

InsSo3D: Inertial Navigation System and 3D Sonar SLAM for turbid environment inspection

Simon Archieri^{1*}, Ahmet Cinar², Shu Pan¹, Jonatan Scharff Willners², Michele Grimaldi¹,
Ignacio Carlucho¹ and Yvan Petillot¹

Abstract— This paper presents InsSo3D, an accurate and efficient method for large-scale 3D Simultaneous Localisation and Mapping (SLAM) using a 3D Sonar and an Inertial Navigation System (INS). Unlike traditional sonar, which produces 2D images containing range and azimuth information but lacks elevation information, 3D Sonar produces a 3D point cloud, which therefore does not suffer from elevation ambiguity. We introduce a robust and modern SLAM framework adapted to the 3D Sonar data using INS as prior, detecting loop closure and performing pose graph optimisation. We evaluated InsSo3D performance inside a test tank with access to ground truth data and in an outdoor flooded quarry. Comparisons to reference trajectories and maps obtained from an underwater motion tracking system and visual Structure From Motion (SfM) demonstrate that InsSo3D efficiently corrects odometry drift. The average trajectory error is below 21cm during a 50-minute-long mission, producing a map of 10m by 20m with a 9cm average reconstruction error, enabling safe inspection of natural or artificial underwater structures even in murky water conditions.

SLAM, Marine Robotics, Mapping, Sonar

I. INTRODUCTION

Underwater 3D perception is an essential component of safe navigation in an unknown environment for Autonomous Underwater Vehicles (AUVs), and critical for many tasks from underwater mapping to offshore asset inspection. Most existing approaches use optical solutions with a single or multiple cameras [1] or based on lidar [2]. Although these solutions are effective, they are very dependent on water conditions and work only at short ranges, generally less than a few meters. Acoustic solutions, such as Sonar, on the other hand, are less dependent on water conditions and can operate at much larger ranges. However, their signal-to-noise ratio and spatial resolution are inferior to those of an optical-based system. Acoustic imaging is also subject to a wide variety of complex noise and multipaths, which makes automatic processing very challenging. Traditional sonar covers a fan-shaped Field of View (FOV) and produces 2D images representing a projection of the volume of water observed, where the elevation angle is lost. Resolving the elevation ambiguity is very challenging and is the main reason why most perception solutions using sonar have been limited to 2D [3]. However, innovations from sonar manufacturers have enabled the creation of a new generation of sonar which does not suffer from this elevation ambiguity; we refer to those sonars as 3D Sonar. Those new sonars generate a depth

¹School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh, UK

²Frontier Robotics, The National Robotarium, Edinburgh, UK

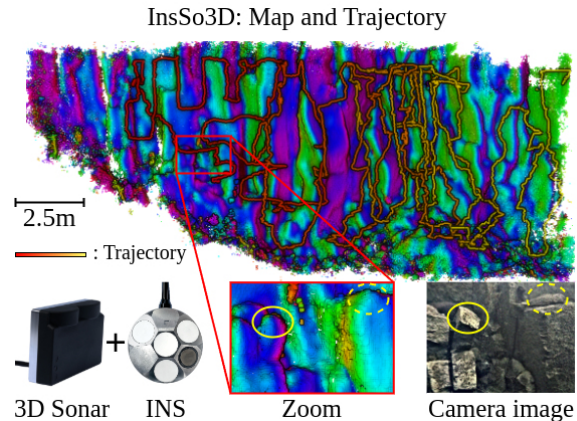


Fig. 1: InsSo3D uses the 3D Sonar and the INS to estimate the robot's trajectory and generate a 3D map of its environment. The trajectory colour indicates the acquisition timestamp of the 3D Sonar frame.

images akin to a lidar where each pixels are characterised by its range, azimuth and elevation angles. 3D Sonar are also affected by the different types of noise mentioned above and generally produces a sparse measurement, making automatic registration of 3D Sonar frames challenging. Fig. 2 shows a comparison of 3D Sonar and Camera data.

In this article, we address the challenge of Simultaneous Localisation and Mapping (SLAM) using a 3D Sonar as shown in Fig. 1. To achieve this, we rely solely on a 3D Sonar and an INS composed of: a Doppler Velocity Log (DVL), an Attitude and Heading Reference System (AHRS) and a pressure sensor.

We evaluated our method in real-world experiments by comparing its results to photogrammetry data obtained from a stereo camera and using a motion tracking system. The average trajectory error relative to the reference remains consistently below 21cm. Furthermore, our method achieves large-scale SLAM, reconstructing a $20m \times 10m$ environment with an average map error of 9cm compared to the reference map during our longest experiment. Those results demonstrate the effectiveness of the method for accurate SLAM, effectively correcting odometry drift and generating a large and consistent map. The key contributions enabling this performance are:

- Use of a 3D Sonar that enables robust, large-scale, and accurate SLAM of an unstructured environment.
- Integration of CFEAR [4] for robust, noise-resistant 3D




Fig. 2: Comparison between (a) Camera image, (b) 3D Sonar depth image, and (c) Camera depth image generated with RoMa [5] thanks to excellent visibility conditions. The red rectangle indicates the overlapping area between 3D Sonar and camera FOV, and the pink ellipse highlights a similar feature. The range colour scale is indicated on the right. Also, the camera FOV is 60° by 45° , and the 3D Sonar FOV is 90° by 40° .

frame and sub-map matching, loop closures detection and global alignment.

- Comprehensive validation in different real-world underwater environments with quantitative performance analysis.

II. RELATED WORK

Underwater 3D mapping and SLAM systems are fundamentally shaped by the sensing modality employed, as underwater environments impose unique challenges not present in terrestrial or aerial domains. These include limited visibility due to turbidity, light absorption with depth, dynamic water currents, and complex acoustic conditions such as multipath and scattering [6], [7]. Consequently, the choice of sensor affects not only the fidelity and scale of generated maps but also the feasibility of real-time localisation, long-term autonomy, and operational robustness. The two dominant classes of sensors in this domain are optical sensors, which leverage light to capture high-resolution visual or depth information, and acoustic sensors, which rely on sound propagation to probe the environment.

Optical sensors are widely used for underwater SLAM due to their ability to capture high-resolution textures and geometric features [8]. Monocular cameras provide lightweight, low-power solutions suitable for small autonomous underwater vehicles (AUVs), with SLAM techniques adapted from ORB-SLAM [1], [9] recovering scene structure through feature tracking and temporal geometry estimation. However, monocular methods suffer from scale ambiguity and rely on external cues, such as pressure sensors, to resolve metric scale. Their performance also degrades in visually challenging environments, such as turbid water or featureless seafloors [6]. Stereo camera systems attempt to overcome some of these limitations by triangulating depth from two calibrated cameras, enabling dense point-cloud mapping and

detailed 3D reconstruction. While successful in underwater inspection, stereo systems remain sensitive to turbidity, lighting variation, and backscatter, require precise calibration, and are limited in effective range. These constraints restrict their applicability for large-scale, deep, or highly dynamic underwater environments.

In contrast, acoustic sensors like sonars remain the backbone of underwater perception, operating effectively where optical systems fail, such as in turbid, low-light, or deep-sea environments. By exploiting long-range sound propagation, sonars enable robust mapping and navigation in visually degraded settings. Depending on type and configuration, acoustic sensors can support 2D SLAM and 3D reconstruction, typically at the cost of lower spatial resolution, slower updates, and higher processing demands. Traditional sonar systems cover a fan-shaped FOV but only provide 2D images with range and bearing information. This lack of elevation information is a central challenge in underwater perception with imaging sonar. Addressing this ambiguity is therefore critical for building an accurate 3D representation. Four main categories of SLAM using 2D sonar stand out. Early work has largely sidestepped elevation ambiguity by projecting the 2D sonar image onto a 2D horizontal plane [10], [11], [3], [12] or projecting the 2D detected features along the 3D elevation arc to help improve the 2D localisation [13], [14]. However, scan matching in 2D projections is fundamentally limited, often leading to significant errors in environments with any elevation variation, a condition almost always met in underwater applications. Second, classical methods address the elevation ambiguity problem primarily through geometric modelling and multi-view integration [15], [16], [17], [18], [19]. These approaches can provide accurate reconstructions when vehicle motion is well-constrained and navigation data is reliable. However, classical methods are highly sensitive to degenerate motion cases such as pure translation or pure rotation, where elevation ambiguity cannot be resolved [17]. This requirement for well-constrained motion makes piloting and trajectory planning more challenging in practice. A promising solution to the elevation ambiguity is the use of multiple orthogonally mounted imaging sonars, enabling direct triangulation of reflective targets in 3D, analogous to stereo vision [20], [21], [22], overcoming the fundamental limitation of single-sonar systems. However, the practical deployment of orthogonal sonar configurations faces several challenges. The overlapping field of view between sensors is often limited, resulting in sparse coverage and incomplete reconstructions. Additionally, the increased size, complexity, and cost associated with multiple sonars may be prohibitive. Finally, learning-based approaches are also used for solving the elevation ambiguity [23], [24], leveraging neural networks to extract 3D information directly from acoustic images. However, such methods depend on large labelled datasets, which are costly to collect underwater, limiting scalability and generalisation. A major leap in sonar 3D reconstruction has come from neural implicit and neural rendering methods [25], [26] inspired by Neural Radiance Fields (NeRF) [27]. However, most neural implicit




Fig. 3: Algorithm overview: new 3D Sonar frames are registered to the current sub-map using frame-to-frame and frame-to-sub-map registration. Then, a graph optimisation within the sub-map is performed to fuse the two registrations and estimate the correct frame pose. Finally, the new frame is integrated into the sub-map TSDF. When the submap completion criterion is met, the submap is closed and sent to the backend. The backend will register it to the previous one and try to detect loop closures with older sub-maps contained in the sub-map database (DB) before including the new sub-map in the factor graph. Finally, a TSDF representation of the global map is generated using each sub-map TSDF representation and the sub-map poses optimised by the factor graph.

approaches are computationally intensive, typically requiring minutes to hours for reconstruction, and are suited for offline processing. Moreover, they remain sensitive to pose accuracy; uncorrected navigation drift can degrade reconstruction quality. Although traditional sonar SLAM has demonstrated strong performance, extending it to full 6DOF estimation remains a significant challenge as it usually requires computationally intensive methods and highly constrained motion to resolve scene geometry, largely due to elevation ambiguity. In contrast, 3D sonar SLAM inherently supports 6DOF pose estimation, as it is not affected by elevation ambiguity, enabling more robust and flexible performance in complex underwater environments. Up to now, very few 3D sonars have been available, and they have all been expensive, bulky, and power-hungry (such as the echoscope), making them unsuitable for 3D SLAM onboard small platforms. Their use in SLAM applications has been limited—primarily due to prohibitive costs. Consequently, the few existing works on 3D sonar SLAM have focused mainly on frame-to-frame matching [28], [29]. These approaches do not incorporate loop closure detection or sub-maps, which are essential for global consistency and long-term mapping. However, recent innovations have significantly reduced the cost of 3D sonar systems, paving the way for more accessible and scalable SLAM implementations, including the development presented in this work.

To overcome these limitations, we propose a SLAM system based on the *Waterlinked Sonar3D-15*, a small and low-power volumetric acoustic sensor capturing azimuth, elevation, and range in a single ping. Unlike conventional approaches, it produces depth images in real time without post-processing, as can be seen in Fig. 2. This direct sensing enables robust, low-latency 3D mapping and navigation in dark, turbid, or turbulent conditions, offering a compact, high-reliability solution for operationally constrained underwater environments.

III. METHOD

In this section, we present the key components of the methodology. As shown in Fig. 3, the algorithm is split into two parts. First, the Sonar 3D frames are accumulated into

a sub-map by using a frame-to-sub-map and frame-to-frame registration. When the sub-map completion criterion is met, the new sub-map is closed and registered to those captured previously to create a consistent global map. We will first present the point cloud registration algorithm that uses the odometry as prior. This algorithm is used for frame-to-frame, frame-to-sub-map and sub-map to sub-map registration. Second, we will introduce the frontend algorithm responsible for sub-map creation and finally, how the backend registers and merges new sub-maps with older ones in order to build a consistent global map using the Truncated signed distance function (TSDF) approach.

A. Notations

Before we start with the description of our method, we outline the notation used in this section to ensure clarity for readers. The odometry frame is depicted as \mathcal{F}_o in NED coordinates, while \mathcal{F}_f represents the position of the robot in the frontend frame after correction by the frontend but before backend correction. The world frame representing the robot position after correction by the frontend and the backend is depicted as \mathcal{F}_w . Thus a 3D Sonar frame f is characterised by its pose in the odometry frame \mathcal{F}_o and in the frontend frame \mathcal{F}_f , denoted as ${}^o\mathbf{F}(f)$ and ${}^f\mathbf{F}(f)$ respectively. Similarly, a sub-map s is characterised by its pose in the odometry frame \mathcal{F}_o and in the corrected world frame \mathcal{F}_w , denoted as ${}^o\mathbf{P}(s)$ and ${}^w\mathbf{P}_n(s)$ respectively. Note that each pose ${}^w\mathbf{P}_n(s)$ of a specific sub-map in \mathcal{F}_w is kept after each graph optimisation and represents the pose of the sub-map after optimisation number n . This way, we keep a history of the different sub-maps pose to adjust the global map TSDF if needed after graph optimisation. In practice, we optimise the graph after acquiring a new sub-map, and therefore n corresponds to the total number of sub-maps.

B. Point cloud registration

A 3D variant of *CFEAR* [4] is used to register two point clouds together. *CFEAR* generates a sparse representation for both the source and target point clouds. This representation aims to model the scene geometries as a set of oriented surface points that are stable in at least one direction. First, as




Fig. 4: Process of computing the *CFEAR* oriented surface point cloud. (a) The point cloud is inserted into a voxel grid of voxel size r . (b) The mean and covariance are computed for each voxel using points within a distance r of the voxel centroid. (c) Final representation.

shown in Fig. 4(a), the point cloud is inserted into a voxel grid of resolution r . For each voxel u , all points within a radius r from the voxel centroid are used to compute the sample mean μ_u and covariance Σ_u of the voxel as shown in Fig. 4(b) and 4(c). By applying the eigen decomposition of Σ_u , the surface normal n_u can be estimated using the eigenvector that corresponds to the smallest eigenvalue. The oriented surface points are therefore defined by (μ_u, n_u) for each voxel u that has more than 6 points within a sphere of radius r .

Unlike *CFEAR*, we use *fast_gcip* [30], an implementation of Generalized ICP [31] to optimise the registration transformation between the source and target oriented surface point clouds. Generalized ICP uses a plane-to-plane distance metrics, which make it suitable for oriented surface point-cloud registration. The maximal correspondence distance of GICP is set to r , the intuition is that for aligned scans, nearby surface points are expected to lie within this distance.

C. Frontend: Sub-map Creation

1) *Registration*: A new incoming 3D Sonar frame i is projected into \mathcal{F}_f using the relative motion from the odometry between frames i and $i - 1$ using the following equation:

$${}^f\mathbf{F}(i)^* = \begin{cases} {}^f\mathbf{F}(i-1) \circ \mathbf{F}(i-1)^{-1} \circ \mathbf{F}(i), & \text{if } i > 0. \\ \mathbf{F}(i), & \text{otherwise.} \end{cases} \quad (1)$$

In order to increase robustness (see Fig. 3), *InsSo3D* performs two types of registration within the frontend. First, a frame-to-frame registration is performed between frames i and $i - 1$. Second, the new frame i is registered either to the TSDF point cloud representation of the sub-map or to the first frame i^* of the sub-map, depending on which one has the highest number of points. After performing both registrations, the new frame and its registrations are added to the frontend factor graph as shown in Figure 5. Frame-to-frame registrations are included using a sequential between-factor between frame i and $i - 1$. Frame-to-sub-map registrations are included using a non-sequential between-factor connecting the new frame i to the first sub-map frame i^* . The pose ${}^f\mathcal{F}(i)$ is therefore estimated during graph optimisation.




Fig. 5: Factor graph representation of the **frontend**, **red** nodes are the frame’s pose, **blue** nodes are sequential frame-to-frame factors, and the **yellow** nodes are frame-to-sub-map factors. Finally, the **green** node is a prior factor.

2) *Sub-map TSDF*: After graph optimisation the new 3D Sonar frames are integrated into the sub-map TSDF using the space carving variation introduced in [32]. Compared to normal TSDF, where only voxels inside the truncation region are updated when integrating a point, TSDF space-carving updates all voxels along the integration ray. With TSDF space-carving, each voxel has a maximum absolute value equal to the truncation threshold. This variation allows for keeping track of regions that have been observed and considered free. We generate a TSDF model for each sub-map independently before any backend correction so that they can be rigidly transformed if needed after backend graph optimisation. Moreover, as the frontend registration needs a real-time point cloud representation of the TSDF, the use of a meshing algorithm like *Marching Cube* [33] is not convenient. Instead, a coarse point cloud representation is generated by converting to a point each voxel that has an absolute signed distance below half of the voxel size [14].

3) *Sub-map completion criterion*: Sub-maps are considered to be complete when the percentage of overlap between the last and first 3D Sonar frame is below a threshold. The intuition is to keep sub-maps small to limit the pose drift inside the sub-map, as no loop closures are being detected within the sub-map. In addition, the sub-map TSDF point cloud representation must contain more points than the first sub-map frame to be considered complete. This avoids closing sub-maps that are too small or in the event of a wrong frame-to-sub-map registration due to insufficient points in the registration target. In this event, the frontend can recover from the bad registration, as we also use frame-to-frame registration.

The pose of a sub-map s in \mathcal{F}_o is defined by the pose of the first 3D Sonar frame i^* contained in the sub-map as ${}^o\mathbf{P}(s) = {}^o\mathbf{F}(i^*)$.

D. Backend: Global Map Generation

1) *Registration*: A pose-graph framework is used within the backend to optimise the sub-map pose [34]. When a new sub-map k is generated it is first projected into \mathcal{F}_w using the relative motion from the odometry between sub-map k and $k - 1$ using the following equation:

$${}^w\mathbf{P}(k)^* = \begin{cases} {}^w\mathbf{P}_{k-1}(k-1) \circ \mathbf{P}(k-1)^{-1} \circ \mathbf{P}(k), & \text{if } k > 0. \\ \mathbf{P}(k), & \text{otherwise.} \end{cases} \quad (2)$$

A sequential sub-map to sub-map registration is then performed between sub-map k and $k - 1$. Therefore the sub-map k position in \mathcal{F}_w is refined as ${}^w\mathbf{P}(k)^{**}$ using the following




Fig. 6: Factor graph representation of the **backend**, **red** nodes are the sub-map pose, **blue** nodes are sequential sub-maps to sub-map factors, and the **yellow** node is a loop closure factor. Finally, the **green** node is a prior factor.

equation:

$${}^w\mathbf{P}(k)^{**} = {}^{k-1}\mathbf{T}_{\text{ICP}_k} {}^w\mathbf{P}(k)^*. \quad (3)$$

Where ${}^{k-1}\mathbf{T}_{\text{ICP}_k}$ is the estimated transformation obtained using the registration algorithm. This refined pose is then used to detect loop-closures as explained in the section III-D.2. Those calculated transformations are then inserted into a factor graph as shown in Fig. 6. The pose ${}^w\mathbf{P}_n(k)$ is therefore obtained after the graph optimisation.

2) *Loop closures*: When the robot revisits a previously explored part of its environment, a loop-closure constraint can be added to the factor graph. In order to detect those loop closures, when a new sub-map k is being registered, a set of sub-maps with potential loop closures $\mathcal{S}_{lc}(k)$ is identified. To do so, the new sub-map k point cloud representation is inserted into a voxel grid of resolution r . Then, every other sub-map point cloud representations are inserted into different voxel grids of the same resolution, and their overlap ratio with the new sub-map k voxel grid is calculated. The corresponding sub-map is added to $\mathcal{S}_{lc}(k)$ if the overlap ratio exceeds 50%. Afterwards, the sub-map k is registered with each sub-map $s \in \mathcal{S}_{lc}(k)$. If the final registration error between sub-map k and sub-map s is below a threshold, a loop closure will be considered during graph optimisation. This loop closure will use the transformation calculated during the registration.

3) *Global Map TSDF*: In order to adapt the global map accordingly to the backend graph optimisations, sub-maps are integrated into TSDF independently and then merged together. Each sub-map that has moved after graph optimisation is deleted from the global map and then re-added with a corrected pose. To achieve this, InsSo3D saves $\phi_n(k)$, the index at which a sub-map k is integrated into the global map, as graph optimisation n is made. In practice for real-time implementation, a sub-map k will be reprocessed only if the rotation and the translation between ${}^w\mathbf{P}_n(k)$ and ${}^w\mathbf{P}_{\phi_{n-1}(k)}(k)$ are above a threshold.

Let $D_n(\mathbf{x})$ and $W_n(\mathbf{x})$ be the SDF and weight function of the global map at graph optimisation n in \mathcal{F}_w . Furthermore, $d_k(\mathbf{x})$ and $w_k(\mathbf{x})$ are the equivalents for the sub-map k in \mathcal{F}_o , and $\mathcal{S}_m(n)$ is a set containing the sub-map whose pose moved during the graph optimisation step n . The updated global map SDF and weight function considering the new sub-map n and the moved sub-map in $\mathcal{S}_m(n)$ are expressed as follows:

$$D_n(\mathbf{x}) = \frac{W_{n-1}(\mathbf{x})D_{n-1}(\mathbf{x}) + \psi_{n,n}(\mathbf{x}) + U(\mathbf{x})}{W_n(\mathbf{x})}, \quad (4)$$

$$W_n(\mathbf{x}) = W_{n-1}(\mathbf{x}) + w_n(\theta_{n,n}(\mathbf{x})) + UW(\mathbf{x}). \quad (5)$$

$D_n(\mathbf{x})$ is the sum of three terms: the existing value $W_{n-1}(\mathbf{x})D_{n-1}(\mathbf{x})$, the new sub-map n , $\psi_{n,n}$, and the changes caused by the sub-maps whose pose have moved $U(\mathbf{x})$. Similarly, $W_n(\mathbf{x})$ is the sum of the three equivalent terms for the new weight calculation. $U(\mathbf{x})$ and $UW(\mathbf{x})$ are calculated for each sub-map $k \in \mathcal{S}_m$ by removing the sub-map using the pose ${}^wP_{\phi_{n-1}(k)}$ and adding it back using the pose ${}^wP_n(k)$, as in the following equations:

$$U(\mathbf{x}) = \sum_{k \in \mathcal{S}_m(n)} \psi_{n,k}(\mathbf{x}) - \psi_{\phi_{n-1}(k),k}(\mathbf{x}), \quad (6)$$

$$UW(\mathbf{x}) = \sum_{k \in \mathcal{S}_m(n)} w_k(\theta_{n,k}(\mathbf{x})) - w_k(\theta_{\phi_{n-1}(k),k}(\mathbf{x})). \quad (7)$$

Where $\psi_{q,k}(\mathbf{x})$ is the weighted signed distance in \mathcal{F}_w of the sub-map k using the pose at graph optimisation q . Finally, the function $\theta_{q,k}(\mathbf{x})$ transforms a point \mathbf{x} of a sub-map k from \mathcal{F}_w to \mathcal{F}_o using the pose at graph optimization q . This transformation is required because the TSDF representations of each sub-map are independently integrated in \mathcal{F}_o .

$$\psi_{q,k}(\mathbf{x}) = w_k(\theta_{q,k}(\mathbf{x})) \times d_k(\theta_{q,k}(\mathbf{x})), \quad (8)$$

$$\theta_{q,k}(\mathbf{x}) = {}^o\mathbf{P}(k) \times {}^w\mathbf{P}_q(k)^{-1}. \quad (9)$$

Finally, as the global map rendering is not needed in real-time, the *Marching Cube* algorithm [33] is used to generate a mesh representation of the global map TSDF.

IV. EVALUATION

InsSo3D have been tested using a *WaterLinked Sonar3D-15* sonar running at 1.2 MHz with a 90° by 40° FOV and a maximal operating range of 15m. It is important to note that our method is independent of the sensors used and would work effectively with other sources of depth images. The system is installed on a *BlueRov 2 Heavy* equipped with a stereo camera [35], [36] from *Frontier Robotics* and a *Nortek Nucleus 1000 DVL* that also includes a built-in AHRS and pressure sensor. The odometry is estimated using *robot_localisation* [37].

We evaluated InsSo3D in two different environments: i) in a water tank, ii) in a water-filled quarry. Fig. 7 shows the different environments used. Excellent visibility conditions during our trials have enabled us to generate a reference trajectory and 3D geometry of the scene using photogrammetry, which allows us to assess the performance of our



(a) Water filled quarry

(b) Water tank

Fig. 7: View of the different environments used for the evaluation.




Fig. 8: Map qualitative comparison for the *Quarry* experiments. InsSo3D results are on the first row, and SFM results are on the second row. The first column shows a front view of the maps and the second a top view. The colour represents the normal azimuth angle calculated using the x and y components of the normal vector.

method. The photogrammetry reconstruction was performed using the visual-based SfM software *Colmap* [38], and the scale was estimated based on the stereo camera’s baseline. Furthermore, the tank is equipped with a *Qualisys* motion capture system, which was also used to obtain trajectory references during experiments in the test tank.

A. Evaluation Scenarios

The first environment evaluated is a water-filled quarry. The experiment was conducted at a short operating distance and consisted of 147 sub-maps during a 47-minute mission over a travelled distance of 230m. The survey was conducted at short range in order to maintain good visibility conditions for the stereo camera and thus evaluate the quality of the map and trajectory using the SFM results.

InsSo3D has also been tested in a test tank using a *Qualisys* motion tracking system, allowing for accurate trajectory estimation. Moreover, a map of the test tank has been obtained by photogrammetry, using the images from the stereo camera for map evaluation. The test tank is a difficult evaluation environment for both sonar and compass due to the rectangular shape and the reinforced concrete construction of the tank. This results in multipath effects on the 3D Sonar point cloud, as well as significant drift on the compass and DVL data. This resulted in significant odometry drift during this experiment. In addition, the test tank only has a few geometric features to align with, which makes it a challenging environment for depth based SLAM.

B. Localisation and 3D Reconstruction Accuracy

As the odometry and InsSo3D coordinate systems are different from the reference coordinate system, we first put all trajectories in the same coordinate system with a common origin. Then, the absolute position error $APE(i)$ with respect to the reference throughout the acquisition of 3D Sonar frames is computed. For each trajectory, we also compute the root-mean-square (RMS) and standard deviation (STD) of the APE, denoted APE_{RMS} and APE_{STD} . Second, the consistency of the output trajectories of InsSo3D is evaluated, the trajectories are aligned to the references




Fig. 9: Map qualitative comparison for the *Tank* experiments. InsSo3D results are on the first row, and SFM results are on the second row. The first column shows a front view of the maps and the second a top view. The colour represents the normal azimuth angle calculated using the x and y components of the normal vector.

using Umeyama alignment [39], and the RMS and STD of the resulting absolute position error $APE_{RMS,align}$, and $APE_{STD,align}$ are computed. Finally, we evaluate the output 3D reconstructions by using the mean distance to the reference mesh e_{map} as well as its standard deviation $e_{STD,map}$ after alignment, using CloudCompare [40].

We first provide an overview of the maps generated by InsSo3D and compare them with the reference maps of each scenario. Fig. 8 shows InsSo3D’s map reconstruction of the *Quarry* Test environments, and the corresponding SFM results. It can be seen that InsSo3D is able to reconstruct the map with high fidelity, matching closely that of the SFM reference. Figures 9 present the *Tank* Test environment reconstruction. Despite this environment’s complexity, due to sonar reflections and compass issues, InsSo3D still achieves accurate 3D scene reconstruction. Furthermore, it should be noted that thanks to the long range of the 3D Sonar compared to the camera, the back wall of the tank is visible on the InsSo3D map but not on the SFM results.

In all scenarios, InsSo3D provides excellent scene reconstruction, enabling robots to operate in environments where cameras are not suitable, such as low-light or murky water conditions. Additionally, the sonar’s larger field of view compared to stereo cameras allows InsSo3D to generate larger maps than SFM, as demonstrated in the *Tank* scenario (Fig. 9).

Table I shows InsSo3D’s trajectories and map errors, and standard deviation for all scenes. It can be seen that map errors are only a few centimetres for all sequences. Such small errors can provide very accurate maps, which can be used for inspection or mapping purposes.

Furthermore, the trajectory errors are given in Fig. 10. In all experiments, InsSo3D reduced position drift and standard deviation compared to odometry, even during the long 47-minute missions. In addition, Table I shows that, $APE_{RMS,align}$ is lower than APE_{RMS} for all scenarios, maintaining a very low $APE_{STD,align}$ value, meaning that the trajectories are even more accurate locally, with respect to the scene, than they are with respect to the vehicle




Fig. 10: Position error over x , y , yaw in and the absolute pose error APE for the different experiments over time t .

starting point. Figure 10 shows that InsSo3D error remains relatively small, unlike the odometry error. Moreover, during the *Tank* scenario, the odometry drifts significantly by nearly 40 degrees on yaw due to magnetic interference from the reinforced concrete structure of the tank, which disturbs the compass. Even in these situations, InsSo3D is still able to correct its trajectory and maintains the yaw error low.

Eventually, in all scenarios, InsSo3D APE_{RMS} remains small compared to motion and mapped object scale, and therefore appears relevant for safe navigation in these test scenarios.

C. Runtime Evaluation

This evaluation considers the processing time of InsSo3D’s frontend, backend and global TSDf map generation. The TSDf pipeline is implemented on the GPU using CUDA, while the frontend and backend are CPU-based. The evaluation is performed on the *Quarry* scenario using an *Nvidia RTX 3070* GPU with 8GB of VRAM and an *Intel Core i7-12700H* with 32GB of RAM. Table II reports the time of the different components. It should be noted that each of these components runs asynchronously; therefore, only the

	Frontend	Backend	Global TSDf
Mean Time [ms]	69.9	593.0	4797.4
STD [ms]	24.3	394.3	3519.8

TABLE II: Runtime of InsSo3D main components. The frontend time is reported per frame, the backend is performed per sub-map (> 100 frames) and the TSDf global map update performed continuously.

frontend performance impacts the SLAM ability to operate in real time. On this dataset, the frontend takes an average of 70ms per frame, or approximately 14Hz, which is faster than the sonar operating frequency at 6Hz.

V. CONCLUSIONS

This work focuses on underwater online localisation and 3D mapping using a 3D Sonar combined with the vehicle’s INS. We developed a new SLAM framework using frontend and backend graph optimisation to perform robust and large-scale 3D SLAM using a 3D Sonar. Our comprehensive evaluation, in extensive field experiments, demonstrates that InsSo3D significantly reduces navigation drift while generating high-fidelity 3D maps even during long missions. Comparison with ground truth trajectories and maps shows that InsSo3D is able to achieve trajectory accuracies consistently below 21cm and successfully reconstruct quality maps of large-scale environments. These results establish InsSo3D’s effectiveness for real-time, long-range acoustic SLAM in challenging underwater conditions. InsSo3D can be used in real-world deployment for initial map generation in poor visibility conditions, as well as safe motion planning for additional closer inspection.

While the results obtained from the developed method are promising, there are still limitations that could be addressed in future work. Just like all feature matching algorithms,

Sequence		<i>Quarry</i>	<i>Tank</i>
Odom	APE_{RMS} [m]	0.942	1.784
	APE_{STD} [m]	0.382	1.021
	$APE_{RMS,align}$ [m]	0.365	1.034
	$APE_{STD,align}$ [m]	0.185	0.430
InsSo3D	APE_{RMS} [m]	0.360	0.089
	APE_{STD} [m]	0.162	0.046
	$APE_{RMS,align}$ [m]	0.213	0.078
	$APE_{STD,align}$ [m]	0.090	0.045
e_{map} [m]		0.086	0.073
$e_{STD,map}$ [m]		0.0787	0.063

TABLE I: Error and standard deviation of odometry and InsSo3D for trajectory and map (smallest in bold).

GICP registration is not robust to environments that contain spatial symmetry or that lack geometric features. Also, in environments prone to sonar multipath effects, the 3D Sonar can produce noisy data and therefore affect the quality of the mapping and localisation. Future work will investigate the fusion of 3D Sonar and camera data to improve robustness and accuracy. Moreover, the camera could also be used to texture the generated map, as well as to improve the localisation.

REFERENCES

- [1] C. Campos, R. Elvira, J. J. Gómez, J. M. M. Montiel, and J. D. Tardós, "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.
- [2] A. Filisetti, A. Marouchos, A. Martini, T. Martin, and S. Collings, "Developments and applications of underwater lidar systems in support of marine science," in *OCEANS 2018 MTS/IEEE Charleston*. IEEE, 2018, pp. 1–10.
- [3] J. Wang, F. Chen, Y. Huang, J. McConnell, T. Shan, and B. Englot, "Virtual maps for autonomous exploration of cluttered underwater environments," *IEEE Journal of Oceanic Engineering*, vol. 47, no. 4, pp. 916–935, 2022.
- [4] D. Adolphsson, M. Magnusson, A. Alhashimi, A. J. Lilienthal, and H. Andreasson, "Lidar-level localization with radar? the cfar approach to accurate, fast, and robust large-scale radar odometry in diverse environments," *IEEE Transactions on Robotics*, vol. 39, no. 2, pp. 1476–1495, 2023.
- [5] J. Edstedt, Q. Sun, G. Bökman, M. Wadenbäck, and M. Felsberg, "RoMa: Robust Dense Feature Matching," *IEEE Conference on Computer Vision and Pattern Recognition*, 2024.
- [6] M. Grimaldi, D. Nakath, M. She, and K. Köser, "Investigation of the challenges of underwater-visual-monocular-slam," *arXiv preprint arXiv:2306.08738*, 2023.
- [7] K. G. Foote, "Using a sonar in a different environment from that of its calibration: Effects of changes in salinity and temperature," in *OCEANS 2018 MTS/IEEE Charleston*. IEEE, 2018, pp. 1–5.
- [8] M. Johnson-Roberson, O. Pizarro, S. B. Williams, and I. Mahon, "Generation and visualization of large-scale three-dimensional reconstructions from underwater robotic surveys," *Journal of Field Robotics*, vol. 27, no. 1, pp. 21–51, 2010.
- [9] F. Hidalgo, C. Kahlefeldt, and T. Bräunl, "Monocular orb-slam application in underwater scenarios," in *2018 OCEANS - MTS/IEEE Kobe Techno-Oceans (OTO)*, 2018, pp. 1–4.
- [10] S. Xu, K. Zhang, Z. Hong, Y. Liu, and S. Wang, "Diso: Direct imaging sonar odometry," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 8573–8579.
- [11] L. Zhuang, X. Chen, W. Lu, and Y. Yan, "Graph matching for underwater simultaneous localization and mapping using multibeam sonar imaging," *Journal of Marine Science and Engineering*, vol. 12, no. 10, p. 1859, 2024.
- [12] S. Pan, Z. Hong, Z. Hu, X. Xu, W. Lu, and L. Hu, "Russo: Robust underwater slam with sonar optimization against visual degradation," *IEEE/ASME Transactions on Mechatronics*, 2025.
- [13] J. Bai, D. Zhu, M. Chen, and C. Luo, "Sio-uv: Rapid and robust sonar inertial odometry for underwater vehicles," *IEEE Transactions on Industrial Electronics*, 2025.
- [14] S. Archieri, J. Drupt, A. Cinar, M. Grimaldi, I. Carlucho, J. Scharff, and Y. R. Petillot, "3dssdf: Underwater 3d sonar reconstruction using signed distance functions," in *2025 IEEE International Conference on Robotics & Automation (ICRA 2025)*, 2025.
- [15] M. D. Aykin and S. Negahdaripour, "On 3-d target reconstruction from multiple 2-d forward-scan sonar views," in *Oceans 2015-Genova*. IEEE, 2015, pp. 1–10.
- [16] P. V. Teixeira, D. Fourie, M. Kaess, and J. J. Leonard, "Dense, sonar-based reconstruction of underwater scenes," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 8060–8066.
- [17] T. A. Huang and M. Kaess, "Towards acoustic structure from motion for imaging sonar," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2015, pp. 758–765.
- [18] T. Guerneve, K. Subr, and Y. Petillot, "Three-dimensional reconstruction of underwater objects using wide-aperture imaging sonar," *Journal of Field Robotics*, vol. 35, no. 6, pp. 890–905, 2018.
- [19] Y. Wang, Y. Ji, H. Woo, Y. Tamura, H. Tsuchiya, A. Yamashita, and H. Asama, "Acoustic camera-based pose graph slam for dense 3-d mapping in underwater environments," *IEEE Journal of Oceanic Engineering*, vol. 46, no. 3, pp. 829–847, 2020.
- [20] N. Sadjoli, Y. Cai, G. Seet, and B. Elhadidi, "Pcd reconstruction, object classification and pose estimation for underwater vehicles using orthogonal multibeam forward looking sonar fusion," *Ocean Engineering*, vol. 288, p. 116019, 2023.
- [21] J. McConnell, I. Collado-Gonzalez, P. Szenher, and B. Englot, "Large-scale dense 3-d mapping using submaps derived from orthogonal imaging sonars," *IEEE Journal of Oceanic Engineering*, 2024.
- [22] W. Liu, Y. Li, L. Li, W. Zhang, and W. Huang, "Target positioning of dual forward looking sonars based on orthogonal detection," *Mechatronics*, vol. 98, p. 103135, 2024.
- [23] R. DeBortoli, F. Li, and G. A. Hollinger, "Elevatenet: A convolutional neural network for estimating the missing dimension in 2d underwater sonar images," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 8040–8047.
- [24] N. Jaber, B. Wehbe, and F. Kirchner, "Sonar2depth: Acoustic-based 3d reconstruction using cgans," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 5828–5835.
- [25] H. Gao, B. Nie, W. Lu, Y. Feng, and M. Hu, "Lambertian-trf: Lambertian tensorial radiance fields for underwater 3d reconstruction using imaging sonar," *IEEE Transactions on Instrumentation and Measurement*, 2025.
- [26] M. Qadri, M. Kaess, and I. Gkioulekas, "Neural implicit surface reconstruction using imaging sonar," 2022. [Online]. Available: <https://arxiv.org/abs/2209.08221>
- [27] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "Nerf: Representing scenes as neural radiance fields for view synthesis," *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, 2021.
- [28] A. Ferreira, J. Almeida, A. Matos, and E. Silva, "Real-time registration of 3d underwater sonar scans," *Robotics*, vol. 14, no. 2, p. 13, 2025.
- [29] R. K. Hansen, U. Castellani, V. Murino, A. Fusiello, E. Puppo, L. Papaleo, M. Pittore, M. Gobbi, L. Bisone, K. Kleppe, *et al.*, "Mosaicing of 3d sonar data sets-techniques and applications," in *Proceedings of OCEANS 2005 MTS/IEEE*. IEEE, 2005, pp. 2326–2333.
- [30] K. Koide, M. Yokozuka, S. Oishi, and A. Banno, "Voxelized gicp for fast and accurate 3d point cloud registration," *EasyChair Preprint 2703*, EasyChair, 2020.
- [31] A. Segal, D. Hähnel, and S. Thrun, "Generalized-icp," 06 2009.
- [32] I. Vizzo, T. Guadagnino, J. Behley, and C. Stachniss, "Vdbfusion: Flexible and efficient tsdf integration of range sensor data," *Sensors*, vol. 22, no. 3, 2022. [Online]. Available: <https://www.mdpi.com/1424-8220/22/3/1296>
- [33] W. E. Lorensen and H. E. Cline, "Marching cubes: A high resolution 3d surface construction algorithm," in *Seminal graphics: pioneering efforts that shaped the field*, 1998, pp. 347–353.
- [34] F. Dellaert and M. Kaess, *Factor Graphs for Robot Perception*. Foundations and Trends in Robotics, Vol. 6, 2017. [Online]. Available: <http://www.cs.cmu.edu/~kaess/pub/Dellaert17fnt.pdf>
- [35] T. Łuczynski, P. Łuczynski, L. Pehle, M. Wirsum, and A. Birk, "Model based design of a stereo vision system for intelligent deep-sea operations," *Measurement*, vol. 144, pp. 298–310, 2019.
- [36] J. S. Willners, I. Carlucho, S. Katagiri, C. Lemoine, J. Roe, D. Stephens, T. Łuczynski, S. Xu, Y. Carreno, E. Pairet, C. Barbalata, Y. Petillot, and S. Wang, "From market-ready rovs to low-cost auvs," in *OCEANS 2021: San Diego – Porto*, 2021, pp. 1–7.
- [37] T. Moore and D. Stouch, "A generalized extended kalman filter implementation for the robot operating system," in *Proceedings of the 13th International Conference on Intelligent Autonomous Systems (IAS-13)*. Springer, July 2014.
- [38] J. L. Schönberger and J.-M. Frahm, "Structure-from-motion revisited," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [39] S. Umeyama, "Least-squares estimation of transformation parameters between two point patterns," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 4, pp. 376–380, 1991.
- [40] "Cloudcompare," <https://www.danielgm.net/cc/>, 2024.