applicable to fully submerged 3-D trajectories. A tightly coupled visual-inertial-acoustic system [23] and AQUA-SLAM [10] further integrate DVL, camera and IMU in ORB-SLAM3-style backends with online extrinsic calibration, and the latter provides

the Tank dataset [2] for benchmarking. However, these factor-graph-based methods entail significant computational cost and are challenging to deploy on resource-constrained platforms for real-time applications.

III. FILTER DESCRIPTION

A. State Definition and Propagation

Following the state formulation in [24] and augmenting it with sensor extrinsic parameters, we define the system state as

$$\mathbf{x}_b = [\mathbf{p}_b^w \ \mathbf{v}_b^w \ \mathbf{R}_b^w \ \mathbf{b}_a \ \mathbf{b}_g \ \mathbf{T}_c^b \ \mathbf{T}_D^b]^\top, \quad (1)$$

where $\mathbf{p}_b^w$, $\mathbf{v}_b^w$, and $\mathbf{R}_b^w$ denote the position, velocity, and orientation of the body frame $\{b\}$ expressed in the world frame $\{w\}$, respectively. The vectors $\mathbf{b}_a$ and $\mathbf{b}_g$ represent the accelerometer and gyroscope biases of the IMU. The terms $\mathbf{T}_c^b$ and $\mathbf{T}_D^b$ denote the extrinsic transformations from the body frame $\{b\}$ to the camera frame $\{c\}$ and the DVL frame $\{D\}$, respectively, each consisting of a rotation and translation (e.g., $\mathbf{T}_c^b = \{\mathbf{R}_c^b, \mathbf{p}_c^b\}$). Accordingly, the associated error-state vector $\delta\mathbf{x}_b$ is defined as

$$\delta\mathbf{x}_b = [\delta\mathbf{p}_b^w \ \delta\mathbf{v}_b^w \ \delta\boldsymbol{\theta}_b^w \ \delta\mathbf{b}_a \ \delta\mathbf{b}_g \ \delta\mathcal{T}_c^b \ \delta\mathcal{T}_D^b]^\top. \quad (2)$$

For non-rotational part, the standard additive error model $\mathbf{x} = \hat{\mathbf{x}} + \delta\mathbf{x}$ is adopted, where $\hat{\mathbf{x}}$ denotes the nominal state. For rotational components, perturbations are defined on SO(3) using the approximation

$$\mathbf{R} = \text{Exp}([\delta\boldsymbol{\theta}]_\times) \hat{\mathbf{R}}, \quad (3)$$

where $\delta\boldsymbol{\theta} \in \mathbb{R}^3$ is minimal rotation error representation, $[\cdot]_\times$ denotes skew-symmetric operator, and $\text{Exp}(\cdot)$ is the matrix exponential operation. $\delta\mathcal{T} \in \mathbb{R}^6$ is the minimal perturbations of the camera and DVL extrinsic transformations, each stacking a translation and a rotation error state.

The continuous-time dynamics of the nominal state are given by (4)

$$\begin{aligned} \dot{\mathbf{p}}_b^w &= \hat{\mathbf{v}}_b^w, \ \dot{\mathbf{b}}_a = \mathbf{0}, \ \dot{\mathbf{b}}_g = \mathbf{0} \\ \dot{\mathcal{T}}_a^b &= \mathbf{0}, \ \dot{\mathcal{T}}_d^b = \mathbf{0} \\ \dot{\mathbf{v}}_b^w &= \hat{\mathbf{R}}_b^w (\tilde{\mathbf{a}} - \hat{\mathbf{b}}_a) + \mathbf{g}^w \\ \dot{\mathbf{R}}_b^w &= \hat{\mathbf{R}}_b^w [(\tilde{\boldsymbol{\omega}} - \hat{\mathbf{b}}_g)]_\times, \end{aligned} \quad (4)$$

where $\tilde{\mathbf{a}}$ and $\tilde{\boldsymbol{\omega}}$ are the raw accelerometer and gyroscope measurements, respectively, and $\mathbf{g}^w$ is the gravity vector in the world frame. The error-state propagation can be derived by linearizing the (4), yielding

$$\dot{\delta\mathbf{x}}_b = \mathbf{F}\delta\mathbf{x}_b + \mathbf{G}\mathbf{n}_b, \quad (5)$$

where $\mathbf{F}$ is the state transition Jacobian, $\mathbf{G}$ is the noise input matrix. $\mathbf{n}_b = [\mathbf{n}_a^\top \ \mathbf{n}_{aw}^\top \ \mathbf{n}_g^\top \ \mathbf{n}_{gw}^\top]^\top$ stacks the processes noise, $\mathbf{n}_a$ and $\mathbf{n}_g$ denote the Gaussian white noise of the accelerometer and gyroscope measurements, respectively, while $\mathbf{n}_{aw}$ and $\mathbf{n}_{gw}$ model the random-walk processes driving the accelerometer and gyroscope biases.

B. Visual Measurement and Equivalent Residual Model

To enable multi-view constraints, we maintain a fixed-size sliding window of N keyframes by cloning the current IMU pose when a new frame arrives; if it becomes a keyframe, the oldest frame is marginalized, otherwise the second-newest is discarded. For each landmark observed across these keyframes, assume that the j -th 3D landmark $\hat{\boldsymbol{\xi}}_j^w \in \mathbb{R}^3$, expressed in the world frame w , is observed by the camera at keyframe i . The projection model can be written as

$$\hat{\mathbf{z}}_{ij} = \pi(\hat{\mathbf{x}}_b, \hat{\boldsymbol{\xi}}_j^w), \quad (6)$$

where $\hat{\mathbf{z}}_{ij} \in \mathbb{R}^2$ is the predicted pixel coordinate of the landmark in the image plane, $\pi(\cdot)$ denotes the camera projection function that maps the 3D landmark in the world frame to 2D pixel coordinates based on the current state estimate $\hat{\mathbf{x}}_b$. Linearizing around the nominal state yields the reprojection residual

$$\mathbf{r}_{ij} = \mathbf{z}_{ij} - \hat{\mathbf{z}}_{ij} \simeq \mathbf{H}_{x,ij} \delta\mathcal{X} + \mathbf{H}_{f,ij} \delta\boldsymbol{\xi}_j^w + \mathbf{n}_{ij}, \quad (7)$$

where $\mathbf{r}_{ij}$ and $\mathbf{z}_{ij}$ are the reprojection residual and visual measurement from the frontend track result, respectively, $\mathbf{n}_{ij}$ is the measurement noise, $\mathbf{H}_{x,ij}$ and $\mathbf{H}_{f,ij}$ are the Jacobians of the projection function with respect to the system state and the landmark position, and $\delta\mathcal{X}$ and $\delta\boldsymbol{\xi}_j^w$ are the state and landmark position perturbations, respectively. Their explicit forms are given by

$$\begin{aligned} \mathbf{H}_{x,ij} &= \mathbf{J}_\pi \left[(\hat{\mathbf{R}}_c^b)^\top [\hat{\boldsymbol{\xi}}_j^b]_\times \mathbf{R}_{b_i}^{w\top} \quad -\mathbf{R}_{c_i}^{w\top} \right], \\ \mathbf{H}_{f,ij} &= [\mathbf{J}_\pi \mathbf{R}_{c_i}^{w\top}], \end{aligned} \quad (8)$$

$\mathbf{J}_\pi$ denotes the Jacobian of the camera pixel projection function with respect to the 3D point in the camera frame, and its explicit form depends on the specific camera model. Stacking all observations across multiple keyframes in sliding window with (7), we obtain

$$\mathbf{r} = \mathbf{H}_x \delta\mathcal{X} + \mathbf{H}_f \delta\boldsymbol{\xi} + \mathbf{n}, \quad (9)$$

where $\mathbf{r}$ and $[\mathbf{H}_x \ \mathbf{H}_f]$ are the stacked residuals and Jacobians, respectively. The stacked noise term is $\mathbf{n} = [u, u, \dots, u]^\top$ with covariance $\mathcal{R} = \text{diag}(u^2, u^2, \dots, u^2)$. Then, we directly project the measurement model of (9) into the Jacobian space $[\mathbf{H}_x \ \mathbf{H}_f]^\top$ to construct an equivalent observation model as

$$\begin{bmatrix} \mathbf{H}_x^\top \\ \mathbf{H}_f^\top \end{bmatrix} \mathbf{r} = \begin{bmatrix} \mathbf{H}_x^\top \\ \mathbf{H}_f^\top \end{bmatrix} [\mathbf{H}_x \ \mathbf{H}_f] \begin{bmatrix} \delta\mathcal{X} \\ \delta\boldsymbol{\xi} \end{bmatrix} + \mathbf{n}', \quad (10)$$

$\mathbf{n}'$ is the equivalent observation noise, the corresponding covariances are given by

$$\mathcal{R}' = \begin{bmatrix} \mathbf{H}_x^\top \\ \mathbf{H}_f^\top \end{bmatrix} \mathcal{R} [\mathbf{H}_x \ \mathbf{H}_f]. \quad (11)$$

Equations (10) and (11) thus define an equivalent observation model and noise model in Jacobian space that fully preserve

the information from all visual measurements, effectively eliminating the original measurement-dimension limitation. For clarity, (10) can be expanded as

$$\begin{bmatrix} \mathbf{H}_x^\top \mathbf{r} \\ \mathbf{H}_f^\top \mathbf{r} \end{bmatrix} = \begin{bmatrix} \mathbf{H}_x^\top \mathbf{H}_x & \mathbf{H}_x^\top \mathbf{H}_f \\ \mathbf{H}_f^\top \mathbf{H}_x & \mathbf{H}_f^\top \mathbf{H}_f \end{bmatrix} \begin{bmatrix} \delta \mathbf{x} \\ \delta \mathbf{f} \end{bmatrix} + \underbrace{\begin{bmatrix} \mathbf{n}' \\ \mathbf{n}_1' \\ \mathbf{n}_2' \end{bmatrix}}_{\mathbf{n}''}. \quad (12)$$

Due to the limitation of the dimension in the system, the filter state does not include the landmark perturbation $\delta \mathbf{f}$, a standard observation model must be constructed by marginalizing $\delta \mathbf{f}$. To this end, adopt a Schur-Complement based elimination [14] to eliminate the landmark state from the observation model in (12), yielding

$$[\mathbf{b}_1 - \mathbf{C}_2 \mathbf{C}_3^{-1} \mathbf{b}_2] = [\mathbf{C}_1 - \mathbf{C}_2 \mathbf{C}_3^{-1} \mathbf{C}_2^\top] \delta \mathbf{x} + \mathbf{n}_1'', \quad (13a)$$

$$\mathbf{R}_1'' = [\mathbf{C}_1 - \mathbf{C}_2 \mathbf{C}_3^{-1} \mathbf{C}_2^\top] \mathbf{u}^2, \quad (13b)$$

where $\mathbf{n}_1''$ is the resulting equivalent observation noise with covariance $\mathbf{R}_1''$. Equations (13a) and (13b) define an equivalent observation model and noise model that depend only on the error state $\delta \mathbf{x}$, fully preserving the information from all visual measurements while successfully marginalizing out the landmark state $\delta \mathbf{f}$. The resulting equivalent residual can then be inserted into the standard EKF update procedure. After the EKF pose update, we perform a lightweight Gauss-Newton step to update each landmark's 3D position ξ_j^w by minimizing its multi-view reprojection residuals across the sliding window. This decoupled point update keeps landmarks geometrically consistent with the updated state without expanding the filter dimension.

C. DVL Measurement Model

1) *Single-Beam Doppler Velocity*: A DVL transducer emits a narrowband acoustic wave of known carrier frequency f_t and receives the echo reflected by the seabed or water-borne scatterers. Let f_r denote the received frequency and $\Delta f = f_r - f_t$ denote measured doppler shift. Under the standard narrowband, small-velocity assumption $|v_r| \ll c_s$, where c_s denotes the speed of sound in water, the classic monostatic doppler relation yields the radial velocity v_r along the acoustic beam as

$$v_r \approx -\frac{c_s}{2f_t} \Delta f. \quad (14)$$

Here the sign convention is chosen such that $v_r > 0$ corresponds to the vehicle moving towards the seabed along the beam direction. For each of the four DVL beams (indexed by $i = 1, \dots, 4$, as shown in Fig. 3), let $\tilde{v}_i$ denote the scalar radial velocity obtained by converting its measured doppler shift via (14). Ideally, $\tilde{v}_i$ coincides with the true radial velocity $v_{r,i}$ along that beam, in practice, it is corrupted by noise and occasional outliers due to low signal-to-noise ratio or loss of bottom lock. We model the single-beam measurement as

$$\tilde{v}_i = v_{r,i} + n_i, \quad n_i \sim \mathcal{N}(0, \sigma_i^2), \quad (15)$$

where n_i denotes zero-mean Gaussian measurement noise with variance σ_i^2 .




Fig. 3. DVL transducer measurements are illustrated in both 2-D and 3-D views. The instrument comprises four transducers oriented in different directions, with Transducer 1 shown as a representative example.

2) *DVL Velocity Measurement Model*: The linear velocity at the DVL acoustic center, expressed in the DVL frame $\{D\}$, as

$${}^D \tilde{\mathbf{v}} = [{}^D \tilde{v}_x, {}^D \tilde{v}_y, {}^D \tilde{v}_z]^\top. \quad (16)$$

Following the DVL geometry in Fig. 3, the orientation of each beam $\mathbf{e}_i$ with respect to the frame $\{D\}$ has been parameterized by a fixed tilt angle α from the horizontal plane $x_D y_D$ and an azimuth angle β_i about the z_D -axis. The azimuth β_i is measured in the $x_D y_D$ -plane from the x_D -axis to the projection of the i -th beam. For a four-beam Janus configuration, the beams share the same tilt α and have azimuths β_i . The example $\mathbf{e}_i$ of transducer 1 can be expressed as

$$\mathbf{e}_1 = [\cos \beta_1 \cos \alpha, \sin \beta_1 \cos \alpha, \sin \alpha]^\top. \quad (17)$$

Then the true radial velocity measured by this beam is simply the projection of the 3D velocity onto the beam direction:

$$v_{r,i} = \mathbf{e}_i^\top {}^D \tilde{\mathbf{v}}. \quad (18)$$

Substituting (18) into (15), the single-beam DVL measurement model can be written as

$$\tilde{v}_i = \mathbf{e}_i^\top {}^D \tilde{\mathbf{v}} + n_i, \quad (19)$$

which states that each transducer provides an observation of the vehicle velocity component along its own acoustic axis.

Stacking the four scalar beam measurements into a vector ${}^b \tilde{\mathbf{v}} = [\tilde{v}_1, \tilde{v}_2, \tilde{v}_3, \tilde{v}_4]^\top$ and collecting the beam direction vectors into a matrix, the single-beam model (19) can be written in compact form as

$${}^b \tilde{\mathbf{v}} = \mathbf{E} {}^D \tilde{\mathbf{v}} + \mathbf{n}_b, \quad (20)$$

where $\mathbf{E} = [\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, \mathbf{e}_4]^\top \in \mathbb{R}^{4 \times 3}$ is the beam direction matrix, and $\mathbf{n}_b = [n_1, n_2, n_3, n_4]^\top$ is the stacked noise vector. The matrix $\mathbf{E}$ is full column rank (rank 3) as long as the beams are not coplanar, which holds for standard DVL configurations. Therefore, the 3D velocity ${}^D \tilde{\mathbf{v}}$ can be uniquely determined from the four beam measurements by solving the overdetermined linear system (20) in a least-squares sense:

$${}^D \tilde{\mathbf{v}} = (\mathbf{E}^\top \mathbf{E})^{-1} \mathbf{E}^\top {}^b \tilde{\mathbf{v}}, \quad (21)$$

where the matrix inverse $(\mathbf{E}^\top \mathbf{E})^{-1}$ depends only on the known beam geometry and can be pre-computed offline.

Since we are assuming independent Gaussian noise on each beam in (15), we have

$$\mathbf{n}_b \sim \mathcal{N}(\mathbf{0}, \Sigma_b), \quad \Sigma_b = \text{diag}(\sigma_1^2, \sigma_2^2, \sigma_3^2, \sigma_4^2). \quad (22)$$

Substituting (20) into (21) and separating the true velocity and noise terms yields

$${}^D\tilde{\mathbf{v}} = {}^D\mathbf{v}_{\text{true}} + \mathbf{A}\mathbf{n}_b, \quad \mathbf{A} = (\mathbf{E}^\top \mathbf{E})^{-1} \mathbf{E}^\top, \quad (23)$$

where ${}^D\mathbf{v}_{\text{true}}$ denotes the true DVL-frame velocity. Since $\mathbf{n}_b$ is zero-mean Gaussian, its linear image $\mathbf{A}\mathbf{n}_b$ is also zero-mean Gaussian. Consequently, the estimated DVL-frame velocity remains Gaussian:

$${}^D\tilde{\mathbf{v}} \sim \mathcal{N}({}^D\mathbf{v}_{\text{true}}, \Sigma_D), \quad \Sigma_D = \mathbf{A} \Sigma_b \mathbf{A}^\top. \quad (24)$$

In the common case of identical per-beam variance $\sigma_i^2 = \sigma^2$, this simplifies to

$$\Sigma_D = \sigma^2 (\mathbf{E}^\top \mathbf{E})^{-1}, \quad (25)$$

which provides a convenient closed-form expression for the covariance of the DVL-frame velocity measurement.

3) *DVL Residual for ESKF State Update*: Given the estimated state, the predicted DVL-frame velocity is computed as

$${}^D\hat{\mathbf{v}} = (\hat{\mathbf{R}}_D^b)^\top \left((\hat{\mathbf{R}}_b^w)^\top \hat{\mathbf{v}}_b^w + [{}^b\hat{\boldsymbol{\omega}}]_\times \hat{\mathbf{p}}_D^b \right), \quad (26)$$

where ${}^b\hat{\boldsymbol{\omega}}$ is the bias-corrected angular velocity in the body frame, $\hat{\mathbf{p}}_D^b$ is the IMU–DVL lever arm expressed in the body frame, and $\hat{\mathbf{R}}_D^b$ is the rotation from DVL to body frame. Thus we can get the DVL residual for the state update, according to the (26) and (21)

$$\mathbf{r}_{\text{DVL}} = {}^D\tilde{\mathbf{v}} - {}^D\hat{\mathbf{v}}. \quad (27)$$

This residual is then linearized with respect to the error state and used in the standard ESKF update with measurement covariance Σ_D .

D. AWARE Module: Reliability Evaluation

Most fusion-based filters assume fixed, time-invariant measurement noise for each sensor, which is rarely valid in practice: visual quality changes with texture, illumination and motion, and DVL measurements degrade under poor bottom lock, scattering or flow disturbances. If such variations are ignored, bursts of bad measurements from one modality can corrupt the entire estimate even when other sensors remain reliable. AWARE addresses this by continuously assessing visual and DVL quality, adapting their effective covariances, and temporarily disabling severely degraded sensors, thus preventing any single faulty source from dominating the fusion.

For each sensor $s \in \{\text{VIS}, \text{DVL}\}$, AWARE maintains a reliability scale σ_s and a fixed-length queue Q_s of recent “unhealthy” events. At every measurement, a sensor-specific quality score $q_s \in [0, 1]$ is computed (q_{VIS} from feature tracking statistics and reprojection error, q_{DVL} from velocity consistency and DVL

Algorithm 1: AWARE Update for a Generic Sensor s .

Input: sensor stream $\{(t, \mathbf{z}_s)\}$, unhealthy events Q_s .
Output: adapted scale σ_s .

```

1 foreach measurement  $(t, \mathbf{z}_s)$  do
2    $q_s \leftarrow \text{QualityScore}_s(\mathbf{z}_s)$ 
3   if  $\text{enabled}_s$  then
4     if  $q_s \geq \tau_s$  then
5       // healthy
6        $\mathbf{R}_s^{\text{eff}} \leftarrow \mathbf{R}_s / \sigma_s^2$ 
7       ESKF_UPDATE( $\mathbf{z}_s, \mathbf{R}_s^{\text{eff}}$ )
8     else
9       // unhealthy
10      append  $(t, q_s)$  to  $Q_s$ 
11       $\sigma_s \leftarrow \gamma_s \sigma_s$ 
12      if  $|Q_s| > N_s$  then
13        remove oldest element from  $Q_s$ 
14      if  $|Q_s| = N_s$  and  $\text{span}(Q_s) < \Delta T_s$  then
15         $\text{enabled}_s \leftarrow \text{false}$ 
16         $\sigma_s \leftarrow 1, \quad Q_s \leftarrow \emptyset$ 
17      ESKF_UPDATE( $\mathbf{z}_s, \sigma_s \mathbf{R}_s$ )
18 else
19   // sensor disabled
20   if  $q_s \geq \tau_s^{\text{rec}}$  then
21      $\text{enabled}_s \leftarrow \text{true}$ 
22      $\sigma_s \leftarrow 1, \quad Q_s \leftarrow \emptyset$ 
```

residuals). These scores drive covariance scaling and sensor gating decisions, as summarized in Algorithm 1.

E. Pose Prior Visual Frontend Tracking

Our visual frontend detects sparse Shi–Tomasi [25] corners on an image pyramid and tracks them via pyramidal Lucas–Kanade [26] optical flow, initialized with IMU-predicted pose priors to improve robustness under fast motion and blur. For stereo matching, candidate correspondences are constrained to a small window along the epipolar line, which is critical for long-baseline setups where unconstrained optical flow fails due to large disparities.

IV. EXPERIMENTS

A. Experiment Settings and Datasets

In this section, we evaluate the proposed FAR-AVIO on the public Tank dataset [2], which provides synchronized stereo, IMU, DVL and depth measurements collected in a wave tank. Accurate ground-truth (GT) camera poses are generated by the TankGT pipeline using AprilTag markers mounted on the underwater structure, enabling quantitative benchmarking under realistic underwater conditions. The eight sequences are grouped into three trajectory types (Structure, HalfTank, and WholeTank) with varying difficulty levels (Easy/Medium/Hard) determined by vehicle speed, lighting and the amount of textureless area.

TABLE I
AVERAGE TRANSLATIONAL RMSE/STD (M) ON TANK SEQUENCES

Sequence	FAR-AVIO	FAR-AIO	AQUA-SLAM	UVA-SLAM	SVIN2	ORB-SLAM3	VINS-Fusion
Structure Easy	0.11 / 0.03	0.13 / 0.04	0.07 / 0.03	0.21 / 0.12	0.09 / 0.03	0.28 / 0.09	0.13 / 0.04
Structure Medium	0.16 / 0.05	0.36 / 0.18	0.18 / 0.08	0.54 / 0.30	2.94 / 1.64	3.30 / 1.08	NaN / NaN
Structure Hard	0.13 / 0.04	0.14 / 0.04	0.50 / 0.24	0.50 / 0.23	3.26 / 1.43	2.73 / 1.45	NaN / NaN
HalfTank Easy	0.19 / 0.10	0.22 / 0.10	0.28 / 0.17	1.69 / 1.10	6.01 / 4.33	2.69 / 1.45	1.24 / 0.30
HalfTank Medium	0.33 / 0.11	0.48 / 0.14	0.29 / 0.14	0.44 / 0.22	3.40 / 1.76	0.74 / 0.38	NaN / NaN
HalfTank Hard	0.25 / 0.10	0.33 / 0.13	0.36 / 0.22	0.58 / 0.37	77.6 / 55.07	1.10 / 0.70	NaN / NaN
WholeTank Medium	0.34 / 0.15	0.53 / 0.23	0.52 / 0.28	1.34 / 0.73	0.72 / 0.41	1.18 / 0.71	13.08 / 7.22
WholeTank Hard	0.57 / 0.27	0.72 / 0.29	0.22 / 0.12	1.11 / 0.83	0.83 / 0.65	2.96 / 2.49	NaN / NaN
Total RMSE	2.08	2.91	2.42	6.41	94.85	14.98	NaN




Fig. 4. Different baseline methods estimated trajectory comparisons.

Fig. 1 illustrates typical visual challenges and running result in the HalfTank-Easy sequence.

B. Localization Performance Comparison

We have benchmarked FAR-AVIO against five representative baselines: AQUA-SLAM, UVA-SLAM, SVIN2, ORB-SLAM3, and VINS-Fusion. For fairness, all methods use the same camera/IMU intrinsic and extrinsic parameters, and are run in stereo-inertial (or stereo-inertial-DVL when available) configurations. Estimated trajectories are aligned with ground truth using the method described in [27], the average root-mean-square error (RMSE) and standard deviation (STD) of the absolute translation error (ATE), summarized in Table I. Entries marked as NaN indicate repeated tracking failure or divergence before sequence completion.

FAR-AVIO achieves the best average translational accuracy among all methods, consistently ranking first or second on individual sequences. On challenging sequences (Structure-Hard, HalfTank-Hard, WholeTank-Medium), FAR-AVIO reduces translational RMSE by up to 75% compared with AQUA-SLAM (e.g., 0.13 m vs. 0.50 m on Structure-Hard) and by over an order of magnitude compared with purely visual-inertial methods. Notably, FAR-AIO (our system with vision disabled) remains functional across all sequences, demonstrating the robustness of the acoustic-inertial fusion backbone even under complete visual failure. Fig. 4 visualizes representative trajectory comparisons, for methods that diverged, only the successfully tracked portion is shown.

In contrast, pure visual-inertial methods (ORB-SLAM3, VINS-Fusion, SVIN2) frequently incur large drift or tracking failure under turbidity and visual degradation, as reflected by




Fig. 5. CPU load and memory usage of FAR-AVIO and baselines on Structure-Easy sequence.

meter-level errors and NaN entries. DVL-aided baselines (UVA-SLAM, AQUA-SLAM) substantially reduce drift and achieve sub-meter accuracy on most sequences, yet still exhibit higher errors than FAR-AVIO on difficult settings.

C. Runtime and Computational Load

We evaluate runtime and computational load for methods with public implementations (ORB-SLAM3, VINS-Fusion, SVIN2, FAR-AVIO) on a desktop CPU (AMD Ryzen 9 7950X, 32 GB RAM) and an embedded platform (NVIDIA Jetson Orin NX, 8 GB RAM). AQUA-SLAM (and UVA-SLAM), which are built on the ORB-SLAM-based backend, are not included in this comparison, but their computational cost is expected to be comparable to or higher than ORB-SLAM3. Fig. 5 shows that FAR-AVIO consistently exhibits the lowest CPU utilization and memory footprint among all baselines.

On the Orin NX, the per-module breakdown in Table II shows that FAR-AVIO processes frames in 28.28 ms (≈ 35 Hz),

TABLE II
RUNTIME COMPARISON BETWEEN VINS-FUSION AND FAR-AVIO ON THE EMBEDDED PLATFORM

Module	VINS-Fusion		FAR-AVIO (Ours)	
	Time (ms)	Share	Time (ms)	Share
Visual frontend	27.69	45%	20.21	73%
Visual measure update	33.76	54%	6.08	21%
DVL measure update	—	—	0.78	2%
Others (I/O, etc.)	0.20	1%	1.21	4%
Total	61.65	100%	28.28	100%

achieving about a $2.2\times$ speed-up over VINS-Fusion (61.65 ms, ≈ 16 Hz). The main gain comes from the backend: VINS-Fusion spends 33.76 ms (54%) in visual optimization, whereas FAR-AVIO visual update requires only 6.08 ms (21%), and the additional DVL update adds merely 0.78 ms (2%). As a result, the visual frontend becomes the dominant cost in FAR-AVIO, confirming that the proposed filter based backend effectively removes the optimization bottleneck and is well suited for embedded deployment.

D. Ablation Study for AWARE Modular and Extrinsic Calibration

We conduct ablation experiments to quantify the contribution of the proposed AWARE module and the online calibration of the IMU–DVL extrinsics. The study includes both real Tank sequences and purely numerical simulations with synthetic IMU [16] and DVL measurements, where ground-truth extrinsics are known.

a) Effect of the AWARE module: To assess whether AWARE behaves as intended, we visualize the instantaneous confidence scales applied to visual and DVL measurements along the estimated trajectory. For two representative sequences, Structure–Easy (SE) and WholeTank–Hard (WH), we plot the trajectory and color each point by the corresponding visual and DVL $\Sigma_{\text{scale}} \in [0, 1]$, where values close to 1 indicate high confidence (nominal weighting) and values approaching 0 indicate strong down-weighting of that sensor, see Fig. 6(a) and Fig. 6(b), on the SE sequence, which features relatively clean water and stable lighting, the visual frontend rarely experiences long-term degradation. Accordingly, both the visual and DVL scales remain close to 1 along almost the entire trajectory, with only minor fluctuations (Fig. 6(a)). This indicates that AWARE does not introduce unnecessary re-weighting when all sensors operate nominally, and effectively reduces to a standard tightly coupled VI-DVL fusion scheme on easy sequences.

In contrast, the WH sequence is considerably more challenging, strong turbidity, non-uniform illumination, and large portions of the trajectory with weak or missing image features lead to extended periods of unreliable visual tracking. In these segments, the visual Σ_{scale} is significantly reduced along the corresponding parts of the trajectory, while the DVL scale remains close to 1 (Fig. 6(b)). This behavior shows that AWARE automatically down-weights visual updates when the frontend reports poor tracking quality, and simultaneously leans more on the DVL constraints to stabilize the state estimate. As visual




Fig. 6. AWARE-estimated sigma scaling factors for visual and DVL measurements on two Tank sequences.

TABLE III
ABLATION OF THE AWARE MODULE ON TANK SEQUENCES

Variant	SE	SM	SH	HE	HM	HH
FAR-AVIO (full)	0.11	0.16	0.13	0.19	0.33	0.25
FAR-AIO	0.24	0.21	0.14	0.39	F ¹	F
FAR-AVIO w/o AWARE	0.13	F	F	F	F	F

¹ F means tracking failure or filter divergence.

conditions recover, the visual scale smoothly returns towards 1, and the fusion reverts to a more balanced VI-DVL weighting.

To quantitatively assess the impact of AWARE on localization performance, we compare FAR-AVIO with and without AWARE on the Tank sequences, as well as FAR-AIO, a variant that fuses only IMU and DVL (no visual measurements) with AWARE enabled. As summarized in Table III, FAR-AVIO with AWARE achieves the lowest RMSE on all sequences and is the only variant that successfully completes every run, while both FAR-AIO and FAR-AVIO without AWARE experience failures on the more challenging cases.

b) Effect of IMU–DVL extrinsic calibration: We first validate the IMU-DVL extrinsic calibration convergence in a numerical simulation. Fig. 7 plots the evolution of the extrinsic errors over time for three different initializations: an identity transform (*Identity*), and two perturbed initial extrinsics with medium (*Noise Middle*) and small (*Noise Small*) noise levels. In all cases, the estimated IMU-DVL extrinsics converge toward the ground truth, with the final errors stabilizing at a small residual level. Larger initial perturbations understandably lead to longer convergence transients and slightly higher steady-state error, but the calibration remains stable and convergent even




Fig. 7. Extrinsic of rotation part calibration convergence from different initializations in numerical simulation.

TABLE IV
ABLATION OF THE EXTRINSIC CALIBRATION MODULE ON TANK AND SIMULATE SEQUENCES

Variant	SE	SH	Identity	NM ²	NS ²	Average
FAR-AIO	0.242	0.212	0.574	0.165	0.124	0.263
w/o calib ¹	0.323	0.217	9.372	3.546	2.848	8.152

¹ w/o calib mean FAR-AIO without extrinsic calibration.

² NM mean middle level noise; NS mean small level noise.

from the coarse identity initialization. This demonstrates that the proposed calibration scheme can reliably recover extrinsics from IMU-DVL motion without requiring a carefully tuned initial guess.

We further quantify the impact on trajectory accuracy, as shown in Table IV, both real Tank sequences and simulated sequences are initialized with perturbed extrinsics. On Tank data, calibration reduces RMSE by 10–25%, in simulation with larger perturbations, corrupted extrinsics cause RMSE to exceed 3–9 m without calibration, whereas online calibration recovers sub-meter accuracy (0.124–0.574 m). On average, enabling calibration reduces RMSE from 8.152 m to 0.263 m, confirming that the proposed module reliably recovers accurate extrinsics from noisy initial guesses. We note that, as observed in our experiments, sufficient rotational excitation in the vehicle motion is important for reliable convergence; under near-stationary or purely translational motion, convergence may be slow, and providing a reasonable offline initial estimate is advisable in such cases.

V. CONCLUSION

This letter presented FAR-AVIO, a fast and robust Schur-Complement based Acoustic-Visual-Inertial fusion odometry framework with online sensor calibration for underwater robots. Extensive evaluations on real world sequences and synthetic scenarios demonstrated that FAR-AVIO achieves competitive or superior localization accuracy compared with state-of-the-art underwater and terrestrial baselines, while requiring substantially lower CPU and memory resources and running comfortably in real time on embedded hardware.

REFERENCES

- [1] M. Ferrera, V. Creuze, J. Moras, and P. Trouvé-Peloux, “AQUALOC: An underwater dataset for visual-inertial-pressure localization,” *Int. J. Robot. Res.*, vol. 38, no. 14, pp. 1549–1559, 2019.
- [2] S. Xu et al., “Tank dataset: An underwater multi-sensor dataset for SLAM evaluation,” *Int. J. Robot. Res.*, vol. 45, no. 4, pp. 541–551, 2025.
- [3] C. Campos, R. Elvira, J. J. G. Rodríguez, J. M. Montiel, and J. D. Tardós, “ORB-SLAM3: An accurate open-source library for visual, visual-inertial, and multi-map SLAM,” *IEEE Trans. Robot.*, vol. 37, no. 6, pp. 1874–1890, Dec. 2021.
- [4] T. Qin, P. Li, and S. Shen, “VINS-Mono: A robust and versatile monocular visual-inertial state estimator,” *IEEE Trans. Robot.*, vol. 34, no. 4, pp. 1004–1020, Aug. 2018.
- [5] L. v. Stumberg and D. Cremers, “DM-VIO: Delayed marginalization visual-inertial odometry,” *IEEE Robot. Automat. Lett.*, vol. 7, no. 2, pp. 1408–1415, Apr. 2022.
- [6] M. Ferrera, A. Eudes, J. Moras, M. Sanfourche, and G. Le Besnerais, “OV²SLAM: A fully online and versatile visual SLAM for real-time applications,” *IEEE Robot. Automat. Lett.*, vol. 6, no. 2, pp. 1399–1406, Apr. 2021.
- [7] S. Boche, J. Jung, S. B. Laina, and S. Leutenegger, “OKVIS2-X: Open keyframe-based visual-inertial SLAM configurable with dense depth or LiDAR, and GNSS,” *IEEE Trans. Robot.*, vol. 41, pp. 6064–6083, 2025.
- [8] C. Hu, S. Zhu, Y. Liang, and W. Song, “Tightly-coupled visual-inertial-pressure fusion using forward and backward IMU preintegration,” *IEEE Robot. Automat. Lett.*, vol. 7, no. 3, pp. 6790–6797, Jul. 2022.
- [9] C. Hu, S. Zhu, Y. Liang, Z. Mu, and W. Song, “Visual-pressure fusion for underwater robot localization with online initialization,” *IEEE Robot. Automat. Lett.*, vol. 6, no. 4, pp. 8426–8433, Aug. 2021.
- [10] S. Xu, K. Zhang, and S. Wang, “AQUA-SLAM: Tightly coupled underwater acoustic-visual-inertial slam with sensor calibration,” *IEEE Trans. Robot.*, vol. 41, pp. 2785–2803, 2025.
- [11] S. Rahman, A. Quattrini Li, and I. Rekleitis, “SVIn2: A multi-sensor fusion-based underwater SLAM system,” *Int. J. Robot. Res.*, vol. 41, no. 11–12, pp. 1022–1042, Sep. 2022.
- [12] L. Zhao, M. Zhou, and B. Loose, “Tightly coupled visual-DVL-inertial odometry for robot-based ice-water boundary exploration,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2023, pp. 7127–7134.
- [13] E. Vargas et al., “Robust underwater visual SLAM fusing acoustic sensing,” in *Proc. IEEE Int. Conf. Robot. Autom.*, 2021, pp. 2140–2146.
- [14] G. Sibley, L. Matthies, and G. Sukhatme, “Sliding window filter with application to planetary landing,” *J. Field Robot.*, vol. 27, no. 5, pp. 587–608, Sep. 2010.
- [15] Y. Fan, T. Zhao, and G. Wang, “SchurVINS: Schur complement-based lightweight visual inertial navigation system,” in *Proc. IEEE Int. Conf. Pattern Recognit.*, 2024, pp. 17964–17973.
- [16] P. Geneva, K. Eckenhoff, W. Lee, Y. Yang, and G. Huang, “OpenVINS: A research platform for visual-inertial estimation,” in *Proc. IEEE Int. Conf. Robot. Autom.*, 2020, pp. 4666–4672.
- [17] K. Sun et al., “Robust stereo visual inertial odometry for fast autonomous flight,” *IEEE Robot. Automat. Lett.*, vol. 3, no. 2, pp. 965–972, Apr. 2018.
- [18] C. Forster, M. Pizzoli, and D. Scaramuzza, “SVO: Fast semi-direct monocular visual odometry,” in *Proc. IEEE Int. Conf. Robot. Autom.*, 2014, pp. 15–22.
- [19] S. Leutenegger, S. Lynen, M. Bosse, R. Siegwart, and P. Furgale, “Keyframe-based visual-inertial odometry using nonlinear optimization,” *Int. J. Robot. Res.*, vol. 34, no. 3, pp. 314–334, 2015.
- [20] S. Xu et al., “Underwater visual acoustic SLAM with extrinsic calibration,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2021, pp. 7647–7652.
- [21] Y. Huang et al., “Tightly-coupled visual-DVL fusion for accurate localization of underwater robots,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2023, pp. 8090–8095.
- [22] A. Thoms, G. Earle, N. Charron, and S. Narasimhan, “Tightly coupled, graph-based DVL/IMU fusion and decoupled mapping for SLAM-Centric maritime infrastructure inspection,” *IEEE J. Ocean. Eng.*, vol. 48, no. 3, pp. 663–676, Jul. 2023.
- [23] Y. Huang et al., “Visual-inertial-acoustic sensor fusion for accurate autonomous localization of underwater vehicles,” *IEEE Trans. Cybern.*, vol. 55, no. 2, pp. 880–896, Feb. 2025.
- [24] J. Solà, “Quaternion kinematics for the error-state Kalman filter,” Institut de Robotica i Informàtica Industrial, Universitat Politècnica de Catalunya, Barcelona, Spain, 2017.
- [25] J. Shi and C. Tomasi, “Good features to track,” in *Proc. IEEE Int. Conf. Pattern Recognit.*, 1994, pp. 593–600.
- [26] B. D. Lucas and T. Kanade, “An iterative image registration technique with an application to stereo vision,” in *Proc. Int. Joint Conf. Artif. Intell.*, 1981, pp. 674–679.
- [27] S. Umeyama, “Least-squares estimation of transformation parameters between two point patterns,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 4, pp. 376–380, Apr. 1991.