Towards autonomous exploration in confined underwater environments

Angelos Mallios[*]

Applied Ocean Physics & Engineering

Woods Hole Oceanographic Institution

Woods Hole, MA 02543, USA

amallios@whoi.edu

Pere Ridao

Computer Vision and Robotics Institute

Universitat de Girona

17003, Girona, Spain

pere@eia.udg.edu

David Ribas

Computer Vision and Robotics Institute

Universitat de Girona

17003, Girona, Spain

dribas@eia.udg.edu

Marc Carreras

Computer Vision and Robotics Institute

Universitat de Girona

17003, Girona, Spain

marcc@eia.udg.edu

Richard Camilli

Applied Ocean Physics & Engineering

Woods Hole Oceanographic Institution

Woods Hole, MA 02543, USA

rcamilli@whoi.edu

Abstract

In this field note we detail the operations and discuss the results of an experiment conducted in the unstructured environment of an underwater cave complex, using an autonomous underwater vehicle (AUV). For this experiment the AUV was equipped with two acoustic sonar to simultaneously map the caves' horizontal and vertical surfaces. Although the caves' spatial complexity required AUV guidance by a diver, this field deployment successfully demonstrates a scan matching algorithm in a simultaneous localization and map-

[*]The author can be also reached at Computer Vision and Robotics Institute, Universitat de Girona, 17003, Girona, Spain, amallios@eia.udg.edu.

ping (SLAM) framework that significantly reduces and bounds the localization error for fully autonomous navigation. These methods are generalizable for AUV exploration in confined underwater environments where surfacing or pre-deployment of localization equipment are not feasible and may provide a useful step toward AUV utilization as a response tool in confined underwater disaster areas.

1   Introduction

Robust localization of autonomous underwater vehicles (AUVs) in constrained unstructured environments remains challenging. In most cases, AUVs are required to survey large areas with precise coverage while operating unattended for extended durations. In order to meet these requirements transponders are often deployed in the vicinity of operation forming a long base line (LBL) system in order to improve navigation accuracy, but this introduces disadvantages of additional ship-based survey operations, thereby increasing overall complexity. Nowadays, a more common technique is to install the transponders on the support vessel, forming an ultra short baseline (USBL) system which significantly reduces the complexity of the system, with the cost of reduced accuracy relative to LBL and requiring the support vessel to follow the AUV.

An AUV's trajectory is prone to drift over time at a rate of 0.1 - 3 % or more of distance traveled when operating without the benefit of an externally supplied absolute position reference (Kinsey et al., 2006). For this reason, most AUV operations are limited to open water so that adequate safety margins are available for the vehicle's drift error, and where the seafloor is relatively smooth, with limited obstacles or constraints.

Nevertheless, numerous applications in confined environments are currently served either by remotely operated vehicles (ROVs) and divers, or not served at all. AUVs hold the promise of being safer, more cost effective alternative technology for these applications, but AUV navigation will require new methods for concurrent localization and perception of the survey environment.

For example, offshore oil platforms contain structured underwater environments where inspection and maintenance operations are typically performed by ROVs or divers. Recent research indicates that AUVs are suitable for automating tasks such as infrastructure or ship hull inspection, and basic intervention such as valve handling (Evans et al., 2003; Hover et al., 2007; Marani et al., 2009; Carrera et al., 2013). Although AUV navigation in proximity to offshore infrastructure is a challenge, anthropogenic objects (pipes, pylons, etc.) can be perceived using conventional robotic sensing modalities. Moreover, the infrastructure itself can

be equipped with acoustic beacons and visual features that aid identification and navigation accuracy.

In some natural underwater environments, however, autonomous navigation can become less tractable. Geological formations (e.g., faults, canyon walls, hydrothermal vents) are highly unstructured and rarely resemble simplified geometric shapes that conventional localization algorithms can distinguish and use to aid AUV navigation. For this reason, close proximity surveys are generally performed by ROVs using human-in-the-loop interpretation. In these cases attending surface support ships are required, increasing the infrastructure requirements and complexity of operations. If AUVs are used, these vehicles can conduct survey operations few meters from seafloor if relatively smooth gradient is expected. If not, they operate from extended stand-off distances that relax obstacle avoidance requirements, but sacrifice observing resolution and intervention capability. Extremely challenging environments for autonomous navigation, such as underwater caves, often require complex vehicle trajectories in close proximity to unstructured hazards. Umbilical entanglement is a major concern that limits ROV operation in cave systems, thus human divers remain the dominant explorers of these environments.

Diver-operated survey technologies increase the accuracy of in-situ human exploration and mapping (Stone et al., 2000). Complexity and scale of caves which can be kilometers in length with multiple pathways make surveys very difficult and slow, with numerous dives conducted over several weeks (Farr, 1991). Unfortunately, the limited endurance of human physiology and life support systems (i.e., breathing gases supply) require that cave divers conduct inherently dangerous explorations and many have lost their lives, making cave diving one of the highest risk exploration activities (Fock, 2013).

Like caves, marine wreckage sites also present extreme challenges for underwater navigation. During or following a shipwreck, an initial assessment is carried out as quickly as possible in order to minimize loss of life, stabilize the vessel, and prevent or mitigate environmental damage. In shallow waters, these tasks are performed by divers with entanglement and entrapment risks similar to cave diving. In deeper water ROVs are used for external wreck inspection and limited penetration within the wreckage because of vehicle entrapment and umbilical entanglement hazards. AUVs have the potential for accident response, particularly in early response phases requiring rapid initial assessment within the wreckage without creating additional risk to humans. A common misperception is that once a vessel capsizes or sinks all of its entrapped occupants are immediately killed. Multiple cases over the past years are documented where survivors have been rescued from within submerged wreckage (Gray, 2003; BBC News, 2013).

An AUV capable of safe navigation in complex physically restrictive environments will be an invaluable tool

in industry, science, and society. Improved localization is, however, a necessary next step for transitioning robotic operations into confined underwater environments. This paper presents results from an experimental deployment in an underwater cave system based on a rigorous approach to simultaneous localization and mapping (SLAM) in confined and unstructured underwater environments.

2 Background and closely related work

During the last 20 years, AUVs have become one of the standard tools for underwater exploration. Modalities using optical sensors (Marks et al., 1995; Eustice et al., 2005) and acoustic sensors (Lucido et al., 1996; Paduan et al., 2009) are applied to dam inspections (Ridao et al., 2010), marine geology (Escartín et al., 2008) and underwater archaeology (Bingham et al., 2010; Henderson et al., 2013). The majority of commercially available autonomous platforms are currently optimized for side scan sonar and bathymetric multibeam survey operations. Other functionalities, particularly two-dimensional (2D) and three-dimensional (3D) optical mapping, are not yet commonplace although these techniques have been demonstrated extensively in field applications and remain an active area of research (Singh et al., 2004; Richmond and Rock, 2006; Ferrer et al., 2007; Pizarro et al., 2009; Nicosevici et al., 2009; Johnson-Roberson et al., 2010; Inglis et al., 2012).

Maps constructed by AUVs can be used to improve vehicle navigation estimation, either in real-time or in post-processing, through the application of SLAM. This technique relies on filtering algorithms which merge the noisy sensor measurements with information from a kinematic or a dynamic motion model of the system. SLAM was first addressed in a probabilistic framework in the early 1990's with the seminal work from Smith et al. (1990). Since then, a significant amount of research has been conducted and a number of algorithms have been proposed to formulate and solve the problem with notable achievements predominantly in land mobile robotics. Filtering algorithms commonly used for SLAM can be separated into Gaussian filters, such as extended Kalman filters (EKFs), extended information filters (EIFs), and non-parametric filters such as particle filters (PFs). A method that has recently gain popularity is the pose-based graph (aka Graph-SLAM), which offers a natural representation for solving the full SLAM problem. Pose-graphs are abstract representations of an optimization problem consisting of pose nodes, which represent the variables to be estimated and edges that represent constraints on the variables they connect. We refer the reader to Durrant-Whyte and Bailey (2006) and Bailey and Durrant-Whyte (2006) for a comprehensive review, and Paull et al. (2013) for a review focused in the underwater domain.

In the majority of SLAM algorithms, features are recognized and extracted from the images (optical or acoustic) represented as points or lines. For example, Ribas et al. (2008) represent the underwater walls of a marina as line features and use a robust voting algorithm that extracts these line features, together with their uncertainty, from the continuous sonar data flow of a mechanically scanned imaging sonar (MSIS). This technique then fuses the resulting information into a feature-based EKF-SLAM algorithm incorporating sub-mapping techniques. Nevertheless, in a natural underwater environment it may be difficult to recognize features via sonar, as their appearance can change significantly, depending on the point of view, making it extremely difficult to extract robust features. For this reason, a number of researchers have focused their efforts on using featureless methods, such as occupancy grids and scan matching.

AUV navigation in open waters often relies upon generally accepted solutions, and research efforts are now focusing on more complex environments such as man-made structures (Ridao et al., 2014) and ship hull inspections (Hover et al., 2007). Walter et al. (2008) describes a SLAM implementation for ship hull inspection using forward looking sonar (FLS) data from a highly maneuverable, hovering AUV. This method applies an exactly sparse extended information filter (ESEIF) algorithm to perform SLAM based upon features manually selected within FLS images. To automate the process, Hurtós et al. (2013) proposed the use of Fourier-based techniques to register 2D images from a multibeam FLS. The estimated registrations are used within a pose-graph optimization framework to obtain a globally consistent trajectory and to render acoustic mosaics of underwater environments.

Under ice regions is another example of complex overhead environment for AUV operations. An AUV must be able to return to its recovery location because aborting and surfacing while under ice is not a vital option. Moreover, in Arctic regions the ice is often moving preventing recovery locations from being determined in advance, so the AUV must be able for re-localize a new recovery location in real-time. To aid navigation, acoustic localization systems, such as LBL, are used (Jakuba et al., 2008; Plueddemann et al., 2012). Since the first reported Arctic AUV deployment in 1972 (Francois and Nodland, 1972), there have been numerous missions with many notable examples. The Theseus AUV was successfully laid autonomously 175 km long fiber-optical cable (Thorleifson et al., 1997). Maridan Martin 150 AUV gathered the first side-scan imagery of the underside of the ice (Wadhams et al., 2004), and the Autosub-II AUV obtain the first under-ice multibeam measurements (Wadhams et al., 2006). In the ALTEX project an AUV capable of ranges between 1500 and 3000 km and with a depth rating of up to 4500 m, surveyed the water column under the Arctic (Bellingham et al., 2008). The Puma and Jaguar AUVs have explored hydrothermal vents in the Arctic ocean (Kunz et al., 2009). Explorer AUV collected under-ice bathymetry for 12 days without

surfacing (Kaminski et al., 2010). On the other side of the earth, ENDURANCE AUV successfully surveyed an ice covered lake in Antarctica (Stone et al., 2010; Gulati et al., 2010).

Despite these advances, very few underwater robots have been operated in confined and overhead environments where external navigation aid is not feasible. One notable exception is a modified REMUS AUV for internal inspection of New York City's 72 km long, 4 m diameter aqueduct (Vision Systems Design, 2004; Stokey et al., 2005). This inspection survey relies on acoustic ranges from two Doppler velocity logs (DVLs) that enable the vehicle to stay in the center of the aqueduct while it travels from one end of the pipe to the other end, in an assumed obstacle free environment. Survey imagery is generated by digital still cameras positioned around the vehicle that provide a 360 degree field of view around the vehicle's main axis. In another application, White et al. (2010) tested different mapping and localization techniques across horizontal distances of less than 30 meters in ancient Maltese cisterns with an MSIS-equipped ROV. Among the tested techniques are: sonar image mosaics using stationary sonar scans, SLAM while the vehicle was in motion, SLAM using stationary sonar scans, and localization using previously created maps. However, the resulting sonar mosaic was produced by manually registering the scan images. In the same expedition, McVicker et al. (2012), used a combination of particle filter and least squares (LS) to estimate the ROV pose and to feed an octree-based occupancy grid map.

Pioneering work in natural confined environments was reported by Stone et al. (2000) mapping an underwater cave system (Wakulla Springs, Florida). Although the vehicle, called digital wall mapper (DWM), was guided from divers and not designed as a robotic system, it successfully mapped 6.4 km of passages in 3D. The DWM was equipped with 32 narrow-beam sonar transducers located in a spiral array around the nose of the 2 m long, 150 kg (in air) instrument. It used a ping rate of 4 Hz and dead-reckoning (DR) trajectory was estimated from an inertial measurement and a ring-laser gyro unit. To bound the inherent DR drift, custom designed geo-referenced radio beacons were placed every few hundred meters along the path traveled. Diver-portable sonar survey systems with updated technology can allow for accurate mapping, but remain fundamentally limited by the diver's risk and operational endurance.

The *DEPTHX* AUV (Stone, 2007), which evolved from the DWM, was designed to explore flooded caverns and tunnels, and to collect water column, wall samples, and environmental data. The vehicle used a full navigation sensor suite and an array of 54 narrow-beam sonar transducers that provided range measurements around the vehicle. For pose estimation and map building, Fairfield et al. (2010) implemented a SLAM method consisting of a Rao-Blackwellized particle filter for which the map is stored within a 3D evidence grid. To reduce the memory requirements, it relied on a deferred reference counting octree data structure.

The algorithm was tested on-board the vehicle in real-time in a series of dives in flooded natural cenotes to 318 m depth. Because of the cenotes nature, most of the AUV's trajectory was in the vertical axis where accurate pressure sensors provide absolute measurements, thus minimizing localization drift errors that can occur in horizontal trajectories. Nevertheless, to the best of our knowledge, this experiment remains the only autonomous exploration example in confined natural environments.

In our previous research (Mallios et al., 2014), we presented a pose-based algorithm to solve the full SLAM problem of an AUV localization in a confined, unknown and possibly unstructured environment. This algorithm does not rely on features or any structural information, which is potentially advantageous when applied in natural environments. The method first estimates the local path traveled by the robot while forming the acoustic image (scan) with range data coming from an MSIS, providing position estimates for correcting the distortions that the vehicle motion introduces in the scans. Each new pose of a scan is maintained in an augmented state EKF (ASEKF) and is compared with previous scans that are in the nearby area using probabilistic scan matching. If sufficient scan points overlap, a match between the scans will introduce a constraint between the poses and update the ASEKF. These constraints help to identify and close the loops which correct the entire trajectory, thus bounding the drift. The algorithm was tested with a dataset from the structured environment of a marina and ground truth validation confirmed that this method could bound the mean error to 1.90 m with a variance of $\pm$ 1.09 m, for a 600 m trajectory path.

Here we report the results of this ASEKF SLAM method in the natural environment of an underwater cave system. AUV compactness and long endurance is an overriding consideration for mapping confined overhead environments which can extend in multiple branches with tight cross sections. For this work we used a small size AUV equipped with low power sensors.

3 Field operations setup

In this section we detail the AUV experimental setup for the survey at the underwater cave system, focusing on the AUV systems configuration including the robotic platform itself, its on-board sensors, and key issues of conducting the survey.
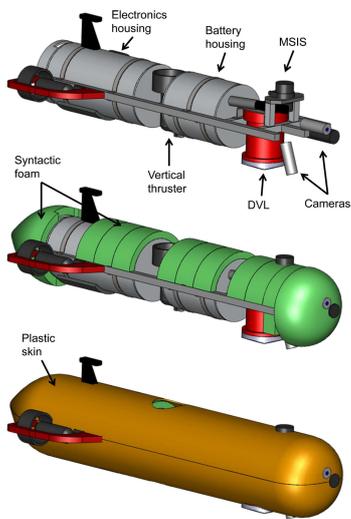
(a) Solid model
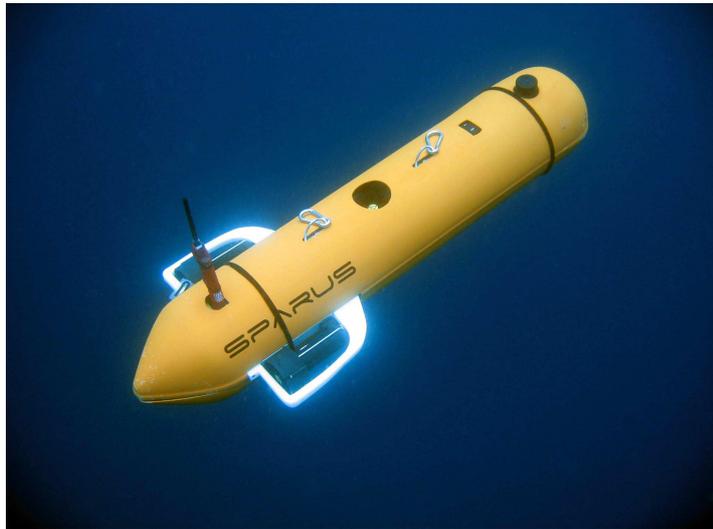
(b) Vehicle being deployed in standard configuration

Figure 1: Solid model of *Sparus* AUV and photo as deployed in standard configuration. The 1.22 m in length by 0.23 m diameter vehicle weighs 25 kg in air and the stored energy allows for 4-5 hours of operations. It is mainly used as research platform for development and testing of new algorithms.

Table 1: Summary of the *Sparus* AUV and sensor suite specifications.

Sensor	Specifications	Value	Description
Vehicle	Depth rating	50 m	UdG in-house design and assembly
	Size	$1.22(L) \times 0.23(D)$ m	
	Mass	25 kg	
	Survey speed	$0.15 - 1.0$ m/s	
	Energy	890 Wh	Li-ion battery pack
	Propulsion	3×80 W	Seabotix BDT150
Navigation	Velocity	$0.2\% \pm 1$ mm/s	LinkQuest NavQuest600Micro
	Pitch/Roll	0.1 deg	Analog Devices ADIS16480
	Heading	0.3 deg	Analog Devices ADIS16480
	Depth	$\pm 0.25\%$	DS2806 HPS-A
Acoustic	**Imaging sonar**		Tritech Micron DST
	Max range:	75 m	
	Horizontal beamwidth:	3 deg	
	Vertical beamwidth:	35 deg	
	Scan rate (360 deg sector):	$5 - 20$ sec	
	Frequency:	Chirped 650 to 750 kHz	
	Profiling sonar		Tritech Super SeaKing DFP
	Frequency:	0.6 | 1.1 MHz	
	Max range:	80 | 40 m	
	Beamwidth:	2 | 1 deg	
	Scan rate (360 deg sector):	$4 - 25$ sec	
Optical	Vision system 1	Forward and down-looking	Analog cameras (PAL)
	Vision system 2	1920×1080 pixels	GoPro Hero2 3D
	Lighting	2×24 W	High intensity discharge (HID)

3.1 Platform: A testbed AUV

The *Sparus* AUV (Fig. 1), was developed in the underwater robotics lab at Universitat de Girona (UdG), Spain, for participating in the Student Autonomous Underwater Challenge - Europe (SAUC-E) 2010. After winning this competition, *Sparus* transitioned to a research platform for development and testing new algorithms as well as for collecting datasets in real environments. It is a small torpedo-shaped vehicle with hovering capabilities and modularity that enables reconfiguration of hardware components. It has three degrees of vehicle motion control and the propulsion consists of three thrusters (two controlling surge and yaw and one controlling heave) which are integrated in its hull. As opposed to most typical single-hull, torpedo-shaped AUVs that rely on forward motion and control surfaces to ascend and descend, the hovering capability of the *Sparus* vehicle allows it to operate close to seafloor and in confined environments. The vehicle dimensions are 1.22 m in length by 0.23 m diameter and weighs 25 kg in air, allowing it to be deployed from a wide variety of vessels, including small boats or from shore.

During the experiment three types of sensors were onboard the *Sparus* AUV: navigation sensors for DR estimation, acoustic sensors for long range perception and mapping, and optical sensors for ground truth validation. The vehicle's sensor suite, as configured for the cave survey, is listed in Table 1. The software architecture of the vehicle is based on the open-source robot operating system (ROS) framework, which provides a structured communication, synchronization, and logging between the vehicle's subsystems and sensors (Quigley et al., 2009).

3.2 Payload: Optical sensors for ground truth

Optical imaging is a common choice for AUV survey operations, but is generally limited to just a few meters (normally <10 m). Although it is practically difficult to capture the full cave scene with vision systems, they can be used for trajectory estimation and ground truth. If the vehicle is keeping constant visual distance from the bottom then visual-odometry algorithms can be applied to verify the trajectory (Furgale and Barfoot, 2010; Scaramuzza and Fraundorfer, 2011). In most cases due to the natural structure of caves, the vehicle will return from the same path and visual features can be identified to close loops in a visual-SLAM framework (Elibol et al., 2010) and estimate the trajectory error.
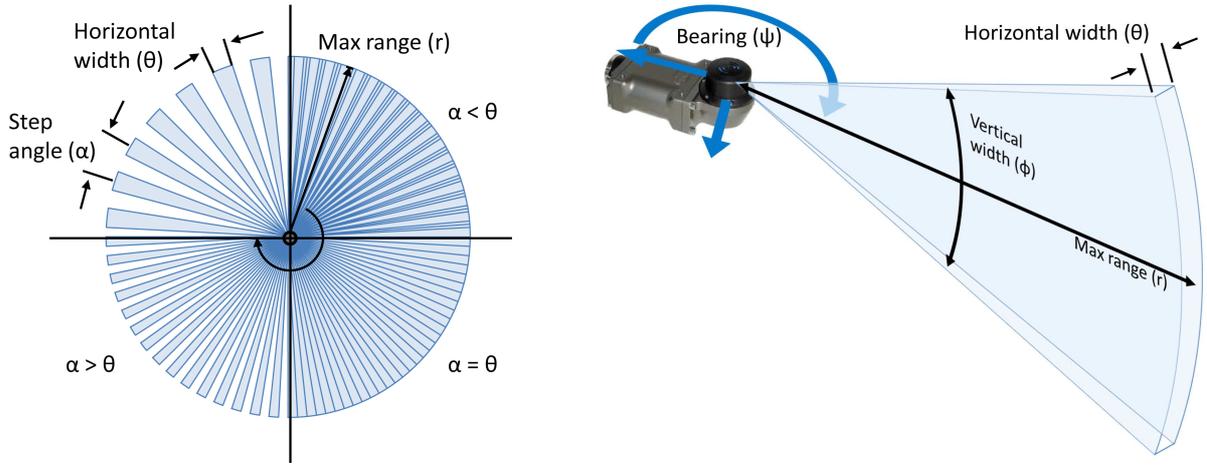
(a) Planar view of beam coverage. Step angle and horizontal width determines beam overlap.

(b) Perspective view of a beam.

Figure 2: MSIS acoustic coverage.

3.3 Payload: Acoustic sensors for long range perception

Because through-water propagation of light is poor (Medwin and Clay, 1998), acoustic sonar sensing (instead of vision or laser sensors) is more appropriate for long range perception and autonomous navigation tasks such as mapping, path planning, and obstacle avoidance. For our research a MSIS was chosen because it enables 360 deg field of view (Fig. 2(a)), wider than commercially available multibeam imaging sonars (typically limited to $\sim$120 deg). Significant drawbacks to MSIS systems are their relatively low resolution and slow rotating mechanical heads.

Although the specifications regarding operating frequency, acoustic beamwidth, frame rate, and the internal beamforming technology depend on the specific sonar model and manufacturer, the principle of operation is the same. First, the sonar insonifies the scene with an acoustic wave, spanning its field of view in the azimuth (θ) and elevation (ϕ) directions, and then the acoustic return is sampled by an array of transducers as a function of range (r) and bearing (ψ) (Fig. 2(b)). Because of the sonar construction, it is not possible to disambiguate the elevation angle of the acoustic return originating at a particular range and bearing. In other words, the reflected echo could have originated anywhere along the corresponding elevation arc. Hence, the 3D information is lost in the projection into a 2D scan image.

For the cave survey presented here, two MSIS were installed on the vehicle, one horizontal and one vertical, which operated simultaneously.

Horizontal scanning sonar   The horizontal scanning sonar is the main sensor used for the SLAM algorithm. As the vehicle moves, it provides the planar views (2D scan points) of the environment that the scan matching algorithm then registers and uses to extract the relative displacement of the vehicle, as described in Mallios et al. (2014). If the scans are sufficiently different, then the scan matching algorithm will be unable to register them. In a natural constraint environment such as the caves, the walls can have sufficiently large variations in vertical relief. For that reason the decreased vertical resolution of a sonar with wide vertical beam-width, can be exploited to aid horizontal scan matching. If the beam is sufficiently wide to overlap at differing vertical displacements, then the intensities per range will be more similar and the resulting scans will be more likely to match. On the other hand, if a sonar with narrow beam had been used, then the scans will highly depend on the vertical position of the vehicle at scan acquisition (Figure 3).

Vertical scanning sonar   In addition to the horizontal scanning sonar, a vertical MSIS was mounted to scan perpendicular to the vehicle's forward motion. The vertical MSIS's high frequency (1.1 MHz) and narrow beam (1 deg) can provide resolution up to 15 mm in order to capture as accurately as possible the 3D shape of the cave. However, because the beam is rotating, the final scan accuracy is correlated with the localization accuracy and the speed of the vehicle.

3.4   Dataset: The underwater caves complex

The dataset was obtained in the underwater cave complex "Coves de Cala Viuda", located in the L' Escala area of Costa Brava, Spain (Lat: 42.103883, Lon: 3.182550) (Fig. 4(a)). The complex consists of three single-branch caves and several tunnels of different sizes and is part of a broader cave complex in the area that is popular among recreational divers (Llamas and Cáceres, 2010).

The survey was performed in June 2012 and was designed to provide quantified acoustic, optical, and inertial measurements for research in autonomous navigation. The dataset was collected with the *Sparus* AUV guided by a diver. It required three dives, all completed in a single day. The first dive was without the AUV for installing ground truth features (described subsequently). The second dive was with the AUV and devoted to the acquisition of an initial dataset in order to test the trajectory and sensor settings, and the third dive was the actual survey. This dataset explores two of the caves and tunnels, closing a loop of approximately 500 m length. The floor of the cave entrances are at approximately 20 m depth and ascend upward, following corridors with diameters varying from 1 to 17 m, to interior surface air chambers at the end of each cave (Fig. 4(b)).
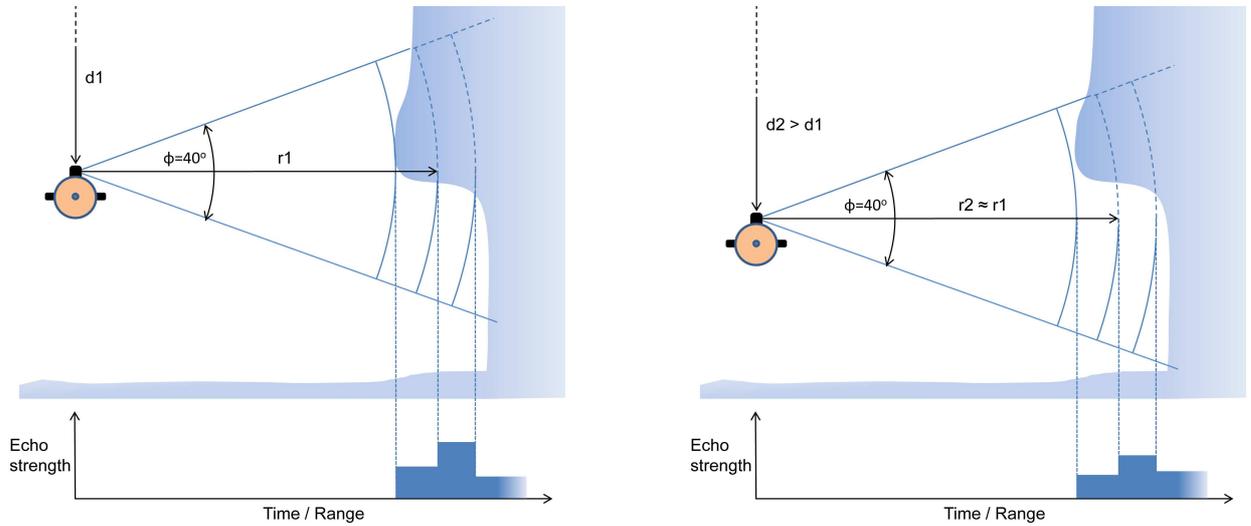
(a) Pencil beam MSIS with 2 deg beam height. Vehicle's depth change can have relatively large effect in output beam pattern.
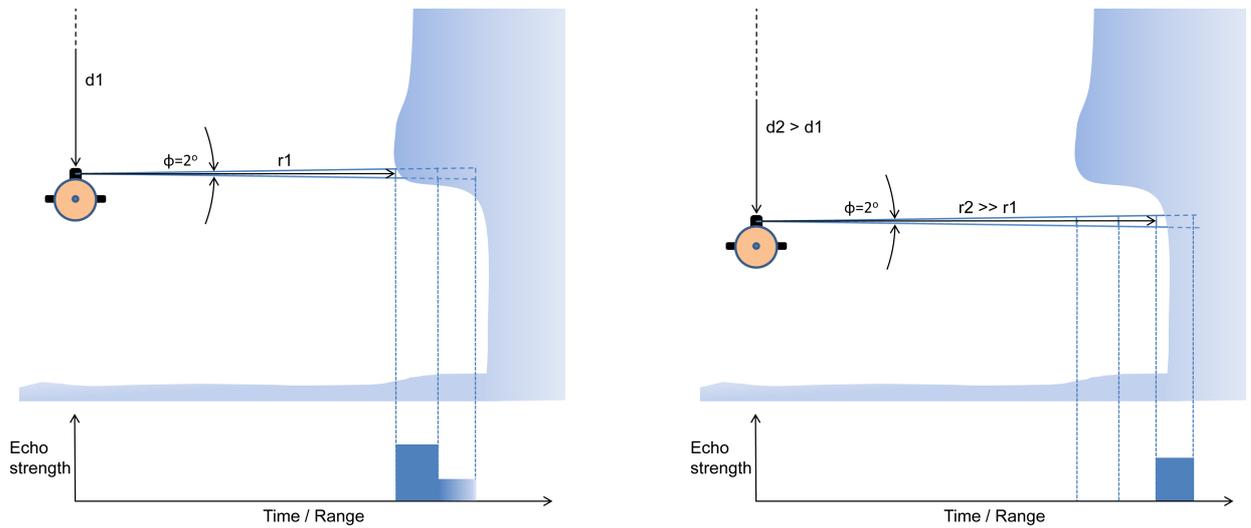
(b) MSIS with 40 deg beam height. Vehicle's depth change has relatively small effect in output beam pattern.

Figure 3: Effect of the sonar beam height in wall with strong relief.

The horizontal MSIS was configured to scan the entire 360 deg sector at 20 m range, with a 0.05 m resolution at 1.8 deg angular step which also defines the sensor's standard deviation. For this survey the payload of the *Sparus* AUV was augmented with additional sensors (Fig. 5). A 3D video system (Hero2, GoPro) together with two high-intensity discharge (HID) lights were mounted downward-looking in front of the main camera, for collecting imagery from the cave floor. The data recording of this unit was uncoupled from the vehicle system, and time synchronization was achieved by comparing specific frames with the main camera which timestamped its frames in the vehicle's software architecture. Finally, in order to capture the 3D shape of the caves the mechanically scanned profiling sonar was mounted in front and parallel to the forward axis of the AUV.The vertical MSIS emits a conical shape beam with 1 deg angle at 1.1 MHz and was configured
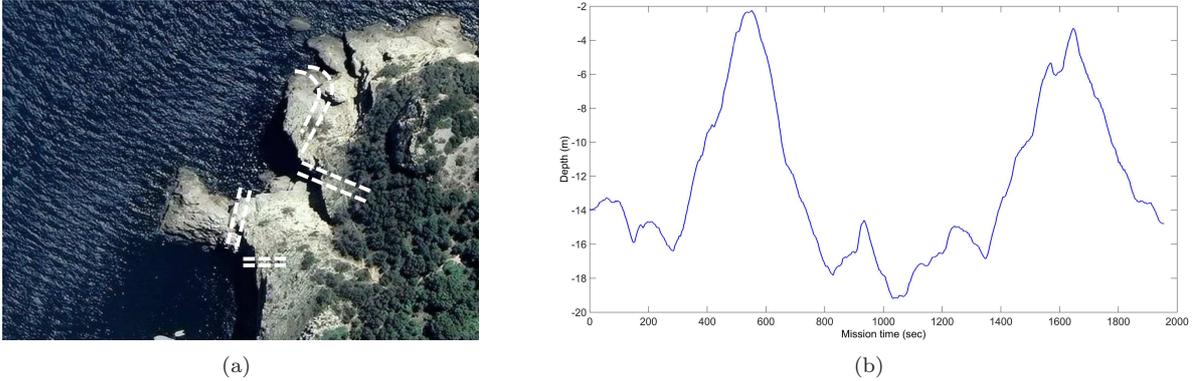
Figure 4: Gathering the dataset: (a) The approximate position of the underwater caves and tunnels are identified by the dotted white lines. (b) The depth profile of the AUV trajectory in this dataset.

to scan the whole 360 deg sector at 10 m range, with a 0.2 m resolution and 1.8 deg angular step. In that configuration, a full sector scan requires approximately 4 sec to complete.

For algorithm validation, the dataset includes six traffic cones as ground truth points (Fig. 6(a)). The cones were placed in strategic locations along the AUV trajectory over which the vehicle passed twice, including the entrances of the caves, and the beginning and end of the survey. The relative position of the cones can be found by extracting the timestamps from the video frames where the cones appear closest to the center of the image and compare them with the timestamps of the trajectory. Then the distance of the cone from the reference frame of the vehicle can be found by using the existing camera calibration (Bouguet, 2004). An additional ground truthing step was included to manually measure and record the distances between the cones with a tape measure. Due to long distances and rock outcrops between the cones, the tape was not always following a straight line and these measurements are the upper bound of unknown variance (Fig. 6(b)). These tape measurements bound scaling error (a problem which is very common with scan matching algorithms in corridor areas) by providing a maximum length scale.

4 Results

Survey data indicate that the AUV's DR estimate, derived from the onboard DVL and the attitude and heading reference system (AHRS), drifted at a mean rate of 2.12 % of distance traveled (based on errors measured at the cones). This drift was sufficient to generate observable inconsistencies in the resulting map (Fig. 7 top), such as the corridor of the top cave and the corridor of the vertical tunnel. Fig. 7 (bottom) shows the trajectory and the map estimated with the proposed SLAM algorithm, where these inconsistencies have been greatly reduced. The SLAM algorithm was able to bound the drift by cross-registering the majority

(a) Scanning planes of the two sonar. No interference was observed due to their different operating frequencies.

(b) Close-up of the sensors.

Figure 5: The *Sparus* AUV extended payload for the cave experiment included: (a) Horizontal MSIS, (b) Vertical MSIS, (c) 3D video camera, and (f) HID lights. Standard payload at the bottom part of the vehicle include: (d) analog camera, and (e) DVL.

of the scans in areas that the vehicle visited more than once. Fig. 8 (top), shows the constraints from the cross-registration between the scan poses. The trajectory and map reconstruction can be seen as a video file in the supplementary material that accompanies this paper (Online Resource 1).

Due to turbidity conditions in most of the trajectory, data from the AUV's video systems couldn't be used for ground truth via visual-odometry or visual-SLAM techniques. Fortunately, the visibility was adequate to record the traffic cones. These ground truth cones provide two metrics for evaluating the performance of the algorithm relative to DR using DVL and AHRS.

The first metric, is the apparent difference in position (based on AUV trajectory estimate) between an initial and a subsequent appearance of a single cone when the vehicle passes over it twice. Table 2 summarizes these errors for the DR and the SLAM trajectory. Both are visualized in Fig. 8 (bottom) and with additional detail in Fig. 9. The maximum error can be observed at cones 1 and 6. Cone 1 is in an open area where there is not enough vertical information for the scan matching algorithm to identify and close the loop. At cone 6 the algorithm successfully cross-registered nearby scans but the relatively large distance from the walls to the vehicle increases the uncertainty of the scans, resulting in larger error. For cones located in well confined areas, the algorithm was able to significantly reduce the error. The overall performance of the SLAM algorithm, compared to DR, decreased the error from 0.5 to 1.0 order of magnitude.

(a)   (b)
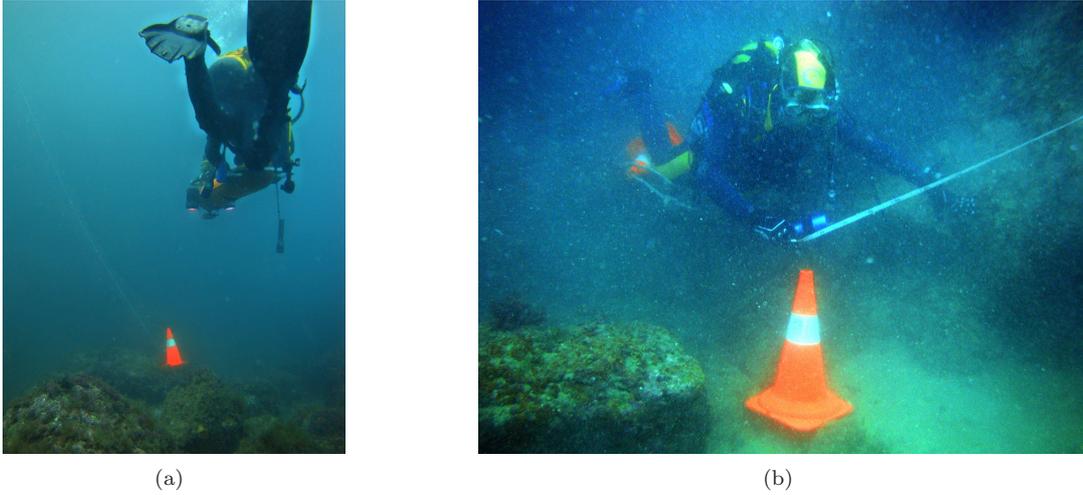
Figure 6: Ground truth points for the cave dataset. In (a) *Sparus* AUV is guided over the cone. (b) Measuring the distances between cones.

Table 2: Cave experiment cones errors (*relatively to distance traveled).

Cone	DR error		SLAM error		Improvement
#	(m)	%*	(m)	%*	%
1	6.60	1.36	2.20	0.45	66.67
2	3.84	1.16	1.68	0.51	56.25
3	2.81	1.05	0.28	0.10	90.04
4	3.53	1.47	0.49	0.20	86.12
5	2.44	3.06	1.40	1.75	42.62
6	4.37	4.59	2.43	2.55	44.39

The second metric is the difference between the cone to cone distances recorded by divers using tape measures, with those estimated from AUV trajectories. As previously described, the diver measurements represent the upper bound of variance due to obstructions preventing the tape from following a straight line. Table 3 summarizes all the measured and estimated distances. All of the AUV trajectory estimated distances are smaller than the diver measured distances, demonstrating that the algorithm does not suffer from a significant scaling problem, which can be very common in corridor areas. Unfortunately, due to the aforementioned tape measure obstruction issues, there is no accurate way of determining which estimated distance is most accurate. However, we can observe that the consistency of measurements between the same pairs using the SLAM algorithm has greatly improved over the DR estimation.

Additional results validating the algorithm are shown in Fig. 10. The top figure is the DR acoustic map which was generated using sonar measurements from the complete dataset. The bottom figure is the acoustic map using the complete dataset and based on the SLAM estimated trajectory. The diffuse appearance of the
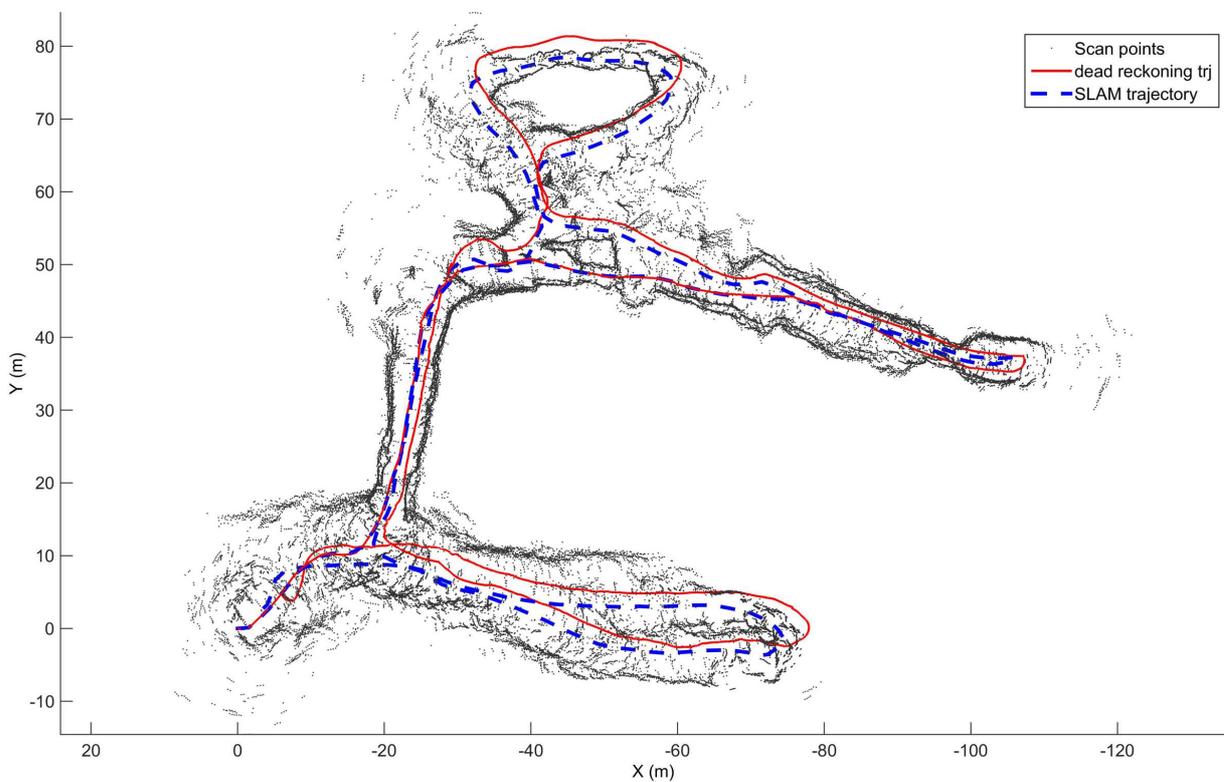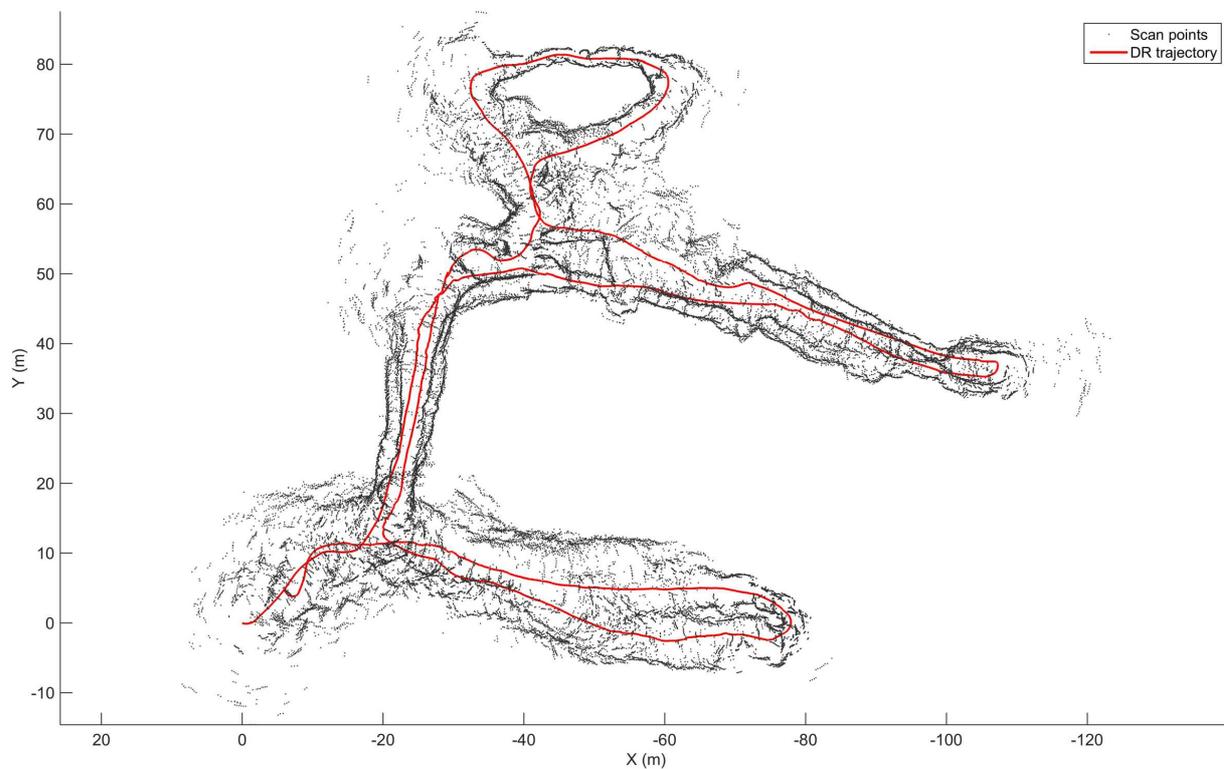
Figure 7: Cave complex trajectories and maps in 2D as seen from above. Top: map based on dead-reckoning estimation. Bottom: map based on SLAM result.

Table 3: Cave experiment error analysis. Cone pairs distance (*see text about ground truth accuracy).

Cone pairs from - to	Ground truth* (m)	DR (m)	SLAM (m)
1 - 2	19	17.07	17.31
2 - 1	19	14.35	16.91
2 - 3	32	31.53	31.00
3 - 2	32	31.38	30.16
3 - 4	16	12.52	12.83
4 - 3	16	13.23	12.87
4 - 5	n.a.	n.a.	n.a.
5 - 4	n.a.	n.a.	n.a.
1 - 6	30	28.36	25.95
6 - 1	30	26.42	26.86

DR image is caused by erroneous trajectory estimates. The SLAM trajectory provides a more self-consistent placement of the measurements which results in a sharper image.

The survey dataset was acquired in approximately 32 min. The SLAM algorithm was implemented as a post-process in MATLAB, requiring approximately 19 min to complete using an Intel Core2 Quad @ 3.00 GHz CPU. The SLAM algorithm's limited computational requirements suggest that it may be appropriate for real time operations under optimized implementation.

As previously described, during the survey the vehicle's vertically mounted MSIS scanned 360 deg perpendicular to the traveling direction. This dataset was combined with the SLAM trajectory estimates to build a 3D map of the cave system.

To accomplish this, the collected scans were first segmented and the range points were extracted using the same process as the horizontal scanning sonar. Next, a cloud of points, (Fig. 11 top), was generated by projecting the scan points over the SLAM trajectory (taking into account depth, roll, and pitch). This cloud of points is the input of the splats distance normalized cut (SDNC) algorithm (Campos et al., 2013) that is used to build the 3D surface model of the cave system. For each input point, a local quadratic surface was constructed using its k-Nearest Neighbors inside a RANSAC procedure (Fig. 11 bottom). These local surfaces were merged in an unsigned distance function evaluated on a tetrahedral grid adapting to the density of the input points. Next, the distance function was signed by using a normalized cut algorithm to segment the volume of the object into its inside and outside. Finally, the surface was extracted at the interface of the two volumes using the restricted Delaunay triangulation surface mesher (Boissonnat and Oudot, 2005).
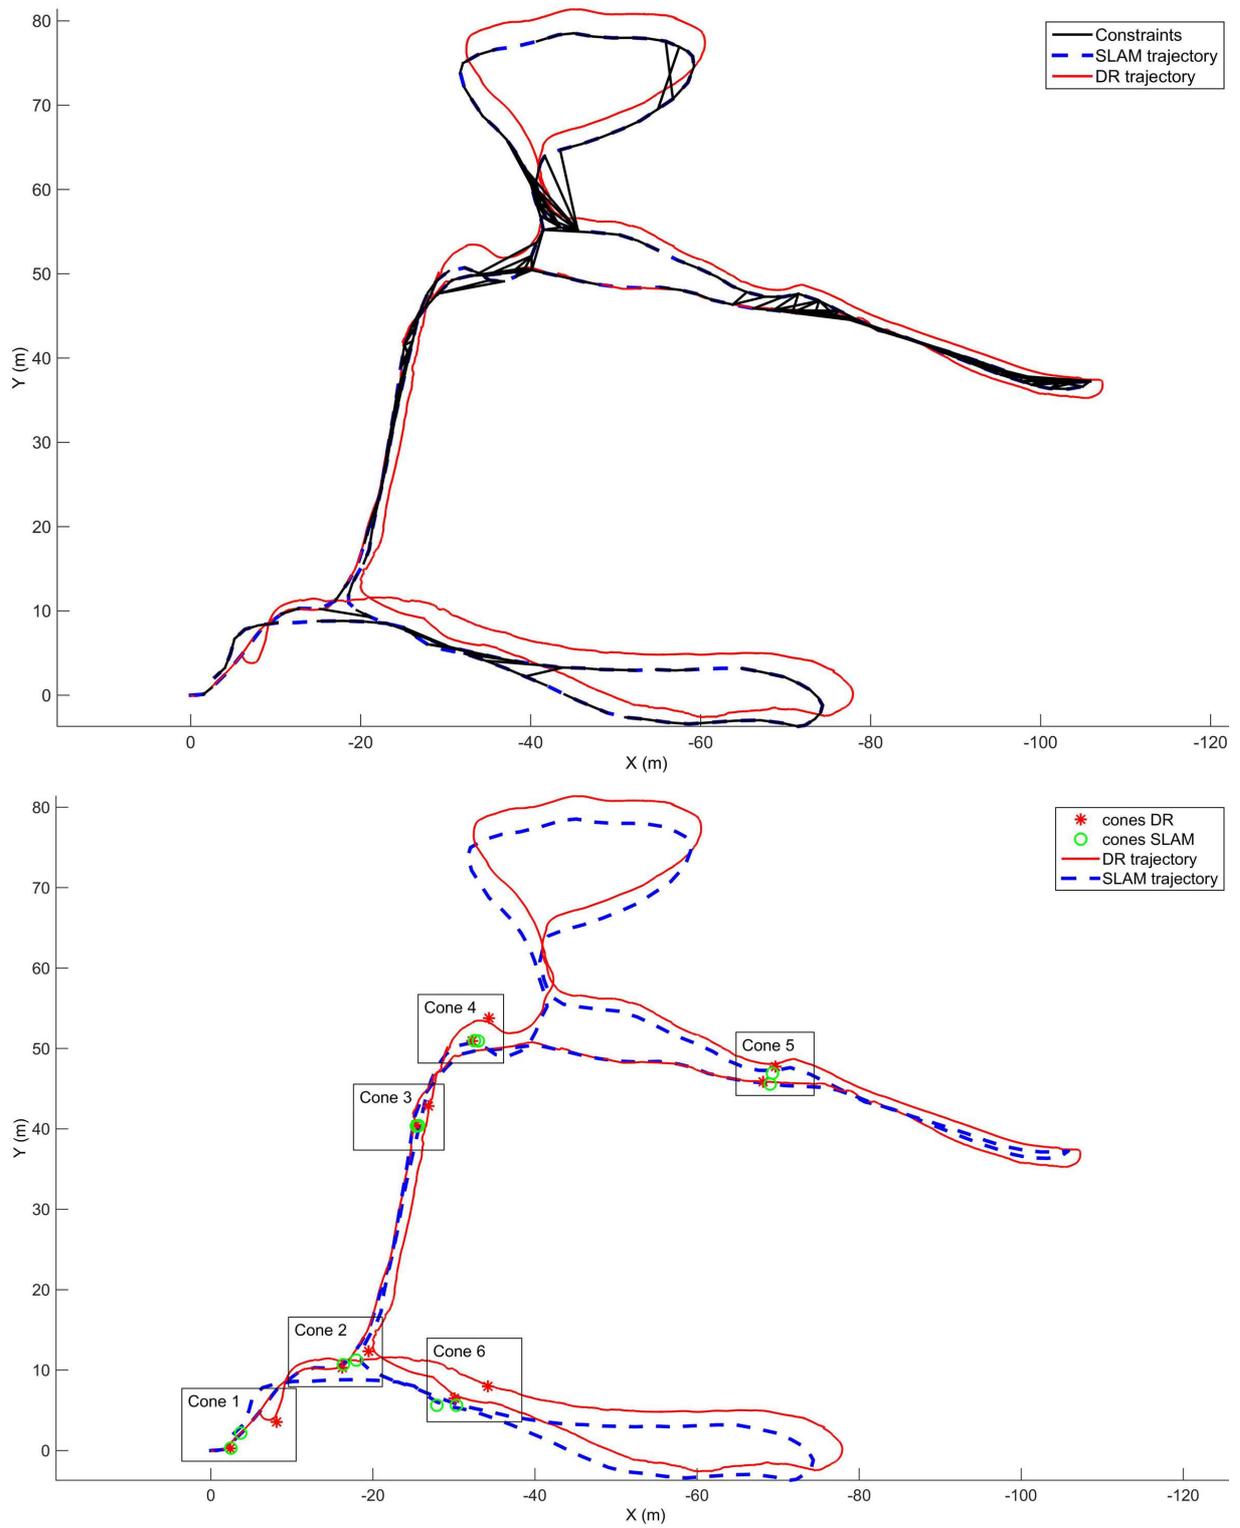
Figure 8: Cave complex map of constraints and ground truth points. Top: map of constraints. Bottom: ground truth points as estimated from dead-reckoning (red asterisks) and from SLAM (green circles).
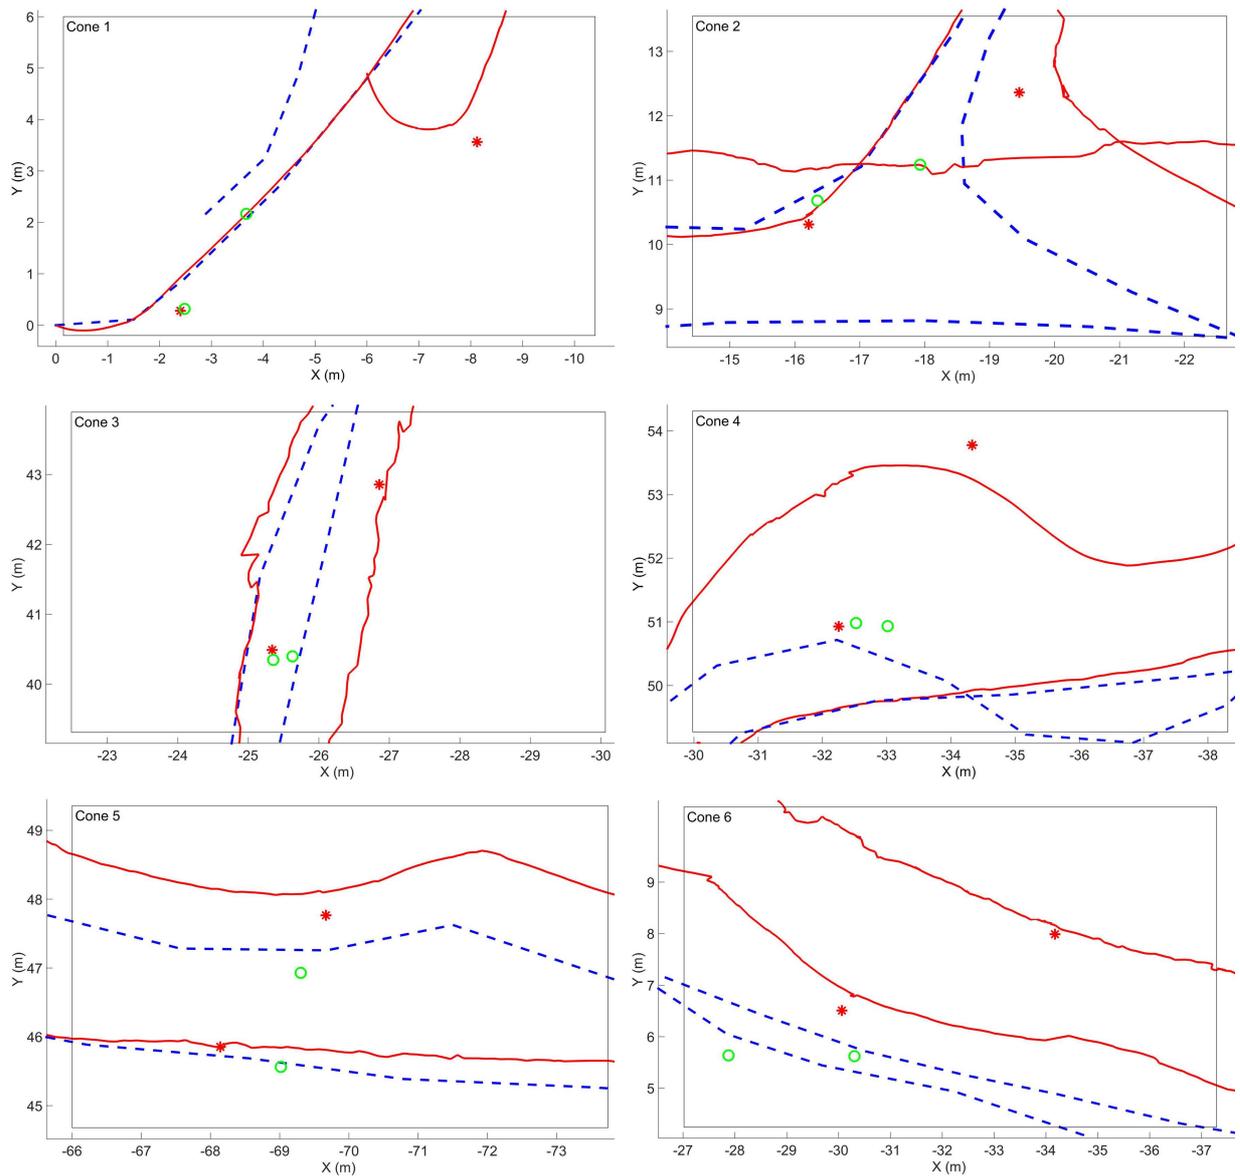
Figure 9: Zoom-in at the cone areas of Figure 8 bottom. Ground truth points as estimated from dead-reckoning (red asterisks) and from SLAM (green circles).
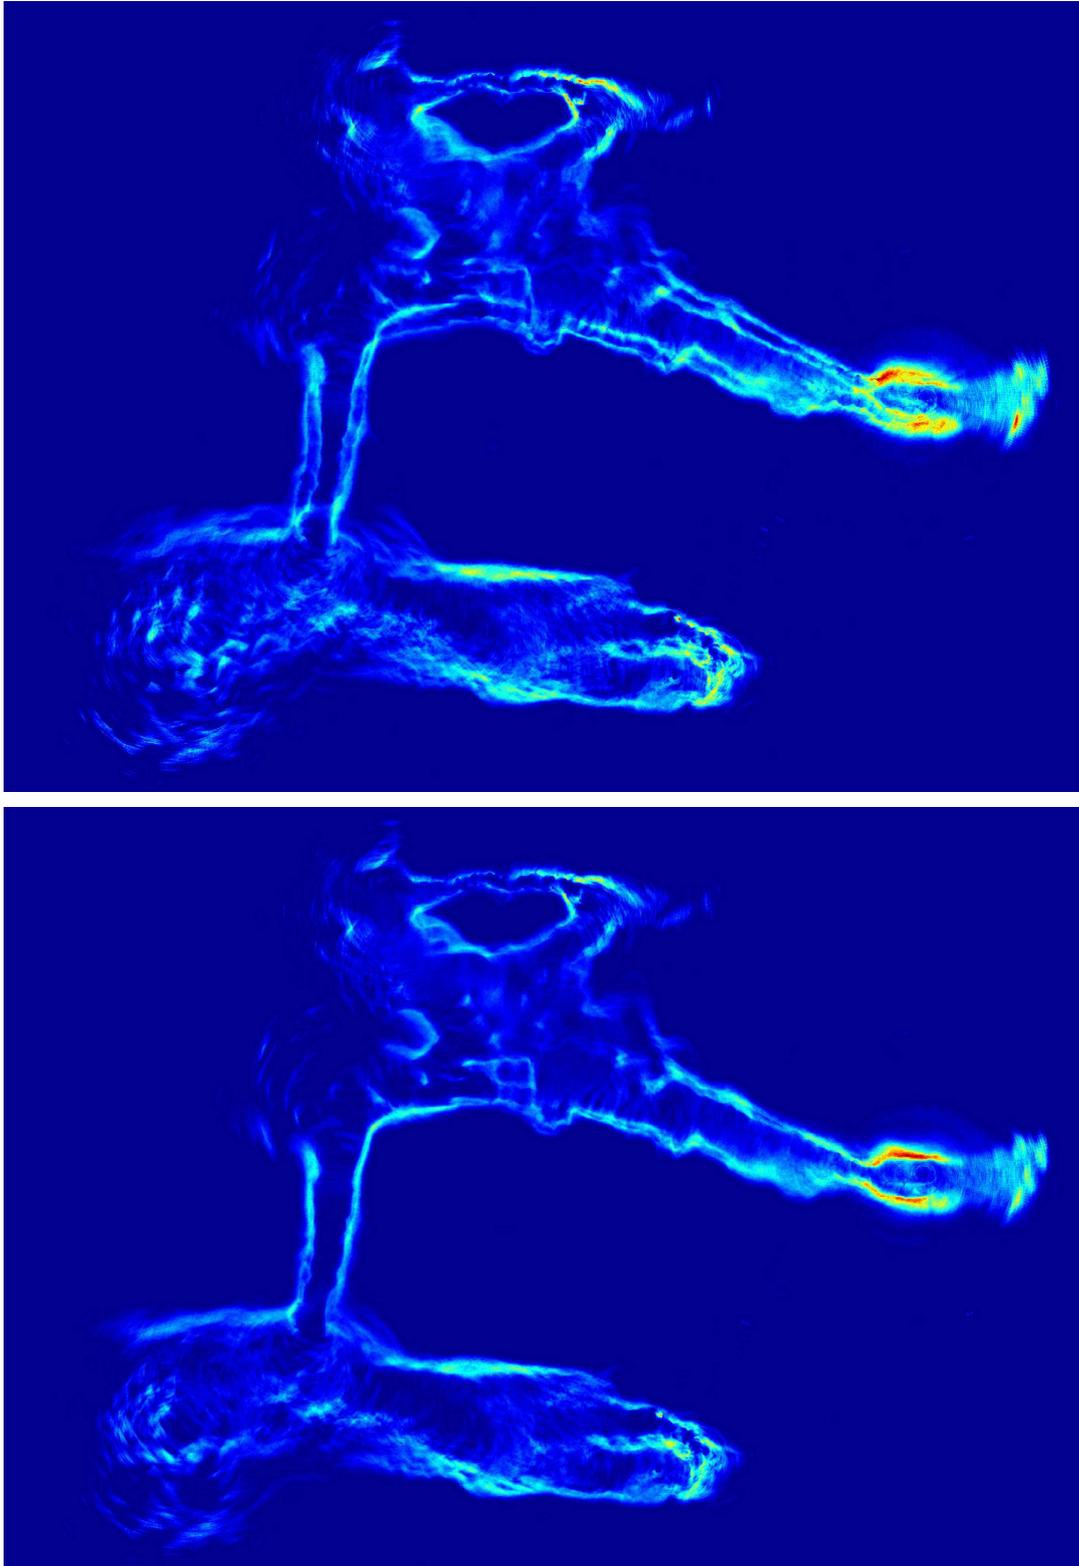
Figure 10: Acoustic maps obtained after an averaged composition of the sonar readings. Top: projected over the dead-reckoning trajectory. Bottom: projected over the SLAM estimated trajectory. Color represents the averaged intensities strength.

It is worth noting that the final three dimensional reconstruction reveals passages that were not visible to the divers due to low water clarity. The cave map can be seen as an interactive 3D figure and video file in the supplementary material that accompanies this paper (Online Resource 2 & 3).

5   Discussion

The experimental methods described here performed satisfactorily as a proof-of-concept in the complex environment of the cave system; however, there remains significant potential for improvement.

Perception and data interpretation.   The horizontal MSIS, the main perception sensor used for localization in this experiment, provides 2D information which can be adequate for obstacle avoidance and path planning in areas where the vehicle explores at constant or gradually changing depths. In more complex environments (e.g. in a dead-end tunnel where the path continues above the vehicle or at an odd angle), the information from a native 3D sensor will be invaluable in order to find the right path. Relevant research in the underwater domain in confined environments is at its infancy, with the lack of low power/size 3D sensors and relevant experiments (Poppinga et al., 2011).

Scans registration.   Scan matching is a well-established technique for registering scans in a wide variety of environments. This technique, in tandem with the horizontal MSIS, performed as intended within the cave system complex where its internal walls provided sufficient data points. However, in areas with insufficient acoustic returns, or self-similar regions scans reconstruction will be inaccurately represented because of the relatively low number of scan points. The result is a high rate of rejected registrations and increased dependence on the DR estimation as an initial displacement step. Thus, the accuracy of the DR can have a direct impact on the final scan matching result. Exploitation of 3D information from the vertical MSIS (or, in principle, a native 3D sonar), can form dense scans that can be used to improve accuracy by minimizing false positive registrations and increasing the scan registration rate, thereby improving overall robustness.

Trajectory estimation.   The pose-based SLAM algorithm used in this experiment is built around an ASEKF stochastic map implementation for estimating scan poses. For our purposes the well-known quadratic expansion problem of the map covariance matrix is partially relaxed because of the relatively low rate of pose additions and short duration of the mission. Recent research in Graph-SLAM (Carlone et al., 2012; Kaess et al., 2012) and octomaps (Hornung et al., 2013) approaches offer promising efficiency gains and data consistency for longer missions.
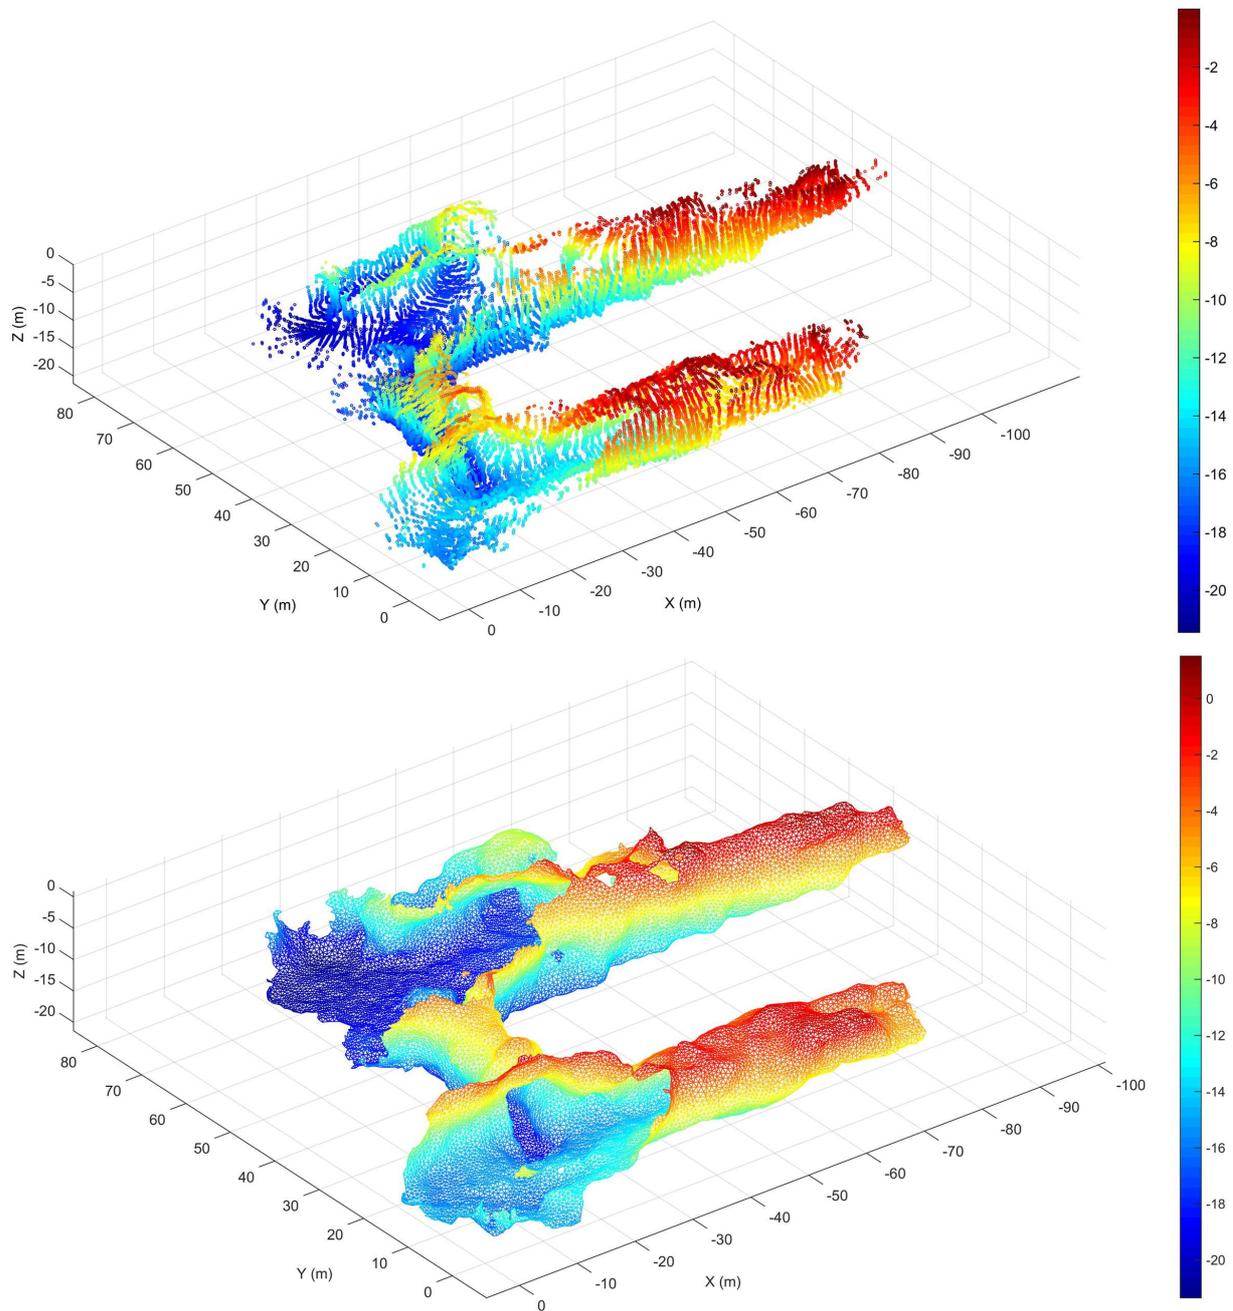
Figure 11: Top: Cloud of points of the cave dataset as extracted from the profiling sonar and projected over the SLAM trajectory. Bottom: Meshed map using the SDNC algorithm. Color map represents depth in meters.

Mapping. In our experimental setup, the map quality is a function of the trajectory estimation quality, vertical scanning sonar resolution, and vehicle speed. Assuming accurate trajectory estimation, speed becomes a limiting factor. Increased sonar resolution and range requires decreased scanning speed (typical rotational rates can be on the order of tens of seconds), hence vehicle speed directly influences point cloud density. Multibeam profiling sonars can provide high resolution with rapid ping rates (on the order of few Hz); however, because of their limited swath angle (∼120 deg) multiple devices operating simultaneously would be needed in order to cover a full sector, significantly increasing the vehicle's dimensions, power consumption, and cost.

Robotic platform. An initial goal for exploration in confined environments would be robotic operations in areas where a human diver can transit. The *Sparus* AUV is well suited for operation at this size scale. However, there will be cases (such as inside a wreck) where the flexibility of a diver's body will permit maneuvering in places difficult or impossible for a rigid AUV. Assuming that navigation and control requirements are solved, there are two possible solutions: a smaller AUV, the size of which may be limited by hardware component requirements; or advancements in an articulating/flexible body AUV.

Ground truth. Acoustic positioning and localization methods are prone to multipath effects and limited to line of sight telemetry. Radio frequency methods have been successfully used in the past (Stone et al., 2000) however it requires specialized hardware, attenuates rapidly over short distances, and is challenging to calibrate. Visual features, such as the traffic cones used here, provide simple point of reference for relative comparison. These markers, however, require prior installation and surveying, and an accurate georeferenced position is not always possible. Accurate ground truthing is an unresolved issue and a unified approach has not yet been developed for all possible deployment cases.

Scene complexity. Navigating an AUV in complex and confined environments represents a step forward in field robotics. Although the cave system in which we contacted our experiment was well explored and mapped over a period of many years (Llamas and Cáceres, 2010), our work has demonstrated the ability to substantially improve the accuracy and spatially resolution of the map in a single 30 min dive. This field experiment offers the possibility to foster research towards autonomous investigation in other confined and potentially hazardous underwater environments, including shipwrecks, debris fields, deep caves, and unexplored natural systems.

6   Summary and conclusions

This AUV experiment represents one of the very few successful field demonstrations in confined underwater environments, and to our knowledge, is the first reported demonstration of horizontal traversal through a cave system, with a trajectory loop closing of more than 500 m. Furthermore, this field experiment demonstrates localization and mapping using a compact AUV equipped with a modest (i.e., low power, small, and low cost) MSIS sonar as the main perception input for posed-based SLAM. This approach using rudimentary AUV hardware is enabled by a localization method that first estimates the local path traveled by the robot while forming the acoustic image (scan) with range data coming from an MSIS. Position estimates are then used to correct the distortions that the vehicle's motion produces in the scans. Next, the corrected scans are cross-registered under a probabilistic scan-matching technique for estimating the displacements of the vehicle, including the uncertainty of the scan-matching result. Successful registrations introduce constraints that an ASEKF use to close loops, estimating a bounded error trajectory. The overall performance of the SLAM algorithm has been evaluated at six ground truth points, and the comparison to DR shown error decrease from 0.5 to 1.0 order of magnitude.

A second MSIS, which scanned perpendicular to the traveling direction, captured cross-sections of the cave system. These scans were projected over the SLAM trajectory estimates and, after meshing, used to build 3D surface model of the cave system.

Results indicate that the trajectory can be estimated in real time and the SLAM algorithm may be useful for autonomous obstacle avoidance and path planning. Moreover, this field experiment provides a new dataset for autonomous underwater navigation research in confined environments where surfacing or pre-deployment of localization equipment is not feasible.

Appendix

Index to electronic supplementary material

Online resource	Media type	Description
1	Video file	SLAM algorithm visualization in the cave dataset
2	Video file	3D cave visualization
3	Interactive PDF	3D cave visualization

Bibliography

Bailey, T. and Durrant-Whyte, H. F. (2006). Simultaneous localization and mapping (SLAM): Part II, state of the art. *IEEE Robotics and Automation Magazine*, 13(3):108–117.

BBC News (2013). Nigerian cook tells how he survived for three days trapped under sunken ship. http://www.bbc.co.uk/news/world-africa-22888828. Retrieved Dec. 18, 2013.

Bellingham, J., Cokelet, E., and Kirkwood, W. (2008). Observation of warm water transport and mixing in the Arctic basin with the ALTEX AUV. In *IEEE/OES Autonomous Underwater Vehicles*, pages 1–5.

Bingham, B., Foley, B., Singh, H., Camilli, R., Dellaporta, K., Eustice, R., Mallios, A., Mindell, D., Roman, C., and Sakellariou, D. (2010). Robotic tools for deep water archaeology: Surveying an ancient shipwreck with an autonomous underwater vehicle. *Journal of Field Robotics*, 27(6):702–714.

Boissonnat, J.-D. and Oudot, S. (2005). Provably good sampling and meshing of surfaces. *Graph. Models*, 67:405–451.

Bouguet, J.-Y. (2004). Camera Calibration Toolbox for Matlab.

Campos, R., Garcia, R., Alliez, P., and Yvinec, M. (2013). Splat-based surface reconstruction from defect-laden point sets. *Graphical Models*, 75(6):346 – 361.

Carlone, L., Aragues, R., Castellanos, J., and Bona, B. (2012). A Linear Approximation for Graph-based Simultaneous Localization and Mapping. *Robotics: Science and Systems VII*, pages 41–48.

Carrera, A., Carreras, M., Kormushev, P., Palomeras, N., and Nagappa, S. (2013). Towards valve turning

with an AUV using Learning by Demonstration. In *Proceedings of the MTS/IEEE Oceans conference, Bergen*, pages 1–7.

Durrant-Whyte, H. F. and Bailey, T. (2006). Simultaneous localization and mapping (SLAM): Part I, the essential algorithms. *IEEE Robotics and Automation Magazine*, 13(2):99–108.

Elibol, A., Gracias, N., and Garcia, R. (2010). Augmented state-extended Kalman filter combined framework for topology estimation in large-area underwater mapping. *Journal of Field Robotics*, 27(5):656–674.

Escartín, J., García, R., Delaunoy, O., Ferrer, J., Gracias, N., Elibol, A., Cufi, X., Neumann, L., Fornari, D. J., Humphris, S. E., and Renard, J. (2008). Globally aligned photomosaic of the Lucky Strike hydrothermal vent field (Mid-Atlantic Ridge, 37°18.5'N): Release of georeferenced data, mosaic construction, and viewing software. *Geochemistry, Geophysics, Geosystems*, 9(12).

Eustice, R., Singh, H., Leonard, J., Walter, M., and Ballard, R. (2005). Visually navigating the RMS Titanic with SLAM information filters. In *Proc. Robotics: Science and Systems, MIT Press, Cambridge, MA*, pages 57–64.

Evans, J., Redmond, P., Plakas, C., Hamilton, K., and Lane, D. (2003). Autonomous docking for Intervention-AUVs using sonar and video-based real-time 3D pose estimation. In *Proceedings of the MTS/IEEE Oceans conference, San Diego*, volume 4, pages 2201–2210.

Fairfield, N., Kantor, G., Jonak, D., and Wettergreen, D. (2010). Autonomous exploration and mapping of flooded sinkholes. *The International Journal of Robotics Research*, 29(6):748–774.

Farr, M. (1991). *The darkness beckons: The history and development of cave diving.* Cave Books.

Ferrer, J., Elibol, A., Delaunoy, O., Gracias, N., and Garcia, R. (2007). Large-area photo-mosaics using global alignment and navigation data. In *Proceedings of the MTS/IEEE Oceans conference, Vancouber*.

Fock, A. W. (2013). Analysis of recreational closed-circuit rebreather deaths 1998–2010. *Diving and Hyperbaric Medicine*, 43(2):78–85.

Francois, R. and Nodland, W. (1972). Unmanned Arctic research submersible (UARS) system development and test report. Technical Report APL-UW 7219, University of Washington.

Furgale, P. and Barfoot, T. D. (2010). Visual teach and repeat for long-range rover autonomy. *Journal of Field Robotics*, 27(5):534–560.

Gray, E. (2003). *Disasters of the Deep: A Comprehensive Survey of Submarine Accidents & Disasters.* Leo Cooper.

Gulati, S., Richmond, K., Flesher, C., Hogan, B. P., Murarka, A., Kuhlmann, G., Sridharan, M., Stone, W. C., and Doran, P. T. (2010). Toward autonomous scientific exploration of ice-covered lakesField experiments with the ENDURANCE AUV in an Antarctic Dry Valley. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 308–315. IEEE.

Henderson, J., Pizarro, O., Johnson-Roberson, M., and Mahon, I. (2013). Mapping submerged archaeological sites using stereo-vision photogrammetry. *International Journal of Nautical Archaeology*, pages n/a–n/a.

Hornung, A., Wurm, K. M., Bennewitz, M., Stachniss, C., and Burgard, W. (2013). OctoMap: an efficient probabilistic 3D mapping framework based on octrees. *Autonomous Robots*, 34(3):189–206.

Hover, F., Vaganay, J., Elkins, M., Willcox, S., Polidoro, V., Morash, J., Damus, R., and Desset, S. (2007). A vehicle system for autonomous relative survey of in-water ships. *Marine Technology Society Journal*, 41(2):44–55.

Hurtós, N., D.Ribas, X.Cufí, Petillot, Y., and Salvi, J. (2013). Fourier-based registration for robust FLS mosaicing in low visibility underwater environments. *Journal of Field Robotics*.

Inglis, G., Smart, C., Vaughn, I., and Roman, C. (2012). A pipeline for structured light bathymetric mapping. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4425–4432. IEEE.

Jakuba, M. V., Roman, C. N., Singh, H., Murphy, C., Kunz, C., Willis, C., Sato, T., and Sohn, R. A. (2008). Long-baseline acoustic navigation for under-ice autonomous underwater vehicle operations. *Journal of Field Robotics*, 25(11-12):861–879.

Johnson-Roberson, M., Pizarro, O., Williams, S., and Mahon, I. (2010). Generation and visualization of large-scale three-dimensional reconstructions from underwater robotic surveys. *Journal of Field Robotics*, 27(1):21–51.

Kaess, M., Johannsson, H., Roberts, R., Ila, V., Leonard, J. J., and Dellaert, F. (2012). iSAM2: Incremental smoothing and mapping using the Bayes tree. *The International Journal of Robotics Research*, 31(2):216–235.

Kaminski, C., Crees, T., Ferguson, J., Forrest, A., Williams, J., Hopkin, D., and Heard, G. (2010). 12 days under ice - an historic AUV deployment in the Canadian High Arctic. In *IEEE/OES Autonomous Underwater Vehicles (AUV)*, pages 1–11.

Kinsey, J., Eustice, R., and Whitcomb, L. (2006). A survey of underwater vehicle navigation: Recent advances and new challenges. In *Proceedings of the 7th Conference on Maneuvering and Control of Marine Craft (MCMC)*. IFAC, Lisbon.

Kunz, C., Murphy, C., Singh, H., Pontbriand, C., Sohn, R., Singh, S., Sato, T., Roman, C., Nakamura, K., Jakuba, M., et al. (2009). Toward extraplanetary under-ice exploration: Robotic steps in the Arctic. *Journal of Field Robotics*, 26(4):411–429.

Llamas, A. and Cáceres, P. (2010). *Guía submarina de les Illes Medes i la Costa del Montgrí.* Anthias.

Lucido, L., Opderbecke, J., Rigaud, V., Deriche, R., and Zhang, Z. (1996). A terrain referenced underwater positioning using sonar bathymetric profiles and multiscale analysis. In *Proceedings of MTS/IEEE OCEANS*, volume 1, pages 417–421.

Mallios, A., Ridao, P., Ribas, D., and Hernández, E. (2014). Scan matching SLAM in underwater environments. *Autonomous Robots*, 36(3):181–198.

Marani, G., Choi, S., and Yuh, J. (2009). Underwater Autonomous Manipulation for Intervention Missions AUVs. *Ocean Engineering. Special Issue: AUV*, 36(1):15–23.

Marks, R. L., Rock, S. M., and Lee, M. J. (1995). Real-time video mosaicking of the ocean floor. *IEEE Journal of Oceanic Engineering*, 20(3):229–241.

McVicker, W., Forrester, J., Gambin, T., Lehr, J., Wood, Z., and Clark, C. (2012). Mapping and visualizing ancient water storage systems with an ROV – an approach based on fusing stationary scans within a particle filter. In *Proceedings of IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pages 538–544.

Medwin, H. and Clay, C. S. (1998). *Fundamentals of Acoustical Oceanography.* Boston: Academic Press.

Nicosevici, T., Gracias, N., Negahdaripour, S., and Garcia, R. (2009). Efficient three-dimensional scene modeling and mosaicing. *Journal of Field Robotics*, 26(10):759–788.

Paduan, J. B., Caress, D. W., Clagueand, D. A., Paull, C. K., and Thomas, H. (2009). High-Resolution Mapping of Mass Wasting, Tectonic, and Volcanic Hazards Using the MBARI Mapping AUV. In *International Conference on seafloor mapping for geohazard assessmen*, pages 181–186.

Paull, L., Saeedi, S., Seto, M., and Li, H. (2013). AUV Navigation and Localization: A Review. *IEEE Journal of Oceanic Engineering*, PP(99):1–19.

Pizarro, O., Eustice, R. M., and Singh, H. (2009). Large area 3-D reconstructions from underwater optical surveys. *IEEE Journal of Oceanic Engineering*, 34(2):150–169.

Plueddemann, A., Kukulya, A., Stokey, R., and Freitag, L. (2012). Autonomous underwater vehicle operations beneath coastal sea ice. *IEEE/ASME Transactions on Mechatronics*, 17(1):54–64.

Poppinga, J., Birk, A., Pathak, K., and Vaskevicius, N. (2011). Fast 6-DOF path planning for Autonomous Underwater Vehicles (AUV) based on 3D plane mapping. In *IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, pages 345–350.

Quigley, M., Conley, K., Gerkey, B. P., Faust, J., Foote, T., Leibs, J., Wheeler, R., and Ng, A. Y. (2009). ROS: an open-source Robot Operating System. In *ICRA Workshop on Open Source Software*.

Ribas, D., Ridao, P., Tardós, J., and Neira, J. (2008). Underwater SLAM in man made structured environments. *Journal of Field Robotics*, 25(11-12):898–921.

Richmond, K. and Rock, S. M. (2006). An operational real-time large-scale visual mosaicking and navigation system. In *Proceedings of the MTS/IEEE Oceans Conference*. IEEE.

Ridao, P., Carreras, M., Ribas, D., and Garcia, R. (2010). Visual inspection of hydroelectric dams using an autonomous underwater vehicle. *Journal of Field Robotics*, 27(6):759–778.

Ridao, P., Carreras, M., Ribas, D., Sanz, P. J., and Oliver, G. (2014). Intervention AUVs: The Next Challenge. In *19th World Congress. The International Federation of Automatic Control.*, pages 12146–12159.

Scaramuzza, D. and Fraundorfer, F. (2011). Visual Odometry: Part I: [Tutorial]. *IEEE Robotics & Automation Magazine*, 18(4):80–92.

Singh, H., Howland, J., and Pizarro, O. (2004). Advances in large-area photomosaicking underwater. *IEEE Journal of Oceanic Engineering*, 29(3):872–886.

Smith, R., Self, M., and Cheeseman, P. (1990). *Estimating uncertain spatial relationships in robotics*, pages 167–193. Springer-Verlag New York, Inc., New York, NY, USA.

Stokey, R., Roup, A., von Alt, C., Allen, B., Forrester, N., Austin, T., Goldsborough, R., Purcell, M., Jaffre, F., Packard, G., and Kukulya, A. (2005). Development of the REMUS 600 autonomous underwater vehicle. In *Proceedings of the MTS/IEEE Oceans Conference*, volume 2, pages 1301–1304.

Stone, W. (2007). Design and deployment of a 3D autonomous subterranean submarine exploration vehicle. In *Conference on Un-manned, Un-tethered Submersable Technology, Durham, NH*.

Stone, W., Hogan, B., Flesher, C., Gulati, S., Richmond, K., Murarka, A., Kuhlman, G., Sridharan, M., Siegel, V., Price, R., et al. (2010). Design and deployment of a four-degrees-of-freedom hovering autonomous underwater vehicle for sub-ice exploration and mapping. *Proceedings of the Institution of Mechanical Engineers, Part M: Journal of Engineering for the Maritime Environment*, 224(4):341–361.

Stone, W. C., Ende, B., Wefer, F., and Jones, N. (2000). Automated 3D mapping of submarine tunnels. In *Robotics 2000: ASCE Conference on Robotics for Challenging Environments*.

Thorleifson, J., Davies, T., Black, M., Hopkin, D., Verrall, R., Pope, A., Monteith, I., Den Hertog, V., and Butler, B. (1997). The Theseus autonomous underwater vehicle. A Canadian success story. In *Proceedings of the MTS/IEEE Oceans Conference*, volume 2, pages 1001–1006 vol.2.

Vision Systems Design (2004). Submersible imaging system inspects NYC water supply. *www.vision-systems.com*, 9(4):9.

Wadhams, P., Wilkinson, J., and Kaletzky, A. (2004). Sidescan sonar imagery of the winter marginal ice zone obtained from an AUV. *Journal of Atmospheric and Oceanic Technology*, 21(9):1462–1470.

Wadhams, P., Wilkinson, J. P., and McPhail, S. D. (2006). A new view of the underside of Arctic sea ice. *Geophysical Research Letters*, 33(4). L04501.

Walter, M., Hover, F., and Leonard, J. (2008). SLAM for ship hull inspection using exactly sparse extended information filters. In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, pages 1463–1470.

White, C., Hiranandani, D., Olstad, C. S., Buhagiar, K., Gambin, T., and Clark, C. M. (2010). The malta cistern mapping project: Underwater robot mapping and localization within ancient tunnel systems. *Journal of Field Robotics*, 27(4):399–411.