

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

# Above and Below: Heterogeneous Multi-robot SLAM Across Surface and Underwater Domains

John McConnell<sup>1</sup>, Armon Shariati<sup>2</sup>, Paul Szenher<sup>3</sup> and Yaxuan Li<sup>3</sup>

**Abstract**—Multi-robot simultaneous localization and mapping (SLAM) is a fundamental task in multi-robot operations. Robots must have a common understanding of their location and that of their team members to complete coordinated actions. However, multi-robot SLAM between Uncrewed Surface Vessels (USVs) and Autonomous Underwater Vehicles (AUVs) has primarily been achieved through acoustic pinging between robots to retrieve range measurements; a measurement technique requires that robots to be in similar locations simultaneously, have an uninterrupted path for signal propagation, and may necessitate synchronized clocks. This is especially challenging in complex, cluttered maritime environments, where structures may impede signals. However, these same structures may be observable above and below the water’s surface, presenting an opportunity for inter-robot SLAM loop closure between USV and AUV data streams. This work builds upon recent research on inter-robot SLAM loop closure between USV and AUV data [4], extending it to propose a centralized multi-robot SLAM system. Each robot performs its state estimation, and we detect loop closures between each AUV and the USV data. These inter-robot loop closures are used to merge each robot’s state estimate into a centralized graph, yielding estimates for the whole time history of the USV and all AUVs in the system. Validation is performed using real-world perceptual data in three different environments. Results show improved errors for AUVs in the multi-robot SLAM system compared to single-robot SLAM over the same trajectories. To our knowledge, this is the first instance of a multi-robot SLAM system with AUVs and USVs built on loop closures rather than acoustic distance measurements. The views expressed in this document are those of the author(s) and do not reflect the official policy or position of the U.S. Naval Academy, Department of the Navy, the Department of Defense, or the U.S. Government.

**Index Terms**—Marine Robotics, SLAM, Range Sensing

## I. INTRODUCTION

**A**UTONOMOUS underwater vehicles (AUVs) have become critical tools in many settings, including Oil & Gas, aquaculture, offshore renewable energy, and defense. AUVs are routinely employed for inspection, repair, and tasks where using human divers might be dangerous or otherwise impossible [1]. However, underwater localization is challenging compared to other domains due to the lack of GPS and low-quality perceptual data.

A notable approach to addressing the difficulties of AUV localization is a multi-robot solution. Uncrewed Surface Vessels (USV) have easy access to high-accuracy localization through GPS or high-quality perceptual systems. In a multi-robot setting, USV localization is leveraged to enhance the quality




Fig. 1: **Multi Robot Map Result.** A sample multi-robot mission from the harbor environment. The blue points show the 2D LiDAR map, with the blue line showing the USV trajectory. Orange points indicate the sonar map, with the orange line representing the AUV’s trajectory. Grid lines are shown on the satellite image at 10-meter cell size.

and robustness of AUV localization. Multi-robot localization typically comes in two forms, *direct* and *indirect* encounters [2], [3]. Direct encounters typically involve acoustic ranging between robots. However, there are limitations to direct encounters, mainly requiring proximity between robots, a clear line of sight for signal propagation, and *may* require robots to have synchronized clocks. Indirect encounters associate targets observed in the environment to estimate measurements between robots, offering several significant advantages: they do not need robots to be simultaneously in the same part of the environment, can be resolved whenever communication is available, and do not require clock synchronization.

In this work, we consider teams of USVs and AUVs operating in complex, cluttered maritime environments, where underwater structures frequently impede the path of acoustic ranging and communication. However, we consider the utility of human-made structures in littoral settings. Such structures are often observable from above and below the water’s surface, including bridges, docks, ships, and pier pylons. This work will leverage the insight that common structures are observable above and below the water’s surface and use this to develop indirect encounters between USVs and AUVs. Therefore, we consider two important sensing models: (1) wide aperture multi-beam imaging sonar for its robustness to water quality, immunity to lighting conditions, and long sensing range. (2) 3D above water LiDAR for its dense, high-quality information, making it simple to render observations *near* the waterline regardless of USV roll and pitch. This paper utilizes recent developments in indirect encounters (loop closure) between AUV and USV data streams [4]. However, unlike previous work, we use indirect encounters in a fully fledged multi-robot simultaneous localization and mapping (SLAM) system.

Our contributions are as follows:

- The first example of multi-robot SLAM for USVs and AUVs using indirect encounters from perceptual data. The proposed system’s novelty is the improvement and

Manuscript received: June 23, 2025; Revised September 21, 2025; Accepted October 27, 2025.

This paper was recommended for publication by Editor Giuseppe Loianno upon evaluation of the Associate Editor and Reviewers’ comments.

<sup>1</sup>J. McConnell is with the United States Naval Academy, Annapolis, MD, USA. [jmcconne@usna.edu](mailto:jmcconne@usna.edu). <sup>2</sup>A. Shariati is with Shield-AI, San Diego CA, USA. [armon.shariati@shield.ai](mailto:armon.shariati@shield.ai). <sup>3</sup>P. Szenher and <sup>3</sup>Y. Li are with the Stevens Institute of Technology, Hoboken, NJ, USA. [pszenher,yli21}@stevens.edu](mailto:{pszenher,yli21}@stevens.edu).

Digital Object Identifier (DOI):

**IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.**

integration of [4] into the broader SLAM problem.

- A compression method for sonar data that maintains cell resolution, making perceptual data exchange amenable with low bandwidth communications systems.
- Validation of the system in three real-world environments using simulated communications. Showing improved AUV localization quality over single robot SLAM. To the best of our knowledge, these are the first experiments of this type and scale that include ground truth. This is the most important, and broadest contribution of this work.
- An open-source repository for the code and datasets used in this paper. Code Link.

## II. RELATED WORK

### A. Above Water Prior Information

Related to our work is the use of prior information above the water's surface to assist in the underwater SLAM problem. [5] develops a similarity scoring system between an underwater sonar image and a sector of a satellite image of the same sensor footprint using deep learning. In follow-up work, the same authors extend this similarity score to serve as the backbone of a particle filter localization system for an underwater robot equipped with imaging sonar and a priori satellite images [6]. Lastly, [7] provides an example of sonar-based pose graph SLAM, incorporating constraints from a satellite image prior and standard place recognitions.

### B. Place Recognition

A fundamental task in robotics is place recognition, which involves detecting when a robot is in a location it has visited before and finding transformations between past and current timesteps. In the LiDAR sensing domain, Scan Context [8], [9], [10] has demonstrated impressive results without the need for machine learning. Scan context converts dense, high-resolution 3D point clouds from LiDAR into coarse context images. Scan context computes single-dimension descriptors to generate place recognition candidates using a KD-tree [11]. Scan context has recently been applied in underwater imaging sonar [12] and is an essential module in DRACo-SLAM [13], a multi-robot SLAM system designed for underwater robots equipped with imaging sonar. Object scene graphs have also been studied for imaging sonar place recognition [14], requiring a high number of objects in view to compute a transform. Machine learning is used in side-scan sonar place recognition [15], [16]; however, there is little publicly available underwater data, making it challenging to apply these methods to new problem settings. Lastly, [17] applies the iterative closest point (ICP) algorithm [18], using dead reckoning as an initial guess, to evaluate place recognition.

It is crucial to note that a single erroneous place recognition can lead to highly inaccurate robot pose estimation; for this reason, we consider robust outlier rejection systems. Pairwise consistent measurement set maximization (PCM) [19] evaluates measurements as a set rather than in isolation, acting as a pre-filter to the SLAM backend. Dynamic covariance scaling (DCS) [20] adjusts measurement covariance during the solver step, and switchable constraints [21] turn measurements on or off. While these robust backends are impressive, they often defer to the odometry chain, which, in the case of multi-robot SLAM, could be challenging given that the measurements between robots may be sparse.

### C. Cooperative Localization

Measurements between robots typically come in two forms: direct and indirect encounters [2]. Direct encounters occur when robots measure the distance between them, often achieved through acoustic pinging. Indirect measurements are developed by associating commonly observed targets in the environment to form transformations between robots [3]. A significant drawback of relying solely on direct encounters is that robots must be in the same part of the environment, with an unbroken path for the signal to pass between them. Moreover, while many methods avoid the need for clock synchronization [22][23], many others require it to develop range measurements between robots [24], [25], [26].

There are several notable examples of collaborative localization using direct encounters. Firstly, [27] shows a team of underwater robots using direct encounters with a surface station to perform collaborative localization. [28] shows an example of a pair of robots operating in a similar pattern. [29] shows a collaborative localization method with distributed processing. [30] introduces collaborative localization using several fixed-position beacons, comparing the performance across a robot team with and without the addition of fixed beacons. [31] evaluates several system variations that include onboard dead reckoning systems and inter-robot pinging. When considering work with indirect encounters, the above water setting has shown numerous impressive contributions [32]. In a maritime setting [2] presents a simulated example of an underwater multi-robot team using GPS at the surface, inter-robot pinging, and a data association method that relies on GPS for initialization. DRACo-SLAM [13] demonstrates an example of multi-robot SLAM using underwater robots with imaging sonar and no direct encounters. [4] finds inter-robot loop closures between underwater imaging sonar and above-water 3D LiDAR; however, this work focuses on the loop closure problem and does not integrate this tool with a SLAM solution, nor does it consider the cost of data exchange.

### D. Sonar Data Compression

When considering transmitting sonar data, compression is crucial due to bandwidth limitations. DRACo-SLAM [13] compresses data using a voxel grid, degrading resolution. [33] proposes a machine learning approach to sonar image compression; however, new datasets may be out of distribution, potentially degrading performance. In this work, we propose a simple method that we are confident will generalize beyond the experiments presented in this paper.

## III. PROBLEM DEFINITION

In this work we consider a team of  $N$  robots. Each robot  $n \in N$  maintains its own state estimate, in its own reference frame  $\mathcal{I}_n$ . Each robot collects observations,  $\mathbf{z}_{n,t}$  at pose  $\mathbf{x}_{n,t}$  at time step  $t$ . Each robot pose is defined in the plane as

$$\mathbf{x}_{n,t} = \begin{pmatrix} x_r, y_r, \theta_r \end{pmatrix}^\top, \quad (1)$$

and each set of observations is a two dimensional point cloud. AUVs use sonar, with returns in spherical coordinates denoted as ranges  $r_s \in \mathbb{R}_+$ , bearings  $\theta_s \subseteq [-\pi, \pi)$ , elevation angles  $\phi_s \subseteq [-\pi, \pi)$ , and associated intensity values  $\gamma_s \in \mathbb{R}_+$ . Imaging sonar however, does not measure  $\phi_s$ , and therefore we treat it as

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

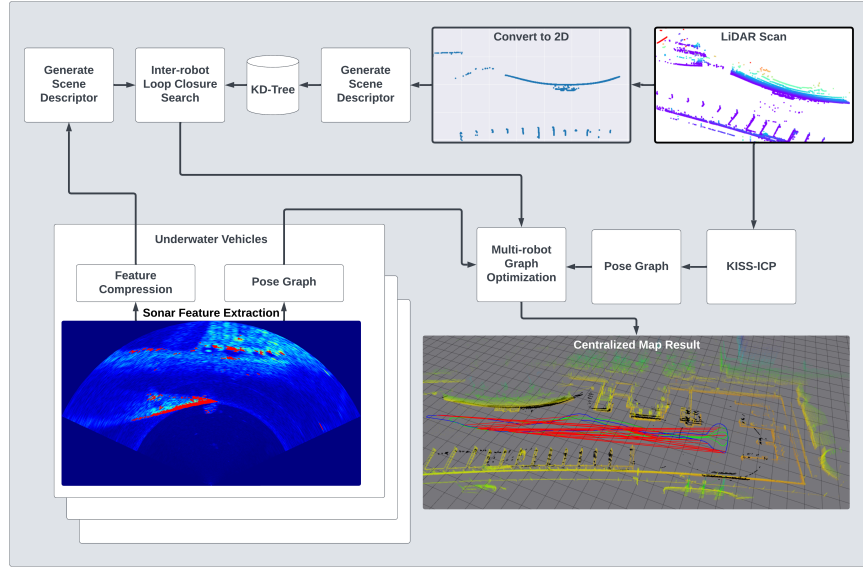


Fig. 2: **System Flow.** At bottom left, we show  $N$  underwater vehicles with imaging sonar; each AUV compresses its features and computes a pose graph, sending results to the USV. The USV converts its LiDAR scans into 2D scans near the water’s surface. The USV performs a loop closure search between the AUV’s sonar data and the USV’s LiDAR data. Loop closures are used to merge USV and AUV pose graphs, resulting in a centralized state estimate and a joint map, as shown in the bottom right. 3D LiDAR point clouds are shown, with color mapped to height. Black points represent the sonar map, the blue line indicates the USV trajectory, the green line represents one of the AUV trajectories, and the red lines denote inter-robot loop closures.

zero, or in-plane. USVs use LiDAR, with returns in spherical coordinates denoted as ranges  $r_L \in \mathbb{R}_+$ , bearings  $\theta_L \subseteq [-\pi, \pi)$ , elevation angles  $\phi_L \subseteq [-\pi, \pi)$ , and associated intensity values  $\gamma_L \in \mathbb{R}_+$ . Each robot moves through the environment according to the dynamics

$$\mathbf{x}_{n,t} = \mathbf{g}(\mathbf{u}_{n,t}, \mathbf{x}_{n,t-1}) + \epsilon_{n,t}, \quad (2)$$

where  $\mathbf{x}_{n,t-1}$  is the previous timestep’s pose,  $\mathbf{u}_{n,t}$  is the control command and  $\epsilon_{n,t}$  is process noise. The posterior probability over the time history of poses is defined as

$$p(\mathbf{x}_{n,1:t}, \mathbf{m}_n | \mathbf{z}_{n,1:t}, \mathbf{u}_{n,1:t}), \quad (3)$$

with map  $\mathbf{m}_n$ . Lastly, while we propose a multi-robot SLAM solution, there is no notion of initial transform between reference frames  $\mathcal{I}_n$ . We assume robots may communicate with one another for the entire mission duration, and only do so at discrete SLAM steps.

#### IV. ALGORITHM

##### A. USV State Estimation

While there are many methods for USV state estimation, we utilize LiDAR odometry. We found it to be simple, effective, provides heading, does not require GPS, yields much higher accuracy than GPS and since the downstream processes use LiDAR, this sensor is present onboard the USV. However, we would like to emphasize that **any** localization system could be used here.

1) *LiDAR Pre-processing:* Maritime environments can introduce noise into LiDAR point clouds due to reflections off the water surface. These reflections can even mirror the above-water structures below the waterline. For this reason, we form a partial state estimate for the USV,

$$\mathbf{x}_{t,p} = [z_I, \phi_I, \psi_I]^T. \quad (4)$$

Where  $z_I$  is the height of the LiDAR, measured from the water surface,  $\phi_I$  is the roll, and  $\psi_I$  is the pitch.  $z_I$  is sourced from

calibrating the height of the LiDAR above the water’s surface using a depth sensor mounted below the waterline with  $\phi_I$  and  $\psi_I$  from the IMU. The LiDAR scan  $\mathcal{P}_L$  is corrected for height, roll, and pitch using this partial state estimate.  $\mathcal{P}_L$  is cropped according to height, where only points above the waterline are retained, denoted as  $\hat{\mathcal{P}}_L$ . This method is simple yet effective for removing erroneous information from LiDAR point clouds.

2) *KISS-ICP:* We select KISS-ICP [34] for its simplicity, robustness, lack of system requirements, and functionality on this maritime data. KISS-ICP aligns each LiDAR scan,  $\hat{\mathcal{P}}_L$ , to the local LiDAR map,  $\mathcal{M}_L$ , a composite of previous scans, and provides a pose estimate for each LiDAR timestep in the USV reference frame.

3) *Pose Graph:* While pose information is available at each LiDAR timestep, we downsample this information to form a more sparse representation. If the USV has moved at least  $\Delta T_{min}$  or rotated  $\Delta R_{min}$ , we formulate a LiDAR odometry factor,  $\mathbf{f}^{LO}$ , which is the transformation from the last time discrete pose. Forming the factor graph,

$$\mathbf{f}(\Theta) = \mathbf{f}^0(\Theta_0) \prod_i \mathbf{f}_i^{LO}(\Theta_i), \quad (5)$$

which does not contain a source of loop closure. We found that LiDAR loop closure was unnecessary in testing, and simple LiDAR odometry provided highly accurate, stable results. However, it can easily be incorporated using similar systems in future work [35].

##### B. AUV State Estimation

In this section, we consider state estimation onboard the AUV. Similar to the USV, the source of this state estimation could be any system. We consider a dead-reckoning pose graph combining Doppler velocity logger (DVL) and IMU data.

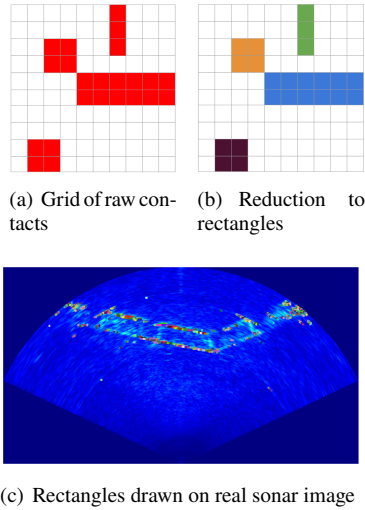


Fig. 3: **Sonar image compression.** (a) shows an illustration of a sonar image with multiple contacts highlighted in red. (b) shows the results of reducing each of these patches to a rectangle represented by the top left and bottom right corners. (c) shows a real-world example sonar image of (b), where all sonar contacts have been grouped into rectangles, each with a different color.

1) *Pose Graph:* We combine DVL+IMU information in an extended Kalman filter (EKF) [36], yielding AUV position estimates. Since AUVs are free floating, the Kalman filter motion model is simply particle motion. Due to space constraints, we direct interested readers to prior work for more details [36]. While DVL+IMU pose information is available at a high rate, we down-sample this information to a more sparse representation. If the AUV has moved at least  $\Delta T_{min}$  or rotated  $\Delta R_{min}$ , we create a DVL odometry factor,  $\mathbf{f}^{DO}$ , the transformation between the current and previous time discrete poses. Yielding the factor graph,

$$\mathbf{f}(\Theta) = \mathbf{f}^0(\Theta_0) \prod_i \mathbf{f}_i^{DO}(\Theta_i). \quad (6)$$

All factor graphs in this paper are implemented using the GTSAM library [37]. By design, this factor graph does not contain a source of loop closure; the factor graph yields a stable, yet drifting estimate to build our system off of. Downstream systems will layer in a source of loop closure, which we will compare in experiments to a more standard, single-robot approach.

### C. Communications Strategy

In this work, we consider a centralized multi-robot SLAM system. The USV acts as the central node, solving the multi-robot SLAM problem using the information provided by AUVs. Each AUV sends a message to the USV consisting of its current pose estimate and associated sonar data at each discrete timestep. This one-way communication of perceptual data is important, as we do not communicate LiDAR point clouds from USV to AUV at any time, avoiding the need to expend limited data exchange capacity on this costly data structure.

1) *Sonar Pre-processing:* Recall that sonar data originates as polar coordinate images, which are costly to transmit. For this reason, we do not propose transmitting raw sonar images from the AUV to the USV. Instead, we select which pixels are

legitimate contacts with the environment, using SOCA-CFAR [38], denoted as  $\mathcal{P}_c$ , contacts. By transmitting only the pixel coordinates of legitimate sonar contacts, represented by a pair of 32-bit integers, we achieve significant savings over raw images.

2) *Sonar Compression:* To further improve communication efficiency, we apply a compression technique to the sonar contacts,  $\mathcal{P}_c$ , with an illustrative example shown in Figure 3(a). We reduce  $\mathcal{P}_c$  to rectangular patches denoted as  $\mathcal{P}_r$  and illustrated in Figure 3(b). Contacts that are not part of a patch are retained using their coordinates. Rather than transmitting the coordinates of all contacts as shown in Figure 3(a), we can transmit the top left and bottom right coordinates of each patch shown in Figure 3(b), with a real-world example of these patches shown in Figure 3(c). While the efficiency of this compression varies with the environment, it maintains the cell resolution, where other methods [13] reduce resolution, and experimental results will demonstrate its effectiveness.

### D. Robot Loop Closure

1) *Loop Closure LiDAR Pre-processing:* In downstream systems, our system will search for possible loop closures between USV LiDAR scans and AUV sonar scans. However, 3D LiDAR scans must be rendered into a standardized representation with sonar data, such as a 2D point cloud. This 2D point cloud consists of points *near* the water's surface. In Section IV-A1, we remove erroneous points and any data presented below the water. This LiDAR scan  $\hat{\mathcal{P}}_L$  is further processed by removing all points with a height above  $h_{max}$ , yielding LiDAR scan  $\hat{\mathcal{P}}^*_L$ . Finally, we drop the height axis of the LiDAR scan altogether, creating a 2D point cloud near the water's surface  $\hat{\mathcal{P}}_L$ .

2) *Loop Closure Sonar Pre-processing:* Sonar data arrives onboard the USV as  $\mathcal{P}_r$ , which are pixel locations in a polar coordinate image, which are converted to cartesian 2D sonar point clouds,  $\mathcal{P}_s$ . These point clouds are in the plane, assuming a zero elevation angle. This detail is important since, with an imaging sonar, this elevation angle is unknown. However, in our work, this zero elevation angle represents a reasonable approach to registering sonar and LiDAR point clouds.

3) *Inter-robot Loop Closure Search:* To find correspondences between USV-LiDAR and AUV-sonar data streams, we form scene descriptors to query for possible matches between scenes. While there are numerous ways to develop these descriptors, we employ a simple and effective method, namely range histograms, similar to many notable loop closure systems [8], [9], [10], [12], [13]. Each 2D point cloud,  $\hat{\mathcal{P}}_L$  and  $\mathcal{P}_s$  is converted from x-y coordinates, to range. We then enforce a maximum range  $r_{max}$  on both data products and discretize the range into a LiDAR scene descriptor,  $k_L$ , and a sonar scene descriptor  $k_s$  as appropriate. Each scene descriptor is encoded as

$$k = (f(r_1), \dots, f(r_N)), f : r_i \rightarrow \mathbb{R} \quad (7)$$

where the contents of each discrete bin,  $f(r_i)$ , is the number of occurrences inside that range bin. Lastly, we normalize the counts in each histogram bin by dividing them by the maximum count in each descriptor. A candidate loop closure, consisting of a pair of discrete time steps,  $i$  from an AUV and  $j$  from the USV, is generated using a KD tree.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

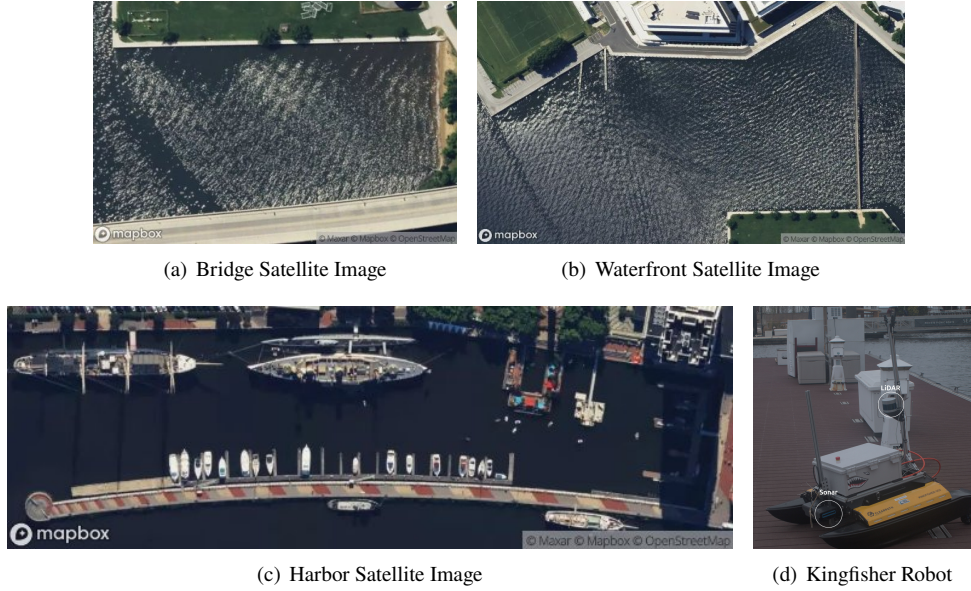


Fig. 4: **Test Environments and real robot.** (a) shows the bridge, (b) shows waterfront and (c) shows harbor. (d) shows our kingfisher robot with sonar and LiDAR used to collect data for the experiments in this work. Satellite images are from Mapbox [41].

4) *Point Cloud Registration*: Given a loop closure candidate, a pair of indexes, the  $i$ th AUV step, and the  $j$ th USV step, we apply point cloud registration between the two. First, we gather a 2D sonar submap, denoted as  $\mathcal{S}_s$ , which spans from time step  $i$  to  $i - w_s$ . Where  $w_s$  is the user specified sonar window size. The submap is formed as,

$$\mathcal{S}_s = \{\mathbf{T}_{i-0}\mathcal{P}_{s,i-0}, \mathbf{T}_{i-1}\mathcal{P}_{s,i-1}, \dots, \mathbf{T}_{i-w_s}\mathcal{P}_{s,i-w_s}\}, \quad (8)$$

by applying the rigid body transform,  $\mathbf{T}_{i-w_s}$ , which denotes the transformation between step  $i$ , and preceding steps inside the window. We apply the same submap process to the 2D LiDAR point clouds as  $\mathcal{S}_L$ , denoted as

$$\mathcal{S}_L = \{\mathbf{T}_{j-0}\mathcal{P}_{L,j-0}, \mathbf{T}_{j-1}\mathcal{P}_{L,j-1}, \dots, \mathbf{T}_{j-w_s}\mathcal{P}_{L,j-w_s}\}, \quad (9)$$

Each submap is voxel down-sampled to the same resolution cell size,  $\mathcal{V}$ . Each voxel down-sampled cloud is denoted as  $\hat{\mathcal{S}}_s$  and  $\hat{\mathcal{S}}_L$ . We then use iterative closest point (ICP) [18] to align the two submaps, specifically Go-ICP [39]. Go-ICP is used to solve for the unknown 3 DoF translation  $T$  and rotation,  $R$  between the two submaps with *no initial guess*. Standard ICP refines Go-ICP results as a secondary registration step.

$$T = \begin{pmatrix} \Delta x & \Delta y \end{pmatrix}^\top, R = \begin{pmatrix} \cos \Delta\theta & -\sin \Delta\theta \\ \sin \Delta\theta & \cos \Delta\theta \end{pmatrix} \quad (10)$$

Go-ICP is used, due to the poor initial overlap between submaps, which is a failure case for standard ICP [18]. Once the submaps are aligned, we compute a registration quality metric, overlap, denoted as  $O_s$ .  $O_s$  is the percentage of points in  $\hat{\mathcal{S}}_s$  that have a nearest neighbor within 1-meter. This simple metric is used to check for outliers, by only retaining loop closures with an  $O_s$  greater than  $O_{min}$ .

5) *Robust Outlier Rejection*: If an inter-robot loop closure has an  $O_s \geq O_{min}$ , then it is further verified by a robust outlier rejection system, Pairwise Consistent Measurement Set Maximization algorithm (PCM) [19]. PCM analyzes loop closures as a group, performing geometric verification and

retaining only the largest set of consistent inter-robot loop closures.

$$\left\| \left( \mathbf{T}_{L_l}^{L_j} \cdot \mathbf{T}_{L_j}^{S_i} \cdot \mathbf{T}_{S_i}^{S_k} \right) \cdot \mathbf{T}_{L_l}^{S_k-1} \right\|_2^2 < \epsilon. \quad (11)$$

A pair of loop closures is consistent if they satisfy the condition in Equation 11. Place recognitions between LiDAR and sonar are denoted as  $\mathbf{T}_{L_j}^{S_i}$  and  $\mathbf{T}_{L_l}^{S_k}$ .  $i, j, k, l$  are the discrete time steps in each data form, with each represented by LiDAR,  $L$ , and sonar,  $s$ . The transformation between each data type are denoted by  $\mathbf{T}_{L_l}^{L_j}$  and  $\mathbf{T}_{S_i}^{S_k}$ .

#### E. Centralized Multi-robot Pose Graph

The centralized multi-robot factor graph is formulated on-board the USV. The USV's pose chain from Section IV-A3 is added as sequential factors between poses, denoted as  $\mathbf{f}^{LO}$ . Next, each AUV's pose chain from Section IV-B1 is added as sequential factors between poses, denoted as  $\mathbf{f}^{DO}$ . Lastly, each inter-robot loop closure is added as  $\mathbf{f}^{IR}$ , completing the pose graph in Equation 12.

$$\mathbf{f}(\Theta) = \mathbf{f}^0(\Theta_0) \prod_i \mathbf{f}_i^{LO}(\Theta_i) \prod_j \mathbf{f}_j^{DO}(\Theta_j) \prod_q \mathbf{f}_q^{IR}(\Theta_q) \quad (12)$$

This graph solves all robot poses jointly, yielding a centralized understanding of the environment. Pose information provides enhanced AUV state estimates and maps of the environment, both underwater and above-water, in a common reference frame. Lastly, these position estimates serve as an initial guess for all subsequent applications of point cloud registration in Section IV-D4.

## V. EXPERIMENTS

### A. Hardware Overview

In this work, we use the heavily modified ClearPath Kingfisher shown in Figure 4(d) to collect datasets for validation. Kingfisher is equipped with a custom flight control system based on the BlueCube ecosystem. For inertial navigation, the Kingfisher uses a Nortek DVL-1000, Vectornav VN100 IMU, and Bluerobotics Bar02 depth sensor. On the perception side,

**IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.**

we utilize a VLP-32C, a 32-beam above-water LiDAR, and, mounted on an underwater pole, a Blueprint Subsea Oculus M750d multi-beam imaging sonar. For dataset ground-truthing, we use the Here3+ RTK-GPS system.

### B. Datasets

When collecting datasets, we operate our Kingfisher in two different modalities. Option one is to collect USV data: above-water LiDAR and IMU. Option two is to collect AUV data: IMU, DVL, and underwater imaging sonar. This method enables us to collect multiple datasets with RTK-GPS ground truth and replay them concurrently with simulated communications to validate the proposed multi-robot system using real-world perceptual data.

We consider three real-world environments, each with a multi-robot team of a single USV and two AUVs. Each environment includes robots with different starting locations and trajectories. **Bridge** is shown in Figure 4(a), consisting of several large seawalls and several sizable bridge pylons. **Waterfront** is shown in Figure 4(b), consisting of seawalls, several piers, and a large footbridge supported by small-diameter pier pilings. **Harbor** is shown in Figure 4(c), consisting of seawalls, piers, floating docks, pier pilings, and three ships. Additionally, all environments feature a small number of moving objects, including pedestrians, cars, and boats, as well as substantial magnetic interference from steel-hulled ships and underwater high-voltage power lines. These three environments represent different, complex littoral settings to validate this work.

### C. Metrics

We use two primary metrics to evaluate results: error and time. For SLAM error, we compare estimated trajectories to RTK-GPS data to develop mean-absolute error (MAE) and root-mean-squared error (RMSE) to capture both the average and size of error distributions. When considering time, we evaluate module runtime in milliseconds (ms).

### D. Multi-robot SLAM Ablation Study

In this section, we compare varying versions of the proposed system to a single-agent underwater SLAM system for robots with imaging sonar [17]. We consider the error of each AUV in the multi-robot system separately and compare it to the same dataset using a single-agent SLAM system.

In our ablation study, we vary the sonar window size and the addition of the PCM robust outlier rejection block. We use a fixed LiDAR window size of three, as the LiDAR sensor range of 200 meters is significantly larger than the sonar range of 30 meters. Due to their comparatively small sensor range, sonar submaps are prone to creating outliers because of their limited context. Therefore, we evaluate the following cases: **Case 1**: Small sonar window size, without PCM, **Case 2**: Small sonar window size, with PCM, **Case 3**: Large sonar window size, without PCM and **Case 4**: Large sonar window size, with PCM.

1) *Bridge Results*: In this section, we consider the bridge dataset with a small window size of 3 and a large window size of 32. Quantitative results are shown at the top of Table I, and qualitative results are shown in Figures 5(a) and 5(b). For case 1 of the ablation study (window size 3, no PCM), there is a high error for AUV 1. Case 2 (window size 3, with PCM) shows a failed SLAM result for AUV-2 due to outliers. Case 3 (window

Case	AUV ID	MAE (meters)	RMSE (meters)
Bridge Case 1	1	15.6	18.02
	2	2.12	2.33
Bridge Case 2	1	Fail	Fail
	2	5.25	5.97
Bridge Case 3	1	15.16	18.02
	2	<b>2.12</b>	<b>2.33</b>
Bridge Case 4	1	<b>2.80</b>	<b>3.37</b>
	2	2.13	2.39
Bridge Single	1	3.35	3.75
Robot SLAM	2	2.80	3.23
Waterfront Case 1	1	Fail	Fail
	2	Fail	Fail
Waterfront Case 2	1	23.95	28.64
	2	24.1	30.58
Waterfront Case 3	1	14.08	20.14
	2	<b>4.13</b>	<b>5.57</b>
Waterfront Case 4	1	<b>9.38</b>	<b>12.06</b>
	2	4.46	4.96
Waterfront Single	1	13.92	18.85
Robot SLAM	2	7.22	7.80
Harbor Case 1	1	Fail	Fail
	2	Fail	Fail
Harbor Case 2	1	Fail	Fail
	2	Fail	Fail
Harbor Case 3	1	<b>4.3</b>	<b>5.31</b>
	2	3.83	<b>4.43</b>
Harbor Case 4	1	5.44	7.02
	2	<b>3.7</b>	4.47
Harbor Single	1	5.92	7.44
Robot SLAM	2	8.92	11.44

TABLE I: **Real world multi-robot SLAM results.** MAE/RMSE is reported for each of the two AUVs in our three robot team. Bridge results are shown at the top, waterfront in the middle and harbor at the bottom. We report data for the ablation study, cases 1-4 and a baseline comparison to single robot SLAM with each AUV dataset.

size 32, no PCM) yields results similar to those of Case 1. Lastly, in case 4 (window size 32 with PCM), there is a low error for both AUVs in the system. The ablation study shows that using a large window and PCM for the bridge environment is well-founded, attenuating any outliers. Moreover, when comparing case 4 to the baseline, single-robot SLAM, case 4 shows improvements in both MAE and RMSE.

2) *Waterfront Results*: In this section, we consider the waterfront dataset with a small window size of 3 and a large window size of 50. Quantitative results are shown in the middle of Table I, and qualitative results are shown in Figure 5(c) and 5(d). Case 1 (window size 3, no PCM) exhibits failure for both AUVs, while Case 2 (window size 3, with PCM) displays significant errors. Cases 3 (window size 50, no PCM) and 4 (window size 50, with PCM) improve these errors. Again, using a large window, both with and without PCM, shows the best MAE/RMSE in the ablation study compared to single-robot SLAM. Any improvement without PCM is most likely because of PCM's extremely judicious nature. PCM can, on occasion, be overly strict when accepting loop closures, resulting in the false rejection of helpful information.

3) *Harbor Results*: In this section, we consider the harbor dataset with a small window size of 3 and a large window size of 30. Quantitative results are shown at the bottom of Table I, and qualitative results are shown in Figure 5(e) and 5(f). Case 1 (window size 3, no PCM) exhibits failure for both AUVs, while Case 2 (window size 3, with PCM) also reports failure.

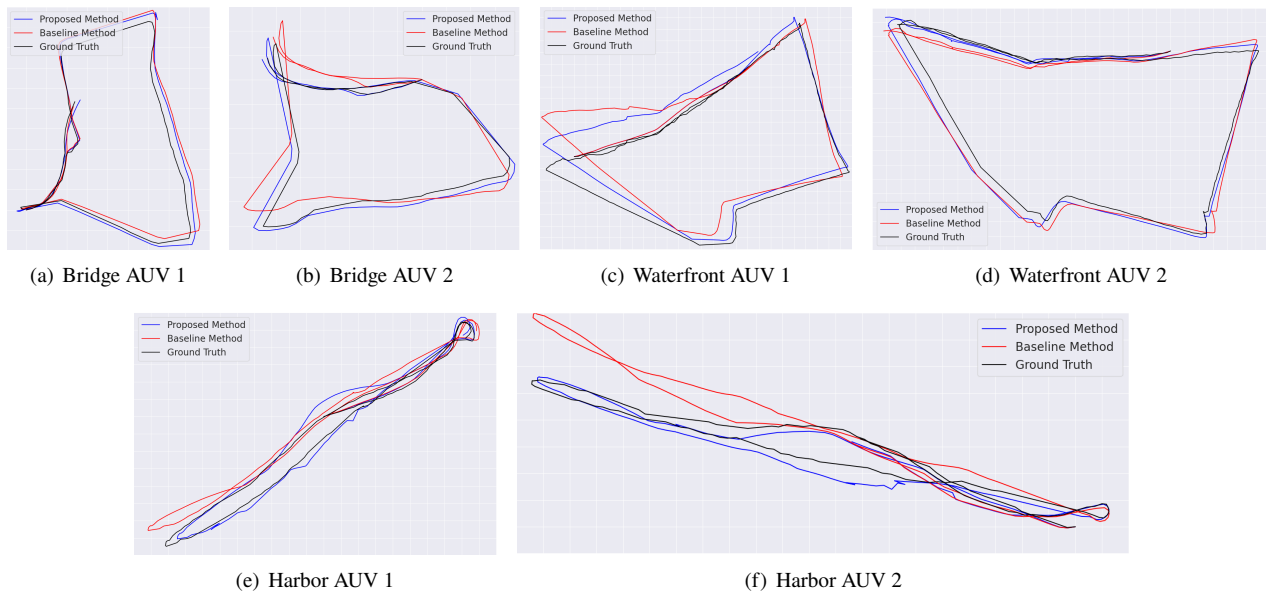


Fig. 5: **Qualitative Results.** Each plot shows a single AUV in the system. Black lines show the ground truth trajectory from RTK-GPS. The red lines indicate the baseline trajectory from single-robot underwater SLAM. Blue lines show the trajectory from the proposed method, centralized multi-robot SLAM. Results from bridge are shown in (a) and (b). Results from waterfront are shown in (c) and (d). Results from harbor are shown in (e) and (f). Grid lines are shown at a 10-meter cell size for all plots.

In case 3 (window size 35, no PCM), this error is attenuated. In case 4 (window size 35, with PCM), again, PCM is a stringent outlier rejection system, which can create higher MAE/RMSE. Similar to the waterfront, there are multiple meter improvements in MAE and RMSE compared to single robot SLAM.

4) *SLAM Results Summary:* For all three environments, there is an improvement when using a larger window size, with mixed results when adding PCM. However, PCM is an important addition due to its robustness against outliers, which is crucial for a SLAM mission to succeed. When comparing the proposed system to the single-robot SLAM baseline, significant improvements are observed in all three environments. Lastly, an example composite map, including sonar and LiDAR from different trajectories, is shown in Figure 1.

#### E. Perceptual Data Analysis

While many complexities are associated with using a system like this, we focus on data rate, noting COTS acoustic modems have rates as high as 62.5 k-bits/second [40]. We evaluate the size of perceptual messages in k-bits in Table II across all three environments. We compare raw sonar images, JPEG, pixels selected by CFAR in Section IV-C1, and the rectangle compression in Section IV-C2. First, we observe the costly nature of exchanging raw and JPEG images. Next, we

Data Type	Mean Size	STD Size	Min Size	Max Size
Raw Sonar Image	1978.37	1978.37	1978.37	1978.37
Jpeg Sonar Image	75.14	15.73	.07	124.0
CFAR Coordinates	14.28	15.5	.10	79.14
CFAR Compressed	5.63	4.76	.06	23.55
DRACo-SLAM [13]	1.12	1.11	.02	6.336

TABLE II: **Perceptual Message Sizes (k-bits).** The first row reports the cost to transmit a sonar image, the second row is the same, with jpeg compression applied. The third row is only transmitting CFAR detection coordinates, and the fourth row is with rectangle compression applied. The last row shows DRACo-SLAM compression [13].

compare to DRACo-SLAM [13] compression, which reduces data resolution from 0.05 to 0.25 meters, and as shown in [13], has a negative impact on SLAM error. Lastly, we observe the savings associated with CFAR down-sampling and the further savings with rectangle compression. Network statistics for bandwidth usage are provided as average kbits/second: Bridge 0.95 kbits/second, Waterfront 1.86 kbits/second, and Harbor 3.38 kbits/second. These results are encouraging. From a data rate perspective, perceptual data may be exchanged over acoustic communications.

#### F. Runtime Analysis

In this section, we consider the runtime of modules in the proposed system. Runtime results are summarized in Table III using an AMD Ryzen 9 9950X 16-Core Processor.

We report the mean and standard deviation of runtime for the point cloud registration system in Section IV-D4, PCM in Section IV-D5, and rectangle compression in Section IV-C2. Firstly, rectangle compression and PCM report low runtimes, with point cloud registration taking more time. It is essential to evaluate these runtime statistics in the context of how often each module is queried. Recall that the proposed method operates over a discrete set of keyframes. Given the vehicle speed for our datasets, these are added at a mean rate of 5469.53 ms with a standard deviation of 1920.61 ms, indicating the total average runtime of the modules in Table III is lower than the rate at which they are required.

Module	Mean Runtime (ms)	STD Runtime (ms)
PCM	6.05	5.23
Point Cloud Registration	766.11	254.60
Rectangle Compression	115.64	40.0

TABLE III: **Runtime statistics in ms.** Point cloud registration is from Section IV-D4, PCM is the outlier rejection system in Section IV-D5, and rectangle compression is the sonar data compression system in Section IV-C2.

**IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.**

## VI. CONCLUSIONS

In this work, we propose a multi-robot SLAM system for USVs equipped with 3D LiDAR and AUVs with imaging sonar. This centralized multi-robot system demonstrates improved AUV state estimation compared to single AUV SLAM. Moreover, this system enables multi-robot mapping of complex littoral environments, with structures above and below the water surface. We performed validation in three real-world environments, using a team of three robots consisting of a single mother-ship USV and two AUVs. When considering failure modes, robots must actually observe the same structures for loop closure to be possible. Additionally, perceptual aliasing is persistent, with many structures appearing identical in sonar imagery, necessitating a strict outlier rejection system. However, this can lead to numerous missed opportunities, resulting in higher underwater vehicle error.

This system represents an exciting development, marking the first multi-robot SLAM system to incorporate USVs and AUVs through indirect encounters. Future work will leverage real acoustic communication systems, both to validate the system in its current form in the field and to add acoustic range measurements between all robots in the system.

## REFERENCES

- [1] J. Vincent *et al.*, "Supervised Multi-Agent Autonomy for Cost-Effective Subsea Operations," *Offshore Technology Conference*, 2020.
- [2] L. Paull, G. Huang, M. Seto and J.J. Leonard, "Communication-constrained multi-AUV cooperative SLAM," *IEEE International Conference on Robotics and Automation*, 2015.
- [3] L. Paull, M. Seto and J.J. Leonard, "Decentralized cooperative trajectory estimation for autonomous underwater vehicles," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014.
- [4] J. McConnell, I. Collado-Gonzalez, P. Szenher and A. Shariati "Maritime Scene Matching for Inter-Robot Localization Across Surface and Underwater Domains," *SSRN Preprint*, 2025.
- [5] M. Machado Dos Santos, G. G. De Giacomo, P. L. J. Drews and S. S. C. Botelho, "Matching Color Aerial Images and Underwater Sonar Images Using Deep Learning for Underwater Localization," in *IEEE Robotics and Automation Letters*, 2022
- [6] M. M. D. Santos, G. G. De Giacomo, P. L. J. Drews-Jr and S. S. C. Botelho, "Cross-View and Cross-Domain Underwater Localization Based on Optical Aerial and Acoustic Underwater Images," in *IEEE Robotics and Automation Letters*, 2022
- [7] J. McConnell, F. Chen and B. Englot, "Overhead Image Factors for Underwater Sonar-Based SLAM," in *IEEE Robotics and Automation Letters*, 2022
- [8] G. Kim and A. Kim, "Scan Context: Egocentric spatial descriptor for place recognition within 3D point cloud map," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018.
- [9] H. Wang, C. Wang and L. Xie, "Intensity Scan Context: Coding Intensity and Geometry Relations for place recognition Detection," *IEEE International Conference on Robotics and Automation*, 2020.
- [10] G. Kim, S. Choi and A. Kim, "Scan Context++: Structural Place Recognition Robust to Rotation and Lateral Variations in Urban Environments," in *IEEE Transactions on Robotics*, 2021.
- [11] J. H. Friedman, J. L. Bentley, and R. A. Finkel. "An Algorithm for Finding Best Matches in Logarithmic Expected Time", *ACM Transactions on Mathematical Software*, 1977.
- [12] H. Kim, G. Kang, S. Jeong, S. Ma and Y. Cho, "Robust Imaging Sonar-based Place Recognition and Localization in Underwater Environments," *IEEE International Conference on Robotics and Automation*, 2023
- [13] 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," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2022
- [14] M.M. Santos, G.B. Zaffari, P. Ribeiro, P.L. J. Drews-Jr. and S.S.C. Botelho, "Underwater place recognition using forward-looking sonar images. A topological approach," in *Journal of Field Robotics*, 2018.
- [15] P.O.C.S. Ribeiro *et al.*, "Underwater Place Recognition in Unknown Environments with Triplet Based Acoustic Image Retrieval," *IEEE International Conference on Machine Learning and Applications*, 2018.
- [16] M. Larsson, N. Bore and J. Folkesson, "Latent Space Metric Learning For Sidescan Sonar Place Recognition," *IEEE/OES Autonomous Underwater Vehicles Symposium*, 2020.
- [17] J. Wang, F. Chen, Y. Huang, J. McConnell, T. Shan and B. Englot, "Virtual Maps for Autonomous Exploration of Cluttered Underwater Environments," in *IEEE Journal of Oceanic Engineering*, 2022.
- [18] P. J. Besl and N. D. McKay, "A method for registration of 3-D shapes," in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1992
- [19] J.G. Mangelson, D. Dominic, R.M. Eustice and R. Vasudevan, "Pairwise Consistent Measurement Set Maximization for Robust Multi-Robot Map Merging," *IEEE International Conference on Robotics and Automation*, 2018.
- [20] P. Agarwal, G. D. Tipaldi, L. Spinello, C. Stachniss and W. Burgard, "Robust map optimization using dynamic covariance scaling," *IEEE International Conference on Robotics and Automation*, 2013.
- [21] N. Sünderhauf and P. Protzel, "Switchable constraints for robust pose graph SLAM," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012.
- [22] M.V. Jakuba, C. N. Roman, H. Singh, C. Murphy, C. Kunz, C. Willis, T. Sato and R.A. Sato "Long-baseline acoustic navigation for under-ice autonomous underwater vehicle operations" *Journal of Field Robotics*, 2008.
- [23] N. H. Kussat, C. D. Chadwell and R. Zimmerman, "Absolute positioning of an autonomous underwater vehicle using GPS and acoustic measurements," in *IEEE Journal of Oceanic Engineering*, 2005.
- [24] N. R. Rypkema, E. M. Fischell and H. Schmidt, "One-way travel-time inverted ultra-short baseline localization for low-cost autonomous underwater vehicles," *IEEE International Conference on Robotics and Automation*, 2017.
- [25] Y. Wang *et al.*, "Passive Inverted Ultra-Short Baseline Positioning for a Disc-Shaped Autonomous Underwater Vehicle: Design and Field Experiments," in *IEEE Robotics and Automation Letters*, 2022.
- [26] E. Fischell, T. Schneider and H. Schmidt, "Design, Implementation, and Characterization of Precision Timing for Bistatic Acoustic Data Acquisition," in *IEEE Journal of Oceanic Engineering*, 2016.
- [27] E.M. Fischell, N.R. Rypkema and H. Schmidt, "Relative Autonomy and Navigation for Command and Control of Low-Cost Autonomous Underwater Vehicles," in *IEEE Robotics and Automation Letters*, 2019.
- [28] Y. Yao, D. Xu and W. Yan, "Cooperative localization with communication delays for MAUVs," *IEEE International Conference on Intelligent Computing and Intelligent Systems*, 2009.
- [29] T. Halsted and M. Schwager, "Distributed multi-robot localization from acoustic pulses using Euclidean distance geometry," *International Symposium on Multi-Robot and Multi-Agent Systems*, 2017.
- [30] Y. Li, Y. Wang, W. Yu and X. Guan, "Multiple Autonomous Underwater Vehicle Cooperative Localization in Anchor-Free Environments," in *IEEE Journal of Oceanic Engineering*, 2019.
- [31] A. Bahr, J. Leonard, M. Fallon, Cooperative Localization for Autonomous Underwater Vehicles. in *The International Journal of Robotics Research*. 2009.
- [32] Y. Huang, T. Shan, F. Chen and B. Englot, "DiSCO-SLAM: Distributed Scan Context-Enabled Multi-Robot LiDAR SLAM With Two-Stage Global-Local Graph Optimization," in *IEEE Robotics and Automation Letters*, 2022.
- [33] I. D. Gerg and V. Monga, "A Learnable Image Compression Scheme for Synthetic Aperture Sonar Imagery," *OCEANS*, 2021.
- [34] I. Vizzo, T. Guadagnino, B. Mersch, L. Wiesmann, J. Behley and C. Stachniss, "KISS-ICP: In Defense of Point-to-Point ICP – Simple, Accurate, and Robust Registration If Done the Right Way," in *IEEE Robotics and Automation Letters*, 2023.
- [35] T. Guadagnino, B. Mersch, S. Gupta, I. Vizzo, G. Grisetti and C. Stachniss, "KISS-SLAM: A Simple, Robust, and Accurate 3D LiDAR SLAM System With Enhanced Generalization Capabilities," *arXiv preprint*, 2023.
- [36] B. A. McElhoe, "An Assessment of the Navigation and Course Corrections for a Manned Flyby of Mars or Venus," *IEEE Transactions on Aerospace and Electronic Systems*, 1966.
- [37] F. Dellaert and GTSAM Contributors, "borglab-gtsam," *Georgia Tech Borg Lab*, 2022.
- [38] M. Richards, *Fundamentals of Radar Signal Processing*, McGraw Hill, 2005.
- [39] J. Yang, H. Li and Y. Jia, "Go-ICP: Solving 3D Registration Efficiently and Globally Optimally," *IEEE International Conference on Computer Vision*, 2013
- [40] EvoLogics, <https://www.evologics.com/acoustic-modem/hs>
- [41] Mapbox, <https://mapbox.com>