

Opti-Acoustic Semantic SLAM with Unknown Objects in Underwater Environments

Kurran Singh¹, Jungseok Hong¹, Nicholas R. Rypkema² and John J. Leonard¹

Abstract—Despite recent advances in semantic Simultaneous Localization and Mapping (SLAM) for terrestrial and aerial applications, underwater semantic SLAM remains an open and largely unaddressed research problem due to the unique sensing modalities and the object classes found underwater. This paper presents an object-based semantic SLAM method for underwater environments that can identify, localize, classify, and map a wide variety of marine objects without *a priori* knowledge of the object classes present in the scene. The method performs unsupervised object segmentation and object-level feature aggregation, and then uses opti-acoustic sensor fusion for object localization. Probabilistic data association is used to determine observation to landmark correspondences. Given such correspondences, the method then jointly optimizes landmark and vehicle position estimates. Indoor and outdoor underwater datasets with a wide variety of objects and challenging acoustic and lighting conditions are collected for evaluation and made publicly available. Quantitative and qualitative results show the proposed method achieves reduced trajectory error compared to baseline methods, and is able to obtain comparable map accuracy to a baseline closed-set method that requires hand-labeled data of all objects in the scene.

SUPPLEMENTARY MATERIAL

Data is made available here: <https://kurransingh.github.io/underwater-semantic-slam/>

I. INTRODUCTION

Recent advances in state estimation and perception methods have significantly improved the capability of terrestrial and aerial robots to map environments using semantically meaningful objects [1], [2], [3]. Robust and reliable semantic mapping systems are crucial for robots to achieve a higher level of autonomy, which is critical for underwater robotics applications given their often long-term and communications-limited nature. Semantic maps can also serve as a compressed map representation for multi-vehicle scenarios [4], and for diver-AUV collaboration [5]. However, underwater robotics has not kept pace with terrestrial and aerial robotics in semantic mapping due to the unique challenges posed by underwater environments.

One of the primary hurdles in underwater semantic Simultaneous Localization and Mapping (SLAM) is the lack of GPS availability and the substantial attenuation of electromagnetic waves underwater. Moreover, underwater visibility

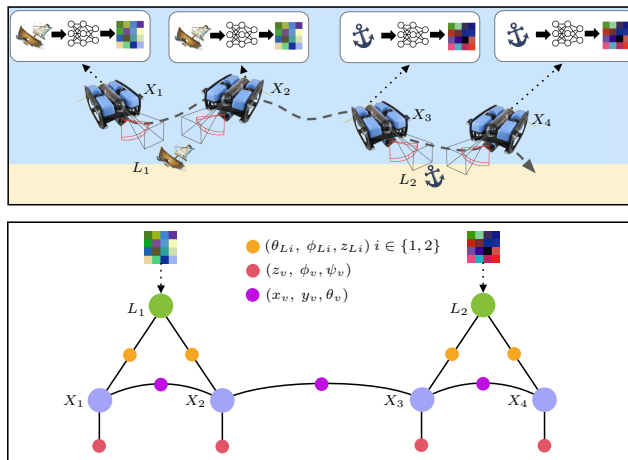


Fig. 1: An example scenario, where an underwater vehicle encounters objects during its navigation, demonstrating the proposed semantic SLAM approach. (Top): The vehicle senses objects using both optical and acoustic sensors. Optical sensing information is used to extract semantic feature embeddings of each object, and the ranges to the objects are obtained from the acoustic sensor. (Bottom): Information obtained from objects and odometry is used for building a factor graph (X_i : vehicle states; L_i : landmark states; $\theta_{L_i}, \phi_{L_i}, z_{L_i}$: bearing, elevation, and ranges to objects; z_v, ϕ_v, ψ_v : depth, pitch, and roll of the vehicle; x_v, y_v, θ_v : location and yaw of the vehicle).

is degraded by light scattering and color absorption, worsening perception challenges.

In an effort to tackle underwater SLAM, existing works primarily use sonar and camera sensors (II-B). However, sonar data is subject to multi-path reflections [6] and other noise-inducing phenomena; underwater camera data contains difficult lighting effects such as caustics [7]. The scarcity of large labeled datasets for both sonar and camera has made the problem more challenging. These limitations have restricted previous underwater semantic SLAM techniques to use a small number of object classes or be developed in simulation. Furthermore, the prohibitive cost of underwater motion capture systems has forced most existing semantic mapping methods to rely on manual measurements and image registration to estimate real-world ground truth trajectories, which are prone to error.

Addressing these challenges, we propose a novel approach to underwater semantic SLAM. This approach involves exploiting unsupervised object segmentation to identify objects and generate corresponding masks for accurate localization of the objects. Figure 1 illustrates how camera and sonar sensors are co-located for fusing each mask with sonar measurements to estimate the 3D position of the object. From segmented masks, feature embeddings of a fixed dimension are

¹K. Singh, J. Hong, and J. Leonard are with the Computer Science and Artificial Intelligence Laboratory (CSAIL) at the Massachusetts Institute of Technology (MIT), 32 Vassar St, Cambridge, MA 02139, USA. ² N. Rypkema is with the Applied Ocean Physics and Engineering Department at the Woods Hole Oceanographic Institute (WHOI), 266 Woods Hole Rd. Woods Hole, MA 02543, USA. Corresponding author: Kurran Singh (singhk@mit.edu)

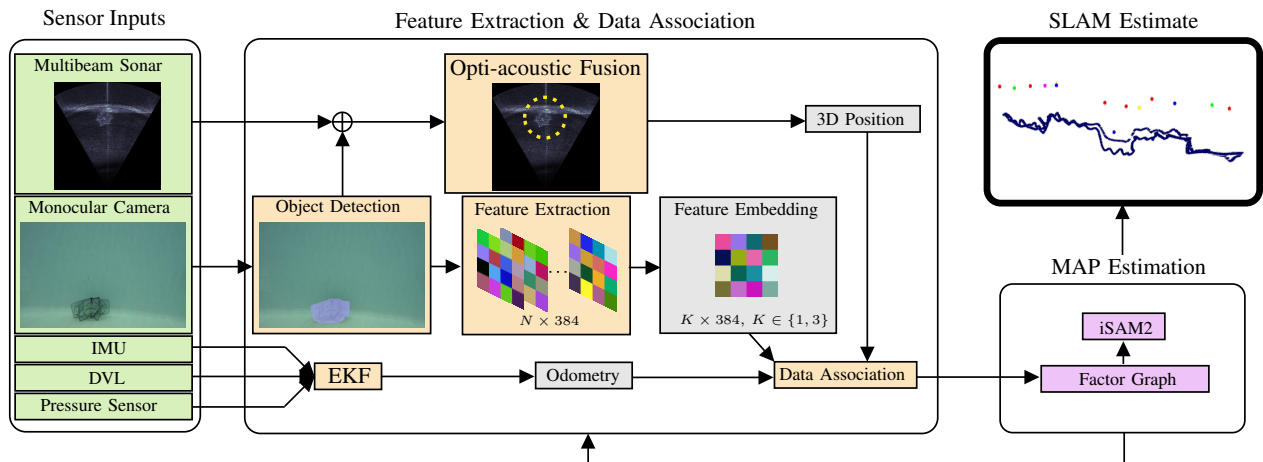


Fig. 2: The pipeline receives sensor inputs from a multibeam sonar, a monocular camera, an IMU, a DVL, and a pressure sensor. Given those sensor inputs, the system does the following: (1) Segment objects from optical images and extract features for each segmentation mask. The dimension of these extracted features is dependent on the corresponding object’s size in the image. These features are projected into a fixed dimension for the data association process. (2) Given the segmentation mask, its pixel centroid is used to find a corresponding range from sonar returns and estimate a 3D position of the object. (3) In parallel, sensor readings from IMU, DVL, and pressure sensor are used to estimate odometry. Outputs from (1)-(3) are used to build a factor graph which is optimized via iSAM2 to produce maximum a posteriori (MAP) estimates. Lastly, we obtain map and trajectory estimates as an output.

produced. Additionally, odometry is estimated by combining an inertial measurement unit (IMU), Doppler velocity log (DVL), and pressure sensor measurements with an Extended Kalman Filter (EKF). By incorporating these measurements into a factor graph framework, as shown in Figure 1, the proposed method detects semantically meaningful objects, uses them as landmarks, and outputs a semantic map, while also reducing the odometric drift of the trajectories. For evaluation, the data from a low-cost Hybrid Long/Inverted Ultra-Short Baseline (iLBL-USBL) acoustic beacon system is used to obtain a ground truth trajectory. This allows for the evaluation of the method quantitatively on both mapping and trajectory accuracy. Qualitative results from various underwater environment scenarios are also presented.

Results demonstrate that the proposed approach enables an underwater vehicle to map underwater environments accurately by identifying and locating semantically meaningful objects. Through this method, the challenges of underwater semantic mapping including the lack of training data for object recognition and the limited visibility for stereo vision are addressed, thereby enhancing autonomy for sophisticated underwater operations. The contributions of this work are as follows:

- A novel underwater semantic SLAM system capable of mapping a wide variety of object classes without additional training, and without any prior knowledge of the scene’s object classes.
- An opti-acoustic data fusion method that provides accurate object localization in 3D.
- Quantitative and qualitative evaluations of the system using real underwater data, which substantiates the map accuracy and trajectory optimization capabilities of the proposed approach.
- A public dataset consisting of camera, sonar, odometry, and ground truth trajectory data, as well as 691 labeled images for common underwater objects across 17 different classes.

II. RELATED WORK

This work ties together concepts from machine learning, sensor fusion, and underwater mapping. We provide a brief overview of relevant related works from the above areas.

A. Object Segmentation

Recent advances in the Transformer architecture [8] have inspired the development of Vision Transformers (ViTs) [9], which outperform convolutional neural networks (CNNs) for multiple computer vision tasks, including object segmentation, detection, and classification. However, understanding the visual representation of an image has remained extremely challenging for a model when the image was not previously seen by the model. To address this issue, Caron *et al.* proposed DINO [10]. It exploits unsupervised learning approaches (e.g. a teacher-student network) to distill information from large datasets without any labels using ViTs. Radford *et al.* [11] presented multi-modal learning, which uses image-text pairs with a contrastive loss. The model, CLIP, learns how to understand images in the context of natural languages. Kirillov *et al.* [12] proposed a promptable segmentation method segment anything (SAM). It uses a pre-trained encoder from CLIP and takes boxes, masks, and points as prompts. These models’ ability to extract contextual information from images has led to the development of several segmentation models. For instance, [13], [14], [15] proposed unsupervised segmentation models without using any explicit labels. [16], [17], [18] are the examples of SAM variants. CLIP- and SAM-based approaches achieve comparable segmentation accuracy to fully supervised models (e.g., [12]) for terrestrial images.

B. Opti-Acoustic Correspondences

Purely sonar-based underwater mapping has a relatively large body of work such as Teixeira *et al.* [6] which proposes a submap-based SLAM method for sonar mapping, and Westman *et al.* [19] which describes a method for dense




Fig. 3: Schematic representation of sensor placements of a camera and sonar. The diagram on the left illustrates the spatial configuration of a camera and sonar sensor with their respective coordinate frames, where C represents a camera frame. The camera’s field of view is depicted in blue, and the sonar’s conical beam is represented in red. The diagram on the right details the observed object’s orientation angles of θ , ϕ_o from the camera and ϕ_s from the sonar, depicting how angular measurements are derived.

sonar-based reconstruction. Purely optical camera-based underwater mapping has also been explored in depth by works such as [20] which performs visual SLAM with a bag-of-words saliency criteria for keyframe selection.

Unfortunately, each sensor modality has a degeneracy stemming from the loss of the elevation angle for sonar, and the loss of the range measurement for monocular cameras. Some works seek to address the degeneracy directly for sonar-only mapping by learning the elevation angle [21], through the fusion of orthogonally mounted sonars [22], or through carefully handling the degeneracy in the back-end [23]. Stereo or multi-view for optical camera-only mapping has been explored as well, but suffers from issues with lighting conditions, reflections, caustics, and featureless areas such as plain walls [7]. However, combining the two sensor modalities to address the degeneracies of each sensor modality is likely the best way to fully utilize their complimentary geometries.

Several previous works exist proposing different methods and use cases of opti-acoustic fusion, with Ferreira *et al.* [24] providing a survey of such opti-acoustic fusion methods. Negahdaripour [25] describes the epipolar geometry for a sonar and camera configuration, and also provides a method for opti-acoustic calibration and dense object reconstruction in [26], while Pecheux *et al.* [27] refine the opti-acoustic calibration process through feature matching which negates the need for a specific calibration structure.

Jang *et al.* [28] attempt to learn opti-acoustic feature correspondences through the use of a convolutional neural network, while Jang *et al.* [29] try to calculate opti-acoustic feature matching by applying a style-transfer method and through integration with SuperGlue for multi-session or multi-vehicle mapping. Lanthier *et al.* [7] combine a stereo camera with sonar and infrared sensors to perform mapping, but note that stereo cameras struggle with obtaining range measurements for underwater vision due to uniform and featureless regions, and difficult lighting effects such as reflections and caustics. Yang *et al.* [30] use sonar to provide an absolute scale to monocular odometry. Kim *et al.* [31] separately perform visual SLAM and sonar space carving in

two different mapping modules, before combining the two for dense reconstruction. Rahman *et al.* [32] combine sonar and monocular camera data in a SLAM framework they refer to as SVIN2, but their mapping technique uses point features rather than objects.

The only work that performs fusion at the *object* level as required by our method, rather than at the *feature* level, is by Raaj *et al.* [33], who use a monocular camera and sonar in a particle filter, albeit only for a single object.

C. Semantic SLAM

While terrestrial semantic SLAM has been thoroughly explored with applications in, for example, dynamic environments (Cui *et al.* [34]), monocular cameras (Civera *et al.* [35]), lidar (Chen *et al.* [36]), probabilistic data association (Bowman *et al.* [37]), and scene graphs (Rosinol *et al.* [1]), far fewer such works exist for underwater scenarios. Arain *et al.* [38] use a stereo camera to perform binary classifications on objects for obstacle avoidance. Himri *et al.* [39] and Vallicrosa *et al.* [40] rely on a database of a priori known objects, and use a laser scanner to detect objects from the database. They demonstrate their method on pipes, valves, elbows, and tees for underwater pipe structures. Similarly, Guerneve *et al.* [41] compare sensor inputs with pre-existing CAD models, albeit by using sonar rather than laser scanners. Machado *et al.* [42] and dos Santos *et al.* [43] detect five classes of objects from sonar data in a harbor environment through training a neural network. Finally, McConnell *et al.* [44] demonstrate a method of predictive sonar-based mapping using Bayesian inference, which is able to identify and thus better reconstruct previously seen underwater objects even from partial views.

Our method differs from the above methods in that it incorporates both camera and sonar to detect, classify, and localize *previously unseen* object classes. The method extends to any object class, and does not require extensive hand-labeled training data of which there is a very little for underwater scenes.

III. METHOD

The proposed approach consists of three primary components: Object segmentation and feature extraction; optical-acoustic fusion; and data association for mapping. Figure 2 shows a block diagram of the approach, illustrating how features, 3D points, and odometry are extracted from sensor inputs and used for data association to perform semantic mapping.

A. Object Segmentation and Feature Extraction

Models leveraging unsupervised learning techniques excel in extracting visual features from unseen images [45]. Consequently, we select an unsupervised segmentation model MaskCut [15] to develop our object segmentation module.

For each image, the segmentation module outputs binary masks $\mathcal{M}_{I,i}$ of detected objects i from the image \mathcal{O}_I and simultaneously provides the masks’ feature embeddings $\mathcal{F}_{I,i}$ extracted by the DINO backbone. The centroids

Algorithm 1 Find Acoustic Range from Image Bearing

```

1: procedure BEARING2RANGE( $\theta$ )
2:    $\theta_s^{\text{res}} = \text{SONAR\_FOV}/\text{NUMBER\_OF\_BEAMS}$ 
       $\triangleright$  Find angular resolution of Sonar
3:   for beam in beams do
4:     if  $\theta_s(\text{beam}) \geq \theta - \theta_s^{\text{res}}/2$  and
5:        $\theta_s(\text{beam}) \leq \theta + \theta_s^{\text{res}}/2$  then
6:         if  $\max(\text{Ping}(\theta_s(\text{beam}))) \geq \tau$  then
               $\triangleright \tau$ : THRESHOLD
7:            $Z = \max(\text{Ping}(\theta_s(\text{beam})))$ 
8:           return  $Z$ 
9:         else
10:          return FALSE
11:        end if
12:      end if
13:    end for
14: end procedure

```

$(u_{\mathcal{M},i}, v_{\mathcal{M},i})$ of each binary mask $\mathcal{M}_{I,i}$ on the image plane are extracted and used for opti-acoustic fusion, which is introduced in the following section (III-B). Next, $\mathcal{F}_{I,i}$ is projected into constant dimensions ($K \times 384$), where $K \in \{1, 3\}$. The original dimension of $\mathcal{F}_{I,i}$ is $(N \times 384)$, with N representing the number of patches for the corresponding binary mask $\mathcal{M}_{I,i}$. N , and consequently the dimension of $\mathcal{F}_{I,i}$, varies based on the size of the corresponding binary mask $\mathcal{M}_{I,i}$.

Patch Averaging (PA) and Uniform Manifold Approximation and Projection (UMAP) [46] are used to accommodate masks of varying sizes while ensuring the output features maintain a constant dimension, effectively condensing the information contained in the larger-size feature embeddings into compact representation. This standardization enables a consistent fusion framework irrespective of the underlying binary mask sizes.

$\mathcal{F}_{I,i}$ is projected into feature embeddings $\mathcal{F}_{I,i}^P$ with constant dimensions ($K \times 384$), where $K \in \{1, 3\}$ and $P \in \{\text{PA}, \text{UMAP}\}$.

- PA: Mean feature embeddings $\mathcal{F}_{I,i}^{\text{PA}}$ of $\mathcal{F}_{I,i}$ are computed by averaging features values across the dimension N , yielding the size of (1×384) .
- UMAP: UMAP is selected to make the feature embeddings more interpretable while preserving the global structure of $\mathcal{F}_{I,i}$ as much as possible. In this case, the size of $\mathcal{F}_{I,i}^{\text{UMAP}}$ is (3×384) .

B. Opti-acoustic Fusion for 3D Projection

$(u_{\mathcal{M},i}, v_{\mathcal{M},i})$ of $\mathcal{M}_{I,i}$, introduced in the previous section and camera parameters (c_x, c_y, f_x, f_y) are used to calculate object bearing θ and elevation ϕ_o as shown Eq 1 and 2.

$$\theta = \arctan\left(\frac{u_{\mathcal{M},i} - c_x}{f_x}\right) = \arctan\left(\frac{X}{Z}\right) \quad (1)$$

$$\phi_o = \arctan\left(\frac{v_{\mathcal{M},i} - c_y}{f_y}\right) = \arctan\left(\frac{Y}{Z}\right) \quad (2)$$

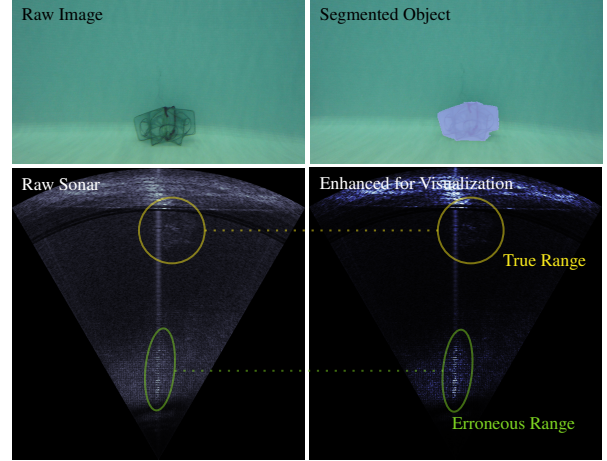


Fig. 4: (Top left): Input image of a lobster cage, (Top right): Detected object, (Bottom left): Corresponding sonar frame for object range estimation, (Bottom right): Enhanced frame for visualization. Sonar data can often be noisy and is subject to multi-path effects that make the detection and range extraction of objects challenging. Our method is largely robust to such issues through the use of a probabilistic data association method that takes into account the range uncertainty that is typical with the use of sonar.

Next, the acoustic beam that has the closest bearing value to θ is identified. Following this selection, the range, Z , along the beam to the target is retrieved from a sonar sensor (Figure 3). Algorithm 1 provides more details on finding Z from θ . After finding Z , X and Y are calculated using the second part of Eq 1 and 2. Through this procedure, the 3D position (X, Y, Z) of each detected object is calculated. As shown in Figure 4, the range extraction method is subject to noisy and erroneous sonar range measurements due to multipath reflections. Such errors are mitigated through the use of a robust data association method, which is described in the next section.

C. Semantic SLAM

The proposed method jointly optimizes both poses X and landmarks L given observations Z through the *maximum a posteriori* (MAP) inference problem:

$$\hat{X}, \hat{L} = \arg \max_{X \in SE(3), L \in \mathbb{R}^3} p(X, L | Z). \quad (3)$$

However, the method must also consider data association i.e. observation to landmark correspondences. Given the object class latent code as outlined in section III-A, and the object position as described in section III-B, as well as an odometry estimate (obtained as explained in Section IV), a keyframe and landmark observation is added to the factor graph. The landmark observation is classified as belonging to a previous landmark, or as the first observation of a new landmark. This classification is determined by first checking the cosine similarity of the object observation's latent space embedding as compared to existing landmarks' embeddings.

$$S_C(A, B) := \frac{A \cdot B}{\|A\|_2 \|B\|_2} \quad (4)$$

For each existing landmark that is determined to meet a heuristically chosen similarity threshold, the Mahalanobis



Fig. 5: (Top): The outdoor tank used for data collection exhibits challenging lighting conditions, including reflections (left) and caustics (right). (Bottom): Selected raw (uncalibrated) frames from the vehicle’s monocular camera. These images demonstrate the caustics (ripple lighting effect due to surface waves), reflections, and sudden lighting changes make processing underwater imagery particularly difficult compared to terrestrial scenes. Feature-based methods fail to maintain tracking due to the dynamic textures; our method uses objects instead and proves to be robust to such lighting effects.

distance is calculated between the observed landmark and the existing landmark following the method described by Kaess *et al.* [47]. The distance follows a six degree-of-freedom chi-square distribution, and thus we use the test below where l is the new landmark, l_j is the existing landmark, and Σ is the recovered covariance matrix.

$$\|l - l_j\|_{\Sigma}^2 < \chi_{6dof}^2 \quad (5)$$

If no landmarks meet both criteria above, then a new landmark is added. Otherwise, the max-likelihood hypothesis is selected given prior estimates X^0, L^0 .

$$\hat{D} = \arg \max_D p(D | X^0, L^0, Z) \quad (6)$$

where $D \triangleq \{d_k : d_k \in \mathbb{N}_{\leq M}, k = 1, \dots, K\}$ and M is the number of existing landmarks. Given the selected data association, the pose and landmark estimates are then updated:

$$\hat{X}, \hat{L} = \arg \max_{X \in SE(3), L \in \mathbb{R}^3} \log p(Z | X, L, \hat{D}). \quad (7)$$


By utilizing a factor graph framework and iSAM2 [48], it is possible to efficiently and accurately solve the above inference problems.

IV. DATASETS


Two different settings with several different configurations were used to collect datasets with a variety of sensors, lighting conditions, and objects to test the proposed method.

A. Indoor Pool

Data was collected in an indoor tank with a large variety of objects (Run 1-4 in Table I), but no available ground truth trajectory. The vehicle used for these experiments is a modified BlueROV2, which is described in detail in [49]. Odometry is obtained by fusing a Teledyne Impulse DVL, a VectorNav IMU, and a Bar30 pressure sensor using an Extended Kalman Filter. The camera is a calibrated Sony Exmor



(a) Run 3: 2 Laps



(b) Run 4: 4 Laps

Fig. 6: Comparison results of trajectory estimates, where the ‘O’ and ‘X’ markers denote the start and end points of each lap, respectively. The start and end points of the true trajectory are the same location, but drift causes the vehicle’s odometric estimate to deviate over time. (a) The two lap trajectory (Run 3) estimated by our patch averaging method (PA) can reduce drift errors from the odometry trajectory. (b) A similar comparison over four laps (Run 4), highlighting corrections made by our method against the odometry. Motion in the z-direction is negligible and thus is not shown.

IMX322 low-light camera with an 80-degree horizontal field-of-view and a 64-degree vertical field-of-view. The sonar is a Oculus m1200d multibeam operating at 2.1MHz, which provides a 2.5 mm range resolution, a 0.4 degree angular resolution, a 60 degree horizontal aperture, and a 12 degree vertical aperture. Due to the vertical aperture, camera data is needed to constrain the objects’ elevation angles.

B. Outdoor Tank

Data was also collected in an outdoor tank on a sunny day with a modified BlueROV2 that has an acoustic array for the iLBL-USBL measurements attached to the top of the vehicle. The vehicle also has a Bar30 pressure sensor for depth measurements, the same camera as in the indoor datasets, and an Oculus m750d multibeam sonar operating at 750kHz frequency with a 5 Hz update rate. In this mode, the sonar has a range resolution of 4 mm, an angular resolution of 1 degree, a horizontal field-of-view of 130 degrees, and a vertical aperture of 20 degrees. The relatively large vertical aperture, which cannot be resolved in the received data, means that the use of the camera for estimating the elevation angle of objects is crucial.

The acoustic beacons are used to provide odometry estimates, and are also free of the magnetic interference that prevented the use of magnetometer-based heading measurements. One-way travel time range measurements from each of the two beacons to the vehicle’s receiver array create a nonlinear least-squares optimization as described below to uniquely determine the vehicle position from overlapping range circles.




Fig. 7: The upper portion of the image illustrates an underwater vehicle navigating through an underwater environment where various objects are placed. The trajectory map below shows the estimated trajectory of the vehicle and the location of objects encountered during the mission. The colors are assigned based on semantic class groupings. Note that the tire and lobster cage are unexpectedly similar in the latent space and thus appear as the same class. Black lines represent correct matches, a red dashed line shows false positives due to an error from feature embedding, and a cyan dashed line is a misaligned object match in the map due to an error from range values from the sonar sensor.

TABLE I: Summary of Datasets

Run	Objects	Laps
1	1 small hand net, 1 large hand net, 2 sets of stairs, 2 recycling bins, 2 types of towfish, 2 trash bins, 2 paddles	2
2	1 model submarine, 1 green seaweed patch, 1 pipe, 1 towfish, 1 star-shaped lobster cage, 1 tire, 1 small hand net	2
3	2 round lobster cages, 1 green seaweed patch, 1 blue seaweed patch, 2 tires, 1 small hand net, 2 types of towfish, 2 pipes with anchors, 1 star-shaped lobster cage	2
4	2 round lobster cages, 1 green seaweed patch, 1 blue seaweed patch, 2 tires, 1 small hand net, 2 types of towfish, 2 pipes with anchors, 1 star-shaped lobster cage	4
5	1 fish cage, 1 ladder, 1 pipe, 1 T-slot extrusion structure	3

$$(\tilde{x}, \tilde{y}) = \min_{x, y} \sum_{k=1}^M [\sqrt{(x - x_k)^2 + (y - y_k)^2} - r_k]^2 \quad (8)$$

where (x_k, y_k) is the position of beacon k , and r_k is the range to beacon k .

Since magnetometer-based heading measurements are adversely affected by the magnetic interference of the tank walls and other nearby equipment, to get a more accurate ground truth trajectory, the vehicle heading is directly inferred from beamforming on the measurements from the vehicle's acoustic array by solving for vehicle heading θ_k :

$$\theta_k = [\phi_k - \arctan 2(\tilde{y} - y_k, \tilde{x} - x_k) - \frac{\pi}{2}] \mod 2\pi. \quad (9)$$

The vehicle position and heading are available at 1 Hz, and are directly used as the absolute (as opposed to relative as in the indoor datasets) odometry input for the semantic SLAM method. The ground truth trajectory is further refined through a pose-graph SLAM method that incorporates ICP-based alignments of the wall and structures. The use of an outdoor tank on a sunny day means that the data has several difficult lighting effects, as seen in Figure 5. The dataset consists of a single run (Run 5 in Table I) of three loops around the circular tank, which is roughly 64 square meters in area with an average depth of 4 meters. Note that while certain markings on the wall due to algae build-up

were detected by the open-set system, those objects are not included in the map accuracy results in Table II.

V. RESULTS

The proposed method is thoroughly tested on a variety of collected datasets, with map accuracy results shown in Table II, and trajectory accuracy along with an ablation study in Table III. Quantitative trajectory evaluations are available for the outdoor tank, but qualitative improvements of the trajectory can be seen in the indoor pool data as well in Figure 6. While the start and end point should be the same, the effects of odometric drift can be clearly seen as the end points of the odometry only trajectories are several meters from the start point. In contrast, our method is able to make loop closures through successfully re-identifying and associating objects as the vehicle returns to the start point, thus allowing the tightly-coupled optimization to correct the odometric drift.

The outdoor tank dataset demonstrates the failure of feature-based methods in such conditions, as the standard implementation for such methods, ORB-SLAM3 [50] is unable to complete more than 27% of the trajectory due to the dynamic movement of caustics throughout the trajectory. Furthermore, our method is able to improve the trajectory accuracy from the odometry through identifying previously seen objects. The results also emphasize the success of the use of the semantic encoding of each object (*e.g.*, Figure 7), as the ablation demonstrates that the trajectory is made worse without the assistance of the semantic latent codes due to incorrect data associations.

The proposed method also achieves strong map accuracy. PA outperforms UMAP consistently, which is surprising since patch averaging only allows for a 384 length encoding, whereas UMAP effectively has a length 1,152 encoding, which should allow for more expressiveness. However, this unexpected outcome is likely attributed to a fundamental discrepancy in the underlying assumptions about data distribution. UMAP's effectiveness is based on the premise

TABLE II: Map accuracy

Dataset	Method	Precision	Recall
Indoor, Run 1	Patch Averaging (ours)	0.82	0.82
	UMAP (ours)	0.58	0.78
	YOLOv8	1.0	1.0
Indoor, Run 2	Patch Averaging (ours)	0.41	1.0
	UMAP (ours)	0.38	0.71
	YOLOv8	1.0	0.75
Indoor, Run 3	Patch Averaging (ours)	0.86	1.0
	UMAP (ours)	0.56	0.82
	YOLOv8	1.0	0.83
Indoor, Run 4	Patch Averaging (ours)	0.73	0.92
	UMAP (ours)	0.48	0.92
	YOLOv8	1.0	0.83
Outdoor Tank, Run 5	Patch Averaging (ours)	0.75	0.60
	UMAP (ours)	0.66	0.75
	YOLOv8	1.0	1.0

that data is uniformly distributed across a manifold, an assumption that may not hold for the features extracted in our pipeline.

A baseline method based on fine-tuning YOLOv8 is also implemented. While off-the-shelf network weights do not identify any of the objects in the collected underwater datasets, 230 hand-labeled images including every object class in the datasets were used to fine-tune a YOLOv8n network over 25 epochs. The 230 images were augmented with rotation and blurring to produce a dataset of 691 training images. The labeled images are included as part of the dataset release. Object detections are encoded as one-hot vectors, with the rest of the system pipeline the same as our method. While the baseline method performs well, it has received training data of the exact objects in the same scene and lighting conditions as the test data, unlike our method which is able to identify and classify objects while receiving no labeled training data or prior knowledge of the scene.

It should be noted that the precision and recall metrics are at the *map* level, not observation level. Thus, each correct object identification requires the correct data association over all observations of the object, which can be over 100 individual observations. However, just one bad data association decision can cause the creation of a new object, which then adversely affects the precision figures. False positives from incorrect data associations occur for two different reasons: 1) very incorrect ranges due to multi-path effects in the sonar data as seen in Figure 4, and 2) difficulty in determining that an observation is of the same class as an existing landmark due to light changes (Figure 5) or viewing angle changes. False negatives occur solely due to the segmentation network, as it often is unable to detect the white pipe object as there is very little contrast with the background. Nevertheless, the overall map accuracy results demonstrate that our system is largely robust to such issues, and is able to achieve comparable results to a closed-set fine-tuned baseline system.

VI. CONCLUSION

A method that combines monocular camera data with sonar data for identifying, localizing, and mapping underwater objects is proposed. A robust data association method is implemented along with a SLAM system using a factor

TABLE III: Outdoor Dataset, Average Pose Error in meters. Note that the ORBSLAM3 Baseline method loses tracking and fails to complete trajectory. The error metric for ORBSLAM3 is shown for completed section only (27% of full trajectory).

Method	APE (m)
Odometry	0.58
Odom + Geometric	1.40
Odom + Geo + Uncertainty	0.60
Odom + Geo + Uncert + Semantics	0.56
Baseline: ORBSLAM3	2.28
Baseline: Closed-Set (YOLOv8)	0.59

graph representation to jointly optimize landmark and vehicle poses. A variety of datasets in challenging conditions unique to underwater scenarios with many different typical underwater objects were collected and open-sourced, and the method is evaluated on these datasets. The method performs well quantitatively and qualitatively on mapping accuracy as well as trajectory accuracy as compared to baseline methods.

ACKNOWLEDGMENTS

The authors would like to thank Thomas Jeongho Song, John Paul Morrison, and Ben Weiss for their assistance in data collection.

This work was supported by the MIT Lincoln Laboratory Autonomous Systems Line which is funded by the Under Secretary of Defense for Research and Engineering through Air Force Contract No. FA8702-15-D-0001, ONR grants N00014-18-1-2832, N00014-23-12164, and N00014-19-1-2571 (Neuroautonomy MURI), and the MIT Portugal Program.

REFERENCES

- [1] A. Rosinol, A. Violette, M. Abate, N. Hughes, Y. Chang, J. Shi, A. Gupta, and L. Carlone, “Kimera: From SLAM to Spatial Perception with 3D Dynamic Scene Graphs,” in *International Journal of Robotics Research*, vol. 40, 2021, pp. 1510–1546.
- [2] L. N. Nicholson, M. Milford, and N. Sünderhauf, “QuadricSLAM: Dual Quadrics as SLAM Landmarks,” in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2018, pp. 426–4261.
- [3] L. Schmid, M. Abate, Y. Chang, and L. Carlone, “Khronos: A Unified Approach for Spatio-Temporal Metric-Semantic SLAM in Dynamic Environments,” *arXiv preprint arXiv:2402.13817*, 2024.
- [4] S. Jamieson, K. Fathian, K. Khosoussi, J. P. How, and Y. Girdhar, “Multi-Robot Distributed Semantic Mapping in Unfamiliar Environments Through Online Matching of Learned Representations,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 8587–8593.
- [5] B. O’Neill, “Adaptive AUV-assisted Diver Navigation for Loosely-Coupled Teaming in Undersea Operations,” Ph.D. dissertation, Massachusetts Institute of Technology, 2023.
- [6] P. V. Teixeira, M. Kaess, F. S. Hover, and J. J. Leonard, “Underwater Inspection Using Sonar-based Volumetric Submaps,” *IEEE International Conference on Intelligent Robots and Systems*, vol. 2016-Novem, pp. 4288–4295, 2016.
- [7] M. A. Lanthier, D. Nussbaum, and A. Sheng, “Improving Vision-Based Maps by Using Sonar and Infrared Data,” *Proceedings of the IASTED International Conference on Robotics and Applications*, pp. 118–123, 2004.
- [8] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, “Attention is All You Need,” in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017.
- [9] A. Kolesnikov, A. Dosovitskiy, D. Weissenborn, G. Heigold, J. Uszkoreit, L. Beyer, M. Minderer, M. Dehghani, N. Houlsby, S. Gelly, T. Unterthiner, and X. Zhai, “An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale,” 2021.
- [10] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin, “Emerging Properties in Self-Supervised Vision Transformers,” in *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021.

- [11] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, *et al.*, “Learning Transferable Visual Models from Natural Language Supervision,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 8748–8763.
- [12] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, P. Dollár, and R. Girshick, “Segment Anything,” *arXiv:2304.02643*, 2023.
- [13] T.-W. Ke, J.-J. Hwang, Y. Guo, X. Wang, and S. X. Yu, “Unsupervised Hierarchical Semantic Segmentation with Multiview Cosegmentation and Clustering Transformers,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 2571–2581.
- [14] D. Niu, X. Wang, X. Han, L. Lian, R. Herzig, and T. Darrell, “Unsupervised Universal Image Segmentation,” 2023.
- [15] X. Wang, R. Girdhar, S. X. Yu, and I. Misra, “Cut and Learn for Unsupervised Object Detection and Instance Segmentation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 3124–3134.
- [16] T. Ren, S. Liu, A. Zeng, J. Lin, K. Li, H. Cao, J. Chen, X. Huang, Y. Chen, F. Yan, Z. Zeng, H. Zhang, F. Li, J. Yang, H. Li, Q. Jiang, and L. Zhang, “Grounded SAM: Assembling Open-World Models for Diverse Visual Tasks,” 2024.
- [17] F. Li, H. Zhang, P. Sun, X. Zou, S. Liu, J. Yang, C. Li, L. Zhang, and J. Gao, “Semantic-SAM: Segment and Recognize Anything at Any Granularity,” *arXiv preprint arXiv:2307.04767*, 2023.
- [18] L. Ke, M. Ye, M. Danelljan, Y. Liu, Y.-W. Tai, C.-K. Tang, and F. Yu, “Segment Anything in High Quality,” in *NeurIPS*, 2023.
- [19] E. Westman, I. Gkioulekas, and M. Kaess, “A Volumetric Albedo Framework for 3D Imaging Sonar Reconstruction,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 9645–9651.
- [20] A. Kim and R. M. Eustice, “Real-time Visual SLAM for Autonomous Underwater Hull Inspection Using Visual Saliency,” *IEEE Transactions on Robotics*, vol. 29, pp. 719–733, 2013.
- [21] Y. Wang, Y. Ji, D. Liu, H. Tsuchiya, A. Yamashita, and H. Asama, “Elevation Angle Estimation in 2D Acoustic Images Using Pseudo Front View,” *IEEE Robotics and Automation Letters*, vol. 6, pp. 1535–1542, 2021.
- [22] J. McConnell, J. D. Martin, and B. Englot, “Fusing Concurrent Orthogonal Wide-Aperture Sonar Images for Dense Underwater 3D Reconstruction,” *IEEE International Conference on Intelligent Robots and Systems*, pp. 1653–1660, 2020.
- [23] E. Westman and M. Kaess, “Degeneracy-Aware Imaging Sonar Simultaneous Localization and Mapping,” *IEEE Journal of Oceanic Engineering*, vol. 45, no. 4, pp. 1280–1294, 2019.
- [24] F. Ferreira, D. Machado, G. Ferri, S. Dugelay, and J. Potter, “Underwater Optical and Acoustic Imaging: A Time for Fusion? A Brief Overview of the State-of-the-Art,” *OCEANS 2016 MTS/IEEE Monterey, OCE 2016*, pp. 19–23, 2016.
- [25] S. Negahdaripour, “Epipolar Geometry of Opti-Acoustic Stereo Imaging,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, pp. 1776–1788, 2007.
- [26] S. Negahdaripour, H. Sekkati, and H. Pirsivash, “Opti-Acoustic Stereo Imaging: On System Calibration and 3-D Target Reconstruction,” *IEEE Transactions on Image Processing*, vol. 18, pp. 1203–1214, 2009.
- [27] N. Pecheux, V. Creuze, F. Comby, and O. Tempier, “Self Calibration of a Sonar-Vision System for Underwater Vehicles: A New Method and a Dataset,” *Sensors*, vol. 23, 2 2023.
- [28] H. Jang, Y. Lee, G. Kim, and A. Kim, “CNN-based Opti-Acoustic Transformation for Underwater Feature Matching,” *The Journal of Korea Robotics Society*, vol. 15, no. 1, pp. 1–7, 2020.
- [29] H. Jang, S. Yoon, and A. Kim, “Multi-session Underwater Pose-graph SLAM using Inter-session Opti-acoustic Two-view Factor,” vol. 2021-May. Institute of Electrical and Electronics Engineers Inc., 2021, pp. 11 668–11 674.
- [30] P. Yang, H. Liu, M. Roznere, and A. Q. Li, “Monocular Camera and Single-Beam Sonar-Based Underwater Collision-Free Navigation with Domain Randomization,” in *The International Symposium of Robotics Research*. Springer, 2022, pp. 85–101.
- [31] J. Kim, M. Lee, S. Song, B. Kim, and S. C. Yu, “3-D Reconstruction of Underwater Objects Using Image Sequences from Optical Camera and Imaging Sonar,” *OCEANS 2019 MTS/IEEE Seattle, OCEANS 2019*, 2019.
- [32] S. Rahman, A. Q. Li, and I. Rekleitis, “SVIn2: An Underwater SLAM System using Sonar, Visual, Inertial, and Depth Sensor,” *The International Journal of Robotics Research*, vol. 41, no. 11-12, pp. 1022–1042, 2022.
- [33] Y. Raaj, A. John, and T. Jin, “3D Object Localization Using Forward Looking Sonar (FLS) and Optical Camera via Particle Filter Based Calibration and Fusion,” *OCEANS 2016 MTS/IEEE Monterey, OCE 2016*, pp. 1–10, 2016.
- [34] L. Cui and C. Ma, “SOF-SLAM: A Semantic Visual SLAM for Dynamic Environments,” *IEEE access*, vol. 7, pp. 166 528–166 539, 2019.
- [35] J. Civera, D. Gálvez-López, L. Riazuelo, J. D. Tardós, and J. M. M. Montiel, “Towards Semantic SLAM Using a Monocular Camera,” in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2011, pp. 1277–1284.
- [36] X. Chen, A. Milioto, E. Palazzolo, P. Giguere, J. Behley, and C. Stachniss, “SUMA++: Efficient Lidar-Based Semantic SLAM,” in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 4530–4537.
- [37] S. L. Bowman, N. Atanasov, K. Daniilidis, and G. J. Pappas, “Probabilistic Data Association for Semantic SLAM,” in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 1722–1729.
- [38] B. Arain, C. McCool, P. Rigby, D. Cagara, and M. Dumbabin, “Improving Underwater Obstacle Detection using Semantic Image Segmentation,” in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 9271–9277.
- [39] K. Himri, P. Ridao, N. Gracias, A. Palomer, N. Palomeras, and R. Pi, “Semantic SLAM for an AUV Using Object Recognition from Point Clouds,” *IFAC-PapersOnLine*, vol. 51, no. 29, pp. 360–365, 2018, 11th IFAC Conference on Control Applications in Marine Systems, Robotics, and Vehicles CAMS 2018.
- [40] G. Vallicrosa, K. Himri, P. Ridao, and N. Gracias, “Semantic Mapping for Autonomous Subsea Intervention,” *Sensors*, vol. 21, no. 20, 2021.
- [41] T. Guerneve, K. Subr, and Y. Petillot, “Underwater 3D Structures as Semantic Landmarks in SONAR Mapping,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 614–619.
- [42] M. Machado, P. Drews, P. Núñez, and S. Botelho, “Semantic Mapping on Underwater Environment Using Sonar Data,” in *2016 XIII Latin American Robotics Symposium and IV Brazilian Robotics Symposium (LARS/SBR)*, 2016, pp. 245–250.
- [43] M. dos Santos, P. Drews, P. Núñez, and S. Botelho, “Object Recognition and Semantic Mapping for Underwater Vehicles Using Sonar Data,” *Journal of Intelligent & Robotic Systems*, vol. 91, pp. 279–289, 2018.
- [44] J. McConnell and B. Englot, “Predictive 3D Sonar Mapping of Underwater Environments via Object-specific Bayesian Inference,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 6761–6767.
- [45] L. Jing and Y. Tian, “Self-Supervised Visual Feature Learning With Deep Neural Networks: A Survey,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 11, pp. 4037–4058, 2021.
- [46] L. McInnes, J. Healy, and J. Melville, “UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction,” 2020.
- [47] M. Kaess and F. Dellaert, “Covariance Recovery from a Square Root Information Matrix for Data Association,” *Robotics and autonomous systems*, vol. 57, no. 12, pp. 1198–1210, 2009.
- [48] 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.
- [49] T. P. Osedach, K. Singh, P. V. Teixeira, J. B. Arber, A. Levesque, C. Chahbazian, D. Jain, B. Englot, J. J. Leonard, S. Vannuffelen, and S. Ossia, “Proxy Platform for Underwater Inspection, Maintenance, and Repair,” *2020 Global Oceans 2020: Singapore - U.S. Gulf Coast*, pp. 3–7, 2020.
- [50] C. Campos, R. Elvira, J. J. Gomez, J. M. M. Montiel, and J. D. Tardos, “ORB-SLAM3: An Accurate Open-Source Library for Visual, Visual-Inertial and Multi-Map SLAM,” *IEEE Transactions on Robotics*, vol. 37, no. 6, pp. 1874–1890, 2021.