

BenthiCat: An opti-acoustic dataset for advancing benthic classification and habitat mapping

Journal Title
XX(X):1–19
©The Author(s) 2024
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/ToBeAssigned
www.sagepub.com/

SAGE

Hayat Rajani¹, Valerio Franchi¹, Borja Martinez-Clavel Valles², Raimon Ramos², Rafael Garcia¹ and Nuno Gracias¹

Abstract

Benthic habitat mapping is fundamental for understanding marine ecosystems, guiding conservation efforts, and supporting sustainable resource management. Yet, the scarcity of large, annotated datasets limits the development and benchmarking of machine learning models in this domain. This paper introduces a thorough multi-modal dataset, comprising about a million side-scan sonar (SSS) tiles collected along the coast of Catalonia (Spain), complemented by bathymetric maps and a set of co-registered optical images from targeted surveys using an autonomous underwater vehicle (AUV). Approximately 36 000 of the SSS tiles have been manually annotated with segmentation masks to enable supervised fine-tuning of classification models. All the raw sensor data, together with mosaics, are also released to support further exploration and algorithm development. To address challenges in multi-sensor data fusion for AUVs, we spatially associate optical images with corresponding SSS tiles, facilitating self-supervised, cross-modal representation learning. Accompanying open-source preprocessing and annotation tools are provided to enhance accessibility and encourage research. This resource aims to establish a standardized benchmark for underwater habitat mapping, promoting advancements in autonomous seafloor classification and multi-sensor integration.

Keywords

side-scan sonar, seabed mapping, seafloor classification, semantic segmentation, self-supervised pretraining, weakly-supervised learning, modality fusion, dataset

1 Introduction

Despite the oceans covering more than 70 % of our planet's surface, a minuscule fraction of the seabed has been thoroughly explored. Benthic maps provide information about the types of habitats present on the seafloor, including details about the composition of the substrate (e.g., sand, mud, rock) and the diversity and distribution of benthic organisms such as corals, sponges, and other marine life. These maps are not only valuable for understanding marine biodiversity and conservation efforts but also for economic activities. Typically, the surveys for creating such maps are conducted from oceanographic vessels, using different types of acoustic sensors, such as the side-scan sonar (SSS), due to their long range of operation and their ability operate effectively even in poor visibility conditions. This is followed by a labour-intensive process of manual interpretation of the surveyed area by a group of experts, which is not only time-consuming but also economically taxing. Additionally, video samples are often collected over certain surveyed areas to allow a visual confirmation of seafloor characteristics, compensating for the low spatial resolution and low signal-to-noise ratio of acoustic sensing, which would otherwise make it difficult to classify the benthos even by a trained human eye. This, however, further adds to the operational costs of creating such detailed benthic maps.

While AUV-based acoustic mapping has evolved dramatically, there is still a clear need to move the ability to interpret these maps from a post-mission expert-based analysis to an on-mission machine-learning capability. The

automation of seafloor classification could revolutionize this domain, offering real-time, precise underwater mapping. In the recent years, there have been various studies on mission-time seafloor classification in SSS imagery (Wang et al. 2019; Zhao et al. 2021; Burguera and Bonin-Font 2020; Yu et al. 2021; Rajani et al. 2023) using Deep Neural Networks (DNNs). The problem is posed as dense pixel-wise categorization, or semantic segmentation, of different seafloor types in patches of acoustic images. However, with the exception of the study by Burguera and Bonin-Font (2020), none of the aforementioned studies have published the data used for their respective experiments. Moreover, due to the high costs and logistical intricacies of data collection in marine environments, compounded by the meticulous and time-consuming nature of data labelling, there is already a severe lack of publicly available labelled datasets that are also large enough for training Machine Learning (ML) models. As such, to the best of our knowledge, there is no standard dataset to date for benchmarking and validation of algorithms or reproduction of results, thereby presenting formidable barriers to stronger statistical analysis and methodological improvements.

¹Computer Vision and Robotics Research Institute, University of Girona, Spain

²Tecnoambiente SLU, Spain

Corresponding author:

Hayat Rajani, Edifici CIRS, Parc Científic i Tecnològic UdG, Carrer Pic de Peguera 13, 17003 Girona, Spain
Email: hayat.rajani@udg.edu

To overcome these challenges, this paper introduces a thorough multi-modal dataset for training and benchmarking ML models on seafloor classification in SSS imagery. Our team's collaboration with the Spanish company TecnoAmbiente has yielded an extensive repository of about a million SSS tiles, encompassing a wide variety of seafloor types, collected as part of a large-scale survey along the entire coast of Catalonia. This huge amount of data can be easily leveraged for self-supervised pre-training of ML models. In order to facilitate subsequent fine-tuning of these models, we also manually annotated about 36 000 images with segmentation masks. We additionally include the raw and mosaiced SSS waterfalls, together with coarse interpretations, as part of this dataset. Although the emphasis of this work lies on seafloor classification in SSS imagery, we also release bathymetric maps of the surveyed area collected via the onboard Multibeam Echosounder (MBES) together with all the associated navigation information to encourage further research. Furthermore, we conducted an independent series of optical surveys to complement the SSS imagery. These surveys solely targeted specific areas identified during the acoustic surveys to collect samples of certain classes of interest, in order to help alleviate inter-class ambiguities in SSS imagery. This resulted in an additional 178 000 images, which were subsequently paired with SSS tiles from the corresponding geographic location, culminating in an opti-acoustic dataset. We believe that such a dataset could be leveraged in modality fusion pipelines to further enhance the quality of seafloor classification.

The remainder of this paper is organised as follows. Section 2 presents an overview of other publicly available datasets for benthic mapping. Section 3 and Section 4 describe the details of the surveys conducted for acoustic and optical data acquisition, including the system setup and sensor characteristics, and outline how the collected data was processed. Section 5, on the other hand, presents details on the additional processing of the collected data to compile datasets for training ML models. Section 6 explains how the datasets are structured. Finally, Section 7 presents conclusions and directions for future work.

2 Related Work

Publicly available datasets containing opti-acoustic data, particularly SSS, are notoriously scarce. A notable contribution in this direction is the AURORA multi-sensor dataset for robotic ocean exploration (Bernardi et al. 2022), which was compiled from multiple missions using the Autosub 6000 AUV by the National Oceanography Centre (NOC) in the Greater Haig Fras Marine Conservation Zone (MCZ). This dataset is designed to facilitate the testing of control, navigation, and sensor processing algorithms for AUVs, as well as to serve as a benchmark for research in robotic ocean exploration. It integrates various modalities such as SSS, MBES, and optical images, along with AUV navigation data and other associated sensor metadata. However, the dataset does not specifically target seabed classification and lacks fine-grained annotations of seabed sediments.

In contrast, the dataset released in conjunction with the work of Burguera and Bonin-Font (2020) is explicitly focused on seabed segmentation in SSS images. The data

was acquired using an EcoMapper AUV in Port de Sóller, Mallorca, Spain traversing over 4 km, resulting in a total of 22 438 swaths. These swaths were translated to overlapping patches of size 83×83 , annotated across three classes: rocks, sand ripples and mud for training.

On the other hand, the recently published SeafloorAI dataset (Nguyen et al. 2024) is designed to support AI-based seafloor mapping. The dataset includes around 696 000 sonar images and 827 000 segmentation masks, collected from 62 different marine survey regions, covering a total area of over 17 300 km². The dataset combines SSS and MBES data and provides pixel-wise labels for five types of geological features: sediment type, physiographic zone, habitat, fault, and fold. In addition, SeafloorAI also includes a vision-language extension called SeafloorGenAI, which provides image-level captions and about seven million question-answer pairs, enabling the development of multimodal models. While SeafloorAI provides annotations for seabed sediments, it only focuses on four broad categories such as rock, gravel, sand and mud, and its objectives quite differ from our work. Where we focus on seabed sediment segmentation with an emphasis on fine-grained, high-resolution distinctions relevant for robot perception, SeafloorAI is designed primarily for large-scale geological interpretation. Its broader goal is to support automated mapping of multiple seafloor attributes (including physiographic zones, structural features like faults and folds, and general habitat types) making it well-suited for regional-scale geoscientific studies, but less so for targeted applications for AUVs such as real-time scene understanding, close-range intervention tasks, navigation and path planning.

To the best of our knowledge these are the only publicly available datasets that align closely with our target application. Other existing datasets in the literature primarily emphasize on object detection (Huo et al. 2020; Du et al. 2023) or binary segmentation of discrete targets such as shipwrecks (Sethuraman et al. 2025), pipelines (Álvarez Tuñón et al. 2024), or mines (Pessanha Santos et al. 2024), rather than continuous benthic habitats.

3 Acoustic Data Acquisition

The acoustic data presented in this work was collected by Tecnoambiente SLU through various geophysical surveys of the entire coast of Catalonia*. These surveys involved over 3500 hours of geophysical (SSS) and hydrographic (MBES) data acquisition, covering more than 1900 km² in area, extending from the coastline up to six nautical miles offshore. In order to ensure accurate characterization of the habitat, additional surveys were carried out collecting about 500 sediment samples, and 150 scientific dives and ROV-enabled missions for acquisition of video samples. It is important to note that the goal of these additional surveys was to provide a visual validation of the habitat in certain locations, and not to conduct an extensive optical survey. As such, there was not an easy way to co-register the acoustic

*This data was used to create an interactive marine habitat map by the Government of Catalonia, available at https://sig.gencat.cat/visors/habitats_marins.html



Figure 1. Overview of the surveyed area, divided into 22 sectors, illustrating a portion of the mosaiced SSS transects from sectors N08 and the corresponding DEM.

and optical data, since the latter comprised much smaller and separate transects. Therefore, the data presented as part of these surveys do not include any optical imagery.

Figure 1 outlines the surveyed area. The entire coast of Catalonia was subdivided into 22 sectors: 16 along the northern and 6 along the southern coast. Each sector comprised multiple overlapping transects, collecting over 50 million swaths, encompassing a wide variety of seabed types, including sand ripples, mud, posidonia (dead and alive), cymodocea, artificial reefs, rocks, corals, detrital funds, etc. Additionally, miscellaneous objects such as buoys, anchors, conduits and cables, shipwrecks, and agricultural infrastructure were also observed. A total of 26 different classes were identified, which were then grouped into 12 categories based on the similarity of visual features or benthic type as summarized in Appendix A.




Figure 2. Schema of the equipment on board the survey vessel.

The remainder of this section provides further details on the system setup, sensor configurations, processing of the raw data and the methodology for creating annotations.

3.1 System Setup

The surveys were conducted using Tecnoambiente's 16 m long survey vessel, *Festina Lente*. The vessel was equipped with a Klein 3000H SSS and an R2Sonic 2022/2024 MBES for acoustic data gathering, together with a Sound Velocity Sensor (SVS) and a Sound Velocity Profile (SVP) for sound velocity corrections, and an Inertial Navigation System (INS), a Motion Reference Unit (MRU) and an Ultrashort Baseline (USBL) to aid with navigation. Figure 2 depicts a schematic of how these sensors interacted with each other during the course of the surveys.

The Klein 3000H (Klein Marine Systems 2021) is a dual frequency (500 kHz/900 kHz) side-scan sonar that was operated as a towfish hauled behind the survey vessel by a cable controlling its operating depth. Only the scans from the 500 kHz frequency domain were utilized to achieve the best balance between resolution and coverage area per transect. Data acquisition was performed at an altitude varying between 5 m and 15 m, relative to about 10 % of the sensor range. The sensor range, on the other hand, was manually adjusted for each zone based on the depth. It varied between 50 m and 150 m for each side (port and starboard). Corrections for navigation errors were applied recursively by the on-board INS during the course of the surveys, together with an online georeferencing of the swaths via the USBL. Initially, a Sonardyne Scout Plus USBL (Sonardyne 2021) was used in conjunction with an Navsight Ekinox INS (SBG Systems 2018), resulting in a position accuracy of 0.2 % of traveled distance and a heading accuracy of 0.02°. This configuration was later replaced by a GAPS Gen 4 USBL (iXblue 2014), which combines high performance USBL and a fiber optic INS in the same housing, resulting in a position accuracy of 0.06 % of the slant range and a heading accuracy of 0.01° secant latitude.


The R2Sonic 2022 (R2SONIC 2022) is a wideband shallow water multibeam echosounder that was mounted to

Table 1. Sensor configurations and characteristics for acoustic surveys.

Parameter	Value
Klein 3000H SSS	
Operating Frequency	500 kHz
Beam Tilt (θ)	15°
Beam Width (α)	40°
Sonic 2022/2024 MBES	
Operating Frequency	400 kHz
Resolution	1 cm
Beam Count	256
Beam Angle	140° (max)
Valeport miniSVS	
Acoustic Frequency	2.5 MHz
Sample Rate	60 Hz
Resolution	0.001 m
Velocity Accuracy	± 0.017 m/s
Valeport SWiFT SVP	
Range	1375–1900 m/s
Resolution	0.001 m/s
Accuracy	± 0.02 m/s
Coda Octopus F-175 MRU	
Heading accuracy	0.1° (with RTK)
Roll and pitch accuracy	0.03° (with RTK)
Navsight Ekinox INS	
Position accuracy	0.2 % of traveled distance
Heading accuracy	0.02°
Roll and pitch accuracy	0.02°
GAPS Gen 4 USBL	
Position accuracy	0.06 % of slant range
Heading accuracy	0.01° secant latitude
Roll and pitch accuracy	0.01°

the hull of the survey vessel. Data acquisition was performed at an operating frequency of 400 kHz, with a maximum ping rate of 60 Hz and 256 beams per ping. The swath angle was manually adjusted for each zone depending on the depth, never exceeding 140°. The R2Sonic 2022 was later replaced by a R2Sonic 2024 (R2SONIC 2024), maintaining the same configurations. For sound velocity corrections, a Valeport miniSVS (Valeport 2016) was mounted at the head of the MBES. The SVS was operated at an acoustic frequency of 2.5 MHz and a sampling rate of 60 Hz. The 100 mm sensor reported a maximum theoretical error of ± 0.017 m/s with a 0.001 m resolution. This was further complemented by sound velocity profiling via a Valeport SWiFT SVP (Valeport 2019), carried out approximately every 3 hours or when changing zones. Real-time corrections for draft and tide variations as well as attitude were performed via the survey vessel's MRU, a Coda Octopus F-175 (Coda Octopus 2013) with RTK corrections, resulting in a heading accuracy of 0.1° and roll and pitch accuracy of 0.03°. This was later replaced by the Navsight Ekinox INS.

Table 1 summarizes the key configurations and characteristics of all sensors.

**Figure 3.** Illustrative example of a portion of an SSS mosaic from sector N08 overlaid with the corresponding benthic habitat interpretation.

3.2 Data Processing

The raw SSS data were recorded in the eXtended Triton Format (XTF) (Kermarrec et al. 2020), a commonly used open source format for recording hydrographic data, which were subsequently gain-corrected and mosaiced using SonarWiz (Chesapeake Technology, Inc. 2010-2024). The mosaiced SSS waterfalls were then georeferenced and interpretations of various sea bottom types were manually drawn by experts using ArcGIS (ESRI 2021). This led to a coarse benthic habitat map of the entire coast.

Figure 3 depicts examples of a portion of one such SSS mosaic overlaid with the corresponding interpretation. Although done by experts, the interpretations are prone to human errors. While certain areas lacked the respective interpretations altogether, those of several others were inaccurate as highlighted in the figure. Moreover, further noise was inevitably introduced during the ground truth generation process as detailed in Section 5.

The raw MBES data were recorded using Qinsy (Quality Positioning Services 2021b), a survey planning and real-time hydrographic data processing software. Qinsy ingests all the data from the acoustic and navigation sensors to compute corrections in real-time and produces a single DB file, a proprietary database format. These raw database files from all the transects were subsequently processed in Qimera (Quality Positioning Services 2021a) and converted to the more standard ASCII gridded XYZ files, which were then rasterized using ArcGIS. The so produced Digitized Elevation Models (DEMs) were exported as GeoTIFFs and Shapefiles (.shp) for each sector, and also as a non-editable PDF file containing all cartographic information.

4 Optical Data Acquisition

As previously mentioned, acoustic sensing suffers from low spatial resolution and low signal-to-noise ratio, particularly in dynamic environments, making it difficult to classify the benthos even by a trained human eye. Video samples are, therefore, typically acquired over the surveyed areas to allow a visual confirmation of seafloor characteristics, and thus have a ground-truth for the manual interpretations. Consequently, there is a strong need to enable conducting integrated surveys by fusing information from different




Figure 4. Overview of the optical transects, depicted as white lines, across the surveyed areas in sectors N07 and N08.

sensors, especially using autonomous robotic platforms. This becomes particularly crucial for distinguishing between seabed types with similar appearance in the acoustic images.

As opposed to the acoustic surveys, optical data acquisition was performed at an altitude varying between 3.5 m and 4.5 m depending on the area and weather conditions. Consequently, the varying fields-of-view and the differences in scale between the acoustic and optical images resulted in the lack of one-to-one correspondence between the two modalities. The methodology thus adopted for aligning the data is discussed in Section 5.

As such, we conducted an independent series of optical surveys targeting certain areas from sectors N7 and N8 which were identified during the acoustic surveys to be abundant in the following six classes of interest: fine sediments, sand ripples, rocks, corals, *Posidonia oceanica* and *Cymodocea nodosa*. Figure 4 depicts a few transects over the surveyed areas. The surveys were conducted from the 22nd until the 26th of April 2024, covering approximately a distance of 27.55 km and an area of 118.2 km² with the optical data acquired at a maximum altitude of 5 m. During the course of the optical surveys, the transects were carried out at three different altitudes: 3.5 m, 4 m, 4.5 m and at two different AUV speeds: 0.5 m/s and 0.8 m/s. Table 2 summarises the footprint of each image on the seabed for each of those altitudes for a downward-looking camera.

Table 2. Altitudes and image footprints during the optical surveys.

Altitude [m]	Image Footprint [m ²]
3.5 m	10.69 m ²
4.0 m	13.95 m ²
4.5 m	17.68 m ²

The remainder of this section provides further details on the system setup, sensor configurations and processing of the raw data.

4.1 System Setup

The surveys were conducted using the Girona1000, a variant of the Girona500 hovering AUV (Ribas et al. 2012) developed at the University of Girona for operations at depths of up to 1000 m. The reconfigurable vehicle was equipped with two cameras to gather the optical data, along with a Doppler-Velocity Logger (DVL), a Global Navigation Satellite System (GNSS), an Inertial Navigation System (INS), and a Sound Velocity Sensor (SVS) to estimate the vehicle's position and orientation during the surveys.




Figure 5. Image of the Girona1000 AUV at port before the start of the optical surveys.

Table 3. Sensor distance offsets (x, y, z) with respect to the Girona1000's centre of gravity for optical surveys.

Sensor	Offset (x,y,z) [cm]
Phins Compact C3	(24.90, 37.70, -28.40)
DVL1000	(-65.0, 0.0, 49.96)
L86	(-43.522, -35.005, -65.681)
miniSVS	(-57.05, 35.0, -42.725)
Blackfly S (left)	(56.0, 0.0, 47.0)
Blackfly S (right)	(56.0, 0.0, 47.0)

Table 4. Sensor orientation offsets (roll ψ , pitch θ , yaw ϕ) with respect to the Girona1000's centre of gravity for optical surveys. The cameras' orientations were changed multiple times during the surveys based on the area.

Sensor	Offset (ψ,θ,ϕ) [rad]
Phins Compact C3	(0, 3.141593, 0)
DVL1000	(0, 0, 3.141593)
L86	(0, 0, 0)
miniSVS	(0, 0, 0)

Figure 5 displays the Girona1000 AUV with all the mounted sensors, while Tables 3 and 4 list the sensor offsets with respect to the Girona1000's centre of gravity. Table 5, in turn, summarizes the key configurations and characteristics of all the sensors.

The Phins Compact C3 (iXblue 2016), based on fiber-optic gyroscope (FOG) technology, is a compact INS designed for customised sub-sea integration. It was used onboard the vehicle as the core Inertial Measurement Unit (IMU) to acquire linear accelerations and angular velocities from the raw sensor measurements. The DVL1000-4000m (Nortek 2015), rated for operational depths of up to 4000 m, was used to measure the vehicle's linear velocity, whereas the Valeport miniSVS (Valeport 2016) supplied low noise, high accuracy and high resolution sound velocity and pressure readings using digital Time-of-Flight processing. On the other hand, a Quectel L86 GNSS module (Quectel 2014), equipped with a patch antenna, was mounted on the antenna mast to provide surface positioning. All sensors were interfaced with the Phins Compact C3, which fused their measurements through its commercial INS algorithm consisting of an Unscented Kalman Filter (UKF) (Wan and Van Der Merwe 2000), providing the real-time pose and velocity estimates of the vehicle throughout the survey.

Two FLIR Blackfly S GigE cameras (FLIR 2016) were mounted on the bottom hull of the Girona1000 AUV. The cameras were specifically arranged to maximise the overlap between the optical data and the near-field region of the sonar data, with each camera oriented toward opposite sides of the vehicle (see Figure 5, bottom-center). The cameras feature a Sony IMX429 sensor model with a 2/3 inch format, a pixel size of 4.5 μm , and a native resolution of 1936 $\times$ 1464 (2.8 MP). They provided colour images over a Gigabit Ethernet interface at a data rate of 1 Gbit/s and supported a maximum frame rate of 43 Hz. The exposure time and the frame-rate were adjusted before the start of the missions in order to achieve the optimal horizontal/vertical overlap

Table 5. Sensor configurations and characteristics for optical surveys.

Parameter	Value
Phins Compact C3 INS	
Position accuracy	0.3 % of traveled distance
Heading accuracy	0.15°
Roll and pitch accuracy	0.05°
DVL1000	
Maximum altitude	75 m
Accuracy	$\pm 0.1\%$ / ± 0.1 cm/s
Velocity resolution	0.01 m/s
Maximum ping rate	8 Hz
miniSVS	
Acoustic Frequency	2.5 kHz
Sample Rate	1 kHz to 32 kHz
Resolution	0.001 m
Velocity Accuracy	± 0.20 m/s
Teledyne FLIR Blackfly S GigE	
Megapixels	2.8 MP
Maximum Frame Rate	43 Hz
Gain Range	0 dB to 47 dB

between the stereo pair, and to minimise the motion blur while simultaneously maintaining a wide range of intensities across the entire image.

The camera intrinsics were calculated by placing a ChArUco board underwater. A ChArUco board is a combination of a chessboard with ArUco markers (Garrido-Jurado et al. 2014) instead of the white squares. The ArUco markers are easily identifiable and their positions inside the chessboard are known beforehand, thus it is not necessary to frame the entire board inside the image bounds in order to perform the calibration. The water introduces additional image distortions to the already present lens distortion, however those distortions are mostly radial in nature, therefore the intrinsic camera calibration will be able to model these distortion with higher radial distortion coefficients. The camera calibration parameters (f_x, f_y, x_0, y_0) and the radial (k_1, k_2, k_3) and tangential lens distortion (p_1, p_2) parameters are listed in Table 6.

The dimensions of the camera's sensor are 8.71 mm $\times$ 6.588 mm. The focal length of the camera (assuming a pinhole camera model) was calculated as 8.11 mm. The horizontal and vertical angles of view based on this focal length were calculated, respectively, as 68.26° and 49.29°.

During the survey, the orientation of the stereo pair was modified to better suit the image overlap and the visible information inside the field of view of the cameras. The extrinsic calibration was performed offline after the missions were completed. Table 7 shows the percentage overlap and angle between the stereo pair for each day of the survey.

Since the extrinsic calibration was not performed using a calibration pattern, and the stereo pair is only offset by a known baseline and angle, it is advisable to optimise them using the provided extrinsics as initial estimates when employing the stereo pairs for any SLAM or localisation task.

Table 6. Camera intrinsic parameters and lens distortion coefficients.

Parameter	Value
Camera Intrinsics	
f_x	1801.899 px
f_y	1801.899 px
x_0	938.805 px
y_0	735.677 px
Lens Distortion	
k_1	-0.19027
k_2	0.10212
p_1	0.00139
p_2	-0.00343
k_3	0.00342

4.2 Data Processing

The navigation and optical data was recorded on a ROS bag (Open Robotics 2022), a file format in ROS (Stanford Artificial Intelligence Laboratory et al. 2018) for storing ROS messages. Bags store serialised message data so it can be replayed back in the exact way they were recorded. The navigation data included the altitude data and the pose of the robot from the Phins Compact C3 INS, while the optical data consisted of the images from the two Blackfly S cameras. The optical data was recorded in a compressed format to reduce the memory size of each bag recorded onboard the Girona1000's computer. The images were saved with a 90 % JPEG compression percentage in order to maintain a very high-quality image while gaining a significant reduction on the original 100 % file. No further post-processing was applied on the optical data. Upon decompressing the images, the average image size was 4.5 MB.

The Phins Compact C3 INS inside the Girona1000 improves the navigation drift by employing all the navigation sensors in its UKF. The Girona1000 navigation, together with the camera extrinsics (Tables 3 and 4), was employed to calculate the exact position of each of the two cameras at every time instance. In order to assign the correct navigational information to the images, the image timestamps were matched to the navigation timestamps, and thus the corresponding navigational data was assigned to the temporally-closest image.

However, the camera images would not have been usable at the altitude at which the SSS data was captured. The effects of the water column on light propagation (refraction, light absorption, light scattering, and colour shift) cause severe image quality degradation and distortion (Shu et al. 2017). As a consequence, we surveyed the same area that had previously been covered during the acoustic surveys, but at a lower altitude. Additionally, to prevent the inclusion of degraded images in the dataset, we filtered out the images that were captured above 5 m of altitude from the seabed during the post-processing. This resulted in approximately 178 000 images acquired from both cameras.

Table 7. Percentage overlap and angle between stereo pair for each day of the surveys.

Date	Overlap	Angle
22/04/2024	26.34 %	26.62°
25/04/2024	22.73 %	28.62°
26/04/2024	25.2 %	27.26°

5 Processing and Compilation of Training Datasets

5.1 SSS Datasets

Mapping the benthic habitat in mission-time using AUVs necessitates direct processing of the raw sensor readings as they are acquired from the SSS. The raw 12-bit waterfalls were first log-scaled and normalized to the [0,1] range to ensure compatibility with ML pipelines, particularly for training DNNs. We used the natural logarithm ($\log_{1p}$) instead of the conventional base-10 logarithm ($\log_{10}$) to preserve numerical precision in low-intensity regions. This transformation improves the dynamic range compression while avoiding undefined values for zero inputs, which are common in sonar returns. The normalized log-transformed waterfalls were stored as NumPy arrays (Harris et al. 2020). Subsequently, slant range correction was applied to the waterfalls, and the available navigation data was used to geocode each bin. This allowed us to fetch the corresponding annotations from the ArcGIS interpretations and automatically generate the ground truth for the waterfalls.

However, the original annotations on the mosaics were done on a much coarser resolution than required for accurate pixel-wise supervision on the waterfalls, which caused the generated segmentation masks to suffer from issues such as unclear inter-class boundaries and absence of annotations for small-scale objects. Additionally, unknown parameters for the delineation process used for mosaicing further introduced spatial discrepancies. Moreover, due to the sheer volume of the data, it was not feasible to produce well curated annotations for the entire dataset. Therefore, a subset of approximately 750 waterfalls, each composed of 1152 pings, was selected across different sectors that provided a suitable distribution of benthic classes. From this subset, 600 waterfalls originating from sectors *N01*, *N02*, *N04*, *N07*, *N08*, *S01*, *S02* and *S03* comprised the training set, while 150 waterfalls from sectors *N03*, *N05*, *N06*, *N09*, *N10*, *N13*, *N15* and *N16* were set aside for the test set. A Python-based annotation tool[†] was developed to support manual refinement of the transferred segmentation masks. Appendix A provides a detailed breakdown of class-wise area coverage, sector-wise distributions, the approximate percentage of annotated pixels per class, etc.

The SSS records intensities of the acoustic backscatter in a series of parallel swaths of data, stacking successive swaths on top of each other as the AUV continues scanning lines along its direction of motion. The so formed acoustic

[†]<https://github.com/CIRS-Girona/MILF>

waterfalls can, however, get too large to be processed in a single pass for each transect of the survey. Therefore, these waterfalls are usually partitioned in batches of certain lines in order to generate smaller tiles, each with a certain degree of overlap along-track and across-track so as to maintain continuity of information. Each tile can then be treated as an independent input to the classification algorithm and the results can be merged back into a coherent waterfall (Wang et al. 2019; Zhao et al. 2021; Burguera and Bonin-Font 2020; Rajani et al. 2023).

In order to generate the datasets presented in this work, the waterfalls from all surveyed sectors were partitioned in batches of 384 pings to generate tiles of size 384×384 , with a 192px overlap along-track and across-track. This yielded over 950 000 tiles. Figure 6 depicts some examples of these tiles. This large volume of unlabelled data can be leveraged for self-supervised pre-training of classification models, thereby serving as a strong proxy to learn semantically meaningful embeddings of the input. For supervised fine-tuning, we similarly partitioned the corrected segmentation masks obtained from the transferred annotations. This resulted in a approximately 36 000 tiles with corresponding pixel-wise annotated ground truth, of which about 30 000 were designated to the training set, while about 7000 were reserved for the test set. Figure 7 illustrates some examples of the data.

5.2 Multi-modal Dataset

As previously discussed, integrated surveys of the seafloor can substantially benefit from the fusion of information from multiple sensors. However, training ML models on such data poses inherent challenges with respect to the alignment between modalities. In our case, the acoustic surveys were more extensive, covering a significantly large area, whereas the optical surveys were more selective,




Figure 6. Samples of SSS tiles generated from sector N07 waterfalls.




Figure 7. Segmentation masks corresponding to the SSS tiles from sector N07 shown in Figure 6.

targeting regions containing certain classes of interest. Consequently, there was no direct spatial alignment between the optical and acoustic data, with many SSS transects lacking corresponding optical imagery. Moreover, even where both modalities were available, the SSS tiles spanned much larger spatial footprints compared to the field of view of individual optical images.

To address this challenge, we adopted a geo-referencing strategy that associates multiple optical images with their corresponding SSS tiles based on spatial proximity. The resulting opti-acoustic groupings do not require strict pixel-level alignment and can be used to facilitate self-supervised learning of cross-modal feature representations via pretext tasks, such as visual similarity matching, wherein the model learns to associate acoustically and optically similar seafloor characteristics without relying on explicit labels. This strategy not only introduces meaningful optical context into SSS classification but also circumvents the need for manual annotation of optical images. Incorporating such visual context can provide an effective mechanism to reduce class ambiguities in acoustic imagery, particularly in scenarios where distinct benthic classes exhibit similar backscatter characteristics but visually distinct textures or colors.

The opti-acoustic registration process involves three main steps:

1. Calculating the geographical coordinates of SSS tiles
2. Calculating the geographical coordinates of optical images
3. Computing the intersection of the optical and SSS footprints on the seabed

The geographical coordinates $(x_j^i, y_j^i) \in \mathbb{R}^2$ of bin j inside SSS ping i were computed using the following set of equations (assuming a flat seabed):




Figure 8. Examples of SSS tiles from sectors N07 and N08, alongside a subset of their corresponding matched optical images.

$$x_j^i = x_i + \Delta x_j \cos \psi_i + \Delta y_i \sin \psi_i \quad (1)$$

$$y_j^i = y_i - \Delta x_j \sin \psi_i + \Delta y_i \cos \psi_i \quad (2)$$

$$\Delta x_j = \sqrt{\left(\frac{j \cdot r}{n}\right)^2 - \left(\frac{h_i}{\cos \theta_i}\right)^2} \quad (3)$$

$$\Delta y_i = h_i \tan \theta_i \quad (4)$$

where $x_i \in \mathbb{R}$ and $y_i \in \mathbb{R}$ denote the easting and northing coordinates at ping i , ψ_i and θ_i are the yaw and pitch rotations at ping i , $h \in \mathbb{R}$ is the altitude at ping i , n is the number of bins inside ping i , and $r \in \mathbb{R}$ is the slant resolution. The pitch θ_i angle helps determine the width and the offset geographical coordinates of the SSS ping, since a pitching movement will cause the SSS pulses to travel further away.

The top-left, top-right, bottom-right and bottom-left bins of an SSS tile (composed of multiple pings) can be used to calculate the bounds of the spatial footprint of the SSS on the seabed via Equations 1 and 2.

Similarly, the geographical coordinates $\mathbf{p}^W \in \mathbb{R}^3$ of the projection of a pixel $\mathbf{p}^I \in \mathbb{R}^2$ of an image on the seabed are determined by:

$$\mathbf{p}^W = {}^w\mathbf{T}_c \mathbf{K}^{-1} \cdot h \cdot \mathbf{p}^I \quad (5)$$

where $\mathbf{K}$ denotes the camera intrinsic matrix, ${}^w\mathbf{T}_c$ represents the transformation matrix from the world to camera frame. Both the intrinsic and extrinsic matrices are required for this operation, and are stored inside a .json file. This camera calibration metadata file contains dictionary objects where each entry corresponds to a distinct ROS Bag recording. Each dictionary entry stores the ROS topic name and the intrinsic and extrinsic parameters of the left and right cameras used during the optical surveys.

The top-left, top-right, bottom-right and bottom-left corners of an image can be used to calculate the bounds

of the optical image footprint on the seabed via Equation 5 (assuming a flat seabed).

A matching procedure was then performed on the seabed footprints of both acoustic and optical data: if more than 40 % of an image lies within the bounds of an SSS patch, the pair is considered a *match*. The process was repeated for all the SSS tiles and optical images until the entire dataset was constructed. The SSS tiles with no optical matches were ignored.

Once the matching procedure was completed, three metadata files were generated:

- SSS tile metadata (.csv file)
- Optical image metadata (.csv file)
- Opti-acoustic registration metadata (.json file)

The opti-acoustic registration metadata consists of dictionary objects, defined as unordered collections of values with unique keys. Each key represents the index of an SSS tile, and the corresponding value is a list of indices identifying the optical images associated with that tile based on the matching criteria. A summary of the SSS and camera metadata files is provided in Appendix B.

For the final dataset, we reduced the size of SSS tiles to 256×256 pixels to better balance the spatial scale between acoustic and optical modalities, minimizing the discrepancy in seabed footprint covered by each data sample. This resulted in a total of 2080 SSS tiles and 35 947 optical images, yielding an average of approximately 18 optical images associated with each SSS tile. Each SSS tile corresponds to a seabed footprint of approximately 25 m^2 , whereas individual optical images cover between 11 m^2 and 18 m^2 , depending on vehicle altitude and camera configuration. Figure 8 depicts an SSS tile and a few of the associated optical images.

6 Repository Structure

The data is provided in both raw and processed formats so as not to constrain users to a specific application. The data has also been synchronised and timestamped to make it easily accessible. The data is available at <https://dataverse.harvard.edu/dataverse/benthicat>. Exemplar code and detailed instructions on how to process the data are available at <https://github.com/CIRS-Girona/opti-acoustic-preprocessing-tools>.

The data is organised as a collection of individual datasets. The collection currently comprises four datasets:

- **BenthiCat - Raw**, which contains all the unprocessed files, including ROS Bag recordings (together with the associated calibration metadata) from the optical imaging surveys, XTF files from the SSS data, and .xyz files from the MBES data. The SSS and MBES folders are further subdivided into sector-specific subdirectories, each containing the corresponding raw data files for that sector. Additionally, the GIS folder contains the georeferenced mosaics and ArcGIS interpretations. Figure 9 presents an overview of the repository structure.
- **BenthiCat - SSS Pre-training** that contains unlabelled SSS tiles for unsupervised representation learning and self-supervised pre-training tasks.
- **BenthiCat - SSS Seafloor Segmentation** containing the manually annotated SSS waterfalls for supervised fine-tuning. The dataset is divided into train and test subsets, with the test set consisting of independent transects that do not spatially overlap with those in the train set.




Figure 9. Structure of the raw dataset.




Figure 10. Structure of the opti-acoustic dataset.

The dataset contains waterfalls instead of pre-cut tiles to avoid constraining users to a fixed tile size; however, a Python script to generate overlapping tiles from the waterfalls is available in the GitHub repository.

- **BenthiCat - SSS-CAM**, which contains the spatially matched optical and SSS samples. It includes two folders: sonar and camera, along with the three metadata files: sonar.csv, camera.csv and correspondences.json. The sonar and camera folders store all the matched SSS tiles and optical images, respectively. An overview of the repository structure is presented in Figure 10.

7 Conclusion

This work presents a comprehensive multi-modal dataset combining SSS and optical imagery for benthic habitat mapping using AUVs. The dataset includes both SSS-only and opti-acoustic subsets, with detailed metadata, curated annotations, and standardized training and testing splits designed to facilitate benchmarking ML algorithms. In addition to the curated datasets, we provide open access to the raw data, including ROSBAG recordings, XTF files, and MBES outputs, enabling researchers to apply alternative processing strategies or generate task-specific datasets. Furthermore, all data preprocessing and annotation scripts are made openly available for reproducible research and to foster rapid extensibility.

Funding

The study was supported by the Spanish government through projects SIREC (PID2020-116736RB-IOO), ASSiST (PID2023-149413OB-IOO) and IURBI (CNS2023-144688).

Acknowledgements

We would like to acknowledge the Ministry of Climate Action, Food and Rural Agenda of the Catalan government for granting permission to publicly disclose specific segments of the georeferenced data.

Declaration of conflicting interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Author contributions

H. Rajani: Methodology, Software, Data Curation, Formal Analysis, Writing - original draft; V. Franchi: Software, Data Curation, Writing - original draft; B. Valles: Data Curation, Writing - review & editing; R. Ramos: Conceptualization, Project Administration, Methodology, Writing - review & editing; R. Garcia: Project Administration, Supervision, Validation, Writing - review & editing; N. Gracias: Project Administration, Supervision, Validation, Writing - review & editing.

References

- Bernardi M, Hosking B, Petrioli C, Bett BJ, Jones D, Huvenne VA, Marlow R, Furlong M, McPhail S and Munafò A (2022) Aurora, a multi-sensor dataset for robotic ocean exploration. *The International Journal of Robotics Research* 41(5): 461–469. DOI:10.1177/02783649221078612. URL <https://journals.sagepub.com/doi/abs/10.1177/02783649221078612>.
- Burguera A and Bonin-Font F (2020) On-line multi-class segmentation of side-scan sonar imagery using an autonomous underwater vehicle. *Journal of Marine Science and Engineering* 8(8): 557.
- Chesapeake Technology, Inc (2010–2024) SonarWiz [software]. URL <https://chesapeake.tech.com/products/sonarwiz-sidescan/>. Accessed: July 9, 2025.
- Coda Octopus (2013) Motion Sensors Product Page. URL <https://www.codaoctopus.com/other-products/motion>. Accessed: July 9, 2025.
- Du X, Sun Y, Song Y, Dong L and Zhao X (2023) Revealing the potential of deep learning for detecting submarine pipelines in side-scan sonar images: An investigation of pre-training datasets. *Remote Sensing* 15(19). DOI:10.3390/rs15194873. URL <https://www.mdpi.com/2072-4292/15/19/4873>.
- ESRI (2021) ArcGIS Desktop [software]. URL <https://www.esri.com/en-us/arcgis/products/arcgis-desktop/overview>. Accessed: July 9, 2025.
- FLIR (2016) Blackfly S GigE Product Page. URL <https://www.flir.eu/products/blackfly-s-gige/?model=BFS-PGE-27S5C-C&vertical=machine+vision&segment=iis>. Accessed: July 9, 2025.
- Garrido-Jurado S, Muñoz-Salinas R, Madrid-Cuevas F and Marín-Jiménez M (2014) Automatic generation and detection of highly reliable fiducial markers under occlusion. *Pattern Recognition* 47(6): 2280–2292. DOI:https://doi.org/10.1016/j.patcog.2014.01.005. URL <https://www.sciencedirect.com/science/article/pii/S0031320314000235>.
- Harris CR, Millman KJ, van der Walt SJ, Gommers R, Virtanen P, Cournapeau D, Wieser E, Taylor J, Berg S, Smith NJ, Kern R, Picus M, Hoyer S, van Kerkwijk MH, Brett M, Haldane A, del Río JF, Wiebe M, Peterson P, Gérard-Marchant P, Sheppard K, Reddy T, Weckesser W, Abbasi H, Gohlke C and Oliphant TE (2020) Array programming with NumPy. *Nature* 585(7825): 357–362. DOI:10.1038/s41586-020-2649-2. URL <https://doi.org/10.1038/s41586-020-2649-2>.
- Huo G, Wu Z and Li J (2020) Underwater object classification in sidescan sonar images using deep transfer learning and semisynthetic training data. *IEEE access* 8: 47407–47418.
- iXblue (2014) GAPS USBL Product Page. URL <https://www.ixblue.com/product-range/ultra-short-baseline-usbl-solutions>. Accessed: July 9, 2025.
- iXblue (2016) Phins Compact C3 Datasheet. URL <https://www.ixblue.com/media-file/9297/ixblue-datasheet-phins-compact-series.pdf>. Accessed: July 9, 2025.
- Kermarrec E, Saman L and Tauvry S (2020) eXtended Triton Format (XTF): File format specification, rev. 42. URL <https://www.ecagroup.com/en/xtf-file-format>. Accessed: July 9, 2025.
- Klein Marine Systems (2021) Side-scan Sonar Product Page. URL <https://www.klein.com/products>. Accessed: July 9, 2025.
- Nguyen KX, Qiao F, Trembanis A and Peng X (2024) SeafloorAI: A large-scale vision-language dataset for seafloor geological survey. In: Globerson A, Mackey L, Belgrave D, Fan A, Paquet U, Tomczak J and Zhang C (eds.) *Advances in Neural Information Processing Systems*, volume 37. Curran Associates, Inc., pp. 22107–22123. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/274de7d60333c0848f42e18ae97f13e3-Paper-Datasets_and_Benchmarks_Track.pdf.
- Nortek (2015) DVL1000/4000 Product Page. URL <https://www.nortekgroup.com/products/dvl1000-6000m>. Accessed: July 9, 2025.
- Open Robotics (2022) Bags. URL <http://wiki.ros.org/Bags>. Accessed: July 9, 2025.
- Pessanha Santos N, Moura R, Sampaio Torgal G, Lobo V and de Castro Neto M (2024) Side-scan sonar imaging data of underwater vehicles for mine detection. *Data in Brief* 53: 110132. DOI:https://doi.org/10.1016/j.dib.2024.110132. URL <https://www.sciencedirect.com/science/article/pii/S2352340924001045>.
- Quality Positioning Services (2021a) Qimera [software]. URL <https://qps.nl/qimera/#>. Accessed: July 9, 2025.
- Quality Positioning Services (2021b) Qinsy [software]. URL <https://qps.nl/qinsy/#>. Accessed: July 9, 2025.
- Quectel (2014) L86 GNSS module. URL <https://www.quectel.com/product/gnss-l86>. Accessed: July 9, 2025.
- R2SONIC (2022) Sonic 2022-V Product Page. URL <https://r2sonic.com/products/sonic-2022/>. Accessed: July 9, 2025.
- R2SONIC (2024) Sonic 2024-V Product Page. URL <https://r2sonic.com/products/sonic-2024/>. Accessed: July 9, 2025.
- Rajani H, Gracias N and Garcia R (2023) A convolutional vision transformer for semantic segmentation of side-scan sonar data. *Ocean Engineering* 286: 115647.
- Ribas D, Palomeras N, Ridao P, Carreras M and Mallios A (2012) Girona 500 auv: From survey to intervention. *IEEE/ASME Transactions on Mechatronics* 17(1): 46–53. DOI:10.1109/TMECH.2011.2174065.
- SBG Systems (2018) Navsight Ekinox INS Product Page. URL <https://www.sbg-systems.com/systems-surveyors/navsight-marine/>. Accessed: July 9, 2025.

- Sethuraman AV, Sheppard A, Bagoren O, Pinnow C, Anderson J, Havens TC and Skinner KA (2025) Machine learning for shipwreck segmentation from side scan sonar imagery: Dataset and benchmark. *The International Journal of Robotics Research* 44(3): 341–354. DOI:10.1177/02783649241266853. URL <https://doi.org/10.1177/02783649241266853>.
- Shu Z, Wang T, Dong J and Yu H (2017) Underwater image enhancement via extended multi-scale retinex. *Neurocomputing* 245. DOI:10.1016/j.neucom.2017.03.029.
- Sonardyne (2021) Scout Plus USBL Datasheet. URL <https://www.sonardyne.com/wp-content/uploads/2021/07/Scout-USBL-Transceiver.pdf>. Accessed: July 9, 2025.
- Stanford Artificial Intelligence Laboratory et al (2018) Robotic operating system. URL <https://www.ros.org>.
- Valeport (2016) miniSVS Product Page. URL <https://www.valeport.co.uk/products/minisvs-sound-velocity-sensor>. Accessed: July 9, 2025.
- Valeport (2019) SWiFT SVP Product Page. URL <https://www.valeport.co.uk/products/swift-svp/>. Accessed: July 9, 2025.
- Wan E and Van Der Merwe R (2000) The unscented kalman filter for nonlinear estimation. In: *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No.00EX373)*. pp. 153–158. DOI: 10.1109/ASSPCC.2000.882463.
- Wang Q, Wu M, Yu F, Feng C, Li K, Zhu Y, Rigall E and He B (2019) Rt-seg: A real-time semantic segmentation network for side-scan sonar images. *Sensors* 19(9): 1985.
- Yu F, He B, Li K, Yan T, Shen Y, Wang Q and Wu M (2021) Side-scan sonar images segmentation for auv with recurrent residual convolutional neural network module and self-guidance module. *Applied Ocean Research* 113: 102608.
- Zhao X, Qin R, Zhang Q, Yu F, Wang Q and He B (2021) Dcnnet: Dilated convolutional neural networks for side-scan sonar image semantic segmentation. *Journal of Ocean University of China* 20(5): 1089–1096.
- Álvarez Tuñón O, Ribeiro Marnet L, Antal L, Aubard M, Costa M and Brodskiy Y (2024) Subpipe: A submarine pipeline inspection dataset for segmentation and visual-inertial localization. DOI:10.5281/zenodo.10808161. URL <https://doi.org/10.5281/zenodo.10808161>.

A Supplementary Dataset Statistics

This appendix provides supplementary statistics and analyses of the datasets presented in this study. The tables included herein provide detailed information on class distributions, spatial coverage, and additional relevant metrics not reported in the main text. These materials are intended to support research and further exploration of the dataset.

Table A.1 lists the classes selected for annotation within the SSS dataset, along with their assigned IDs, the RGB triplets used for visualizing the annotations, and the percentage of total annotated pixels.

Table A.1. Summary of annotated classes in the SSS dataset.

ID	Class	RGB	% Pixels
0	undefined	(255,255,255)	—
1	breakwaters	(0,0,0)	—
2	artificial_reefs	(57,48,138)	—
3	rocks	(107,90,58)	—
4	corals	(127,14,138)	—
5	sand_ripples	(252,140,4)	—
6	fine_sediment	(235,186,101)	—
7	cymodocea	(5,247,215)	—
8	posidonia	(0,136,0)	—
9	cables	(255,0,0)	—
10	trawl_scars	(223,116,255)	—
11	misc_objects	(140,129,247)	—

Figures A.1-A.3 summarize the distribution of benthic classes across all surveyed sectors. Figure A.2 shows the distribution of each class across all sectors, expressed as the percentage of the total area of each class contributed by every sector, whereas Figure A.3 shows the relative contribution of each class within individual sectors, expressed as a percentage of the total annotated area per sector. Figure A.1, on the other hand, presents the absolute area (in hectares) of each class within every sector.

Figures A.4 and A.5 illustrate the altitudes at which SSS surveys were conducted as box plots. The northern sectors are shown in Figure A.4 and the southern sectors in Figure A.5. Similarly, figures A.6 and A.7 present the corresponding distribution of survey depths, with Figure A.6 for the southern sectors and Figure A.7 for the northern sectors.

Figure A.8 summarizes the class distribution within the multimodal dataset in terms of the percentage of total area covered per class and the number of samples per class, for both optical images and SSS tiles.

Total Area Breakdown (ha)																								
N01	N02	N03	N04	N05	N06	N07	N08	N09	N10	N11	N12	N13	N14	N15	N16	S01	S02	S03	S04	S05	S06	TOTAL		
algae	0	16.12	0	0	0.02	0	0	0	0	0	0	0	0	0	0	0	0	0	2693.87	7.61	19.55	2737.18		
artificial reefs	0.31	0	0.39	0	8.4e-03	0.05	0.02	0.36	0.07	0.21	0.27	0.26	0.43	0.48	0.41	0.06	0.51	0.36	0.31	2.3e-03	2.69	7.17		
breakwaters	21.10	0	22.71	4.7e-03	2.5e-03	9.59	6.78	33.38	17.04	6.0e-04	52.03	179.55	6.9e-03	0.15	7.0e-04	32.04	35.28	4.08	1.27	0	2.27	417.30		
cables	0.02	0	0	0.17	0	0.07	0.06	0.14	0.21	0.33	0.14	0.55	2.84	0.62	0.45	1.36	0.59	1.05	0	0	2.31	11.31		
caves	0	0	0	0	0.49	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.49		
corals	66.94	0	1.41	4.0e-04	40.70	57.04	62.19	8.04	7.03	0.11	0	0	0	0	0	0	0	0	0	0	0	244.85		
cynodoclea	14.20	0	0	684.74	95.72	5.68	25.46	9.96	3.09	1.03	36.69	0	7.78	50.66	100.18	151.76	66.25	28.86	208.37	55.04	43.14	1594.51		
detritus	7.25	1.12	4.87	31.39	75.94	69.14	64.71	4.25	213.35	0.01	0.36	8076.89	5.33	47.92	247.75	7.93	17.63	772.24	139.37	2693.35	3503.02	14742.07		
gravel	16.27	3.55	23.17	1.36	30.30	11.84	2.97	0.06	1.39	0	0	0	0	0	0	0.10	32.34	6.74	0.41	4.32	0	134.80		
maerl	0	3.7e-03	0.08	0	1.81	307.90	373.98	3376.88	62.92	14.91	0	0	0	0	0.06	0	0	0.61	0	0	0	4139.15		
misc. objects	0	0	0	12.16	0	2.0e-04	0	8.0e-04	0	0.47	0.08	0.05	0	0.02	0	0.04	136.80	30.09	0	94.84	274.55			
mud	77.60	106.82	8104.49	1684.23	172.17	386.11	580.23	461.66	376.76	1066.99	7578.90	179.02	11717.01	27.51	11452.80	64.23	8471.75	223.03	49393.78	19.50	20.02	1147.50	103312.10	
posidonia A	128.27	23.99	130.07	150.22	21.21	43.30	31.94	87.49	82.37	0	434.54	0.91	0	576.02	132.27	52.02	1445.51	1095.78	0	0	0	4435.90		
posidonia D	0	0	1.30	0	0	0	0	0.03	0	0	0	0	0	0.02	0	0	5.70	0.70	0	0	0	7.76		
rock	407.98	227.21	297.20	163.64	162.75	230.16	199.68	140.28	112.04	27.39	150.57	124.16	0.08	9.12	25.23	19.30	44.69	67.64	500.20	31.38	14.36	0	2955.05	
sand	466.01	80.47	1475.01	1239.03	848.24	750.14	1295.95	14542.36	4.65	34.70	114.48	529.76	3371.57	1312.45	4814.90	92.90	3691.64	665.03	12484.20	2457.30	925.39	672.81	51868.38	
shipwrecks	0.04	3.6e-03	2.1e-03	0.52	0.18	9.0e-04	0.12	0	0.06	6.1e-03	0.07	0	0.03	0.03	3.9e-03	0.16	0.11	0.05	0.16	0.23	7.6e-03	1.78		
TOTAL	1205.97	443.18	10053.73	3990.56	1447.20	1555.81	2580.23	15628.40	4215.33	1219.62	8295.73	9000.44	15276.69	1405.87	17168.49	417.70	12463.21	3314.27	63792.58	8136.00	4529.98	3413.36	189554.35	

Figure A.1. Absolute area (in hectares) of each class within each sector.

Class Distribution across Sectors (%)

	N01	N02	N03	N04	N05	N06	N07	N08	N09	N10	N11	N12	N13	N14	N15	N16	S01	S02	S03	S04	S05	S06
algae	0	0	0.59	0	0	<0.01	0	0	0	0	0	0	0	0	0	0	0	0	0	98.42	0.28	0.71
artificial reefs	4.30	0	0	5.42	0	0.12	0.64	0.22	4.96	0.99	2.92	3.73	3.67	5.93	6.65	5.67	0.82	7.05	4.96	4.39	0.03	37.58
breakwaters	5.06	0	0	5.44	<0.01	<0.01	2.30	1.63	8.00	4.08	<0.01	12.47	43.03	<0.01	0.04	<0.01	7.68	8.46	0.98	0.30	0	0.54
cables	0.14	0	0	1.47	0	0.66	0.53	1.24	1.85	2.90	1.24	4.82	25.07	5.48	3.96	3.69	12.01	5.21	9.32	0	0	20.40
caves	0	0	0	0	0	99.69	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
corals	27.34	0	0.58	<0.01	16.62	23.30	25.40	0.57	3.28	2.87	0.04	0	0	0	0	0	0	0	0	0	0	0
cymodocea	0.89	0	0	42.94	6.00	0.36	1.60	0.62	0.37	0.19	0.06	2.30	0	0.49	3.18	6.28	9.52	4.15	1.81	13.07	3.45	2.71
detritus	0.04	<0.01	0.03	0.18	0.44	0.40	0.37	0.02	1.23	<0.01	<0.01	46.39	0.03	0.28	1.42	0.05	0.10	4.44	0.80	15.47	20.12	8.20
gravel	12.07	2.63	17.19	1.01	22.48	8.78	2.20	0.04	1.03	0	0	0	0	0	0	0	0.07	23.99	5.00	0.30	3.20	0
maerl	0	<0.01	<0.01	0	0	0.04	7.44	9.04	81.58	1.52	0.36	0	0	0	0	<0.01	0	0	0.01	0	0	0
misc. objects	0	0	0	4.43	0	<0.01	0	0	<0.01	0	0.17	0.03	0.02	0	<0.01	0	0	0.01	49.83	10.96	0	34.54
mud	0.08	0.10	7.84	1.63	0.17	0.37	0.56	0.45	0.36	1.03	7.34	0.17	11.34	0.03	11.09	0.06	8.20	0.22	47.81	0.02	0.02	1.11
posidonia	1.45	0.27	9.85	1.70	0.24	0.49	0.36	1.20	0.93	0	4.90	0.01	0	0	6.61	1.49	0.58	53.05	16.87	0	0	0
rock	13.81	7.69	10.06	5.54	5.51	7.79	6.76	4.75	3.79	0.93	5.10	4.20	<0.01	0.31	0.85	0.65	1.51	2.29	16.93	1.06	0.49	0
sand	0.90	0.16	2.84	2.39	1.64	1.45	2.50	28.04	<0.01	0.07	0.22	1.02	6.50	2.53	9.28	0.18	7.12	1.28	24.07	4.74	1.78	1.30
shipwrecks	2.08	0.20	0.12	29.15	10.33	0.05	6.70	0	0	3.16	0.34	3.90	0	1.60	1.47	0.22	9.02	6.31	2.83	8.93	13.04	0.43

Figure A.2. Distribution of each class across all sectors. Values are expressed as the percentage of the total area of each class contributed by every sector.

Class Composition per Sector (%)

	N01	N02	N03	N04	N05	N06	N07	N08	N09	N10	N11	N12	N13	N14	N15	N16	S01	S02	S03	S04	S05	S06
algae	0	0	0.16	0	0	<0.01	0	0	0	0	0	0	0	0	0	0	0	0	0	33.11	0.17	0.57
artificial reefs	0.03	0	0	<0.01	0	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	0.03	<0.01	0.10	<0.01	0.02	<0.01	<0.01	<0.01	0.08
breakwaters	1.75	0	0	0.57	<0.01	<0.01	0.37	0.04	0.79	1.40	<0.01	0.58	1.18	<0.01	<0.01	<0.01	0.26	1.06	<0.01	0.02	0	0.07
cables	<0.01	0	0	<0.01	0	<0.01	<0.01	<0.01	<0.01	0.03	<0.01	<0.01	0.02	0.04	<0.01	0.10	0.01	0.02	<0.01	0	0	0.07
caves	0	0	0	0	0	0.03	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
corals	5.55	0	0.01	<0.01	2.81	3.67	2.41	<0.01	0.19	0.58	<0.01	0	0	0	0	0	0	0	0	0	0	0
cymodocea	1.18	0	0	17.16	6.61	0.36	0.99	0.06	0.14	0.25	0.01	0.41	0	0.55	0.30	23.97	1.22	2.00	0.05	2.56	1.21	1.26
detritus	0.60	0.25	0.05	0.79	5.25	4.44	2.51	0.03	5.06	<0.01	<0.01	89.74	0.03	3.41	1.44	1.90	0.14	23.30	0.22	33.10	77.33	41.85
gravel	1.35	0.80	0.23	0.03	2.09	0.76	0.12	<0.01	0.03	0	0	0	0	0	0	0	<0.01	0.98	0.01	<0.01	0.10	0
maerl	0	<0.01	<0.01	0	0	0.12	11.93	2.39	80.12	5.16	0.18	0	0	0	0	0.01	0	0	<0.01	0	0	0
misc. objects	0	0	0	0.30	0	<0.01	0	0	<0.01	0	<0.01	<0.01	<0.01	0	<0.01	0	0	<0.01	0.21	0.37	0	2.78
mud	6.43	24.11	80.61	42.20	11.90	24.81	22.49	2.95	8.94	87.46	91.36	1.99	76.70	1.96	66.71	15.37	67.98	6.73	77.43	0.24	0.44	33.62
posidonia	10.64	5.41	1.30	3.76	1.47	2.78	1.24	0.56	1.95	0	5.24	0.01	0	0	3.36	31.64	0.42	43.79	1.72	0	0	0
rock	33.83	51.29	2.96	4.10	11.25	14.79	7.74	0.90	2.66	2.24	1.81	1.38	<0.01	0.65	0.15	4.62	0.36	2.04	0.78	0.39	0.32	0
sand	38.64	18.17	14.67	31.05	58.62	48.21	50.21	93.05	0.11	2.84	1.38	5.89	22.07	93.35	28.05	22.22	29.62	20.07	19.57	30.20	20.43	19.71
shipwrecks	<0.01	<0.01	<0.01	0.01	0.01	<0.01	<0.01	0	0	<0.01	<0.01	<0.01	0	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01

Figure A.3. Relative composition of classes within each sector. Values are expressed as the percentage of the total annotated area per sector.




Figure A.4. Box plots of survey altitudes for the SSS campaigns in the northern sectors.




Figure A.5. Box plots of survey altitudes for the SSS campaigns in the southern sectors.




Figure A.6. Box plots of survey depths for the SSS campaigns in the southern sectors.




Figure A.7. Box plots of survey depths for the SSS campaigns in the northern sectors.




Figure A.8. Class distribution within the multimodal dataset for (a) SSS tiles and (b) optical images.

B Metadata for Opti-acoustic Dataset

This appendix provides a summary of metadata corresponding to the SSS tiles and camera images generated during the opti-acoustic matching procedure.

Table B.1. Metadata summary for SSS tiles.

Parameter	Unit	Data Type
Timestamp	s	float
Patch file path		string
XTF file path		string
Start ping index		int
End ping index		int
Start bin index		int
End bin index		int
Easting	m	float
Northing	m	float
Roll	rad	float
Pitch	rad	float
Yaw	rad	float
Longitude	°	float
Latitude	°	float
EPSG code of UTM zone		int
Altitude	m	float
Top-left corner easting	m	float
Top-left corner northing	m	float
Top-right corner easting	m	float
Top-right corner northing	m	float
Bottom-right corner easting	m	float
Bottom-right corner northing	m	float
Bottom-left corner easting	m	float
Bottom-left corner northing	m	float

Table B.2. Metadata summary for camera images.

Parameter	Unit	Data Type
Timestamp	s	float
Image file path		string
ROS Bag file path		string
ROS Topic name		string
ROS Message sequence index		int
Easting	m	float
Northing	m	float
Roll	rad	float
Pitch	rad	float
Yaw	rad	float
Longitude	°	float
Latitude	°	float
EPSG code of UTM zone		int
Altitude	m	float
Depth	m	float
Top-left corner easting	m	float
Top-left corner northing	m	float
Top-right corner easting	m	float
Top-right corner northing	m	float
Bottom-right corner easting	m	float
Bottom-right corner northing	m	float
Bottom-left corner easting	m	float
Bottom-left corner northing	m	float