

GeVI-SLAM: Gravity-Enhanced Stereo Visual-Inertial SLAM for Underwater Robots

Yuan Shen^{1*}, Yuze Hong^{1*}, Guangyang Zeng^{1,2†}, Tengfei Zhang¹, Pui Yi Chui³, Ziyang Hong¹, Junfeng Wu¹

Abstract—Accurate visual-inertial simultaneous localization and mapping (VI SLAM) for underwater robots remains a significant challenge due to frequent visual degeneracy and insufficient inertial measurement unit (IMU) motion excitation. In this paper, we present GeVI-SLAM, a gravity-enhanced stereo VI SLAM system designed to address these issues. By leveraging the stereo camera’s direct depth estimation ability, we eliminate the need to estimate scale during IMU initialization, enabling stable operation even under low-acceleration dynamics. With precise gravity initialization, we decouple the pitch and roll from the pose estimation and solve a 4 degrees of freedom (DOF) Perspective-n-Point (PnP) problem for pose tracking. This allows the use of a minimal 3-point solver, which significantly reduces computational time to reject outliers within a Random Sample Consensus framework. We further propose a bias-eliminated 4-DOF PnP estimator with provable consistency, ensuring the relative pose converges to the true value as the feature number increases. To handle dynamic motion, we refine the full 6-DOF pose while jointly estimating the IMU covariance, enabling adaptive weighting of the gravity prior. Extensive experiments on simulated and real-world data demonstrate that GeVI-SLAM achieves higher accuracy and greater stability compared to state-of-the-art methods.

I. INTRODUCTION

Underwater robotic perception challenges conventional algorithms [1], [2]. For visual-inertial (VI) systems, beyond well-documented issues like poor illumination and turbidity [3], underwater scenes introduce a set of compounding difficulties that are detrimental to perception performance.

First, natural terrains and man-made structures often exhibit repetitive or weak textures, inflating feature ambiguity and yielding high outlier ratios in feature matching and tracking [4]; a frontend must therefore be substantially more robust than in typical terrestrial scenes. Second, scenes are often sparse and near-planar (e.g., flat seabeds and bridge piers), which results in geometric degeneracy and leads to unstable or incorrect state estimation [5]. Third, low accelerations due to high underwater resistance undermine system initialization: VINS-Mono [6] suffers a scale-gravity-bias ambiguity under near-constant velocity [7]; the stereo-inertial measurement unit (IMU) mode in ORB-SLAM3 [8] can recover scale but lacks excitation to accurately estimate IMU biases and gravity.

Contributions. To tackle the above challenges, we propose a gravity-enhanced underwater VI SLAM system. We

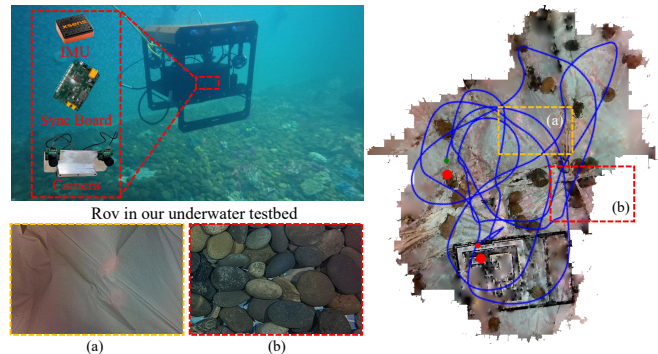


Fig. 1: Estimated trajectory (blue line) using the proposed GeVI-SLAM system, visualized on a dense 3D reconstruction. It demonstrates our algorithm’s robustness against common underwater challenges: (a) **feature-sparse, coplanar scenes** and (b) **highly repetitive patterns**.

note that although low accelerations hinder initialization yet offer a stable and precise gravity prior once the gravity has been well initialized. Hence, we first design a stereo-IMU initialization with known scale, minimizing preintegration residuals to robustly estimate gravity and IMU biases. Then, our main focus is to leverage IMU gravity prior to decouple 6-degrees of freedom (DOF) pose and build a 4-DOF (yaw angle and 3D translation) visual frontend, which is resilient to high outlier rates and geometric degeneracy. Furthermore, adaptive VI fusion gives rise to accurate full 6-DOF refinement. In summary, our contributions are three-fold:

- **Gravity-enhanced 4-DOF Perspective-n-Point (PnP):** By using gravity prior to reduce pose to 4-DOF, a minimal 3-point Random Sample Consensus (RANSAC) delivers faster and more robust outlier rejection (our required iteration number, relative to 5-point RANSAC, is on the order of the squared inlier ratio). In addition, our 4-DOF PnP estimator is proven to be consistent and performs stably in near-degenerate scenes.
- **Adaptive VI fusion:** We adopt a block coordinate descent (BCD) strategy to jointly refine 6-DOF pose and estimate gravity prior covariance, enabling adaptive weighting of the gravity constraint and effectively avoiding the drift of pitch and roll angles.
- **Comprehensive experiment validation:** Simulations demonstrate Cramér-Rao Lower Bound (CRLB)-tight accuracy, robust minimal-set consensus under up to 30% outliers, and drift-free pitch/roll estimation. Real water pool experiments show consistently lower Absolute Tra-

¹School of Data Science, Chinese University of Hong Kong, Shenzhen, Shenzhen, China. {yuanshen@link., junfengwu@cuhk.edu.cn

²School of Intelligence Science and Engineering, Harbin Institute of Technology, Shenzhen, Shenzhen, China.

³School of Life Sciences, Chinese University of Hong Kong, Hong Kong.

[†]Corresponding author: zgy180574@gmail.com

*These authors contributed equally to this work.

jectory Error (ATE) and Relative Pose Error (RPE) than ORB-SLAM3 [8], VINS-Fusion [6], and SVIN2 [4], and stable tracking where these baselines may fail.

Notations: Scalars use italic lowercase (a), vectors bold lowercase (\mathbf{a}), and matrices bold uppercase (\mathbf{A}); unless stated otherwise, vectors are columns. Estimated quantities are marked with a hat $\hat{\mathbf{x}}$, while true values have a superscript circle \mathbf{x}^o . Reference frames are denoted by calligraphic letters: \mathcal{W} (world), \mathcal{C} (camera), \mathcal{I} (IMU), and \mathcal{K} (keyframe). A rigid body transformation from frame \mathcal{A} to \mathcal{B} is denoted by ${}^{\mathcal{B}}_{\mathcal{A}}\mathbf{T} \triangleq [{}^{\mathcal{B}}_{\mathcal{A}}\mathbf{R} | {}^{\mathcal{B}}_{\mathcal{A}}\mathbf{t}] \in \text{SE}(3)$, where ${}^{\mathcal{B}}_{\mathcal{A}}\mathbf{R} \in \text{SO}(3)$ is the rotation and ${}^{\mathcal{B}}_{\mathcal{A}}\mathbf{t} \in \mathbb{R}^3$ is the translation. A point ${}^{\mathcal{A}}\mathbf{p}$ is transformed to frame \mathcal{B} via ${}^{\mathcal{B}}\mathbf{p} = {}^{\mathcal{B}}_{\mathcal{A}}\mathbf{R}{}^{\mathcal{A}}\mathbf{p} + {}^{\mathcal{B}}_{\mathcal{A}}\mathbf{t}$. The world frame \mathcal{W} is defined such that its z -axis is aligned with the gravity vector.

II. RELATED WORK

A. Underwater Visual SLAM

Underwater VI SLAM can be broadly categorized into two groups. Feature-based methods match sparse features and estimate motion via epipolar geometry or PnP; they are robust to viewpoint and illumination changes but can be computationally heavy and vulnerable in low-texture underwater scenes [9], [10]. In contrast, direct methods minimize photometric error over image intensities and are more efficient and tolerant to weak texture; however, they require good initialization and are sensitive to illumination changes [11], [12]. Recent tightly coupled underwater VI SLAM systems mainly adopt feature-based frontends fused with IMU in joint optimization. For example, Wang *et al.* [13] fused ORB points, diagonal segments, IMU preintegration, and loop closure to achieve low Root Mean Square Error (RMSE) in real time. Meanwhile, unified variants leverage both methods, e.g., UniVIO [14], which jointly optimize feature-based reprojection and direct photometric errors, along with inertial terms, in a unified stereo-IMU backend. Finally, some tightly coupled multi-sensor estimators combine cameras and IMU with more sensors, such as acoustic and depth sensors [4].

B. Leveraging Gravity Priors in SLAM

Gravity provides a global prior fixing roll and pitch, leaving yaw as the only rotational DOF. Lupton *et al.* [15] preintegrated IMU measurements to recover gravity and velocity, enabling constant-time VI SLAM. Li *et al.* [16] refined gravity on a unit sphere using vertical edges to decouple accelerometer bias and speed up initialization. Usenko *et al.* [17] tightly coupled IMU with a stereo camera such that roll and pitch are observable and alignment is robust under fast motion and low texture. GaRLIO [18] fused radar Doppler with LiDAR and IMU for a velocity-aware gravity estimate that suppresses vertical drift. Kubelka *et al.* [19] enforced gravity in LiDAR iterative closest point by fixing roll and pitch and solving only yaw and translation, reducing median normalized ATE by around 30%.

The existing work mainly uses gravity as a cost term in iterative optimization. In this paper, we not only incorporate the gravity to achieve reduced-degree visual tracking, which

allows a minimal-set closed-form solution, but also adaptively weight the gravity prior to account for time-varying motion acceleration.

III. GRAVITY-ENHANCED 4-DOF CAMERA POSE ESTIMATION

A. Euler Angles and Characterization of a 4-DOF Pose

Unless otherwise stated, a camera frame refers to the left camera frame of a stereo camera. We aim to estimate the relative pose ${}^{\mathcal{C}}_{\mathcal{K}}\mathbf{T} \in \text{SE}(3)$ from the camera keyframe \mathcal{K} to the current camera frame \mathcal{C} . This transformation is composed of a rotation matrix ${}^{\mathcal{C}}_{\mathcal{K}}\mathbf{R} \in \text{SO}(3)$ and a translation vector ${}^{\mathcal{C}}_{\mathcal{K}}\mathbf{t} \in \mathbb{R}^3$. The rotation ${}^{\mathcal{C}}_{\mathcal{K}}\mathbf{R}$ can be parameterized using three Euler angles: roll (ϕ), pitch (θ), and yaw (ψ). For our application, we decompose the rotation into two components: a rotation representing yaw, \mathbf{R}_{ψ} , and a rotation representing the combined roll and pitch, $\mathbf{R}_{\theta\phi}$. The full rotation is expressed as ${}^{\mathcal{C}}_{\mathcal{K}}\mathbf{R} = \mathbf{R}_{\psi}\mathbf{R}_{\theta\phi}$, where $\mathbf{R}_{\psi} = \begin{bmatrix} \cos \psi & \sin \psi & 0 \\ -\sin \psi & \cos \psi & 0 \\ 0 & 0 & 1 \end{bmatrix}$ and $\mathbf{R}_{\theta\phi} = \begin{bmatrix} \cos \theta & -\sin \theta \sin \phi & \sin \theta \cos \phi \\ 0 & \cos \phi & \sin \phi \\ -\sin \theta & -\cos \theta \sin \phi & \cos \theta \cos \phi \end{bmatrix}$.

This factorization is particularly useful because in a world frame \mathcal{W} , the xy -plane is aligned with the ground, and an IMU can estimate roll and pitch relative to gravity [20]. With roll and pitch known, only yaw remains, motivating a 4-DOF pose consisting of the yaw angle ψ and ${}^{\mathcal{C}}_{\mathcal{K}}\mathbf{t}$.

B. Acquisition of Drift-Free Roll and Pitch

To provide a reliable roll and pitch for our 4-DOF PnP solver, we fuse measurements from the IMU's gyroscope and accelerometer. Our approach is inspired by the complementary filtering [20], which optimally combines the strengths of both sensors. Gyroscopes yield accurate high-frequency attitude updates over short intervals but drift over time. Accelerometers provide a drift-free, low-frequency gravity reference in quasi-static motion (stationary or near-constant velocity) but are sensitive to linear acceleration. By integrating gyroscope measurements for dynamic tracking and using accelerometer data within a probabilistic optimization framework (detailed in Section IV) for drift correction, our system can maintain accurate pitch and roll angles.

C. Problem Formulation

As illustrated in Fig. 2, after feature tracking and outlier removal, we obtain 2D point correspondences $\{\mathbf{q}_i, \mathbf{z}_i, \mathbf{y}_i\}_{i \in \mathcal{I}}$, where \mathcal{I} is the inlier set, \mathbf{q}_i 's denote the points in the left image of the current frame, and \mathbf{z}_i 's and \mathbf{y}_i 's are the points in the left and right images of the keyframe, respectively. Unless otherwise stated, image 2D points are expressed using the normalized image coordinates. We assume the feature tracking errors are independent and identically distributed (i.i.d.) zero-mean Gaussian noises, i.e., $\mathbf{z}_i \sim \mathcal{N}(\mathbf{z}_i^o, \sigma^2 \mathbf{I}_2)$ and $\mathbf{y}_i \sim \mathcal{N}(\mathbf{y}_i^o, \sigma^2 \mathbf{I}_2)$. The zero-mean and isotropic covariance is reasonable because, after outlier rejection, the errors from intensity-based trackers tend to be unbiased and inherit the spatially uniform nature of physical sensor noise. This model is not only a good approximation

for pinhole cameras [21] but also practical, as its variance σ^2 can be consistently estimated from data [22].

Leveraging the 4-DOF parameterization introduced previously, our goal is to estimate the state vector $\chi = [\psi, {}^c_{\mathcal{K}}\mathbf{t}^\top]^\top$. From this state, the full relative rotation is constructed as ${}^c_{\mathcal{K}}\mathbf{R}(\psi) = \mathbf{R}_\psi \mathbf{R}_{\theta_\phi}$, where \mathbf{R}_{θ_ϕ} is the roll and pitch estimate obtained from IMU information. Here, we utilize point-to-epipolar-line distances as costs and construct a maximum likelihood (ML) estimation problem. For the i -th correspondence, we formulate two geometric costs. The first relates to the current frame observation \mathbf{q}_i and the keyframe's left image observation \mathbf{z}_i . The epipolar line $\mathbf{l}_{i,\mathcal{K}_L}$ in the left keyframe image is generated by back-projecting \mathbf{q}_i using the pose ${}^c_{\mathcal{K}}\mathbf{T} = [{}^c_{\mathcal{K}}\mathbf{R}(\psi) \mid {}^c_{\mathcal{K}}\mathbf{t}]$: $\mathbf{l}_{i,\mathcal{K}_L} = ([{}^c_{\mathcal{K}}\mathbf{t}] \times {}^c_{\mathcal{K}}\mathbf{R}(\psi))^\top \mathbf{q}_i^h$, where \mathbf{q}_i^h is the homogeneous coordinates of \mathbf{q}_i , and $[\cdot]_\times$ is the operator that maps a vector in \mathbb{R}^3 to its corresponding 3×3 skew-symmetric matrix, such that for any two vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^3$, $[\mathbf{a}]_\times \mathbf{b} = \mathbf{a} \times \mathbf{b}$. Then, the distance of point \mathbf{z}_i to epipolar line $\mathbf{l}_{i,\mathcal{K}_L}$ is $d_{i,\mathcal{K}_L} = \frac{\mathbf{l}_{i,\mathcal{K}_L}^\top \mathbf{z}_i^h}{\|\mathbf{E} \mathbf{l}_{i,\mathcal{K}_L}\|}$, where $\mathbf{E} = [\mathbf{e}_1, \mathbf{e}_2]^\top$, and \mathbf{e}_i denotes the unit vector in \mathbb{R}^3 whose i -th element is 1. The second cost, defined in the keyframe's right image, uses the pose ${}^c_{\mathcal{K}_R}\mathbf{T}$ obtained via the stereo extrinsic: ${}^c_{\mathcal{K}_R}\mathbf{T} = {}^c_{\mathcal{K}}\mathbf{T} ({}^{\mathcal{K}_R}_{\mathcal{K}}\mathbf{T})^{-1} = [{}^c_{\mathcal{K}_R}\mathbf{R} \mid {}^c_{\mathcal{K}_R}\mathbf{t}]$. The corresponding epipolar line is $\mathbf{l}_{i,\mathcal{K}_R} = ([{}^c_{\mathcal{K}_R}\mathbf{t}] \times {}^c_{\mathcal{K}_R}\mathbf{R})^\top \mathbf{q}_i^h$, and the distance is $d_{i,\mathcal{K}_R} = \frac{\mathbf{l}_{i,\mathcal{K}_R}^\top \mathbf{y}_i^h}{\|\mathbf{E} \mathbf{l}_{i,\mathcal{K}_R}\|}$.

Under the aforementioned noise assumptions on \mathbf{z}_i and \mathbf{y}_i , one can verify that the resulting point-to-epipolar-line distance errors follow the Gaussian distribution $\mathcal{N}(0, \sigma^2)$. Hence, the ML estimate for the 4-DOF relative pose χ is obtained by solving the following non-linear least-squares problem over all inlier correspondences \mathcal{I} :

$$\min_{\chi \in \mathbb{R}^4} \sum_{i \in \mathcal{I}} (d_{i,\mathcal{K}_L}^2 + d_{i,\mathcal{K}_R}^2). \quad (1)$$

The optimal solution to the ML problem (1) is called the ML estimate, denoted as $\hat{\chi}^{\text{ML}}$. Solving this non-linear optimization is challenging as the objective function is non-convex, making iterative solvers highly dependent on a good initial guess. To circumvent this issue, we propose linearizing the model to derive an efficient and consistent initial estimator.

D. Consistent 4-DOF Estimator

To initialize the non-linear optimization (1), we develop a consistent estimator for the 4-DOF relative pose. The process involves three main steps: uncertainty-aware 3D point triangulation, construction of a linear model, and a closed-form bias-eliminated (BE) estimator.

1) *Uncertainty-Aware 3D Point Triangulation*: We first use the method in [22] to obtain a consistent estimate $\hat{\sigma}^2$ of 2D noise variance σ^2 . It involves solving a generalized eigenvalue problem. For the details, one can refer to Appendix A of [22]. Then, following equations (1) and (2) in [21], we can triangulate the 3D points ${}^{\mathcal{K}}\mathbf{p}_i$'s from stereo observations in the previous keyframe along with their uncertainties $\Sigma_{\mathbf{p}_i}$'s. As a result, we establish 3D-2D point correspondences $\{\mathcal{K}\mathbf{p}_i, \mathbf{q}_i\}_{i \in \mathcal{I}}$, where the 3D points




Fig. 2: 4-DOF pose estimation based on point-to-epipolar-line distances.

${}^{\mathcal{K}}\mathbf{p}_i$'s are noisy with known covariance $\Sigma_{\mathbf{p}_i}$, while their corresponding 2D matches \mathbf{q}_i 's in the current frame are noise-free.

2) *Model Linearization*: Our goal is to linearize the camera projection model to enable an efficient closed-form solution. We begin with the ideal projection of a true 3D point ${}^{\mathcal{K}}\mathbf{p}_i^o$, expressed in the left keyframe, into the current frame's normalized image plane:

$$\mathbf{q}_i = h({}^{\mathcal{K}}\mathbf{p}_i^o; \chi) \triangleq \frac{\mathbf{E} ({}^c_{\mathcal{K}}\mathbf{R}(\psi) {}^{\mathcal{K}}\mathbf{p}_i^o + {}^c_{\mathcal{K}}\mathbf{t})}{\mathbf{e}_3^\top ({}^c_{\mathcal{K}}\mathbf{R}(\psi) {}^{\mathcal{K}}\mathbf{p}_i^o + {}^c_{\mathcal{K}}\mathbf{t})}, \quad (2)$$

where h is the pinhole camera projection function. Recall that the 2D point \mathbf{q}_i is noise-free, and the triangulated 3D point ${}^{\mathcal{K}}\mathbf{p}_i$ is noisy, modeled as ${}^{\mathcal{K}}\mathbf{p}_i = {}^{\mathcal{K}}\mathbf{p}_i^o + \epsilon_{\mathbf{p}_i}$, with $\epsilon_{\mathbf{p}_i} \sim \mathcal{N}(\mathbf{0}, \Sigma_{\mathbf{p}_i})$.

To linearize the model, we multiply both sides of equation (2) by $\mathbf{e}_3^\top ({}^c_{\mathcal{K}}\mathbf{R}(\psi) {}^{\mathcal{K}}\mathbf{p}_i^o + {}^c_{\mathcal{K}}\mathbf{t})$:

$$\mathbf{e}_3^\top ({}^c_{\mathcal{K}}\mathbf{R}(\psi) {}^{\mathcal{K}}\mathbf{p}_i^o + {}^c_{\mathcal{K}}\mathbf{t}) \mathbf{q}_i = \mathbf{E} ({}^c_{\mathcal{K}}\mathbf{R}(\psi) {}^{\mathcal{K}}\mathbf{p}_i^o + {}^c_{\mathcal{K}}\mathbf{t}). \quad (3)$$

We now define the state vector to be estimated as $\mathbf{x} \triangleq [\cos(\psi), \sin(\psi), {}^c_{\mathcal{K}}\mathbf{t}^\top]^\top$. Let $\rho_i \triangleq \mathbf{R}_{\theta_\phi} {}^{\mathcal{K}}\mathbf{p}_i$ be the pre-rotated point and $\epsilon_{\rho_i} \triangleq \mathbf{R}_{\theta_\phi} \epsilon_{\mathbf{p}_i}$ be its propagated noise. By substituting ${}^{\mathcal{K}}\mathbf{p}_i^o = {}^{\mathcal{K}}\mathbf{p}_i - \epsilon_{\mathbf{p}_i}$ and rearranging equation (3), we obtain a linear equation with respect to \mathbf{x} :

$$\mathbf{A}_i \mathbf{x} = \mathbf{b}_i + \epsilon_i, \quad (4)$$

where $\mathbf{b}_i = \rho_i(3) \mathbf{q}_i$, $\mathbf{A}_i = \begin{bmatrix} \rho_i(1) - \rho_i(2) & 1 & 0 & -\mathbf{q}_i(1) \\ \rho_i(2) & \rho_i(1) & 0 & 1 - \mathbf{q}_i(2) \end{bmatrix}$, and $\epsilon_i = \begin{bmatrix} \epsilon_{\rho_i}(1) \cos(\psi) - \epsilon_{\rho_i}(2) \sin(\psi) - \epsilon_{\rho_i}(3) \mathbf{q}_i(1) \\ \epsilon_{\rho_i}(2) \cos(\psi) + \epsilon_{\rho_i}(1) \sin(\psi) - \epsilon_{\rho_i}(3) \mathbf{q}_i(2) \end{bmatrix}$. Stacking (4) for all point correspondences yields the full linear system

$$\mathbf{A} \mathbf{x} = \mathbf{b} + \epsilon. \quad (5)$$

3) *Bias-Eliminated Consistent Estimator*: Based on (5), an ordinary least-squares (LS) solution can be obtained as

$$\hat{\mathbf{x}}^{\text{LS}} = (\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{b}. \quad (6)$$

However, the regressor term \mathbf{A} contains noise and is correlated with the noise in the regressand term \mathbf{b} . As a result, this ordinary LS solution is generally statistically biased [22]. To obtain a consistent estimate, we employ a BE approach that corrects for this correlation. The BE estimator is given by:

$$\hat{\mathbf{x}}^{\text{BE}} = \left(\frac{\mathbf{A}^\top \mathbf{A}}{n} - \mathbf{G}_1 \right)^{-1} \left(\frac{\mathbf{A}^\top \mathbf{b}}{n} - \mathbf{G}_2 \right), \quad (7)$$

where the bias correction terms \mathbf{G}_1 and \mathbf{G}_2 are derived from the known covariances $\Sigma_{\rho_i} = \mathbf{R}_{\theta_\phi} \Sigma_{\mathbf{p}_i} \mathbf{R}_{\theta_\phi}^\top$. Specifically, let

$\Delta \mathbf{A} = \mathbf{A} - \mathbf{A}^o$ and $\Delta \mathbf{b} = \mathbf{b} - \mathbf{b}^o$. The first correction term \mathbf{G}_1 is given by

$$\begin{aligned} \mathbf{G}_1 &= \mathbb{E} \left[\frac{1}{n} \Delta \mathbf{A}^\top \Delta \mathbf{A} \right] \\ &= (\bar{\Sigma}_\rho(1, 1) + \bar{\Sigma}_\rho(2, 2)) (\mathbf{e}_1 \mathbf{e}_1^\top + \mathbf{e}_2 \mathbf{e}_2^\top), \end{aligned}$$

where $\bar{\Sigma}_\rho = \frac{1}{n} \sum_{i=1}^n \Sigma_{\rho_i}$ is the average covariance matrix. The second correction term, $\mathbf{G}_2 = \mathbb{E}[\frac{1}{n} \Delta \mathbf{A}^\top \Delta \mathbf{b}]$, is a five-dimensional vector whose last three components are zero. Its first two components are given by $\frac{1}{n} \sum_{i=1}^n (\Sigma_{\rho_i}(1, 3) \mathbf{I}_2 + \Sigma_{\rho_i}(2, 3) \mathbf{J}) [\mathbf{q}_i(1), \mathbf{q}_i(2)]^\top$, where \mathbf{J} is the skew-symmetric matrix $\begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$.

Definition 1 (Big O_p): An estimator $\hat{\gamma}$ is called a \sqrt{n} -consistent estimator of γ^o if $\hat{\gamma} - \gamma^o = O_p(1/\sqrt{n})$. That is, for any $\varepsilon > 0$, there exist finite constants M and N such that for all $n > N$, $\mathbb{P}(\|\sqrt{n}(\hat{\gamma} - \gamma^o)\| > M) < \varepsilon$.

Theorem 1: The BE estimator $\hat{\mathbf{x}}^{\text{BE}}$ is a \sqrt{n} -consistent estimator for the true state vector \mathbf{x}^o .

Proof: The proof is based mainly on the following lemma:

Lemma 1: ([21], Lemma 1): Let $\{X_i\}$ be a sequence of independent random variables with $\mathbb{E}[X_i] = 0$ and bounded $\mathbb{E}[X_i^2]$. Then, $\sum_{i=1}^n X_i/n = O_p(1/\sqrt{n})$.

In the noise-free case, $(\mathbf{A}^{o\top} \mathbf{A}^o)^{-1} \mathbf{A}^{o\top} \mathbf{b}^o$, or equivalently $(\frac{\mathbf{A}^{o\top} \mathbf{A}^o}{n})^{-1} \frac{\mathbf{A}^{o\top} \mathbf{b}^o}{n}$ results in the ground truth \mathbf{x}^o . The idea of the proof is to show that $\frac{\mathbf{A}^\top \mathbf{A}}{n} - \mathbf{G}_1$ and $\frac{\mathbf{A}^\top \mathbf{b}}{n} - \mathbf{G}_2$ converge to $\frac{\mathbf{A}^{o\top} \mathbf{A}^o}{n}$ and $\frac{\mathbf{A}^{o\top} \mathbf{b}^o}{n}$, respectively. First, $\Delta \mathbf{A}^\top \Delta \mathbf{A}$ contains both the first and second order terms of ϵ_{ρ_i} . For the second-order term $\epsilon_{\rho_i} \epsilon_{\rho_i}^\top$, we can subtract it by Σ_{ρ_i} to achieve the zero mean. It can be verified that $-\mathbf{G}_1$ actually performs this procedure. Hence, according to Lemma 1, it holds that

$$\frac{\mathbf{A}^\top \mathbf{A}}{n} - \frac{\mathbf{A}^{o\top} \mathbf{A}^o}{n} = \frac{\Delta \mathbf{A}^\top \Delta \mathbf{A}}{n} + O_p\left(\frac{1}{\sqrt{n}}\right) = \mathbf{G}_1 + O_p\left(\frac{1}{\sqrt{n}}\right). \quad (8)$$

Second, similarly, for $\mathbf{A}^\top \mathbf{b}/n$, based on Lemma 1, we can obtain

$$\frac{\mathbf{A}^\top \mathbf{b}}{n} - \frac{\mathbf{A}^{o\top} \mathbf{b}^o}{n} = \mathbf{G}_2 + O_p\left(\frac{1}{\sqrt{n}}\right). \quad (9)$$

Finally, by combining (9) and (8), we obtain

$$\begin{aligned} \hat{\mathbf{x}}^{\text{BE}} &= \left(\frac{\mathbf{A}^{o\top} \mathbf{A}^o}{n} + O_p\left(\frac{1}{\sqrt{n}}\right) \right)^{-1} \left(\frac{\mathbf{A}^{o\top} \mathbf{b}^o}{n} + O_p\left(\frac{1}{\sqrt{n}}\right) \right) \\ &= \left(\frac{\mathbf{A}^{o\top} \mathbf{A}^o}{n} \right)^{-1} \frac{\mathbf{A}^{o\top} \mathbf{b}^o}{n} + O_p\left(\frac{1}{\sqrt{n}}\right) = \mathbf{x}^o + O_p\left(\frac{1}{\sqrt{n}}\right), \end{aligned}$$

which completes the proof. \blacksquare

E. Epipolar-Based Gauss-Newton Refinement

The previous stage provides a robust and consistent pose estimate. Using the consistent estimate as the initial value, we further refine the pose by implementing the Gauss-Newton (GN) algorithm over the epipolar ML problem (1) using the inlier set \mathcal{I} . Thanks to the \sqrt{n} -consistent initial estimator (Theorem 1), only a one-step GN iteration is sufficient to achieve the asymptotic efficiency, which is formally stated

in the following lemma. Denote the one-step GN iteration of $\hat{\mathbf{x}}^{\text{BE}}$ by $\hat{\mathbf{x}}^{\text{GN}}$.

Lemma 2 ([23, Theorem 2]): We have

$$\hat{\mathbf{x}}^{\text{GN}} - \hat{\mathbf{x}}^{\text{ML}} = o_p\left(\frac{1}{\sqrt{n}}\right).$$

The small o_p implies that $\hat{\mathbf{x}}^{\text{GN}}$ has the same asymptotic property as $\hat{\mathbf{x}}^{\text{ML}}$ [23], i.e., when the point number is large, $\hat{\mathbf{x}}^{\text{GN}}$ achieves the theoretical lower bound-CRLB.

IV. SIMULTANEOUS COVARIANCE ESTIMATION AND POSE REFINEMENT

Recall that we use IMU preintegration and visual PnP for relative 2-DOF (pitch and roll) and 4-DOF (yaw and translation) estimation, respectively. Without the gravity prior, the pitch and roll inevitably drift just like the remaining 4-DOF components of the pose. In this section, we further fuse visual information with the gravity prior to refine the 6-DOF pose and prevent the drift of pitch and roll angles. To handle time-varying accelerations and adaptively weight the gravity prior, we follow [24] and jointly optimize the pose and the gravity prior covariance online.

The state vector is expanded to include not only the relative 4-DOF pose between the current frame and keyframe but also the absolute orientation at each intermediate IMU measurement time. Specifically, the state is defined as $\chi = [\psi, \mathcal{C}_k \mathbf{t}^\top, \phi_1, \theta_1, \dots, \phi_K, \theta_K]^\top$. Here, ψ and $\mathcal{C}_k \mathbf{t}$ are the relative yaw and translation from the previous keyframe, while $(\phi_k, \theta_k)_{k=1}^K$ represents the absolute roll and pitch of the IMU with respect to the world frame \mathcal{W} at each of the K intermediate IMU measurement instants between the keyframe and current frame. The corresponding residuals are:

1) Visual residuals: For each inlier $i \in \mathcal{I}$, this is the 2D reprojection error $\mathbf{r}_{\text{vis},i} = \mathbf{q}_i - h(\mathcal{C}_k \mathbf{R}^K \mathbf{p}_i + \mathcal{C}_k \mathbf{t})$. Here, the relative rotation $\mathcal{C}_k \mathbf{R}$ is fully determined by $\{\psi, \phi_1, \theta_1, \phi_K, \theta_K, \mathcal{I} \mathbf{R}\}$.

2) Gravity prior residuals: Under the assumption of low motion acceleration, the accelerometer reading $\bar{\mathbf{a}}_k$ should align with the gravity vector ${}^{\mathcal{W}}\mathbf{g}$ transformed into the IMU frame. We define the residual as the directional difference: $\mathbf{r}_{\text{imu},k}(\chi) = \bar{\mathbf{a}}_k - \mathcal{I} \mathbf{R}(\phi_k, \theta_k) {}^{\mathcal{W}}\mathbf{g}$. Here, $\mathcal{I} \mathbf{R}(\phi_k, \theta_k)$ is the rotation from world to IMU constructed from the state variables ϕ_k and θ_k , and ${}^{\mathcal{W}}\mathbf{g}$ is the unit gravity vector. This residual measures the non-gravitational component of acceleration. A non-zero value indicates either the violation of the low motion acceleration assumption or a misalignment in the estimated orientation $\{\phi_k, \theta_k\}$.

The underlying estimation problem is heteroscedastic due to multi-sensor fusion with time-varying noise. We jointly estimate the state χ and the IMU measurement covariance Σ_{imu} by solving:

$$\min_{\chi, \Sigma_{\text{imu}}} \sum_{i \in \mathcal{I}} \|\mathbf{r}_{\text{vis},i}(\chi)\|_{\Sigma_{\text{vis}}}^2 + \sum_{k=1}^K \|\mathbf{r}_{\text{imu},k}(\chi)\|_{\Sigma_{\text{imu}}}^2, \quad (10)$$

where $\|\mathbf{r}\|_{\Sigma} \triangleq (\mathbf{r}^\top \Sigma^{-1} \mathbf{r})^{\frac{1}{2}}$. The visual covariance matrix Σ_{vis} is computed once at the beginning from the consistent 2D noise variance estimate obtained in Section III-D and

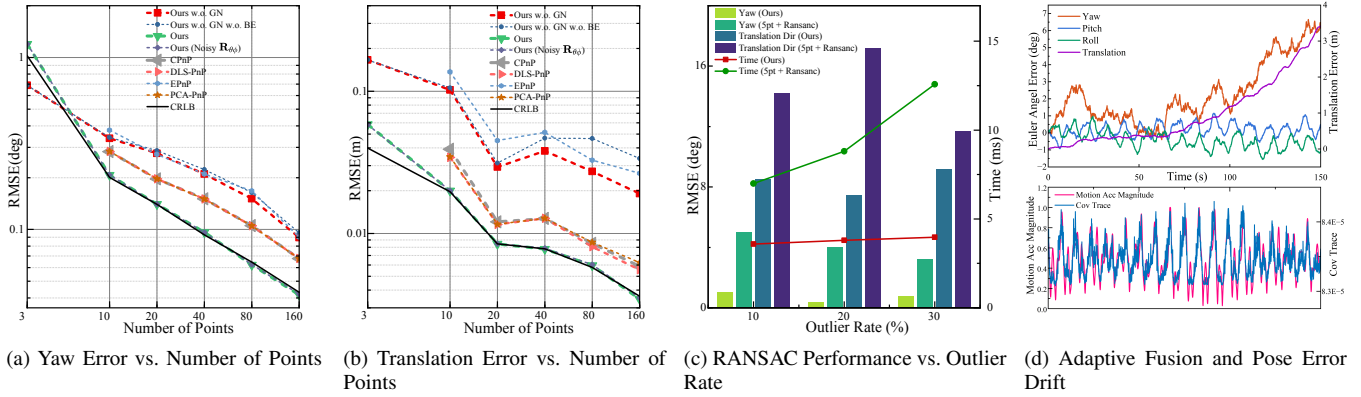


Fig. 4: **(a)-(b)**: Our 4-DOF PnP estimator achieves CRLB accuracy, outperforming other PnP methods. **(c)**: Our 4-DOF minimal-set consensus filter achieves better accuracy at a lower computational cost than the 5-point method. **(d)**: Our adaptive fusion framework maintains drift-free roll and pitch estimates.

the ability to adaptively fuse the IMU gravity prior and maintain drift-free pitch and roll estimates. In our simulations, the baseline of the stereo camera is set as $\mathbf{t} = [0.2, 0, 0]^T$ m and $\mathbf{R} = \mathbf{I}$. The camera has a focal length of $f = 1100$ pixels, an image size of 800×800 pixels, and captures 3D points at depths ranging from 1 to 10 meters. Unless otherwise specified, Gaussian noises with a standard deviation of 2.5 pixels are added to 2D features.

1) *Estimator Consistency and Accuracy*: We evaluate the statistical properties of our 4-DOF PnP estimator via a 700-run Monte Carlo simulation. The results, presented in Fig. 4a and Fig. 4b, benchmark the RMSE of our estimators against several widely used and state-of-the-art solvers, including CPnP [28], EPnP [29], DLS-PnP [30], and PCA-PnP [31]. The theoretical lower bound, CRLB, is also included. As depicted, our initial estimator (*Ours w.o. GN*) is consistent, as the RMSEs of both yaw and translation decrease linearly in the double-log plot as the point count increases. This property stems from our bias elimination method, which guarantees the estimate converges to the true pose. Leveraging this consistent estimate, a single-step GN iteration (*Ours*) is sufficient to asymptotically find the global minimum of the ML problem (1) and achieve the CRLB accuracy. The compared methods are either inconsistent or unable to reach the CRLB. We also note that our 4-DOF estimator remains operational with as few as three points and consistently provides the most accurate estimates across all point numbers. To further assess our estimator’s robustness in a more practical context, we introduce roll and pitch noise with a standard deviation of 0.01° to simulate typical IMU pre-integration errors. Even with this perturbation, our method (*Ours (Noisy $\mathbf{R}_{\theta\phi}$)*) still coincides with the CRLB, highlighting the framework’s resilience.

2) *Effectiveness of Minimal-Set Outlier Rejection*: To evaluate practical robustness and efficiency with outliers, we vary the outlier ratio from 10% to 30% and run 400 Monte Carlo trials within a RANSAC framework. Our 4-DOF minimal-set consensus filter, in conjunction with the

consistent estimator, uses a noisy roll and pitch prior with a 0.2° standard deviation and is benchmarked against MATLAB’s 5-point RANSAC. As shown in Fig. 4c, our method delivers stable yaw and translation-direction estimates across all outlier rates and significantly outperforms the 5-point baseline. Moreover, with 30% outliers, our method reduces 68.5% of the computation time compared to the 5-point essential-matrix RANSAC. Together, these results highlight that our minimal-set consensus filter achieves both higher accuracy and lower computational complexity in practical scenarios with outliers.

3) *Validation of Adaptive Fusion for Drift-Free Pitch/Roll*: We simulate a 150-second trajectory, where the motion acceleration is time-varying and upper-bounded by 1 m/s^2 to account for water resistance. During the simulation, the gyroscope is corrupted by both white noise and a slowly drifting bias. The pose error curves in Fig. 4d show the expected yaw and translation drift, whereas the roll and pitch errors are drift-free and stay near zero, consistent with our claim on the gravity-enhanced joint optimization in Section IV. The figure also shows that as the motion acceleration increases, the estimated IMU gravity prior covariance grows, reducing the weight for the gravity prior residual. Hence, the estimator relies more on visual measurements. In low-acceleration periods, the gravity prior regains influence, and the attitude is re-anchored, suppressing pitch and roll drift. These results confirm the effectiveness of our adaptive fusion for accurate and drift-free roll/pitch estimation over long horizons.

B. Real Water Pool Experiments

To validate real-world performance, we conducted underwater trials with a custom 60kg robot in a $10 \text{ m} \times 5 \text{ m}$, 2m-deep artificial pool with 10 underwater motion-capture cameras for ground truth. The robot carries a hardware-synchronized downward-facing stereo camera (1280×720 , 10 Hz) and an MTI-630 IMU (200 Hz). Two scenarios are covered: (i) scenes with numerous, highly similar feature

points where we collect the **POOL-EASY** dataset and (ii) feature-sparse, largely coplanar scenes where we collect the **POOL-HARD** dataset. The **POOL-HARD** dataset features low texture and higher average accelerations.

We benchmark against state-of-the-art open-source methods: ORB-SLAM3 [8], VINS-Fusion [6], and SVIN2 [4], using ATE and RPE as metrics. For fairness, ORB-SLAM3 and VINS-Fusion are run with their EuRoC Micro Aerial Vehicle (MAV) dataset configuration files, while SVIN2 is run in VI-only mode with sonar and depth sensors disabled. Qualitative results are shown in Fig. 6, which shows that our method achieves better consistency with ground truth.

The quantitative results in Table I show superior accuracy and robustness of our method. On the **POOL-EASY** dataset, our method consistently achieves the best performance in both ATE and RPE, except for the ATE of Easy6 sequence. We note that SVIN2 fails to complete the full trajectory estimation in all sequences. This is because in these repetitive scenes, the BRISK feature descriptor used by SVIN2 yields a large portion of outliers that will be filtered out, leaving too few inliers to reliably initialize and track, as illustrated in the top right in Fig. 5.

The performance gap widens significantly on the more demanding **POOL-HARD** dataset. In this dataset, SVIN2 only fails in one sequence, as a calibration board is placed in the environment to aid initialization. However, ORB-SLAM3 fails in all sequences. This is because the ORB descriptor detects fewer stable corners in these texture-less scenes and is prone to early tracking loss, as shown in the bottom left in Fig. 5. Compared to our methods, VINS-Fusion requires additional scale initialization, which needs more motion excitation and occasionally leads to failure. Notably, our system is the only method that successfully completes all sequences on the **POOL-HARD** dataset. Meanwhile, it exhibits significantly better RPE and consistently delivers the best or second-best ATE. The RPE superiority of our system is mainly attributed to the consistent 4-DOF PnP estimator and adaptive 6-DOF relative pose refinement.




Fig. 5: Illustration of feature matching: ORB-SLAM3 (ORB) vs. SVIN2 (BRISK). The green denotes inliers, while the yellow indicates outliers.

TABLE I: Comparison of ATE and RPE on the **POOL-EASY** and **POOL-HARD** datasets. Values in **blue bold** are the smallest and **blue** are the second smallest.

Seq	ATE (m)				RPE (m)			
	Ours	ORB*	VINS	SVIN2	Ours	ORB*	VINS	SVIN2
POOL-EASY								
Easy1	0.0116	0.0128	0.0127	failed	0.0100	0.0101	0.0264	failed
Easy2	0.0173	0.0196	0.0190	failed	0.0188	0.0188	0.0487	failed
Easy3	0.0319	0.0339	failed	failed	0.0106	0.0106	failed	failed
Easy4	0.0211	0.0218	0.0657	failed	0.0111	0.0111	0.0492	failed
Easy5	0.0147	0.0157	0.0158	failed	0.0134	0.0135	0.0354	failed
Easy6	0.0219	0.0217	0.0408	failed	0.0104	0.0105	0.0262	failed
Easy7	0.0270	0.0485	0.0344	failed	0.0095	0.0095	0.0238	failed
POOL-HARD								
Hard1	0.0246	failed	failed	0.0282	0.0151	failed	failed	0.0631
Hard2	0.0463	failed	0.0513	0.0312	0.0256	failed	0.0456	0.0710
Hard3	0.0406	failed	0.0353	0.0343	0.0190	failed	0.0350	0.0640
Hard4	0.0486	failed	0.0538	0.0301	0.0207	failed	0.0344	0.0731
Hard5	0.0333	failed	failed	failed	0.0158	failed	failed	failed
Hard6	0.0280	failed	0.0338	0.0365	0.0149	failed	0.0260	0.0598
Hard7	0.1054	failed	failed	0.0536	0.0078	failed	failed	0.0485

* denotes the stereo-only mode. The stereo-IMU mode of ORB-SLAM3 is not used as it is difficult to successfully initialize underwater due to the robot's inability to provide sufficient acceleration for IMU excitation.

VII. CONCLUSION AND FUTURE WORK

We presented GeVI-SLAM, a gravity-enhanced stereo visual-inertial SLAM that reduces 6-DOF tracking to a 4-DOF PnP problem. A BE initializer with one-step Gauss-Newton achieves CRLB-level accuracy, and a 4-DOF minimal-set consensus filter efficiently rejects outliers. An adaptive 6-DOF refinement jointly estimates pose and IMU prior covariance, keeping roll/pitch drift-free. Simulations and real pool experiments show superior accuracy and robustness over ORB-SLAM3, VINS-Fusion, and SVIN2.

Limitations include yaw/translation drift without loop closure, the lack of refractive modeling, and uncharacterized roll/pitch uncertainty within the 4-DOF formulation. Future work will add robust loop closure and place recognition for underwater scenes, refractive camera modeling, fusion with DVL/pressure/range measurements, and field validation.

ACKNOWLEDGMENT

This work was supported in part by NSFC under Grant 62273288; in part by Shenzhen Science and Technology Program JCYJ20240813113609013, JCYJ20241202124010014, Guangdong Basic and Applied Basic Research Foundation No. 2024A1515240009.

REFERENCES

- [1] A. Mallios, E. Vidal, R. Campos, and M. Carreras, "Underwater caves sonar data set," *The International Journal of Robotics Research*, vol. 36, no. 12, pp. 1247–1251, 2017.
- [2] B. Joshi, S. Rahman, M. Kalaitzakis, B. Cain, J. Johnson, M. Xanthidis, N. Karapetyan, A. Hernandez, A. Q. Li, N. Vitzilaios *et al.*, "Experimental comparison of open source visual-inertial-based state estimation algorithms in the underwater domain," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 7227–7233.

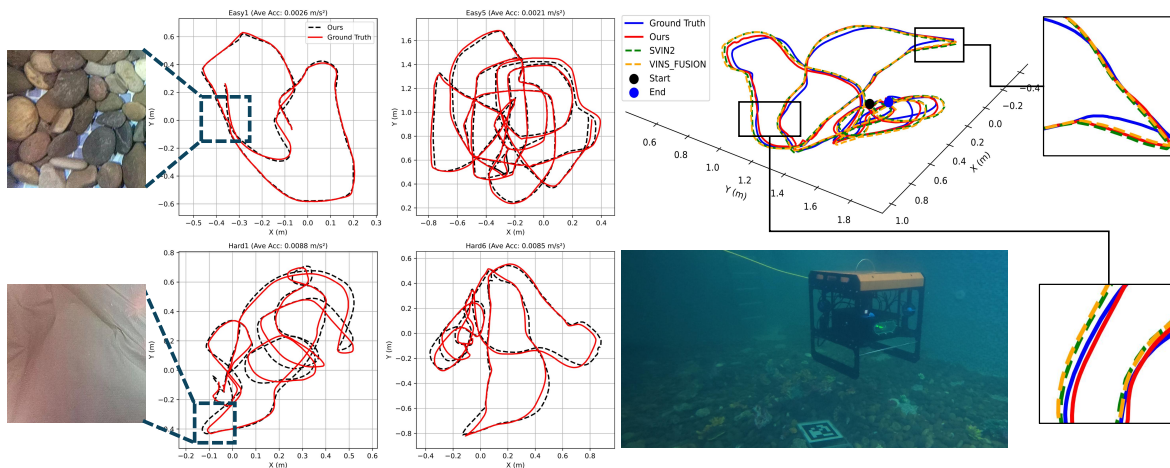


Fig. 6: Trajectory evaluation on POOL-EASY/POOL-HARD. Left: two EASY and two HARD sequences with motion-capture ground truth and our estimates. Top right: a 3D overlay on an EASY sequence comparing GeVI-SLAM with SVIN2, ORB-SLAM3, and VINS-Fusion. Bottom right: the underwater robot during data collection.

- [3] S. Xu, K. Zhang, and S. Wang, "AQUA-SLAM: Tightly-coupled underwater acoustic-visual-inertial SLAM with sensor calibration," *IEEE Transactions on Robotics*, 2025.
- [4] S. Rahman, A. Q. Li, and I. Rekleitis, "SVIn2: An underwater SLAM system using sonar, visual, inertial, and depth sensor," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 1861–1868.
- [5] H. Yao, Y. Ma, P. Li, C. Zhai, J. Song, M. Ouyang, Z. Dai, and X. Zhu, "SG-VIO: Monocular visual-inertial odometry with tightly coupled structural lines and gravity to avoid degeneracy," *IEEE Internet of Things Journal*, 2024.
- [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, pp. 1004–1020, 2018.
- [7] C. Hu, S. Zhu, Y. Liang, Z. Mu, and W. Song, "Visual-pressure fusion for underwater robot localization with online initialization," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 8426–8433, 2021.
- [8] 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 multimap SLAM," *IEEE Transactions on Robotics*, vol. 37, no. 6, pp. 1874–1890, 2021.
- [9] M. Hildebrandt and F. Kirchner, "IMU-aided stereo visual odometry for ground-tracking AUV applications," in *OCEANS'10 IEEE SYDNEY*. IEEE, 2010, pp. 1–8.
- [10] F. Bellavia, M. Fanfani, and C. Colombo, "Selective visual odometry for accurate AUV localization," *Autonomous Robots*, vol. 41, no. 1, pp. 133–143, 2017.
- [11] N. Yang, R. Wang, and D. Cremers, "Feature-based or direct: An evaluation of monocular visual odometry," *arXiv preprint arXiv:1705.04300*, 2017.
- [12] A. Quattrini Li, A. Coskun, S. M. Doherty, S. Ghasemlou, A. S. Jagtap, M. Modasshir, S. Rahman, A. Singh, M. Xanthidis, J. M. O'Kane *et al.*, "Experimental comparison of open source vision-based state estimation algorithms," in *International Symposium on Experimental Robotics (ISER)*. Springer, 2016, pp. 775–786.
- [13] Y. Wang, D. Gu, X. Ma, J. Wang, and H. Wang, "Robust real-time AUV self-localization based on stereo vision-inertia," *IEEE Transactions on Vehicular Technology*, vol. 72, no. 6, pp. 7160–7170, 2023.
- [14] R. Miao, J. Qian, Y. Song, R. Ying, and P. Liu, "UniVIO: Unified direct and feature-based underwater stereo visual-inertial odometry," *IEEE Transactions on Instrumentation and Measurement*, vol. 71, pp. 1–14, 2021.
- [15] T. Lupton and S. Sukkarieh, "Visual-inertial-aided navigation for high-dynamic motion in built environments without initial conditions," *IEEE Transactions on Robotics*, vol. 28, no. 1, pp. 61–76, 2011.
- [16] J. Li, H. Bao, and G. Zhang, "Rapid and robust monocular visual-inertial initialization with gravity estimation via vertical edges," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 6230–6236.
- [17] V. Usenko, J. Engel, J. Stückler, and D. Cremers, "Direct visual-inertial odometry with stereo cameras," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 1885–1892.
- [18] C. Noh, W. Yang, M. Jung, S. Jung, and A. Kim, "GaRLIO: Gravity enhanced Radar-LiDAR-Inertial odometry," *arXiv preprint arXiv:2502.07703*, 2025.
- [19] V. Kubelka, M. Vaidis, and F. Pomerleau, "Gravity-constrained point cloud registration," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 4873–4879.
- [20] R. Mahony, T. Hamel, and J.-M. Pfimlin, "Nonlinear complementary filters on the special orthogonal group," *IEEE Transactions on Automatic Control*, vol. 53, no. 5, pp. 1203–1218, 2008.
- [21] G. Zeng, Y. Shen, Z. Hong, Y. Hong, V. Ila, G. Shi, and J. Wu, "Bias-eliminated PnP for stereo visual odometry: Provably consistent and large-scale localization," *arXiv preprint arXiv:2504.17410*, 2025.
- [22] G. Zeng, Q. Zeng, X. Li, B. Mu, J. Chen, L. Shi, and J. Wu, "Consistent and asymptotically statistically-efficient solution to camera motion estimation," *arXiv preprint arXiv:2403.01174*, 2024.
- [23] B. Mu, E.-W. Bai, W. X. Zheng, and Q. Zhu, "A globally consistent nonlinear least squares estimator for identification of nonlinear rational systems," *Automatica*, vol. 77, pp. 322–335, 2017.
- [24] K. Khosoussi and I. Shames, "Joint state and noise covariance estimation," *arXiv preprint arXiv:2502.04584*, 2025.
- [25] 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 Robotics and Automation Letters*, vol. 6, no. 2, pp. 1399–1406, 2021.
- [26] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.
- [27] F. Dellaert and G. Contributors, "borglab/gtsam," May 2022. [Online]. Available: <https://github.com/borglab/gtsam>
- [28] G. Zeng, S. Chen, B. Mu, G. Shi, and J. Wu, "CPnP: Consistent pose estimator for Perspective- n -Point problem with bias elimination," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 1940–1946.
- [29] V. Lepetit, F. Moreno-Noguer, and P. Fua, "EPnP: An accurate $\mathcal{O}(n)$ solution to the PnP problem," *International Journal of Computer Vision*, vol. 81, no. 2, pp. 155–166, 2009.
- [30] J. A. Hesch and S. I. Roumeliotis, "A direct least-squares (DLS) method for PnP," in *2011 International Conference on Computer Vision (ICCV)*. IEEE, 2011, pp. 383–390.
- [31] L. Zhou, Z. Wei, and X. Wang, "A novel iterative solution to the Perspective- n -Point problem via cost function approximation," *IEEE Transactions on Robotics*, 2025.