

Mapping the Catacombs: An Underwater Cave Segment of the Devil’s Eye System

Michalis Chatzisprou^{a*}, Luke Horgan^{b*}, Hyunkil Hwang^{a*}, Harish Sathishchandra^{b*},
Monika Roznere^c, Alberto Quattrini Li^d, Philippos Mordohai^b, Ioannis Rekleitis^a

Abstract—This paper presents a framework for mapping underwater caves. Underwater caves are crucial for fresh water resource management, underwater archaeology, and hydrogeology. Mapping the cave’s outline and dimensions, as well as creating photorealistic 3D maps, is critical for enabling a better understanding of this underwater domain. In this paper, we present the mapping of an underwater cave segment (the catacombs) of the Devil’s Eye cave system at Ginnie Springs, FL. We utilized a set of inexpensive action cameras in conjunction with a dive computer to estimate the trajectories of the cameras together with a sparse point cloud. The resulting reconstructions are utilized to produce a one-dimensional retract of the cave passages in the form of the average trajectory together with the boundaries (top, bottom, left, and right). The use of the dive computer enables the observability of the z-dimension in addition to the roll and pitch in a visual/inertial framework (SVIn2). In addition, the keyframes generated by SVIn2 together with the estimated camera poses for select areas are used as input to a global optimization (bundle adjustment) framework – COLMAP – in order to produce a dense reconstruction of those areas. The same cave segment is manually surveyed using the MNemo V2 instrument, providing an additional set of measurements validating the proposed approach. It is worth noting that with the use of action cameras, the primary components of a cave map can be constructed. Furthermore, with the utilization of a global optimization framework guided by the results of VI-SLAM package SVIn2, photorealistic dense 3D representations of selected areas can be reconstructed.

I. INTRODUCTION

Underwater caves are crucial for groundwater flow in karst topographies, supplying freshwater to almost 25% of the world’s population [1]. In the US, 23% of freshwater withdrawals in 2000 came from bedrock aquifers, with Florida’s

* The first four authors have contributed equally to this work and are listed in alphabetical order.

^aUniversity of Delaware, Newark, DE, USA,
{michalis,hkhwang,yiannisr}@udel.edu.

^bStevens Institute of Technology, Hoboken, NJ, USA,
{lhorgan,hsathish,pmordoha}@stevens.edu

^cBinghamton University, Binghamton, NY, USA
mrozner1@binghamton.edu

^dDartmouth College, Hanover, NH, USA
alberto.quattrini.li@dartmouth.edu

This research has been supported in part by the National Science Foundation under grants 1943205, 2024541, 2024653, and 2024741. The authors are grateful for the assistance of Greg St Amand, Per Normark, and Evelyn Unti during the data collection. Furthermore, the authors would like to acknowledge the help from Woodville Karst Plain Project (WKPP), El Centro Investigador del Sistema Acuifero de Quintana Roo A.C. (CINDAQ), Global Underwater Explorers (GUE), Ricardo Constantino, and Project Baseline over the years in providing access to challenging underwater caves and mentoring us in underwater cave mapping. The authors are also grateful for equipment support by Halyon Dive Systems, Teledyne FLIR LLC, and KELDAN GmbH lights.



Fig. 1: GoPro setup deployed inside a cave, Ginnie Springs, FL. A fiducial tag can be seen, used for estimating ground truth.

aquifer having the highest withdrawal rate [2]. Given its importance, a project started in 1987 in the Woodville Karst Plain (WKPP), which resulted in the exploration of over 34 miles of cave systems in Florida [3], [4], proving the cave system to be the longest in USA [5]. Furthermore, caves provide a unique preservation environment for archaeological secrets [6], [7], [8], yielding to major historical finds across the globe, such as Mexico [9], [10], [11], [12], Florida [13], Greece [14], and Philippines [15].

Thus, mapping underwater caves is important [16]. Currently, cave divers dive in these extreme environments and map caves with basic tools, such as knotted lines and compass, to obtain the topology of a cave [17]. It is very dangerous even for highly skilled people [18], with more than 350 deaths since 1969 in the USA alone. These environments are still quite challenging for an underwater robot to explore given the absence of any localization or communication infrastructure. Hence, the community would benefit from progress in automated technology tools that can accurately map such environments, but are also easy to use to reduce the cognitive load and preserve the safety of divers.

Mapping underwater caves presents several challenges. LiDAR and RGB-D sensors that provide dense point cloud cannot be used underwater because of the light absorption. The most common mapping sensor used underwater is sonar [19], however, measurements can be inaccurate due to comparatively large beam width and multipath effects. Cameras are promising due to their low cost and amount of scene

information it can capture, but to use them effectively, it is fundamental to overcome limited visibility, color absorption, hazing, absence of natural light, etc. [20].

Some work produced good mapping results with cameras, e.g., using light to infer the surrounding structures [21]; integrating camera, sonar, and IMU to leverage their complementarity [22]. Still, a large dense map of an underwater cave has not been fully demonstrated yet, especially using only tools available to cave divers.

This paper presents a framework for mapping underwater caves using inexpensive action cameras and a dive computer, carried by the divers. We integrate a visual-inertial SLAM framework, which provides a reduced set of images that are informative from the state estimation perspective, together with a structure-from-motion global bundle adjustment to improve geometric consistency and achieve a denser reconstruction. We use datasets collected by cave divers – see, e.g., Fig. 1 – over multiple passes to verify the quality of the reconstruction of the proposed framework and compare it with the manual survey that cave divers perform. The dataset has a long duration of 81 minutes and a covered distance of more than 1km.

Besides presenting this framework, this paper contribution is twofold. First, we provide a unique dataset on a cave system that contains several branches and loops and that we will make public. Second, we discuss open problems from the lessons learned in reconstructing a segment of that cave system, towards advancing technology for divers and developing robust underwater robots.

II. RELATED WORK

While the literature is vast on 3D mapping, ranging from photogrammetry [23] to simultaneous localization and mapping [24] and more recently Neural Radiance Field/Gaussian Splatting [25], we discuss here work specifically related to underwater cave mapping, highlighting the current state of the art and main challenges and gaps.

Traditionally, scuba divers performed the topological mapping of underwater caves using knotted lines or measuring tape for distance, a compass for orientation, and a (water) depth sensor/dive computer for recording the depth. Cartographers would then use the survey information to create depictions of the caves. This process is valuable, as it provides an understanding of how caves expand, but is prone to error. The cave diving community looks forward to technological progress making these surveys easier [17]. In this paper, we integrate surveying methods used by the cave diving community with 3D reconstruction pipelines.

There has long been robot technology that can support cave surveying, even if they are expensive and require expertise in processing the collected data. In 2007, Stone Aerospace presented DEPTHX (DEep Phreatic THERmal eXplorer) [26], a teleoperated robot, equipped with an IMU, two depth sensors, a DVL, and 54 narrow single-beam echosounders. DEPTHX was used to map a vertical shaft filled with water [27]. Later, the same company developed Sunfish [28], an autonomous underwater vehicle with

a multibeam sonar, an underwater dead-reckoning system based on a fiber-optic gyroscope (FOG) IMU, DVL, and pressure-depth sensors. Sunfish was deployed inside some caves in Florida. However, sonar-based reconstruction can be inaccurate in caves due to wide beam width and multipath effects [29].

The advancements in camera technology and their miniaturization resulted in a broadening of their use during cave exploration by divers. However, full 3D reconstruction of underwater caves remains an open problem. Some attempts that specifically considered ease of use for divers included a solution that used the boundary between the scuba dive light and shadow to enhance feature matching, tracking, and 3D reconstruction [21]. A visual-inertial SLAM system, called SVIn2 [30], [31], [22], integrated stereo cameras, an IMU, a depth sensor, and a profiling sonar in a tightly-coupled optimization framework to improve the state estimation reliability in underwater scenarios. Experiments were carried out using a custom stereo rig with such sensors [32], as well as GoPro cameras that have an IMU [33]. The integration of the two systems, light and visual-inertial SLAM, system resulted in a denser and accurate contour-based reconstruction [34]. SVIn2 is the base of our reconstruction pipeline, given its robustness.

Photogrammetry can help in recording marine cave assemblages [35] and achieving high-resolution 3D models [36], but it presents several challenges in achieving highly accurate results due to low visibility, light degradation, and absence of well-defined landmarks. Recently, an approach that combines VIO and a structure-from-motion pipeline demonstrated abilities to reconstruct sections of a cave [37]. Similarly, we integrate the visual SLAM pipeline with a global bundle adjustment system, though we also feed reliable keyframes to achieve more accurate reconstructions.

Generally, compared to other domains, there is a lack of underwater cave datasets. Only one is currently available for general use [38], though there is another for benchmarking VIO/SLAM systems [39]. A dataset for cave visual segmentation has also been recently released [40]. We will contribute a new cave dataset that comprises a large system with several branches and loops.

III. PROPOSED FRAMEWORK

A. Experimental setup

Three GoPro™ Hero9 black action cameras were used. The cameras were rigidly mounted on an aluminum frame with handles; see Fig. 1, and they were carefully positioned to produce a wide, partially overlapping field of view. In particular, the left and right cameras were aligned horizontally but facing approximately thirty degrees off center in opposing directions, while the center camera was pointed approximately 30 degrees upward. Two Keldan™ video lights were connected to the handles to generate the required illumination, as there is no natural light inside the cave. One Shearwater™ Perdix2 AI dive computer (carried often by divers) was utilized to record the depth of the water at the position of the diver, who operated the camera rig.

The video clips from the camera are converted into a ROS1 bagfile using the approach proposed by Joshi *et al.* [33], [41]. This framework encodes the video (at 30 fps) and the IMU (100 Hz) datastream as ROS1 messages with synchronized timestamps. The resulting bagfile can be played back and used as input to a number of VIO/VI-SLAM packages; in our case SVIn2 [22] is used.

In a separate deployment, the MNemo survey data acquisition system [42] (see Fig. 2 left) is used to record data pertaining to the deployed caveline. MNemo records depth and azimuth at each endpoint of a segment and also the length of the segment. The right image in Fig. 2 presents the raw data as recorded from MNemo. Please note that, during deployment, the data acquisition was interrupted due to operator error (blue line) and re-started (red line). The software Ariane [43] was used to calculate loops (light grey dashed lines) and then produce a stick map of the caveline; see Fig. 4 for the processed map.




Fig. 2: MNemo [42]: underwater cave survey data acquisition device. Raw data collected by MNemo with loops manually selected for least square loop closure in the Ariane software [43].

B. Visual-Inertial SLAM – SVIn2

We utilize the Visual Inertial SLAM framework SVIn2 [22] for extracting the camera trajectory and a sparse point cloud. Although the detailed description is beyond the scope of this paper, the main points of this framework are described next.

SVIn2 is a tightly-coupled optimization framework that integrates IMU, camera, water pressure sensor, and sonar, by jointly minimizing error terms for each sensor. SVIn2 operates with a frontend, which provides a trajectory based on the optimization in a local map, and with a backend that supports loop closure for global trajectory optimization of keyframes. A new keyframe is one that maintains a certain overlap with a previous keyframe. A subsequent extension [33] uses the optimized camera poses from the backend to update, in the frontend, the point clouds visible from each camera pose.

At the end, our proposed pipeline extracts the camera trajectory composed of keyframes and corresponding point clouds used in the next steps of the pipeline.

C. Absolute Water Depth Correction

The dive computer (Perdix) provides precise depth measurements throughout the entire dive. In contrast, the SVIn2

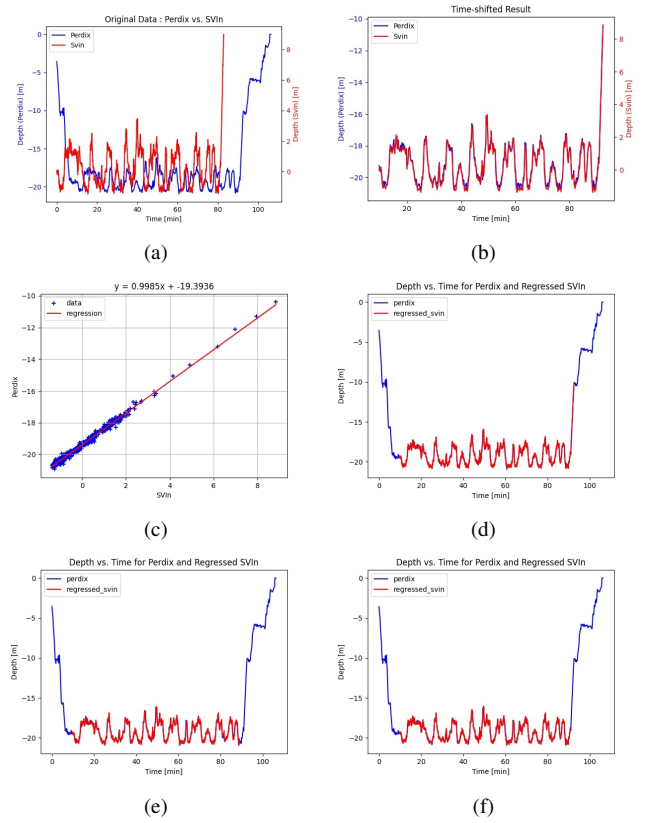


Fig. 3: (a) The SVIn2 z-coordinates (red) and the dive computer depth (blue) before synchronization for the left GoPro dataset. (b) The two time series shifted to a common time but at different scales. (c) Linear regression between the SVIn2 and the Perdix data. (d) The depth estimates time-shifted and with adjusted depth; in red, the SVIn2 produced trajectory; in blue, the data from the dive computer, for the left GoPro dataset. (e),(f) The depth estimates time-shifted and with adjusted depth; in red, the SVIn2 produced trajectory; in blue, the data from the dive computer, for the right and center GoPro datasets respectively.

trajectory starts from a random location and provides relative positioning based on the starting point, with absolute positions remaining unknown [44], [45]. In addition, the two signals are sometimes acquired in different time zones, resulting in significant time drift. Therefore, we adjusted the SVIn2 start time to match the Perdix start time to reduce the time drift to some extent; see Fig. 3(a) for the asynchronous data in the Catacombs-Left dataset. Moreover, the dive computer saves depth data every 10 seconds at a frequency of 0.1 Hz, which is much slower than the frequency of SVIn2 based on the selected keyframes; in our case 2.8 Hz (Catacombs-Left), 3.8 Hz (Catacombs-Right), and 4 Hz (Catacombs-Center).

For every dataset, we already aligned the SVIn2 and Perdix data, thus they have the same start time. In the first step, the depth data is interpolated to 100 Hz to match the same frequency. These interpolated signals were used to determine the time shift values between them by calculating the cross-correlation of the two datasets; see Fig. 3(b) where the Catacombs-Left time sequence is shifted to align with

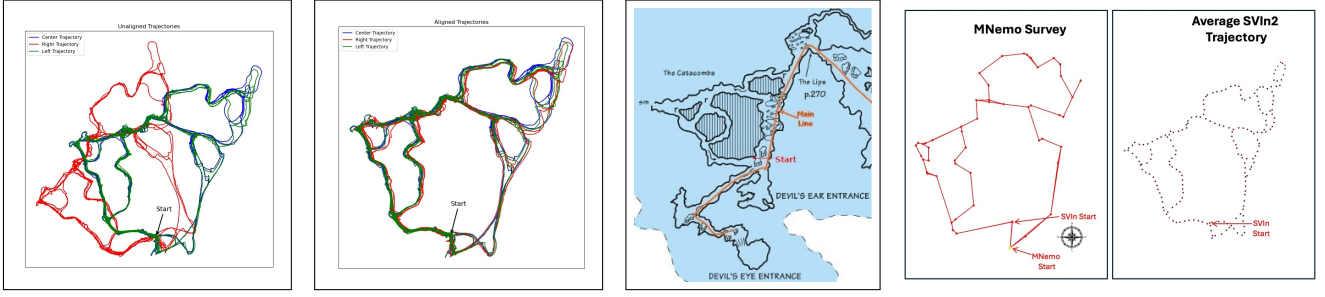


Fig. 4: From the left: The un-aligned trajectories produced from SVIn2 for the left (green), center (blue), and right (red) GoPro. The aligned trajectories. Segment of the cave map with the main line overlayed and the start of the deployment marked. The results from the manual survey using MNemo2 [42] survey tool and Ariane [43] software. Last on the right, the average camera trajectory from SVIn2.

the Perdix time sequence, whereas they still have different scales (Perdix in blue and SVIn2 in red). The following step is to calculate the depth shift values between the two data using linear regression; see Fig. 3(c), which indicates the input data (blue) and the regression result (red). Finally, the SVIn2 depth data are shifted to the real depth based on the values calculated by regression; see Fig. 3(d) where the SVIn2 and Perdix data are aligned to the same depth scale and time in the Catacombs-Left dataset. The same approach is implemented for the Catacombs-Right (see Fig. 3(e)) and Catacombs-Center (see Fig. 3(f)) datasets. For all three datasets, the GoPro (SVIn2) time was shifted by 9.64 min relative to the initial Perdix time. The depth offsets were 19.39 m (Catacombs-Left), 19.4 m (Catacombs-Right), and 19.36 m (Catacombs-Center), respectively. These results reflect the experimental setup, as the three GoPro cameras were aligned in the rig, had similar depths and were set to the same time zone.

D. Trajectory Alignment

For transforming the three trajectories to the same coordinate frame, we utilize a calibration target that was placed at the beginning of the deployment; see the first image of Fig. 8 for a reconstruction of the target inside the cave passage. The aprilgrid package [46] is employed to detect the target in each of the datasets and the relative pose $T_m^{C_i}$ between the camera at time i (C_i) and the target (m). Consequently, by employing the pose of the camera ($P_{C_i}^W$) in world coordinates, we can estimate the pose of the target in world coordinates as ${}^{C_i}P_m^W = P_{C_i}^W \cdot T_m^{C_i}$.

First we estimate the average pose for the target as $\bar{P} = \sum_{i=1 \dots N} \frac{{}^{C_i}P_m^W}{N}$. Then, the distance (D_i) is calculated between the target ${}^{C_i}P_m^W$ for each camera i that viewed the target and the average pose of the target (\bar{P}) as $D_i = \| {}^{C_i}P_m^W - \bar{P} \|^2$. We calculate the standard deviation σ_D of all the distances and then remove all targets with distances more than one σ_D from the mean. Then, we filter outliers based on the orientation estimates, which are more noisy. From the rotation matrix of the pose of the target in the global coordinate system we extract the angles around the principal axes (roll, pitch, yaw) then we remove all the targets with at least one angle (roll, pitch, or yaw) that falls outside one standard deviation. From the remaining (inlier) targets we

estimate an average pose. If there are two trajectories (A and B) the result is two poses $P_m^{W_A}$ and $P_m^{W_B}$ for the same target in two different coordinate systems. The coordinate transformation from World frame B to world frame A is $T_{W_A}^{W_B} = P_m^{W_A} \cdot (P_m^{W_B})^{-1}$ and brings trajectory B in the coordinate frame of trajectory A.

E. Dense Reconstruction – COLMAP

Our pipeline includes a structure-from-motion system, COLMAP [47], run on the keyframes using as prior corresponding poses, identified by SVIn2. This step refines the geometric consistency of the camera poses given that all collections of images are used. The output will be a sparse point cloud. We then densify that point cloud with a depth map estimation using patch-match stereo with per-pixel view selection and fusion.

To obtain the mesh reconstruction, we use a screened Poisson reconstruction, where an implicit function indicator is computed to determine whether a point is inside or outside the model [48].

IV. EXPERIMENTAL RESULTS

A. Dataset

The catacombs section of the Devil's Eye cave system is a narrow passage to the left of the main line as we are entering the cave. The area is characterized by several interconnected narrow tunnels, some of which are accessible only in side-mount configuration¹; see Fig. 4 center for a depiction of the area. Please note that the cave map presented in the center of Fig. 4 is not metrically accurate but provides a generic outline of the area. The catacombs are not marked with a permanent line, however, for the experiments presented here, one was temporarily installed. During a single deployment of one hour, forty-six minutes at an average depth of seventeen meters, the area was mapped using the three cameras described in the experimental setup section. Multiple loops were performed as can be seen in Fig. 4, covering all accessible parts of the catacombs. At

¹In cave diving the passages are characterized as back-mount, where the diver carries double scuba tanks on their back; side-mount, where the tanks are slung on the diver's side, thus creating a narrower vertical profile; and no-mount, where the diver pushes their tanks in front of them, while breathing from them.

the entrance of the passage a calibration target was placed and at the start of the deployment a short calibration motion sequence was performed. In total, data collection lasted for 81 minutes and the total estimated trajectory was more than 1km.

A second dive was performed for manually surveying the cave line deployed inside the catacombs. We started at the mainline (see Fig. 4 center) and followed the line into the catacombs; see the second from right of Fig. 4 where the MNemo start is marked and also, where the VI-SLAM process started. The data collection process was hampered by the current (flow) that is ever-present in the Devil's Eye cave system. Small orientation errors resulted in deviation of the line segments; see Fig. 2, however, the Ariane [43] software optimized the marked loop closures resulting in the map presented in Fig. 4.

B. Sparse reconstruction, Camera Trajectory Generation, Cave Mapping

The SVIn2 [22], [33] package was utilized to extract the camera pose trajectories and sparse point clouds for each of the three cameras (left, center, and right). The un-aligned trajectories can be seen in Fig. 4 on the left. By utilizing the calibration target placed at the beginning of the catacombs section, the three trajectories are aligned in a common frame of reference and adjusted to the correct depth (z-dimension); see Fig. 4. The resulting sparse point cloud representations can be seen in Fig. 5 for the left, center, and right GoPro respectively, while the combined dataset, aligned in the same coordinate frame can be seen in the right image of Fig. 5. Please note that the center camera covers more of the ceiling, while the left and right capture more of the walls. The segment of the main passage near the main line (see Fig. 4 center) is characterized by a high ceiling and wide passage, which was reflected in the lack of detected points in all three cameras.

The three trajectories are used as a guideline for the center of the passage. A pose of the center trajectory is selected and every pose in all three trajectories is flagged. The flagged poses are used to calculate an average pose, and then the next pose of the center camera that has not been flagged is selected. The average trajectory produced can be seen in Fig. 4. The Minimum Spanning Tree (MST) algorithm is used to connect these average poses and from each one the closest point on the left and right of the average pose (at the same depth – z-coordinate) is selected. Fig. 6 presents the average trajectory in green, the left wall in red and the right wall in blue. As can be seen, this naive method is sensitive in the order of traversal for designating left and right boundaries, and the MST does not capture the full connectivity of the passages. We cleaned up manually the data – see the center image of Fig. 6 – and the variability in the passage width is captured. Inside the catacombs, the passages are narrower, while at the main line the passage is wide. In addition, the point above and below the average pose (green) are selected, marking the ceiling (red) and the floor (blue) of the passage; see Fig. 6, where the floor is mapped at an approximately

constant distance from the average camera pose while the ceiling displayed greater variability. Over the years, materials brought by the water accumulate on the floor, smoothing it and filling any holes, while the ceiling erodes over time, generating a lot more 3D structure.

C. Dense reconstruction

From the three cameras, more than forty thousand keyframes were generated. In this paper, we selected six representative areas; see Fig. 7, for dense reconstruction. For each area, we picked one central pose, taken as the pose for a particular image from the center camera. We then found all keyframes, from all three cameras, with poses inside a radius of 2.5 meters, as measured from the central pose, except area 5 where a radius of 5 meters was used. Reconstruction was performed using the global optimization package COLMAP [47], where the camera poses from SVIn2 were used as constraints, as described in the previous section. From the resulting dense Poisson mesh, representative views are displayed in Fig. 8. More specifically, the calibration target is clearly visible at the right side of the passage, while the entry to the catacombs is visible to the left. Areas 2-4 and 6 are inside the catacombs and are characterized by narrow passages. It is worth noting that the floor is relatively flat, while the ceiling presents a more detailed structure; see Area 4 and 6. Area 5 was selected outside the catacombs, in the main passage of the cave. In this relatively open area, distant points are visible in our keyframes. A limited number of pixels cover an extended area, resulting in a very noisy reconstruction. The dense point cloud captures the rocks that exist in the middle of the main passage, as can be seen in Fig. 8.

V. CONCLUSIONS AND FUTURE WORK

We presented a framework for mapping underwater caves, utilizing inexpensive action cameras and a dive computer, available to scuba divers. The proposed pipeline based on VI SLAM and SfM allows for sparse and dense reconstruction of parts of the caves on a cave dataset (which will be publicly released) with several branches and loops extending for more than 1km.

The results showed the ability of the proposed framework to reliably estimate the camera trajectory and achieve dense reconstructions in sections of the cave.

We plan several research directions in the immediate short term based on some of the lessons learned from this paper. While the trajectory looks accurate, achieving a dense map of the whole cave system was unattainable, because there were many outliers preventing a good reconstruction. In addition, the occupied memory raised significantly, making it challenging to manage even on a powerful workstation with 256GB. We plan to investigate strategies to iteratively perform dense reconstruction in a window. We also plan to add to the dataset a 3D sonar released by WaterLinked that has a narrow beam and provides a 3D point cloud.




Fig. 5: The left, center, and right GoPro cameras, sparse point cloud produced by SVIn2. On the right, the combined point cloud from all three cameras.

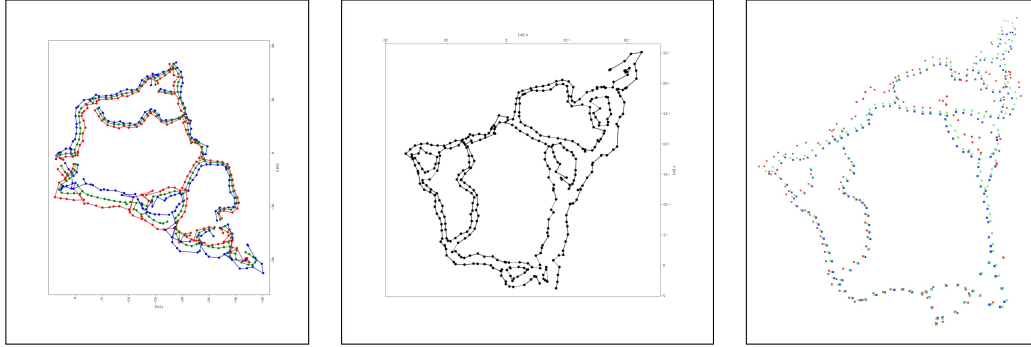


Fig. 6: Left: the average trajectory (in green) connected using the minimum spanning tree (MST) together with the two walls, left (blue) and right (red) extracted from the combined sparse point cloud. Center: the left and right walls (adjusted manually) showing the width of the cave.

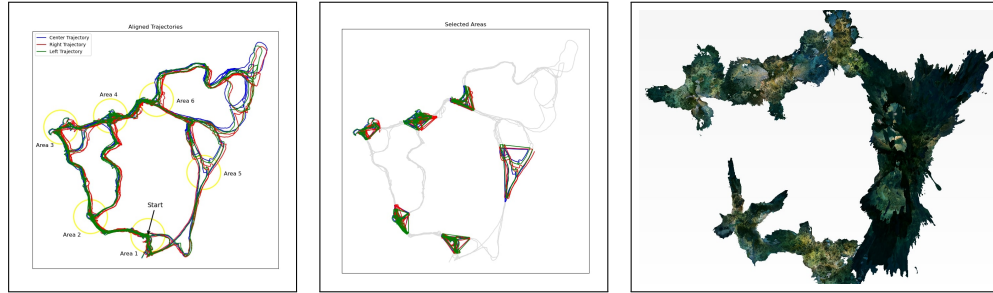


Fig. 7: The aligned trajectories were used as a guideline to select interesting areas for dense reconstruction. On the left, the six areas selected. In the middle, the parts of the trajectories used for the dense reconstruction. On the right, the Poisson reconstruction of the six areas.

In the longer term, we plan to integrate the pipeline for real-time 3D reconstruction for eventual robot navigation, towards autonomous mapping of underwater caves.

REFERENCES

- [1] D. Ford and P. Williams, *Introduction to Karst*. John Wiley & Sons, Ltd, 2007, ch. 1, pp. 1–8.
- [2] M. D. Covington, D. H. Doctor, J. N. King, and C. M. Wicks, “Research in Karst: A Model for Future Directions in Hydrologic Science,” *AGU Hydrology Section Newsletter*, July 2011.
- [3] C. McKinlay, “Woodville Karst Plain Project (WKPP),” URL: <http://www.wkpp.org>, Apr. 2015.
- [4] Global Underwater Explorers, “Woodville Karst Plain Project (WKPP),” URL: <https://www.globalunderwaterexplorers.org/Projects/WKPP>, Apr. 2015.
- [5] B. Gulden, “World longest underwater caves,” URL: <http://www.caverbob.com/uwcaves.htm>, Apr. 2015.
- [6] P. B. Campbell, Ed., *The Archaeology of Underwater Caves*. The Highfield Press, Nov. 2017.
- [7] P. B. Campbell, “An introduction to archaeology in underwater caves,” *The archaeology of underwater caves*, pp. 5–25, 2018.
- [8] A. H. G. González, C. R. Sandoval, A. T. Mata, M. B. Sanvicente, and E. Acevez, “The arrival of humans on the Yucatan Peninsula: Evidence from submerged caves in the state of Quintana Roo, Mexico,” *Current Research in the Pleistocene*, vol. 25, pp. 1–24, 2008.
- [9] J. Arroyo-Cabrales, P. Luna, J. C. Chatters, D. Rissolo, R. C. Arce, A. N. Blank, and H. Barba, “Underwater archaeology and prehistory: the case of the cenotes in Mexico,” *WORLD HERITAGE HEADS* 5, p. 54, 2015.
- [10] D. Rissolo, A. N. Blank, V. Petrovic, R. C. Arce, C. Jaskolski, P. L. Erreguerena, and J. C. Chatters, “Novel application of 3d documentation techniques at a submerged late pleistocene cave site in Quintana Roo, Mexico,” in *2015 Digital Heritage*, vol. 1. IEEE, 2015, pp. 181–182.
- [11] S. de Azevedo, M. C. Bortolini, S. L. Bonatto, T. Hünemeier, F. R. Santos, and R. González-José, “Ancient remains and the first peopling of the Americas: Reassessing the Hoyo Negro skull,” *American Journal of Physical Anthropology*, vol. 158, no. 3, pp. 514–521, 2015.
- [12] B. L. MacDonald, J. C. Chatters, E. G. Reinhardt, F. Devos, S. Meacham, D. Rissolo, B. Rock, C. Le Maillot, D. Stalla, M. D.




Fig. 8: Dense (Poisson) reconstructions from the six areas depicted in Fig. 7. Area 1: the start of the recordings, the calibration target is clearly visible. Area 2: inside the catacombs, see the boundaries of the tunnel. Area 3: as the images for the dense reconstruction are limited, the tunnel is interrupted. Area 4: the floor and walls of the passage are clearly visible of special interest are the formations from the ceiling. Area 5: this was an open area with the walls and ceiling of the tunnel too far to be illuminated. Dense reconstruction point cloud of the floor and some rock formations. Area 6: another narrow passage, with a flat floor.

- Marino *et al.*, "Paleoindian ochre mines in the submerged caves of the yucatán peninsula, quintana roo, mexico," *Science advances*, vol. 6, no. 27, p. eaba1219, 2020.
- [13] M. K. Faught, "The underwater archaeology of paleolandscapes, apalachee bay, florida," *American Antiquity*, vol. 69, no. 2, pp. 275–289, 2004.
- [14] N. Galanidou, K. Dellaporta, and D. Sakellariou, "Greece: Unstable landscapes and underwater archaeology," *The Archaeology of Europe's Drowned Landscapes*, pp. 371–392, 2020.
- [15] J. A. Peterson, A. Jalandoni, and C. Rocha, "An archaeological survey of an underwater cave in marigondon, philippines," *The archaeology of underwater caves*, 2017.
- [16] P. Kambesis, "The importance of cave exploration to scientific research," *Journal of cave and karst studies*, vol. 69, no. 1, pp. 46–58, 2007.
- [17] J. Heinerth, "Exploration of caves: Underwater exploration," in *Encyclopedia of Caves*. Elsevier, 2019, pp. 414–419.
- [18] S. Exley, *Basic Cave Diving: A Blueprint for Survival*. National Speleological Society Cave Diving Section, 1977.
- [19] J. McConnell, I. Collado-Gonzalez, and B. Englot, "Perception for underwater robots," *Current Robotics Reports*, vol. 3, no. 4, pp. 177–186, 2022.
- [20] M. J. Islam, A. Quattrini Li, Y. A. Girdhar, and I. Rekleitis, "Computer vision applications in underwater robotics and oceanography," in *Computer Vision: Challenges, Trends, and Opportunities*, M. A. R. Ahad, U. Mahbub, M. Turk, and R. Hartley, Eds. CRC Press, 2024, pp. 173–204.
- [21] N. Weidner, S. Rahman, A. Quattrini Li, and I. Rekleitis, "Underwater cave mapping using stereo vision," in *IEEE International Conference on Robotics and Automation (ICRA)*, Singapore, May 2017, pp. 5709 – 5715.
- [22] S. Rahman, A. Quattrini Li, and I. Rekleitis, "SVIn2: A Multi-sensor Fusion-based Underwater SLAM System," *International Journal of Robotics Research*, vol. 41, no. 11-12, pp. 1022–1042, July 2022.
- [23] I. Aicardi, F. Chiabrand, A. M. Lingua, and F. Noardo, "Recent trends in cultural heritage 3d survey: The photogrammetric computer vision approach," *Journal of Cultural Heritage*, vol. 32, pp. 257–266, 2018.
- [24] J. A. Placed, J. Strader, H. Carrillo, N. Atanasov, V. Indelman, L. Carlone, and J. A. Castellanos, "A survey on active simultaneous localization and mapping: State of the art and new frontiers," *IEEE Transactions on Robotics*, vol. 39, no. 3, pp. 1686–1705, 2023.
- [25] N. Muallem, R. Amoyal, O. Freifeld, and D. Akkaynak, "Gaussian splashing: Direct volumetric rendering underwater," *arXiv*, 2024.
- [26] W. C. Stone, "Design and Deployment of a 3-D Autonomous Subterranean Submarine Exploration Vehicle," in *International Symposium on Unmanned Untethered Submersible Technologies (UUST)*, no. 512, 2007.
- [27] M. Gary, N. Fairfield, W. C. Stone, D. Wettergreen, G. Kantor, and J. M. Sharp Jr, "3D mapping and characterization of sistema Zacatón from DEPTHX (DEep Phreatic Thermal eXplorer)," in *Proc. of KARST: Sinkhole Conference ASCE*, 2008.
- [28] K. Richmond, C. Flesher, L. Lindzey, N. Tanner, and W. C. Stone, "SUNFISH@: A human-portable exploration AUV for complex 3D environments," in *MTS/IEEE OCEANS Charleston*, 2018, pp. 1–9.
- [29] T. Guerneve, K. Subr, and Y. Petillot, "Three-dimensional reconstruction of underwater objects using wide-aperture imaging sonar," *Journal of Field Robotics*, vol. 35, no. 6, pp. 890–905, 2018.
- [30] S. Rahman, A. Quattrini Li, and I. Rekleitis, "Sonar Visual Inertial SLAM of Underwater Structures," in *IEEE International Conference on Robotics and Automation*, Brisbane, Australia, May 2018, pp. 5190–5196.
- [31] —, "An Underwater SLAM System using Sonar, Visual, Inertial, and Depth Sensor," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Macau, (IROS ICROS Best Application Paper Award. Finalist), Nov. 2019, pp. 1861–1868.
- [32] S. Rahman, N. Karapetyan, A. Quattrini Li, and I. Rekleitis, "A modular sensor suite for underwater reconstruction," in *MTS/IEEE OCEANS - Charleston*. IEEE, 2018, pp. 1–6.
- [33] B. Joshi, M. Xanthidis, S. Rahman, and I. Rekleitis, "High definition, inexpensive, underwater mapping," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2022, pp. 1113–1121.
- [34] S. Rahman, A. Quattrini Li, and I. Rekleitis, "Contour based reconstruction of underwater structures using sonar, visual, inertial, and depth sensor," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Macau, Nov. 2019, pp. 8048–8053.
- [35] T. Pulido Mantas, C. Roveta, B. Calcinai, M. Coppari, C. G. Di Camillo, V. Marchesi, T. Marrocco, S. Puce, and C. Cerrano, "Photogrammetry as a promising tool to unveil marine caves' benthic assemblages," *Scientific reports*, vol. 13, no. 1, p. 7587, 2023.
- [36] A. Calantropio and F. Chiabrand, "Underwater cultural heritage documentation using photogrammetry," *Journal of Marine Science and Engineering*, vol. 12, no. 3, p. 413, 2024.
- [37] A. Lago, N. Neves, R. Ventura, and J. Gaspar, "Visual-inertial odometry for metric-scale mapping of underwater caves," in *2024 IEEE International Conference on Autonomous Robot Systems and Competitions (ICARSC)*. IEEE, 2024, pp. 183–188.
- [38] A. Mallios, E. Vidal, R. Campos, and M. Carreras, "Underwater caves sonar data set," *The International Journal of Robotics Research*, vol. 36, no. 12, pp. 1247–1251, 2017.
- [39] B. Joshi, S. Rahman, M. Kalaitzakis, B. Cain, J. Johnson, M. Xanthidis, N. Karapetyan, A. Hernandez, A. Quattrini Li, N. Vitzilaios, and I. Rekleitis, "Experimental Comparison of Open Source Visual-Inertial-Based State Estimation Algorithms in the Underwater Do-

- main,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 7221–7227.
- [40] A. Abdullah, T. Barua, R. Tibbetts, Z. Chen, M. J. Islam, and I. Rekleitis, “CaveSeg: Deep Semantic Segmentation and Scene Parsing for Autonomous Underwater Cave Exploration,” in *IEEE International Conference on Robotics and Automation (ICRA)*, Yokohama, Japan, May 2024, pp. 3781–3788.
 - [41] B. Joshi. (2022) gopro_ros: Software for converting a GoPro video into a ROS 1 bagfile. [Online]. Available: https://github.com/AutonomousFieldRoboticsLab/gopro_ros
 - [42] S. Kister. (2025) MNemo: Cave Survey Data Acquisition. [Online]. Available: <https://www.arianesline.com/mnemo/>
 - [43] —. (2025) Ariane: Cave Survey Software. [Online]. Available: <https://www.arianesline.com/mnemo/>
 - [44] J. A. Hesch, D. G. Kottas, S. L. Bowman, and S. I. Roumeliotis, “Camera-imu-based localization: Observability analysis and consistency improvement,” *The International Journal of Robotics Research*, vol. 33, no. 1, pp. 182–201, 2014.
 - [45] Y. Yang and G. Huang, “Observability analysis of aided ins with heterogeneous features of points, lines, and planes,” *IEEE Transactions on Robotics*, vol. 35, no. 6, pp. 1399–1418, 2019.
 - [46] P. Lin. (2023) Pure python version of aprilgrid. [Online]. Available: <https://github.com/powei-lin/aprilgrid>
 - [47] J. L. Schönberger and J.-M. Frahm, “Structure-from-motion revisited,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
 - [48] M. Kazhdan and H. Hoppe, “Screened Poisson Surface Reconstruction,” *ACM Transactions on Graphics (ToG)*, pp. 1–13, 2013.