

Sonar–GPS Fusion for Seabed Mapping in Turbid Shallow Waters with an Autonomous Surface Vehicle

Yisheng Zhang¹, Michael Xu¹, Alan Williams², Matthew Gray², Nare Karapetyan^{3,†}, Miao Yu^{1,†}

Abstract—Accurate seabed mapping is essential for habitat monitoring and infrastructure inspection. In turbid, shallow coastal waters, such as shellfish aquaculture farms, the effectiveness of traditional optical methods is limited. Autonomous surface vehicles (ASVs) equipped with forward-looking sonar (FLS) offer a promising alternative. However, existing sonar-based systems face challenges in achieving fine resolution mapping over long trajectories due to low-resolution positioning measurements and accumulated drift over long trajectories. In this paper, we present a drift-resilient seabed mapping framework that integrates local FLS frame alignment using the Fourier–Mellin transform (FMT) with global trajectory optimization based on an extended Kalman filter (EKF) that fuses global positioning system (GPS), inertial measurement unit (IMU), and compass data. A variance-based image blending strategy is used to further reduce visual artifacts in overlapping regions. Field trials on a structured oyster farm site show that our framework helps reduce drift in RMSE by 9.5% relative to the FMT-only baseline. This framework also enables sub-meter reconstruction accuracy and preservation of high-resolution textures needed for oyster inventory estimation within the mapped areas.

I. INTRODUCTION

Reliable seabed mapping is critical for a range of applications including marine habitat monitoring, resource exploration, infrastructure inspection, and search-and-rescue missions. Traditional optical imaging methods, however, are severely constrained in low-visibility underwater environments due to high turbidity, low illumination, and strong light scattering. Under these conditions, forward-looking sonar (FLS) has been widely adopted due to its robustness. However, registration of overlapping sonar frames remains challenging due to low signal-to-noise ratio of sonar images, partial overlaps, and shadow cast artifacts [1], [2]. Traditional registration approaches such as feature-based, template-based, region-based, and Fourier-based methods struggle to maintain robustness due to noisy sonar imagery and the simplified 2-D similarity transforms used in these methods [3]–[5]. Furthermore, although satisfactory local alignment can be achieved, due to the inevitably accumulated residual errors over long sequences, these methods suffer

This work was supported by the United States Department of Agriculture (USDA) National Institute of Food and Agriculture (NIFA) Sustainable Agricultural Systems (SAS) Program under Grant 2020-68012-31805 and in part by Independent Research & Development Program (IR&D) at Woods Hole Oceanographic Institution (WHOI).

¹University of Maryland, College Park, MD 20742, USA, {mixu1235, yiszhang, mmyu}@umd.edu

²University of Maryland Center for Environmental Science, Cambridge, MD 21613, USA, {awilliams, mgray}@umces.edu

³Woods Hole Oceanographic Institution, Woods Hole, MA 02543, USA. nare@whoi.edu

[†] Equal Contribution.




Fig. 1: Field experimental setup. An ASV equipped with a multi-beam FLS was deployed to survey a farm site with pre-arranged on-bottom quadrats. Oysters are distributed within the quadrats. The top right inset shows a photo of a sample PVC quadrat from up close.

from increasingly pronounced drift in large-scale and/or long-range mapping tasks.

Autonomous surface vehicles (ASVs) equipped with a FLS and a global positioning system (GPS) offer a promising platform for seabed mapping [6]. While the on-board GPS can provide global references to correct sonar-based registration errors, it lacks the fine-grained accuracy required for high-resolution mapping. This gap is especially evident in mapping of on-bottom oyster farms in shallow, highly-turbid coastal waters, where surveying large areas of oyster inventory is needed to enable precision oyster farming. The current inventory monitoring approaches rely on divers performing manual counting or dredge sampling.

Although individual oysters are hardly identifiable in FLS imagery, this work aims to generate accurate georeferenced seabed mosaics on a sub-meter scale and improve texture quality to support subsequent mapping of oyster inventory. These mosaics establish a spatial framework that enables the integration of machine-learning based classification algorithms to estimate oyster density and size within a mapped area [7]. Our work addresses a critical bottleneck: the lack of a low-cost, drift-resilient mapping baseline in turbid environments.

We propose a unified pipeline that integrates local sonar alignment, global trajectory optimization, and adaptive image blending. Our main contributions are summarized as follows.

- **Drift-resilient registration:** We introduce a framework that fuses Fourier–Mellin transform (FMT)-based local alignment with global priors from an extended Kalman filter (EKF) (GPS, compass, and inertial measurement unit (IMU)) through a global least-squares optimiza-

tion. This enforces global consistency in sonar image registration and significantly mitigates cumulative drift, which is a key limitation of prior FMT-only approaches.

- **Detail-preserving robust blending:** We develop a variance-based blending strategy that combines exponential moving averages (EMA)-smoothed short- and long-term variance statistics with adaptive weighted averaging. This enables reliable pixel fusion in overlapping sonar frames, effectively suppressing striping and noise while preserving fine seabed details—addressing the artifact-prone results of existing blending methods.
- **Field validation with low-cost ASVs:** We demonstrate the full pipeline in shallow, turbid coastal waters using a low-cost ASV platform (Fig. 1). The system achieves sub-meter accuracy with a 9.5% drift RMSE reduction over the FMT-only baseline, outperforms other representative registration methods, and delivers clearer, more detailed textures than existing blending strategies, demonstrating practical viability for habitat mapping.

II. RELATED WORK

Underwater mapping and navigation have been widely studied, spanning sonar image registration and sensor fusion with iterative optimization. While most prior work rely on autonomous underwater vehicles (AUVs), ASVs can directly access GPS and follow near-planar trajectories, making them well suited for constructing textured 2-D overhead mosaics. Many existing pipelines are rooted in simultaneous localization and mapping (SLAM), which emphasizes landmark tracking and sparse point-cloud mapping for navigation and 3-D reconstruction. Although these pipelines are effective for localization, they cannot be used to generate dense, interpretable 2-D seabed maps needed for ecological monitoring. Our evaluation emphasizes mosaicing fidelity and georeferencing consistency, which are not directly captured by SLAM-centric localization and loop-closure metrics.

A. Sonar-Based Mapping and Sensor Fusion

Early studies on FLS imagery focused on mosaicing through direct frame-to-frame registration. Hurtós et al. [2] applied Fourier-based phase correlation, which was proven to be robust but was limited to 2-D similarity assumptions. To advance this method, Aykin et al. [3] proposed a feature-based method incorporating parallax and elevation cues. While these approaches achieve reliable local alignment, they lack mechanisms to suppress drift over long trajectories and often rely on high-frequency, costly sensors [5], [8], limiting their applicability for large-area mosaicing.

To mitigate drift, external references were incorporated. Li et al. [9] proposed a real-time pipeline for a 3-D phased-array imaging sonar, initialized with GPS and aligned by using a variant of the iterative closest point (ICP) algorithm [10], [11]. While effective, this method resulted in significant computational overhead on large datasets. Dos Santos et al. [12] combined sonar imagery with aerial photography, by applying ICP to register sonar with aerial maps to enhance localization in structured environments. While these methods demonstrated the benefits of external anchoring, they incurred increased sensing and computational requirements.

SLAM-based methods that integrate sonar with inertial and other onboard sensors were also explored and typically framed as pose-graph or factor-graph optimization [13]. Early work fused dense sonar features with Doppler velocity log (DVL)/IMU in a graph-based framework [14], and iSAM2 enabled scalable incremental optimization [15]. ICP was widely adopted in underwater SLAM as a fundamental alignment method [16], [17]. Building on modality-specific advances, Cheung et al. [18] extended SLAM to synthetic aperture sonar with a non-Gaussian formulation, achieving long-mission robustness without expensive inertial aids. For multi-robot operations, DRACo-SLAM and its successor DRACo-SLAM2 represented sonar maps as object graphs and enabled efficient distributed loop closure [19], [20]. McConnell et al. explored cross-modal integration: one study employed FLS observations in a particle filter with the adaptive Monte Carlo localization (AMCL) against aerial imagery and odometry [21], while another used an orthographic imaging sonar within a submap-based framework to achieve dense 3-D reconstructions [22]. Although these frameworks provide strong drift resilience and scalability, their reliance on landmark tracking and sparse or volumetric representations aligns them with navigation and 3-D reconstruction tasks rather than the creation of dense, textured 2-D mosaics.

In summary, direct registration and external anchoring methods offer lightweight pipelines better suited for constructing detailed 2-D overhead maps, whereas SLAM-centric approaches prioritize localization and structural mapping. Motivated by these distinctions, we develop a compact, drift-resilient 2-D seabed mapping framework that leverages robust local registration cues while incorporating GPS constraints, without resorting to full SLAM formulations.

B. Challenges in ASV-Based Seabed Mapping

Unlike AUVs, ASVs can directly access GPS for global positioning [23]. Early surveys, such as the United States Geological Survey (USGS) studies in Long Island Sound [24], registered sidescan sonar imagery to GPS tracks with manual corrections, illustrating both the utility and limitations of GPS anchoring. While GPS provides reliable global references, its spatial resolution is not enough for high-resolution mapping. On the other hand, sonar-based registration preserves local details but suffers from accumulated drift over long trajectories. Therefore, there is a critical need for methods that can help achieve drift-resilient, high-accuracy seabed mapping with ASVs to overcome the aforementioned constraints.

C. Image Blending in Mapping

Blending overlapping sonar images into a unified map is essential but often underexplored. In the broader field of image processing, image fusion techniques have been extensively surveyed [25], which can be divided broadly into spatial-domain (e.g., averaging, principal component analysis (PCA), Brovey transforms) and transform-domain (e.g., pyramid- or wavelet-based) approaches. Spatial domain approaches are simple but prone to artifacts, while transform-based strategies generally preserve more detail and remain widely adopted across domains.

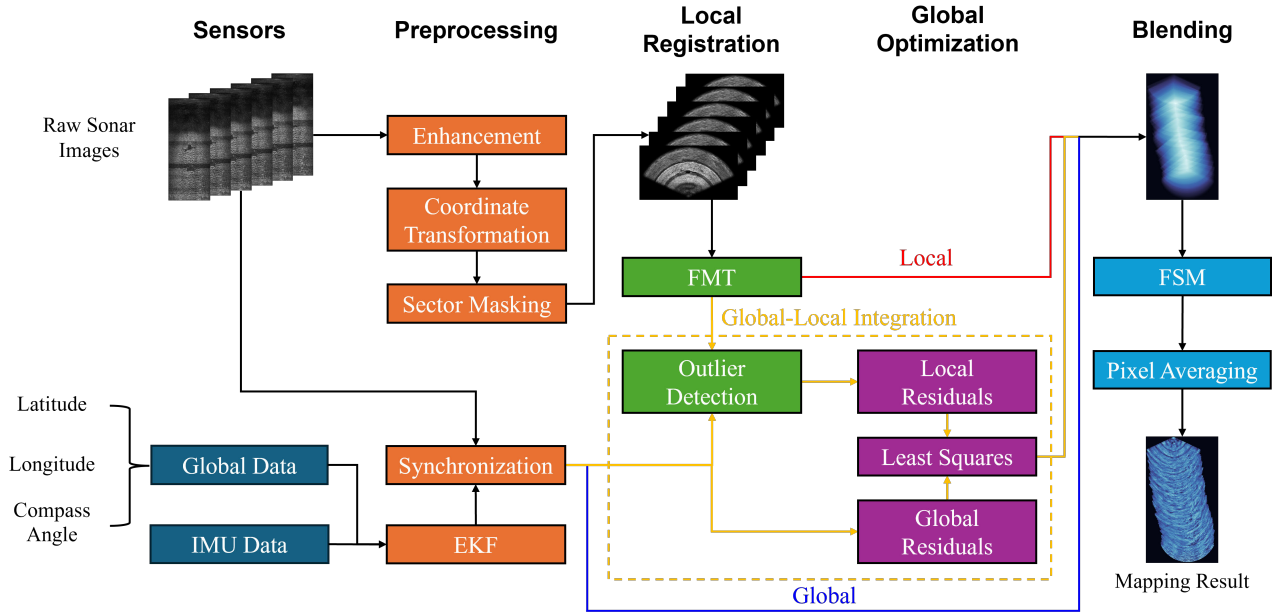


Fig. 2: The architecture of the proposed system. The pipeline integrates global priors and sonar images through preprocessing, local registration, global optimization, and blending to construct a globally consistent 2-D map. Preprocessing modules enhance sonar images before registration. Local alignment is performed with FMT, while global consistency is enforced by least-squares optimization. Variance-based pixel blending resolves conflicts in overlapping regions. In the diagram, **black arrows** denote shared preprocessing and blending steps; **yellow arrows and box** represent the joint local–global optimization pipeline (used as the single-stage scheme and as Stage 2 of the two-stage scheme); **red arrow** indicates Stage 1 of the two-stage scheme, where submaps are produced from local registration and blending without global correction; and **blue arrow** represents the for-reference scheme that builds maps solely from prior global estimates (without local registration).

For sonar-specific applications, Hurtós et al. [26] proposed a blending pipeline tailored to FLS mosaicing that showed improved robustness to photometric irregularities and sonar-specific artifacts arising from the imaging geometry, although residual artifacts may still appear in overlapping regions.

Su et al. [27] proposed a variance-based blending scheme that uses short- and long-term variance statistics to adaptively weight overlapping pixels to improve visual consistency in sonar mosaics. However, the robustness was limited due to the reliance on fixed temporal windows and simple averaging. Our approach is intended to enhance temporal modeling and fusion robustness, yielding smoother transitions and more reliable texture preservation.

As such, existing sonar mosaicing methods lacks drift correction, GPS fusion alone is too coarse, and no prior work integrates these into a lightweight, field-deployable ASV pipeline. Our framework integrates local FMT-based alignment, GPS/compass-constrained global optimization, and variance-based blending into a single pipeline specifically designed for low-cost ASVs operating in shallow, turbid waters. This framework enables accurate sub-meter level mapping compare to sonar-only or GPS-only approaches.

III. APPROACH

We propose a drift-resilient global optimization framework for seabed mapping using FLS imagery. The framework integrates pairwise FLS-based local registrations with EKF-smoothed GPS/compass priors and refines them through nonlinear optimization to construct a globally consistent

2-D seabed maps. Our formulation employs a 2-D registration model, which is appropriate for near-planar seabed conditions. For seabeds with large slopes or complex 3-D topologies, mosaic quality may degrade. In this sections, we discuss the problem formulation followed by the detailed system architecture and components.

A. Problem Formulation

Our objective is to estimate globally consistent poses for a sequence of sonar images and fuse them into a single georeferenced mosaic. The inclusion of EKF-smoothed global priors reduces long-horizon drift and provides additional tolerance to small roll/pitch disturbances.

Input: A sequence of sonar images $\{I_1, I_2, \dots, I_n\}$ and corresponding ASV pose observations $\{G_1, G_2, \dots, G_n\}$, where each $G_i = (G_i^{x,y}, G_i^\theta)$ comprises a global position $G_i^{x,y}$ (from GPS) and a heading G_i^θ (from compass), both refined through EKF integration with IMU observations.

Output: Optimized global positions $\{P_1, P_2, \dots, P_n\}$, orientations $\{\Theta_1, \Theta_2, \dots, \Theta_n\}$, and a fused 2D map \mathcal{M} .

Formally, let $\mathcal{P} = \{P_i, \Theta_i \mid i = 1, \dots, n\}$ denote the global poses to be optimized. The optimization problem is:

$$\mathcal{P}^* = \arg \min_{\mathcal{P}} \mathcal{E}(\mathcal{P}), \quad (1)$$

where \mathcal{E} is a residual function combining consistency with local registration and agreement with global priors.




Fig. 3: Comparison of raw and processed sonar images (5 m range, 130° field of view (FOV)): (a) and (b) are raw and processed images in polar coordinates, respectively. (c) is processed image in Cartesian coordinates. A quadrat (0.62 m side length) is enclosed by the yellow bounding box.

Algorithm 1 FMT-GPS-OPT WITH ROBUST BLENDING

Input: sonar images $\{I_1, \dots, I_n\}$, pose measurements $\{G_1, \dots, G_n\}$, scale s , thresholds τ_p, τ_a , weights λ, α , EMA factors β_s, β_l , ratio ρ , Top- N , MAD factor η , exponent γ

Output: 2D map \mathcal{M}

```

1: Local Registration
2:  $\mathcal{V} \leftarrow \{1\}$  ▷ valid frame indices
3: for  $i \leftarrow 1$  to  $n - 1$  do
4:    $(\Delta p_i, \Delta \theta_i) \leftarrow \text{FMTREGISTRATION}(I_i, I_{i+1}, s)$ 
5:    $(\delta_{p,i}, \delta_{\theta,i}) \leftarrow \text{VAL}((\Delta p_i, \Delta \theta_i), (\Delta p_{\text{glob},i}, \Delta \theta_{\text{glob},i}))$ 
6:   if  $\delta_{p,i} > \tau_p$  or  $\delta_{\theta,i} > \tau_a$  then
7:     discard  $i+1$ 
8:   else
9:      $\mathcal{V} \leftarrow \mathcal{V} \cup \{i+1\}$ 
10:  end if
11: end for
12: Global Optimization
13:  $\mathcal{R} \leftarrow \text{GETRESIDUALS}(\mathcal{P}, (\Delta p_i, \Delta \theta_i), (G_i^{(x,y)}, G_i^\theta))$ 
14:  $\mathcal{P}^* \leftarrow \text{OPTIMIZE}(\mathcal{R}; \lambda, \alpha)$ 
15: Robust Blending
16: for  $i \in \mathcal{V}$  do
17:   Warp  $I_i$  onto the canvas using  $(\Theta_i, P_i) \in \mathcal{P}^*$ 
18:    $(v_{i,s}, v_{i,l}) \leftarrow \text{UPDATEVARIANCEMAPS}(I_i, \beta_s, \beta_l)$ 
19:    $s_i \leftarrow \text{COMPUTEFEATURESCORE}(I_i, v_{i,s}, v_{i,l})$ 
20: end for
21: Initialize  $\mathcal{M}$ 
22: for each canvas pixel  $(x, y)$  do
23:    $\mathcal{K}'(x, y) \leftarrow \text{SELECTCANDIDATES}(x, y, \rho, N, \eta)$ 
24:    $\mathcal{M}(x, y) \leftarrow \text{ROBUSTBLEND}(\mathcal{K}'(x, y), \gamma)$ 
25: end for
26: return  $\mathcal{M}$ 

```

B. Architecture

Our proposed system follows a modular architecture consisting of image preprocessing, local registration, global optimization, and blending (as shown in Fig. 2, also summarized in Algorithm 1). To handle varying dataset scales, we implement three mapping strategies:

Baseline: Frames are aligned solely based on the *Local Registration* block of Algorithm 1, skipping *Global Opti-*

mization and proceeding directly to *Robust Blending*. This configuration is used for the ablation study.

Single-stage: The complete pipeline of Algorithm 1 is executed in this design, which is effective for small to medium-scale datasets.

Two-stage: The sequence is divided into subsets with 50% overlap, balancing computational cost and registration success. In Stage 1, each subset is processed using the *Baseline* procedure. In Stage 2, the resulting submaps are aligned and fused using the complete pipeline of Algorithm 1, identical to the *Single-stage* setting. This hierarchical design enables better preservation of local details and reduces computational complexity for large-scale datasets.

C. Preprocessing

Prior to registration, the following preprocessing steps are performed to reduce sonar-specific artifacts and enhance local textures (Fig. 3): (1) Row-wise intensity equalization on polar frames to suppress radial ring artifacts via mean subtraction per scanline. (2) Optional median filtering with kernel size 5, or non-local means (NLM) [28] filtering with filter strength $h = 15$, template window size 7, and search window size 21, to remove background noise. (3) Global histogram matching with a reference image to normalize intensity, followed by contrast limited adaptive histogram equalization (CLAHE) [29] with clip limit 1.5 and tile grid size 16 to enhance local texture. (4) Optional Gaussian blurring with kernel size 3 and automatically determined standard deviation to prevent over-sharpening induced by CLAHE. (5) Normalization and conversion from polar-to-Cartesian coordinates, with edge smoothing using a Hamming window to suppress boundary discontinuities.

D. EKF-based Sensor Fusion for Enhanced Positioning

GPS receivers provide absolute localization but have limitations in accuracy and update rates. In contrast, IMUs deliver high-frequency measurements of angular rate and acceleration but suffer from drifts. To leverage the complementary strengths of these sensors, we employ an EKF that fuses GPS, IMU, and compass data to yield smooth, accurate estimates of 2-D position and heading.

The EKF maintains an 8-dimensional state vector

$$\mathbf{x}_k = [p_x, p_y, v_x, v_y, \psi, b_\omega, b_{a_x}, b_{a_y}]^\top, \quad (2)$$

where (p_x, p_y) are local positions, (v_x, v_y) are velocities, ψ is heading, and $b_\omega, b_{a_x}, b_{a_y}$ represent IMU biases.

State propagation uses IMU angular rates and accelerations through a nonlinear kinematic model, while GPS provides direct position updates and the compass supplies heading observations. The corresponding measurement matrices are linear in the state. At each step, the EKF performs the standard *predict* stage using IMU inputs, and when GPS or compass data are available, an *update* is performed using the Kalman gain:

$$\hat{\mathbf{x}}_{k|k} = \hat{\mathbf{x}}_{k|k-1} + K_k(z_k - H\hat{\mathbf{x}}_{k|k-1}), \quad (3)$$

with covariance updated accordingly.

After processing all N measurements, the corrected positions and headings are sampled and linearly interpolated to provide a smooth trajectory estimate.

E. Local Registration

Pairwise registration between consecutive frames is performed using the FMT, a widely used phase correlation method [4], [5]. For each consecutive pair (I_i, I_{i+1}) , the relative displacement Δp_i and rotation $\Delta \theta_i$ are estimated:

$$(\Delta p_i, \Delta \theta_i) = \text{FMTREGISTRATION}(I_i, I_{i+1}, s), \quad (6)$$

where s is the pixel scale in meters.

From global priors, relative motion is computed as

$$\Delta p_{\text{glob},i} = G_{i+1}^{(x,y)} - G_i^{(x,y)}, \quad \Delta \theta_{\text{glob},i} = G_{i+1}^\theta - G_i^\theta. \quad (7)$$

Motion differences relative to global priors are then evaluated:

$$\delta_{p,i} = \|\Delta p_i - \Delta p_{\text{glob},i}\|, \quad \delta_{\theta,i} = |\Delta \theta_i - \Delta \theta_{\text{glob},i}|. \quad (8)$$

The pair (I_i, I_{i+1}) is rejected if $\delta_{p,i} > \tau_p$ or $\delta_{\theta,i} > \tau_\theta$. This procedure yields accurate local alignment while constraining drift through global priors, and corresponds to the *Local Registration* block in Algorithm 1 (lines 1–11).

F. Global Optimization

The global optimization module refines the ASV trajectory by combining local registration residuals with global constraints. Specifically, for each consecutive pair, the local and global residuals, respectively, are defined as:

$$\begin{cases} r_{\text{local,pos},i} = (P_{i+1} - P_i) - \Delta p_i, \\ r_{\text{local,angle},i} = (\Theta_{i+1} - \Theta_i) - \Delta \theta_i, \\ r_{\text{glob,pos},i} = P_i - G_i^{(x,y)}, \\ r_{\text{glob,angle},i} = \Theta_i - G_i^\theta. \end{cases} \quad (9)$$

The optimization problem is then defined as:

$$\mathcal{P}^* = \arg \min_{\mathcal{P}} \sum_i \left(\lambda (\|r_{\text{local,pos},i}\|^2 + \alpha \|r_{\text{local,angle},i}\|^2) + (\|r_{\text{glob,pos},i}\|^2 + \alpha \|r_{\text{glob,angle},i}\|^2) \right). \quad (10)$$

Solved via nonlinear least squares initialized from global measurements, in Algorithm 1, these residuals are encapsulated by GETRESIDUALS and optimized in the *Global Optimization* block (lines 12–14).

G. Image Blending

After optimization, each image is rotated and aligned according to its refined orientation Θ_i and position P_i . To resolve conflicts in overlapping regions, we adopt a variance-based blending strategy derived from Su et al. [27], with modifications to improve robustness. For each image, a local variance map is computed, and both short-term and long-term statistics are smoothed using EMA. Specifically, for frame i , the EMA updates are defined as

$$v_{i,s} = \beta_s v_{i-1,s} + (1 - \beta_s) \tilde{v}_i, \quad (11)$$

$$v_{i,l} = \beta_l v_{i-1,l} + (1 - \beta_l) \tilde{v}_i, \quad (12)$$

where \tilde{v}_i denotes the instantaneous local variance map, and $\beta_s, \beta_l \in (0, 1)$ are the decay factors for the short-term and long-term statistics, respectively. Larger β values

yield smoother but more inertial statistics, replacing fixed-size temporal windows and improving feature stability near sequence boundaries.

The feature score at pixel (x, y) is then computed as

$$s_i = v_{i,s}^{\text{norm}} \cdot e^{-v_{i,l}^{\text{norm}}}, \quad (13)$$

where $v_{i,s}$ and $v_{i,l}$ are normalized within the valid fan-shaped region. The resulting scores form the feature score map (FSM) for frame i . For each canvas pixel (x, y) , candidate intensities from overlapping frames are filtered by an adaptive threshold relative to the maximum score. Specifically, the retained set is defined as

$$\mathcal{K}(x, y) = \left\{ i : s_i(x, y) \geq \rho \cdot \max_j s_j(x, y) \right\}, \quad (14)$$

where ρ is the relative threshold parameter, and at most N highest-scoring candidates are retained. To further suppress outliers, median absolute deviation (MAD)-based trimming with factor η is applied to the retained intensity set, yielding a trimmed set $\mathcal{K}'(x, y)$.

A weighted trimmed mean is then computed as

$$\mathcal{M}(x, y) = \frac{\sum_{i \in \mathcal{K}'(x,y)} w_i I_i}{\sum_{i \in \mathcal{K}'(x,y)} w_i}, \quad w_i = s_i^\gamma, \quad (15)$$

where $\mathcal{M}(x, y)$ is the final blended intensity at canvas location (x, y) , I_i is the intensity from the i -th overlapping frame, and γ controls the sharpness of score weighting. This strategy ensures seamless transitions, preserves fine texture, and suppresses striping and outlier artifacts. These steps correspond to the *Robust Blending* block in Algorithm 1 (lines 15–26).

IV. EXPERIMENTAL RESULTS

We conducted a series of field experiments with an ASV equipped with FLS to evaluate our mapping framework. This section describes the experimental setup, followed by the results from three complementary performance evaluations.

A. Experimental Setup

The ASV surveyed a test site with a pre-deployed quadrat array following lawnmower trajectories. To assess robustness under varying motion conditions, a total of 15 trials were conducted with different survey paths. The collected dataset provides a controlled yet diverse basis for evaluating our mapping framework against the known quadrat geometry.

1) *ASV Platform*: The ASV was built on the chassis of a pontoon boat and driven by BlueRobotics T200 thrusters. Navigation and control relied on a Pixhawk PX4 sensor suite with an M8N GPS module and a magnetic compass. The IMU provided angular rate and acceleration at 55 Hz, while GPS and compass supplied absolute position and heading at 2 Hz. An Oculus M1200d FLS was mounted at a 15° grazing angle on the back of the ASV. The FLS (130° FOV and 5 m range) performed scans at 2.1 MHz and 10 FPS.

TABLE I: Ablation of preprocessing.

Eq	Filter	CLAHE	GB	Ent.	Var.	RMS	Ten.	VoL	SS
<i>No CLAHE (GB = -)</i>									
-	-	-	-	0.905	0.036	0.189	0.333	0.687	0.184
✓	-	-	-	0.921	0.035	0.187	0.358	0.743	0.178
✓	Median	-	-	0.864	0.035	0.188	0.099	0.060	0.143
✓	NLM	-	-	0.892	0.035	0.188	0.268	0.538	0.180
<i>With CLAHE (GB = -)</i>									
-	-	✓	-	0.966	0.056	0.237	0.742	1.550	0.174
✓	-	✓	-	0.963	0.052	0.228	0.804	1.701	0.166
✓	Median	✓	-	0.953	0.048	0.218	0.180	0.108	0.140
✓	NLM	✓	-	0.955	0.055	0.234	0.631	1.281	0.172
<i>GB effect (only when CLAHE = ✓)</i>									
✓	-	✓	✓	0.941	0.044	0.210	0.308	0.154	0.178

TABLE II: Relative-pose consistency across methods (evaluated against FMT-only and global-prior-only alignments).

Method	FMT only		Global prior only	
	MAE (m)	RMSE (m)	MAE (m)	RMSE (m)
ORB + RANSAC	0.316	1.173	0.304	1.161
SIFT + RANSAC	1.058	2.807	1.048	2.800
ECC maximization	0.069	0.106	0.066	0.085
FMT only	-	-	0.042	0.064
Global prior only	0.042	0.064	-	-
Ours (single-stage)	0.005	0.006	0.038	0.062

2) *Deployment Environment*: Field trials were conducted in 3-4 feet deep saltwater test site using a 2×4 array of square quadrats arranged in two parallel rows (see Fig. 1). One row consisted of aluminum frames with 17 in (≈ 0.43 m) side length, and the other of PVC frames with 24.5 in (≈ 0.62 m) side length. The two rows were separated by 60 in (≈ 1.52 m) between their nearest edges, and quadrats within each row were spaced 1 m apart center-to-center. Oysters were distributed inside the frames to introduce texture variation.

3) *Implementation and Runtime*: All modules were executed on a desktop workstation CPU (Intel(R) Core(TM) i7-14700K @ 3.40 GHz, 64 GB RAM). For a representative sequence of 200 frames, the average runtimes for frame-to-frame registration and validation, global trajectory refinement (nonlinear optimization), and warping/rendering and blending are 109.3 s, 12.7 s, and 104.8 s, respectively.

B. Image Processing Ablation

Each preprocessing configuration was evaluated over 15 trials, with each trial consisting of more than 400 FLS images. The reported results are averaged across all trials, covering combinations of row-wise brightness equalization, median filtering vs. NLM denoising, CLAHE contrast enhancement, and Gaussian blurring.

Representative before/after patches are shown in Fig. 3. The ablation results are summarized in Table I. We report six indicators: entropy (Ent.) [30] reflecting information richness, variance (Var.) and RMS contrast (RMS) [31] characterizing intensity variation, Tenengrad (Ten.) [32] and variance of Laplacian (VoL) [33] quantifying sharpness, and the similarity score (SS), defined as the peak of normalized cross-correlation (NCC) from FMT registration [34], indicating registration consistency. The results show that CLAHE substantially enhances structural metrics, often achieving the




Fig. 4: Qualitative comparison of map registration on 180 FLS images, all using our proposed blending method. (a) FMT-only baseline, map size $13.7 \text{ m} \times 13.0 \text{ m}$; (b) single-stage, map size $12.6 \text{ m} \times 14.1 \text{ m}$; (c) two-stage, map size $11.5 \text{ m} \times 14.4 \text{ m}$ (both with $\lambda = 10^4$, $\alpha = \frac{1}{225}$). Orange and yellow bounding boxes enclose 4 aluminum and 4 PVC quadrats, respectively.

TABLE III: Quadrat spacing errors across methods.

Method	MAE (m)	RMSE (m)
FMT only	0.189	0.221
Ours (single-stage)	0.173	0.200
Ours (two-stage)	0.176	0.207

highest scores, but it produces ring artifacts around bright regions that hinder blending. Combining row-wise equalization with CLAHE suppresses such artifacts and yields consistently strong performance across all metrics, providing the best option for addressing the trade-off between detail preservation and artifact suppression. Therefore, this approach is adopted as the most effective preprocessing scheme for the mapping pipeline.

C. Registration Accuracy and Ablation

Seven registration strategies were benchmarked: oriented FAST and rotated BRIEF (ORB) [35] + random sample consensus (RANSAC) [36], scale-invariant feature transform (SIFT) [37] + RANSAC, enhanced correlation coefficient (ECC) maximization [38], FMT-only, global-prior-only alignment, and our pipeline in both single- and two-stage variants. ORB and SIFT are classical feature detectors/descriptors, while RANSAC is a robust model fitting scheme. ECC maximization is an intensity-based (direct) parametric registration method that aligns images by maximizing the ECC.

Table II reports relative-pose consistency errors. Feature-based methods perform poorly on sonar imagery: ORB+RANSAC yields large errors, and SIFT+RANSAC fails outright with many outliers. ECC achieves low step-to-step errors but collapses on long sequences due to lacking reliable global constraints. Global-prior-only appears locally accurate but produces distorted maps with indistinguishable quadrats. FMT-only and our single-stage pipeline achieve strong local consistency, with the latter giving the lowest residuals. The two-stage pipeline is excluded here since trajectory-based consistency cannot be directly evaluated after global re-optimization.

Table III evaluates quadrat spacing errors against the ground-truth 2×4 layout. Our single-stage method per-

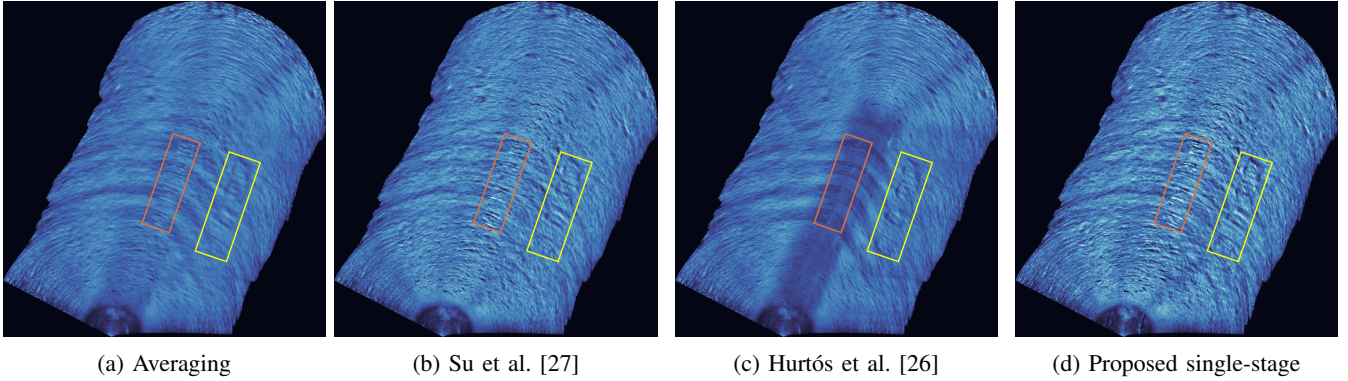


Fig. 5: Qualitative comparison of map blending methods on a $13.2\text{m} \times 13.7\text{m}$ area with 180 images, processed using the single-stage pipeline with $\lambda = 10^4$ and $\alpha = \frac{1}{225}$: (a) simple intensity averaging, (b) variance-based, (c) distance-based, and (d) proposed blending approaches. Orange and yellow bounding boxes enclose 4 aluminum and 4 PVC quadrats, respectively.

TABLE IV: Effect of local registration weight λ on quadrat spacing accuracy (with $\alpha = \frac{1}{225}$).

λ	MAE (m)	RMSE (m)
10^2	0.476	0.606
10^3	0.429	0.503
10^4	0.173	0.200
10^5	0.201	0.248
10^6	0.225	0.262

forms the best. Compared with the FMT-only method, the MAE/RMSE errors reduce by 8.5% and 9.5%. The two-stage variant also outperforms the FMT-only method, with 6.9% and 6.3% reductions in MAE/RMSE errors. Global-prior-only is omitted here as it fails to recover quadrat geometry.

We further evaluated the effect of the local registration weight λ , which balances global constraints and frame-to-frame FMT alignment. The orientation-to-distance residual ratio was fixed at $\alpha = \frac{1}{225}$. Small values ($\lambda \leq 10^1$) were omitted since quadrats were not identifiable. Table IV shows that accuracy improves as λ increases to 10^4 , at which the errors are lowest. Larger values (10^5 – 10^6) slightly degrade performance, indicating overfitting to noisy local estimates.

The three approaches are compared qualitatively in Fig. 4. The single- and two-stage pipelines both reconstruct the global structure and the obtained quadrat spacing aligns well to the ground truth, while the FMT-only baseline, though retaining some local detail, shows evident drift. Between our two variants, the single-stage yields smoother stitching transitions and slightly lower spacing errors, whereas the two-stage strategy—by constructing overlapping submaps before global alignment—preserves sharper textures and clearer structural details. These results suggest that the single-stage pipeline is advantageous for seamless blending, while the two-stage pipeline enhances visual quality and robustness for longer sequences. Overall, our fusion methods achieve a balance of local consistency, global accuracy, and detail preservation, which outperforms the FMT-only baseline.

D. Map Blending Performance

We compared our proposed blending approach against three baselines under the single-stage pipeline: simple in-

TABLE V: Map blending quality metrics using different blending methods.

Method	Ent.	Var.	RMS	Ten.	VoL
Averaging	0.619	0.036	0.189	0.042	0.039
Hurtós et al.	0.634	0.035	0.186	0.047	0.047
Su et al.	0.656	0.043	0.208	0.084	0.077
Ours	0.680	0.048	0.219	0.142	0.145

tensity averaging, the distance-based method of Hurtós et al. [26], and the variance-based method of Su et al. [27]. For our method, the blending parameters were set to $\beta_s = 0.60$, $\beta_l = 0.97$, $\rho = 0.68$, $N = 16$, $\eta = 2.7$, and $\gamma = 1.4$.

As shown in Table V, compared with Su et al., our method shows moderate improvements (3–12%) in Ent., Var., and RMS, while Ten. and VoL are nearly doubled (from 0.084/0.077 to 0.142/0.145). Relative to simple averaging, Ten. and VoL increase more than threefold, highlighting substantial gains in edge sharpness and local detail. Fig. 5 illustrates these differences. Our approach best preserves structural detail, ensures smooth transitions, and suppresses noise and artifacts. Su’s method is less sharp, Hurtós’s method suffers from severe blind-region artifacts, and averaging produces blurred, indistinct structures.

Overall, through both quantitative and qualitative evaluations, our method is confirmed to achieve the most effective blending, offering improved balance between artifact suppression and detail preservation.

V. CONCLUSION

We introduced a drift-resilient seabed mapping framework that couples local sonar registration, globally consistent trajectory refinement, and detail-preserving blending. By integrating GPS-based priors with refined FMT alignment, the system reduces long-term drift and attains sub-meter georeferenced accuracy, with a 9.5% RMSE improvement over the FMT-only baseline and consistently better performance than representative registration baselines. The proposed variance-driven blending further enhances map quality by suppressing artifacts and preserving fine textures, yielding clearer mosaics than existing blending methods. Field trials in structured coastal environments confirm the practicality of

this low-cost ASV-based solution for large-area autonomous seabed surveys.

Limitations and Future Work: While the proposed system produces accurate quadrat-level mosaics, the current formulation assumes near-planar seabed conditions with 2-D registration, which may degrade under strong 3-D relief or large slope conditions. In addition, the current implementation requires offline processing. As next steps, we will extend quadrat-level mapping to inventory surveys and perform substrate classification within each quadrat to estimate oyster distribution, size, and condition. Future work will also enable onboard online mapping to support informative path planning and efficient surveys of regions of interest.

REFERENCES

- [1] S. Negahdaripour, "On 3-d motion estimation from feature tracks in 2-d fs sonar video," *IEEE Transactions on Robotics*, vol. 29, no. 4, pp. 1016–1030, 2013.
- [2] N. Hurtós, S. Nagappa, X. Cufí, Y. Petillot, and J. Salvi, "Evaluation of registration methods on two-dimensional forward-looking sonar imagery," in *2013 MTS/IEEE OCEANS - Bergen*, 2013, pp. 1–8.
- [3] M. D. Aykin and S. Negahdaripour, "On feature matching and image registration for two-dimensional forward-scan sonar imaging," *Journal of Field Robotics*, vol. 30, no. 4, pp. 602–623, 2013.
- [4] N. Hurtós, Y. Petillot, J. Salvi *et al.*, "Fourier-based registrations for two-dimensional forward-looking sonar image mosaicing," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. Ieee, 2012, pp. 5298–5305.
- [5] N. Hurtós, D. Ribas, X. Cufí, Y. Petillot, and J. Salvi, "Fourier-based registration for robust forward-looking sonar mosaicing in low-visibility underwater environments," *Journal of Field Robotics*, vol. 32, no. 1, pp. 123–151, 2015.
- [6] N. Karapetyan, A. Braude, J. Moulton, J. A. Burstein, S. White, J. M. O'Kane, and I. Rekleitis, "Riverine coverage with an autonomous surface vehicle over known environments," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 3098–3104.
- [7] M. A. Xu, N. Karapetyan, K. Rajasekaran, A. Williams, D. A. Pattillo, P. Tokekar, B. Hudson, Y. Tao, D. Webster, M. Gray, and M. Yu, "Machine learning-based classification of on-bottom oyster farms using a mechanical scanning sonar," *IEEE Sensors Journal*, vol. 25, no. 12, pp. 22 068–22 079, 2025.
- [8] X. Liu, J. Yang, C. Lu, E. Zhang, and W. Xu, "Robust forward-looking sonar-image mosaicking without external sensors for autonomous deep-sea mining," *Journal of Marine Science and Engineering*, vol. 13, no. 7, p. 1291, 2025.
- [9] J. Li, P. Wang, H. Zhang, Y. Yang, and B. Liu, "Realtime registration for phased-array submarine sonar image," *OCEANS 2018 MTS/IEEE Charleston*, pp. 1–6, 2018.
- [10] P. J. Besl and N. D. McKay, "Method for registration of 3-d shapes," in *Sensor fusion IV: control paradigms and data structures*, vol. 1611. Spie, 1992, pp. 586–606.
- [11] Z. Zhang, "Iterative point matching for registration of free-form curves and surfaces," *International journal of computer vision*, vol. 13, no. 2, pp. 119–152, 1994.
- [12] W. Dos Santos, T. Oliverira, and E. Brunskill, "Underwater sonar and aerial images data fusion for robot localization," *Robotics and Autonomous Systems*, vol. 121, pp. 103–115, 2019.
- [13] J. Li, M. Kaess, R. M. Eustice, and M. Johnson-Roberson, "Pose-graph slam using forward-looking sonar," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 2330–2337, 2018.
- [14] H. Johannsson, M. Kaess, B. Englot, F. Hover, and J. Leonard, "Imaging sonar-aided navigation for autonomous underwater harbor surveillance," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2010, pp. 4396–4403.
- [15] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. J. Leonard, and F. Dellaert, "isam2: Incremental smoothing and mapping using the bayes tree," *The International Journal of Robotics Research*, vol. 31, no. 2, pp. 216–235, 2012.
- [16] H. Kim, G. Kang, S. Jeong, S. Ma, and Y. Cho, "Robust imaging sonar-based place recognition and localization in underwater environments," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 1083–1089.
- [17] Q. Wu, G. He, C. Liu, J. Li, B. Zhang, and C. Tao, "Research on a high-precision real-time 3d sonar image registration algorithm," in *2023 IEEE International Conference on Control, Electronics and Computer Technology (ICCECT)*, 2023, pp. 413–416.
- [18] M. Y. Cheung, D. Fourie, N. R. Rypkema, P. V. Teixeira, H. Schmidt, and J. Leonard, "Non-gaussian slam utilizing synthetic aperture sonar," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 3457–3463.
- [19] J. McConnell, Y. Huang, P. Szenher, I. Collado-Gonzalez, and B. Englot, "Draco-slam: Distributed robust acoustic communication-efficient slam for imaging sonar equipped underwater robot teams," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 8457–8464.
- [20] Y. Huang, J. McConnell, X. Lin, and B. Englot, "Draco-slam2: Distributed robust acoustic communication-efficient slam for imaging sonar equipped underwater robot teams with object graph matching," in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025, pp. 3587–3594.
- [21] J. McConnell, F. Chen, and B. Englot, "Overhead image factors for underwater sonar-based slam," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4901–4908, 2022.
- [22] J. McConnell, I. Collado-Gonzalez, P. Szenher, and B. Englot, "Large-scale dense 3-d mapping using submaps derived from orthogonal imaging sonars," *IEEE Journal of Oceanic Engineering*, 2024.
- [23] N. Karapetyan, M. Xu, Y. Zhang, K. Joshi, N. Anil, S. Mayenkar, X. Lin, Y. Aloimonos, N. Chopra, P. Malone *et al.*, "Oysterbot: Autonomous surface vehicle based system for aquaculture monitoring operations," in *OCEANS 2025-Great Lakes*. IEEE, 2025, pp. 1–7.
- [24] L. J. Poppe, R. Lewis, H. J. Knebel, E. A. Haase, K. Parolski, and M. DiGiacomo-Cohen, "Sidescan sonar images, surficial geologic interpretations, and bathymetry of the long island sound sea floor in new haven harbor and new haven dumping ground, connecticut," US Geological Survey, Tech. Rep., 2001.
- [25] S. Masood, M. Sharif, M. A. Yasmin, M. A. Shahid, and A. Rehman, "Image fusion methods: A survey," *Journal of Engineering Science & Technology Review*, vol. 10, no. 6, 2017.
- [26] N. Hurtós, X. Cufí, and J. Salvi, "A novel blending technique for two-dimensional forward-looking sonar mosaicing," in *2013 OCEANS - San Diego*, 2013, pp. 1–7.
- [27] J. Su, H. Li, J. Qian, X. An, F. Qu, and Y. Wei, "A blending method for forward-looking sonar mosaicing handling intra- and inter-frame artifacts," *Ocean Engineering*, vol. 298, p. 117249, 2024.
- [28] A. Buades, B. Coll, and J.-M. Morel, "A non-local algorithm for image denoising," in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, vol. 2, 2005, pp. 60–65.
- [29] K. Zuiderveld, *Contrast limited adaptive histogram equalization*. USA: Academic Press Professional, Inc., 1994, p. 474–485.
- [30] C. E. Shannon, "A mathematical theory of communication," *The Bell System Technical Journal*, vol. 27, no. 3, pp. 379–423, 1948.
- [31] E. Peli, "Contrast in complex images," *Journal of the Optical Society of America A*, vol. 7, no. 10, pp. 2032–2040, 1990.
- [32] E. Krotkov, "Focusing," *International Journal of Computer Vision*, vol. 1, no. 3, pp. 223–237, 1988.
- [33] J. Pech-Pacheco, G. Cristobal, J. Chamorro-Martinez, and J. Fernandez-Valdivia, "Diatom autofocusing in brightfield microscopy: a comparative study," in *Proceedings 15th International Conference on Pattern Recognition. ICPR-2000*, vol. 3, 2000, pp. 314–317.
- [34] B. Reddy and B. Chatterji, "An fft-based technique for translation, rotation, and scale-invariant image registration," *IEEE Transactions on Image Processing*, vol. 5, no. 8, pp. 1266–1271, 1996.
- [35] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "Orb: An efficient alternative to sift or surf," in *2011 International Conference on Computer Vision*, 2011, pp. 2564–2571.
- [36] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, p. 381–395, Jun. 1981.
- [37] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International journal of computer vision*, vol. 60, no. 2, pp. 91–110, 2004.
- [38] G. D. Evangelidis and E. Z. Psarakis, "Parametric image alignment using enhanced correlation coefficient maximization," *IEEE transactions on pattern analysis and machine intelligence*, vol. 30, no. 10, pp. 1858–1865, 2008.