

Demonstrating CavePI: Autonomous Exploration of Underwater Caves by Semantic Guidance

Alankrit Gupta^{*†}, Adnan Abdullah^{*†}, Xianyao Li^{*}, Vaishnav Ramesh^{*}, Ioannis Rekleitis[†] and Md Jahidul Islam^{*}

^{*}RoboPI laboratory, Department of ECE, University of Florida, FL 32611, USA,

Email: {gupta.alankrit@, adnanabdullah@, xianyao.li@, vaishnavramesh@, jahid@ece.}ufl.edu.

[†]Department of ME, University of Delaware, Newark, DE 19716, USA, Email: yiannisr@udel.edu.

[‡]These two authors have contributed equally.

Abstract—Enabling autonomous robots to safely and efficiently navigate, explore, and map underwater caves is of significant importance to water resource management, hydrogeology, archaeology, and marine robotics. In this work, we demonstrate the system design and algorithmic integration of a visual servoing framework for semantically guided autonomous underwater cave exploration. We present the hardware and edge-AI design considerations to deploy this framework on a novel AUV (Autonomous Underwater Vehicle) named CavePI. The guided navigation is driven by a computationally light yet robust deep visual perception module, delivering a rich semantic understanding of the environment. Subsequently, a robust control mechanism enables CavePI to track the semantic guides and navigate within complex cave structures. We evaluate the CavePI system through field experiments in natural underwater caves and spring-water sites, and further validate its ROS (Robot Operating System)-based digital twin in a simulation environment. Our results highlight how these integrated design choices facilitate reliable navigation under feature-deprived, GPS-denied, and low-visibility conditions. The system design, code, and data are available on the project website: https://github.com/uf-robopi/CavePI_AUV/.

I. INTRODUCTION

Autonomous underwater cave exploration is an important area of research, providing valuable insights into some of the planet’s least explored ecosystems to understand and manage marine resources. Underwater caves, in particular, hold immense potential for advancing our knowledge of archaeology, hydrology, and geology [26]. The formations and sediments within these caves serve as critical records of historical climate and geology events while playing a key role in monitoring groundwater flow within Karst topographies, which supply freshwater to nearly a quarter of the global population [15]. However, these environments are difficult to access and often hazardous for human divers due to their confined spaces, complex geometries, and lack of natural light [40, 5]. To this end, Autonomous Underwater Vehicles (AUVs) present a promising solution, offering safe and efficient exploration while minimizing risks and improving precision [1].

However, fully autonomous exploration inside underwater caves presents significant challenges: navigation is difficult without GPS, while poor visibility, backscattering, and silt disturbances hinder perception. Thus, uninformed exploration, as often done by AUVs in open-water ecological surveying and mapping tasks [31], is quite challenging. For autonomous navigation inside underwater caves [47, 66], it is important




Fig. 1: The CavePI AUV navigates by leveraging the semantic guidance of a caveline from its down-facing camera. A deep visual perception module extracts the semantic cues, which are processed by an onboard planner to make visual servoing decisions.

to be aware of the entry and exit directions as well as a semantic understanding of existing navigational markers, when available. More specifically, underwater caves explored by humans are augmented by a line (string), referred to as the *caveline* [14], which extends from the entrance through all the major sections of the cave. Other, navigational markers such as *arrows* and *cookies* [1] indicate the closest path to the exit (arrows’ directions) and the presence of landmarks and/or divers (cookies); real-life examples and demonstrations are provided in the the supplementary video. The caveline, together with the navigational markers, provides a 1D retraction of the 3D environment, indicating the orientation and distance from the entrance of the primary passages.

In this paper, we demonstrate the system design and algorithmic integration of CavePI, a novel AUV specifically developed for navigating underwater caves using semantic

guidance from cavelines and other navigational markers; see Fig. 1. Designed for single-person deployments, CavePI’s compact design and 4-degree-of-freedom (4-DOF) motion model enables safe traversal through narrow cave passages with minimal silt disturbance. It features a forward-facing camera for visual servoing and diver-robot interaction, and a downward-facing camera for caveline detection and tracking, effectively minimizing blind spots as low as 30° below the robot, that converges at a distance of 38 cm. A Ping sonar provides altitude data to maintain a safe distance from the ground plane. Additionally, two onboard lights enhance visibility in low-light conditions, complemented by on-device software visual filters [22] for improved perception. The computational framework of CavePI includes two single-board computers (SBCs): a Jetson Nano for perception, and a Raspberry Pi-5 for planning and control.

For visual servoing, we integrate a caveline detector that performs real-time pixel-level segmentation of the caveline using the robot’s down-facing camera stream. While prior studies have proposed vision transformer (ViT)-based frameworks such as CL-ViT [66] and CaveSeg [1] for caveline detection, these approaches are computationally intensive and unsuitable for the resource-constrained SBCs. To address this, we adopt more efficient feature-extractor networks: MobileNetV3 [20] and ResNet101 [18] – combined with a lightweight DeeplabV3 [8] as the segmentation head. By leveraging GPU-accelerated computation on a Jetson Nano SBC, our system achieves a maximum inference rate of 18.2 frames per second (FPS) while ensuring robust segmentation performance. The resulting segmentation map is post-processed to extract caveline contours, which are used to guide a proportional-integral-derivative (PID)-based Pure Pursuit controller [45] for precise heading adjustments for smooth tracking-by-detection [52]. Note that CavePI leverages the caveline position and directions for semantic guidance, while the ground plane is ranged by sonar. Other markers: arrows (exit direction) and cookies (divers’ presence) trigger momentary decisions such as taking u-turns, invoking logging threads, etc. Thus, in this paper, we mainly focused on demonstrating robust caveline following performances of CavePI.

In addition, we develop a digital twin (DT) of CavePI to support pre-mission planning and testing, providing a cost-effective platform for validating mission concepts. Built using the Robot Operating System (ROS) and simulated in an underwater environment via Gazebo, it addresses the challenges posed by expensive and delicate underwater technologies that require stringent safety standards for testing. Even verified systems can exhibit unforeseen behavior due to new technological components; for example, a novel caveline detection algorithm might erroneously identify unrelated objects as cavelines, leading CavePI to follow incorrect paths and potentially creating hazardous conditions. The DT allows for prototyping, rigorous testing, and refinement of new underwater technologies, algorithms, or system modifications before physical deployment, ensuring safer and more economical evaluations. Moreover, as marine exploration missions require

substantial logistical planning and time investment, initial mission planning and performance assessments conducted on the DT facilitate rapid feedback and iterative improvements, enhancing overall mission efficiency.

CavePI’s guided navigation capabilities are initially evaluated by line-following experiments conducted in a controlled environment: a $2\text{m} \times 3\text{m}$ water tank with a maximum depth of 1.5 m. Cavelines are arranged in both regular geometric patterns (*e.g.*, circles, rectangles) and irregular configurations such as curves, spirals, and vertical slopes. The *line-following* accuracy is quantified by measuring the *tracking error*, defined as the distance between the line and the optical center of the down-facing camera on the plane of the line. Similar evaluations are performed in the Gazebo simulation environment using the DT. Additionally, the delay in decision-making and its impact on an overshoot at sharp corners are analyzed. Specifically, the proportional and differential gains of the PID controller are fine-tuned via exhaustive search, and overshoots are minimized for subsequent runs.

Following extensive in-house testing and refinement, CavePI is deployed in diverse real-world environments, including 1.5-6 m deep spring waters and 12-30 m deep natural underwater grottos and caves. These outdoor settings pose unique perception and navigation challenges, such as strong currents and turbidity in open waters, as well as noisy, low-light conditions within overhead environments. These deployments validate the robustness of CavePI, demonstrating its capability to perform long-term autonomous missions in complex underwater environments. Our long-term goal is to make it more generalizable for exploring any overhead structures (*e.g.*, inside ship hulls, pipelines, dams) with no predefined markers.

Need for a new AUV. The existing AUVs, such as SUNFISH [46], CUREE [17] and UX1 [32], offer a broad range of capabilities for marine exploration tasks. We designed CavePI specifically with (i) a small footprint: one person can carry and deploy it; (ii) a down camera dedicated for tracking navigation markers; (iii) a single-enclosure, low-cost, and low-power design with 4+ hours of endurance, as cave missions take longer than open-water tasks. The size, weight, cost, and power statistics of existing robots are not the ideal fit. For instance, SUNFISH (CUREE) weighs about 48 kg (25 kg) with $151,340\text{ cm}^3$ ($63,713\text{ cm}^3$) in volume; CavePI weighs only 8.6 kg with $40,432\text{ cm}^3$ volume (maximum width: 38 cm). Furthermore, the cost of SUNFISH is hundreds of thousands of dollars; in comparison, the overall cost of CavePI is less than \$3,500. On the other hand, UX1 requires on-land logistic support and relies on a slow pendulum system for pitch control – making it unsuitable for turbulent cave environments. BlueROV2 or other ROVs are not autonomous, and thus require significant modifications (additional cameras, autopilot, etc.) to be deployed autonomously.

II. BACKGROUND & RELATED WORK

A. Autonomous Underwater Cave Exploration

While SOTA AUVs are proficient in executing fully autonomous missions, their deployments have primarily been




Fig. 2: The proposed CavePI system design is shown; (a) isometric 3D view of the robot; (b) side-view and top-view displaying the outer shell, sonar, and thrusters' positions; (c) cross-sectional view presenting the assembly of the electronic components inside the computational enclosure; (d) the fully assembled system. CavePI is one-person deployable, weighs 8.8kg, and has a depth rating of 65 meters (213 ft).

constrained to open water environments [48], where obstacles are minimal. Underwater caves, however, present unique and formidable navigation challenges, including overhead and surrounding obstacles, narrow passageways, and complex topologies. Early efforts to address these challenges include the work of Mallios *et al.* [30], where manually operated AUVs collected acoustic data for offline mapping. Similarly, Weidner *et al.* [63, 62] utilized stereo camera systems to generate high-resolution 3D reconstructions of cave walls, floors, and ceilings.

Vision-based state estimation and navigation in underwater caves pose significant challenges due to factors such as lighting variability, backscattering effects, and image degradation [24]. To overcome these obstacles, Rahman *et al.* [42, 44] proposed a data fusion framework that integrates visual, acoustic, inertial, and water depth measurements, enabling robust trajectory estimation and sparse representations of underwater cave environments. Advances such as shadow-based mapping [43], contour extraction [33], and real-time stereo reconstruction [61] have further enhanced the density and fidelity of cave mapping techniques. More recently, visual learning-based approaches have shown promise for autonomous visual servoing within underwater caves [1, 66]. Other contemporary works with the SUNFISH AUV [47] have demonstrated effective navigation capabilities in these environments by employing Doppler Velocity Log (DVL)-based dead reckoning and sonar-based SLAM algorithms. However, it is not one-person portable, and requires considerable logistics for successful deployment.

B. Robot Navigation by Semantic Guidance

Terrestrial and aerial robotic systems benefit from their feature-rich surroundings, leveraging semantic knowledge for navigation; examples include detecting road lanes [10], traffic signs [4], power lines [7, 2] riverbanks [65], and sea horizon lines [16]. Traditional approaches for extracting such features include edge detection [27] and line segment analysis [65]. Advanced learning-based methods employ conditional random fields (CRFs) [68], convolutional neural networks (CNNs) [55], and Vision Transformers (ViTs) [13, 37] for robust semantic scene parsing. More recently, vision-language models, such as CLIP [41] and its variants [51, 54], have been

utilized to generate semantically meaningful embeddings for high-level scene comprehension and robot navigation [21, 12].

However, underwater environments are typically low-light, turbid, cluttered, and unstructured, thereby deprived of distinct semantic features [23]. Underwater robots rely on multi-modal sensing *e.g.*, scanning sonars for additional cues and navigation guidance [6, 59, 67]. To this end, semantic knowledge representation of targets [38] has been proven effective for path planning and real-time decision-making of AUVs in partially-known dynamic environments [39]. Additionally, semantic mapping techniques using laser scanners are employed for subsea pipeline following, inspection, and intervention [60]. Recent studies also utilize 3D laser/sonar point clouds for underwater landmark recognition [19] and eventually integrate them in semantic SLAM pipelines [53].

For navigating underwater caves safely, learning-based frameworks have been developed for detecting and following cavelines [66] and other navigation markers [1]. However, these computationally intensive models require further optimization to achieve real-time deployments [34]. Additionally, their onboard semantic segmentation performance and decision-making capabilities remain underexplored. We address this gap by conducting comprehensive real-world evaluations using our CavePI robotic platform, demonstrating its effectiveness for real-time semantic segmentation and navigation in complex underwater cave environments.

III. CAVEPI SYSTEM DESIGN

The CavePI AUV design includes three major subsystems: sensory bay, computational bay, and locomotion bay. Our proposed system and its components are shown in Fig. 2.

A. Sensor Bay: Acoustic-Optic Perception Subsystem

The CavePI platform includes visual and acoustic sensors: a front-facing fisheye camera, a downward-facing low-light camera, and a Ping2 active sonar. The fisheye camera, housed within a transparent dome at the *head* of the AUV, captures forward-facing visuals with a 160° field-of-view (FOV) and outputs a video feed at 1920×1080 resolution. This camera detects navigation markers and visual commands (QR codes)

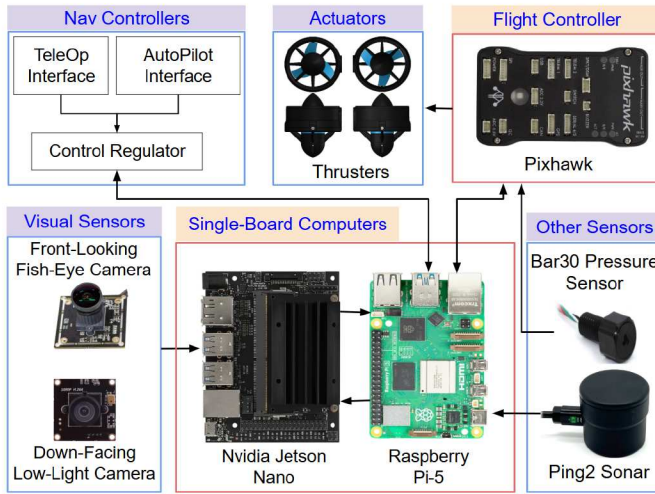


Fig. 3: Major electronics and sensor-actuator connections of CavePI.

from divers, and will assist in obstacle avoidance in the next iteration. It is worth noting that the cylindrical enclosure is a 6" tube while the dome has a 4" diameter. A custom interface is built to connect the two; see section V-A for more details. The low-light camera, mounted inside the computational enclosure, captures downward-facing visuals with an $80^\circ \times 64^\circ$ FOV, also at the same resolution and frame rate. Additionally, a Ping2 sonar altimeter-echosounder from Blue Robotics™ is mounted on the underside of the robot; the sonar has a range of 100 meters, a depth rating of 300 meters, and a resolution of 0.5% of the range, allowing it to detect obstacles directly beneath CavePI at an output frequency of 10 Hz. These sensory components collectively provide robust environmental awareness for autonomous navigation in challenging underwater caves.

B. Computational Bay

As illustrated in Fig. 2, the computational and electronic components of CavePI are housed within an acrylic cylindrical enclosure. This enclosure, with a thickness of 6.35 mm and a depth rating of 65 meters, forms the *main body* of the robot, providing mechanical stability, buoyancy, and waterproof protection for the electronics. The computational elements include a Raspberry Pi-5, a Nvidia™ Jetson Nano, and a Pixhawk™ flight controller. The Jetson Nano is dedicated to processing visual data from the cameras, performing image processing tasks critical for scene perception and state estimation. The Raspberry Pi-5 manages planning and control modules, ensuring real-time underwater navigation. The Pixhawk flight controller acts as a bridge between hardware and software, receiving actuation commands from the Raspberry Pi-5 and transmitting them to the thrusters and lights via the MAVLink communication protocol. Additionally, the Pixhawk integrates a 9-DOF IMU, offering 3-axis gyroscope, accelerometer, and magnetometer measurements, which are used to calculate the attitude of CavePI during underwater operations.

The enclosure also contains the battery compartment, voltage regulators, electronic speed controllers (ESCs), and a Bar-30 pressure sensor. The battery compartment holds a 14.8 V (18 Ah) battery pack, regulated to power internal components




Fig. 4: Data flow among major computational modules of CavePI is shown in the form of ROS Topics: red and blue arrows represent *subscribed* and *published* topics in the ROS graph, respectively.

(e.g., cameras, computers) and external components (e.g., thrusters, sonar). The battery pack is capable of supplying sufficient power to sustain over 6 hours of operation of CavePI. Each thruster is controlled by an ESC, which drives the three-phase brushless motor using PWM signals from the Pixhawk. The Bar-30 sensor provides high-precision pressure readings with a resolution of 0.2 mbar and an accuracy of 2 mm, with a working depth of up to 300 meters. This pressure data is processed to determine underwater depth, ensuring reliable and accurate interoceptive perception during operations.

C. Locomotion Bay: Middleware Integration

The end-to-end integration of CavePI ensures that each computational component operates in sync, tied to a ROS2 Humble-based middleware backbone. The modular design also allows for future upgrades, ensuring that the CavePI can be tailored to meet evolving research in marine ecosystem exploration and monitoring. The sensor-actuator signal communication graph is illustrated in Fig. 3.

The CavePI AUV is designed for low-power operation and integrates a modular ROS2 framework to support application-specific perception, planning, and control methods. As depicted in Fig. 4, the *detector* node acquires visual data from the two cameras to identify the caveline for navigation. The *mission planner* node then integrates the caveline information with the estimated position data to generate subsequent waypoints for the mission. Finally, the *autopilot controller* node utilizes these waypoints, along with the detected caveline, positional data, and depth readings from the Bar-30 sensor, to generate precise actuation signals for the thrusters, enabling accurate movements and depth control. Additionally, CavePI can function as an ROV through an optional tether-based teleoperation module. This module transmits user commands from a joystick to the onboard Raspberry Pi-5, which processes the inputs and relays them to the thrusters for manual control.

Power footprint and synchronization. As shown in Table I, CavePI offers an endurance of 6+ hours at maximum capacity. While it can operate much longer, the battery profile [11] recommends running it above 20% power; the remaining

capacity tends to deplete rapidly, which is a safety issue for recovery. While running in full capacity, the integrated system ensures an overall system throughput of 3.6 Hz. Specifically, both the cameras run at 30 Hz; the IMU and sonar frequencies are 100 Hz and 10 Hz, respectively. In our implementation, software synchronization is achieved via a joint ROS2 publisher that shares a clock across all sensory topics.

TABLE I: The battery power consumption characteristics of CavePI (at maximum capacity) over time, showing 6+ hours of endurance.

Time (hours)	1	2	3	4	5	6	7
Battery %	85	73	61	50	39	29	19

D. CavePI Digital Twin

We develop a digital twin (DT) model of CavePI by using the Unified Robot Description Format (URDF), with links and joints carefully assigned to represent the various CAD components designed in SolidWorks. To replicate the sensor suite of the physical CavePI, Gazebo plugins are integrated to simulate the front-facing camera, down-facing camera, IMU, pressure sensor, and sonar. Additional plugins are employed to simulate environmental forces, including buoyancy, thrust, and hydrodynamic drag, thereby enhancing the physical realism.

A controlled open-water scenario is created in Gazebo to simulate realistic missions, featuring a thin line arranged in a rectangular loop to mimic a caveline. Since the simulated environment lacks real-world perception challenges such as low light or turbid water conditions, the perception subsystem remains simplified. Instead of deploying computationally intensive deep visual learning models, simpler edge detection and contour extraction techniques [56] are used to identify the caveline from the down-facing camera feed. The remaining navigation and control subsystems mirror the real-world implementation and operate via two ROS nodes. The first node processes the extracted contours to make navigation decisions and publishes high-level control commands (e.g., yaw angle). The second node subscribes to these commands, computes the required thrust and hydrodynamic drag forces, and publishes them as ROS topics to control the simulated robot model.

Beyond replicating caveline following experiments, we utilize the DT system for preliminary testing and fine-tuning of new control algorithms. It also enables the simulation of complex cave environments, such as narrow passages, dead ends, and sharp turns. Conducting repeated real-world experiments in such scenarios to improve the control system can be logistically demanding where the simulation offers an efficient alternative for extensive evaluation and fine-tuning.

IV. CAVEPI NAVIGATION PIPELINE

The CavePI AUV is designed to autonomously navigate underwater by following a caveline and other navigation markers. However, a caveline appears a few pixels wide in the bottom-facing camera's FOV and is significantly challenging to detect in noisy conditions with almost no ambient light. Moreover, the caveline is often obscured or blends into the

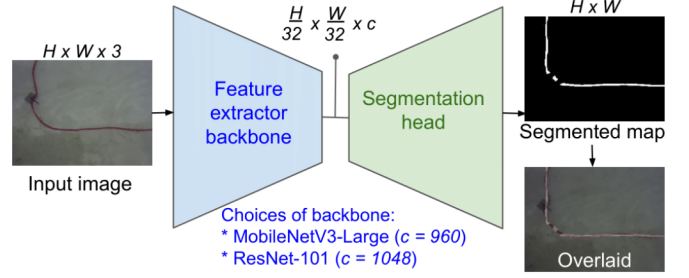


Fig. 5: Simplified model architecture for caveline segmentation is shown; we use a DeepLabV3 [9] head with two choices for backbone network: MobileNetV3 [20] and ResNet101 [18].

background, requiring a robust and fine-tuned visual learning pipeline for reliable operation.

A. Semantic Guidance: Tracking by Detection

The model selection process plays a crucial role in optimizing accuracy and efficiency on resource-constrained edge devices. While object detection and semantic segmentation are two established techniques in visual learning, they produce different interpretations of the target RoI. Object detection models generate bounding boxes around the target, which, given the caveline's irregular shape and orientation, encompass significant background regions, complicating the heading angle estimation. To address this, we opt for semantic segmentation, which provides pixel-level contours of the caveline as a more precise marker for tracking-by-detection.

TABLE II: Edge performances (on a Jetson Nano device) for the two model configurations (Fig. 5) available in CavePI are compared.

Backbone	mIoU	# Params	Inference rate	GPU usage
MobileNetV3 [20]	48.95%	11.02 M	18.2 fps	314 MB
ResNet101 [18]	50.48%	60.99 M	1.4 fps	747 MB

Model architecture and design choices. Onboard resource constraints of CavePI significantly influence the choice of model architecture. The Jetson Nano, with its shared memory architecture, allocates GPU memory dynamically from a limited 2 GB pool, which must also accommodate the dual camera feeds and Nano-Pi control signal communication. Considering these limitations, we initially selected the MobileNetV3 (Large) backbone [20] for feature extraction. This lightweight architecture feeds extracted features into a DeepLabV3 [8] head to generate binary segmentation maps. While the MobileNetV3-DeepLabV3 combination is computationally efficient, requiring only about 314 MB GPU memory for 11 million parameters, it proved insufficient for accurately segmenting the thin caveline in turbid water conditions. To address this limitation, we adopted a heavier backbone, ResNet101 [18], while retaining the DeepLabV3 head. This model achieved slightly better performance, as highlighted in Table II, albeit at the cost of reduced inference speed. Both architectures were initialized with pre-trained weights and fine-tuned on a custom dataset to enhance caveline detection accuracy in challenging underwater environments.

Algorithm 1 Visual Servoing Controller of CavePI

```

1: Input: Segmentation map  $\mathcal{I}_{w \times h}$ , center pixel  $(u_i, v_i)$ 
2: Output: Thruster commands  $\mathbf{v}$ 
3: while robot is active do
4:   Extract set of contours  $\{C\}$  from  $\mathcal{I}$ 
5:   if  $C = \emptyset$  then
6:     State  $\leftarrow$  lost //no caveline detected
7:     Rotate  $360^\circ$  //small loops for recovery
8:   else
9:     if  $|C| > 1$  then
10:      Sort  $\{C\}$ :  $\{C_0, C_1, \dots, C_n, \dots\}$  //toward heading
11:       $(u_n, v_n) \leftarrow \text{Centroid}(C_n)$ 
12:       $C_{\text{next}} \leftarrow \arg \max_{C_n \in \{C\}} ((u_n - u_i)^2 + (v_n - v_i)^2)$ 
        //farthest contour from  $(u_i, v_i)$ 
13:    else
14:       $C_{\text{next}} \leftarrow C$ 
15:    end if
16:     $(u_c, v_c) \leftarrow \text{Centroid}(C_{\text{next}})$  //contour center
17:     $\psi \leftarrow \text{slope}((u_i, v_i), (u_c, v_c))$  //next heading angle
18:     $\mathbf{v} \leftarrow f(\psi)$  //thruster commands
19:  end if
20:  Send  $\mathbf{v}$  to thrusters
21: end while

```

Dataset and model training. For this study, we utilized the open-source CL-ViT dataset, which includes pixel-level annotations of cavelines [66]. It comprises 3,150 RGB images collected from three distinct underwater cave systems. To enhance the dataset’s diversity and variance, we manually annotated an additional 150 images from our laboratory testbed, resulting in a combined dataset of 3,300 images. The dataset was divided into a training set of 3,200 images and a validation set of 100 images. Both models were fine-tuned on this dataset over a maximum of 50 epochs using the following training configurations: the Adam optimizer with an initial learning rate of 0.001, a cosine learning rate scheduler with $T_{\text{max}} = 20$, and a cross-entropy loss function. This training configuration was designed to optimize these models’ performances in diverse underwater environments.

Benchmark comparison. The fine-tuned models are evaluated on the open-source CL-Challenge dataset [66] that contains caveline annotations for 200 noisy low-light images from different cave systems. Using the lighter MobileNetV3 backbone, our model achieves a mean intersection-over-union (mIoU) of 48.95%, whereas the considerably heavier ResNet101 model demonstrates only a marginal improvement with a mIoU of 50.48%. For reference, the highest mIoU score reported on this dataset is 58.3% [66]. Since smooth real-time operation requires a minimum control loop of 3.6Hz, we opt for the MobileNetV3 backbone that exceeds this requirement with an inference speed of 18.2fps (frames per second).

Model compression and integration. The two aforementioned models are optimized to deploy on the Jetson Nano device of CavePI. First, the models are converted into Open

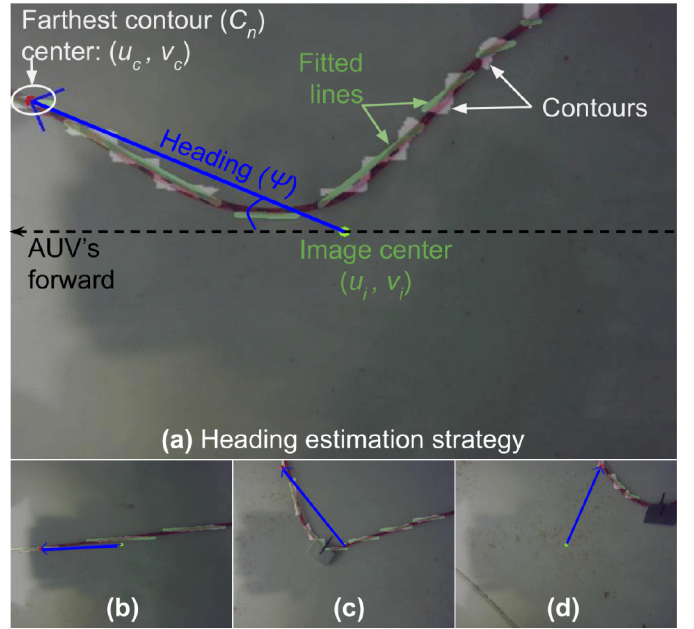


Fig. 6: A few visual servoing test cases are shown; (a) AUV heading is estimated based on its farthest caveline contour; (b,c) scenarios for straight and turn decisions, respectively; and (d) an *overshoot* scenario occurred due to delayed decision-making.

Neural Network Exchange (ONNX) [36] format that uses computational graphs (nodes and edges) to represent operations and data flow within the pipeline. Subsequently, a serialized data structure is created by NVIDIA’s TensorRT SDK [35], which stores highly optimized deep learning models ready for fast inference on NVIDIA GPUs. Table II compares the onboard performance of the two models used in this work.

B. Autonomous Control

Algorithm 1 outlines the visual servoing control algorithm employed by CavePI. The process begins by extracting caveline contours, \mathcal{C} , from the segmentation map \mathcal{I} , for semantic guidance. As the caveline is often detected as fragmented contours (see Fig. 6), the center of the farthest contour (u_c, v_c) is identified and selected as the next waypoint in CavePI’s path planning. The heading angle ψ of this waypoint is then calculated with respect to the image center (u_i, v_i) , about the x-axis of the image frame, following the right-hand rule. This heading angle serves as a high-level navigation command, which is transmitted to the Raspberry Pi-5 for execution within the computational subsystem.

The centroid of all detected contours in the set \mathcal{C} is calculated, with the centroid of the n^{th} contour denoted as (u_n, v_n) . The farthest contour is determined by:

$$(u_c, v_c) = \arg \max_{n \in \mathcal{C}} [(u_n - u_i)^2 + (v_n - v_i)^2]. \quad (1)$$

The instantaneous heading vector is defined as the vector from the image center (u_i, v_i) to the centroid of the farthest detected contour (u_c, v_c) , represented as $[u_c - u_i \quad v_c - v_i]^T$. The equation of the instantaneous line of motion in the image

frame is given by:

$$\begin{bmatrix} u - u_i \\ v - v_i \end{bmatrix} \times \begin{bmatrix} u_c - u_i \\ v_c - v_i \end{bmatrix} = 0, \quad (2)$$

where (u, v) represents the free variable on the line.

Ideally, the forward x-axis of the image frame aligns with the robot's heading direction, enabling the computed heading angle to function as the primary error signal for a PID controller on the Raspberry Pi-5. The PID controller is calibrated to minimize this error, maintaining CavePI's heading toward the designated waypoint. The controller-generated signals are transmitted via the MAVLink communication protocol to the electronic speed controllers (ESCs), which regulate the speed and rotational direction of the thrusters. Upon failure to detect any caveline, CavePI executes a circular search pattern to re-acquire tracking. If a line is not detected, it signals flashing lights (helps if divers are nearby); subsequently, it starts navigating towards the (last seen) arrow direction (toward the cave exit) to rediscover the caveline and/or exit the cave. Additionally, CavePI maintains a user-defined target depth by utilizing a PID controller, which processes the error between the current and desired depths and generates appropriate control signals to achieve stable depth regulation.

V. SYSTEM EVALUATION

A. Mechanical Stability under Hydrostatic Pressure

The CavePI system is evaluated for structural robustness and functional durability to support long-term autonomous operation at depths. Critical components of AUVs generally include the interfacing joints of the actuators (thrusters), sensors (*e.g.*, cameras and sonar), and enclosure doors [29]. As shown earlier in Fig. 2, CavePI's *head* section is connected to the *main body* via a custom-designed aluminum plate, referred to as the dome connector, which ensures a watertight interface for the forward-facing camera. Similarly, the clamps used in the thruster and sonar subassemblies, as well as the plates for electronics mounting within the computational enclosure, are custom-designed to enhance mechanical efficiency and dynamic stability. Components exposed directly to the underwater environment, including thrusters, the Ping2 sonar, the dome connector, and external cables, are identified as sensitive elements requiring careful design and validation to ensure system reliability under challenging conditions.

The thrusters and Ping2 sonar have depth ratings of more than 500 meters and 300 meters, respectively, along with their cables. To evaluate the structural strength of the dome connector of CavePI, we use finite element analysis (FEA) methods [3]. Tetrahedral elements are utilized to generate the finite element mesh from the 3D model of the dome connector, while the Von Mises yield criterion is applied to assess the **maximum stresses**, **total deformation**, and **factor of safety** (FoS). The analysis considered hydrostatic pressures [29, 28] at the maximum water depth of 65 m.

The number of nodes and finite elements in the mesh indicates the mathematical model's quality: a higher count typically reflects a more accurate model. In the analysis of




Fig. 7: (a) FEA mesh of the dome connector; (b) its total deformation due to hydrostatic forces at 65 m water depth; (c) equivalent stresses as per Von-Mises yield criterion; and (d) maximum stresses at the zoomed-in area (best viewed digitally at $2\times$ zoom).

the dome connector, the mesh contained 4,636,236 nodes and 3,300,641 elements; see Fig. 7. Based on the Von-Mises criterion, the maximum equivalent stress was calculated to be 15.3 MPa. The maximum deformation, observed at the inner most circumference of the dome connector, was 0.005 mm. The minimum FoS is calculated using the yield strength (σ_y) of the material (Aluminium 6061-T6), and maximum stress (σ_{max}) as follows:

$$FoS = \frac{\sigma_y}{\sigma_{max}} = \frac{240 \text{ MPa}}{15.3 \text{ MPa}} \approx 15.76.$$

Note that a FoS greater than 2 for hydrostatic applications and 4 for hydrodynamics applications is considered sufficient [25, 58]. Thus, an FoS over 15 indicates that the dome connector is sufficiently stable for the intended operation pressure.

B. Caveline Following Performance

Setup. The line following experiments are initially conducted in a $2\text{ m} \times 3\text{ m}$ laboratory water tank with a maximum depth of 1.5 m. As a first setup, a line is laid on the tank floor in a rectangular loop, without any depth variation; see Fig. 8. Subsequently, we add vertical slopes in the tank to simulate the rough underwater terrains; see Fig. 15 (a). Upon detecting the line via the down-facing camera, CavePI autonomously follows the loop while maintaining a specified depth. The system's line-following accuracy, depth-holding capabilities, and control smoothness are evaluated and fine-tuned in preparation for outdoor deployments.

Evaluation process. CavePI's caveline-following performance is evaluated by measuring the *tracking error*, defined as the perpendicular distance between the optical center of the




Fig. 8: The laboratory setup used for tracking accuracy evaluation is shown. CavePI detects and follows the line laid on the tank floor, completing a loop around the 6 meter perimeter in approximately 2 minutes.

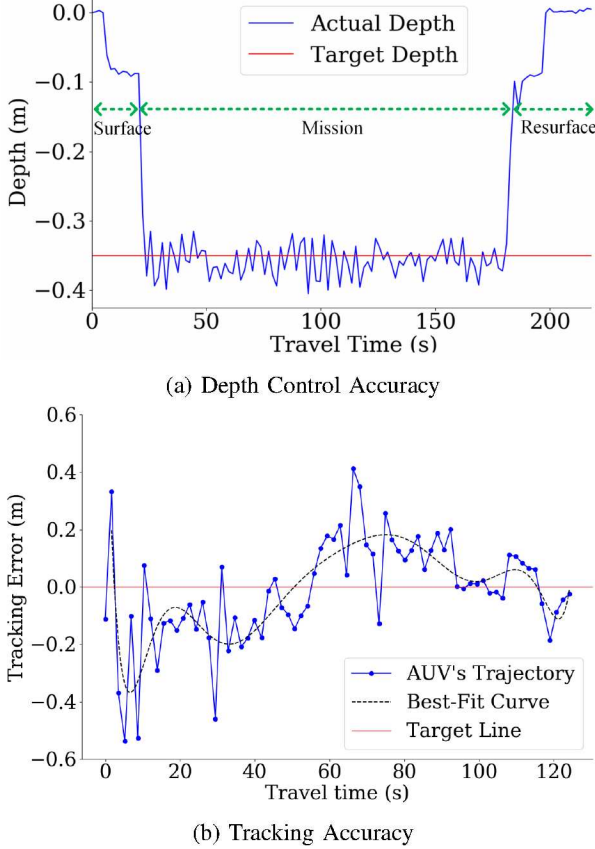


Fig. 9: Line-following and depth-holding accuracy are reported for a laboratory experiment: (a) The target depth was set to 0.35 m; (b) the tracking error (δ), *i.e.*, the offset of the optical center of the camera from the detected line is plotted for a rectangular loop following task.

downward-facing camera and the nearest point on the caveline in the plane of caveline; see Fig. 9. This measurement leverages pose information from CavePI's IMU, depth readings from the sonar, and visual information from the segmentation mask of the caveline captured by the down-facing camera.

To formalize the calculation, we define several reference frames; see Fig. 10. The initial pose of the robot assumes its front axis is aligned with the caveline and its downward axis perpendicular to the caveline plane. This reference frame, denoted as $\{N\}$, has its origin at the camera. A corresponding frame $\{S\}$, with the same orientation as $\{N\}$ but with its origin




Fig. 10: Configurations of the coordinate frames are shown at: (a) initial time t_0 , and (b) an arbitrary time t . Note that the $\{C\}$ is a (rigid) body-fixed frame attached to the robot while $\{N\}$ and $\{S\}$ are non-rotating (fixed orientation) frames, even though their origins translate along with the robot.

at the sonar, is also defined. Additionally, we define the global reference frame $\{G\}$, and the CavePI body frame $\{C\}$ with its origin at the down-facing camera. The transformation matrix from any frame a to b is defined as ${}^b_a\mathbf{T} = [{}^b_a\mathbf{R}_{3 \times 3} \mid {}^b_a\mathbf{t}_{3 \times 1}]$. The roll, pitch, and yaw angle measurements from IMU are expressed in the global frame $\{G\}$ at any time instance.

For a given image I , we extract the caveline contour closest to the camera center and calculate the distance between the camera center and the pixel ${}^I P$ on the contour edge, then convert it to a 3D point ${}^C P$ using:

$${}^I P = [u_P \quad v_P \quad 1]^T, \quad {}^C P = \lambda K^{-1}({}^I P); \quad (3)$$

where K is the camera intrinsic matrix and λ is the depth scale factor. The rotation of the CavePI's frame with respect to the sonar frame is calculated as follows:

$${}^C_S \mathbf{R} = ({}^S_G \mathbf{R})^{-1} {}^C_G \mathbf{R} = ({}^N_G \mathbf{R})^{-1} {}^C_G \mathbf{R}; \quad (4)$$

where ${}^C_G \mathbf{R}$ and ${}^N_G \mathbf{R}$ are computed from instantaneous and initial IMU measurements, respectively. Subsequently, ${}^C P$ is transformed to the sonar frame using:

$$\begin{bmatrix} {}^S P \\ 1 \end{bmatrix} = \begin{bmatrix} {}^C_S \mathbf{R} & {}^C_S \mathbf{t} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} {}^C P \\ 1 \end{bmatrix}, \quad (5)$$

$${}^S P = \lambda [a \quad b \quad c]^T; \quad (6)$$

where the scalars a , b , and c are obtained from the above equation. The raw depth measurement from sonar ${}^C d$ is converted to the sonar frame ${}^S d$ and λ is calculated as:

$${}^S d = {}^C_S \mathbf{R}_{(3,3)} \cdot {}^C d, \text{ and } \lambda = \frac{{}^S d}{c}. \quad (7)$$

TABLE III: Tracking errors from the line following experiments are compared for various choices of K_p and K_d values of the heading controller. Here, δ : mean tracking error (cm), σ : standard deviation.

K_p	1	2	3	3	3.5	3.5	3.4	3.4	3.5	3.4
K_d	0	0	0	0.5	0.5	0.7	0.7	0.9	0.9	1.0
δ	19.8	17.2	15.9	14.7	14.1	13.9	13.6	13.1	13.4	13.3
σ	26.5	24.1	22.3	21.6	20.9	21.2	20.3	19.8	21.5	20.0

TABLE IV: Depth errors from the line following experiments are compared for various choices of K_p and K_d of the depth controller. Here, μ : mean depth error (cm), σ : standard deviation.

K_p	500	500	550	600	600	600	600	600	650	720
K_d	0	10	10	10	20	30	50	100	200	300
μ	3.09	2.98	2.47	2.39	2.06	2.61	1.91	1.80	2.98	1.92
σ	2.14	1.97	1.72	1.63	1.38	1.84	1.38	1.70	2.02	1.44

The 3D coordinates of ${}^N\mathbf{P}$ is calculated using:

$${}^N\mathbf{P} = -\frac{C}{S}t + {}^S\mathbf{P}. \quad (8)$$

Finally, δ is obtained from the y -coordinate of ${}^N\mathbf{P}$.

C. PID Controller Tuning

To ensure stable operation in turbulent water conditions, CavePI's PID-based *Pure Pursuit* controller [45, 50] is designed and fine-tuned through extensive experimentation. Pure Pursuit is a widely used geometric tracking algorithm that guides a vehicle toward a target point on a path [49], calculated at a fixed *lookahead distance*. By steering the vehicle to continuously align with this moving target point, the controller ensures smooth convergence to and following of the intended trajectory, which we tuned and tested empirically.

Our initial trials reveal challenges, including high overshoot, difficulty maintaining depth, and instability at sharp corners. These issues are iteratively calibrated to achieve robust and reliable motion control. A grid search technique is employed to optimize the proportional gain (K_p) and differential gain (K_d) for the two onboard controllers – depth and heading controllers. For each tested parameter pair, the robot traverses a complete 6-meter loop in the testbed at a depth of 0.35 meters. Tracking (δ) and depth errors for each trial are recorded at 0.6 Hz. Table III and IV summarize the observations of the tuning parameters combinations for heading and depth controllers, respectively. We identified that a controller tuned with $K_p = 3.4$ and $K_d = 0.9$ minimized the tracking error, and therefore employed these values as the heading controller gains in all subsequent experiments. Similarly, $K_p = 600$ and $K_d = 50$ are chosen for the depth controller.

D. ROS Prototyping and Gazebo Simulation

The digital twin of CavePI is developed using ROS and simulated within Gazebo's underwater environment to evaluate its line-tracking performance; see Fig. 11. As a simplified representation of the physical platform, it enables extensive pre-deployment testing of mission-critical functionalities. The digital twin is designed to be positively buoyant, consistent with the actual CavePI, and receives depth-control inputs to maintain a certain depth. Simulating the digital twin under

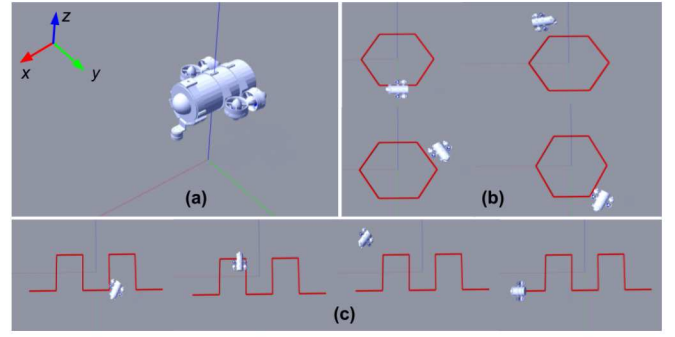


Fig. 11: The digital twin (DT) of CavePI, modeled in ROS, is used for virtual testing in Gazebo underwater environment. (a) An isometric view of the DT in Gazebo; (b,c) Line following experiments in a hexagonal loop and a lawn-mower pattern, respectively.

realistic underwater conditions allows for detailed analysis of hydrostatic and hydrodynamic forces, informing component placement for stable buoyancy and achieving near-zero roll and pitch angles when stationary. A complete sensor suite is integrated to confirm optimal camera placement, particularly for the downward-facing camera essential for caveline detection. Additionally, it incorporates a PID controller to simulate CavePI's line-tracking functionality, serving as a platform for preliminary control algorithm design and validation. Nevertheless, final controller tuning is performed on the physical system through experimental trials due to factors such as non-linear drag forces, caveline detection dependencies, and design parameter variations between the CAD model and the actual hardware, which cannot be replicated in simulation. We also developed a VR interface for live interactive simulations and immersive experiments. In this interface, users wear a head-mounted display (HMD) to engage in an immersive virtual mission experience; they experience an egocentric view (*i.e.*, robot's perspective) as well as an external global perspective in real-time while the DT model of the robot conducts a mission. Additionally, they can switch from autonomous mode to teleoperation mode and directly control the CavePI in




Fig. 12: Interactive demonstration setup for CavePI's virtual missions is shown. A user wears an HMD to experience various camera perspectives while using a handheld controller.

different virtual worlds; see Fig. 12.

VI. FIELD EXPERIMENTAL DEMONSTRATION

A. Field Trials and Experimental Setups

The guided navigation capability of CavePI was demonstrated in two distinct real-world environments: (1) shallow riverine areas (2 m - 6 m depth) near springs' outlets; and (2) deep natural underwater grottos and caves (15 m - 30 m depth). The experiments included 15 open-water trials in spring areas and 10 trials inside underwater grottos. In the open-water trials, a 20-meter rope line was laid in both irregular loop patterns and linear configurations along the uneven riverbed. In contrast, cave trials used an actual caveline, which varied in color, texture, and thickness. Each environment posed unique challenges. Strong currents near the spring outlets caused significant drift, particularly when navigating across or against the flow. Additionally, low-light conditions inside underwater caves hindered the accurate detection of the semantic markers.

Field deployments were conducted under the supervision of two support divers responsible for initiating and concluding the missions. Divers carried QR-code tags to issue visual commands to CavePI via its front-facing camera. Upon receiving the `start` command, the robot activated depth-hold mode and began line-following while maintaining the specified depth. A surface station operator monitored each mission remotely. For cave deployments, a tether connected the robot to the surface station to enable remote emergency intervention if necessary (see Fig. 13). The trials highlighted both the strengths and potential areas for improvement in CavePI's perception, planning, and control mechanisms.






Fig. 13: (a) A support diver is placing the robot on a caveline to initiate tracking; (b) two divers are following the robot once tracking is initiated. Note that the tether is used only for remotely aborting the mission in case of emergency.

B. Results and Observations

Laboratory tank tests. Sec. V-B details the setup and evaluation procedures for assessing CavePI's line-following performance in laboratory tests. During these tests, CavePI operated at a depth of 0.35 meters below the water surface. The depth controller demonstrated robust performance, maintaining the desired depth with a mean error of approximately 2 cm. At the beginning of each trial, heading control signals were disabled for 10 seconds to allow CavePI to stabilize at the target depth after its descent. Consequently, higher tracking errors were observed during this initial period, as shown earlier in Fig. 9.



(a) Depth Control Accuracy



(b) Tracking Accuracy

Fig. 14: Line-following and depth-holding accuracy are reported for an open-water experiment: (a) the target depth was set to 0.525 m; (b) the higher tracking error is caused by strong currents, leading to loss of tracking on one occasion.

Additionally, occasional perception failures were encountered, where the tank edges were misidentified as caveline segments (see Fig. 16 a), leading to *overshooting*. Despite these transient deviations, CavePI successfully completed multiple loops, demonstrating the robustness of its control strategy.

Open-water tests: Spring riverine environments. These trials were evaluated using the same methodology described in Sec. V-B, focusing on both line-tracking and depth-regulation accuracy. Fig. 14 presents a representative result from a 10-minute open-water experiment. As expected, the field results exhibit greater variation compared to the controlled laboratory experiments due to environmental disturbances. Despite these challenges, the depth controller demonstrated robust performance, maintaining a depth within ± 10 cm without requiring trial-specific parameter tuning.

In contrast, the heading controller proved less resilient under strong currents, resulting in much higher tracking errors. As illustrated in Fig. 14 b, the system experienced significant lateral drift, occasionally causing CavePI to lose the caveline. This limitation became particularly pronounced when traveling



Fig. 15: A few snapshots from our field trials for caveline tracking and following experiments with CavePI are shown. The setups include: (a) rectangular loops and slopes in a water tank; (b) irregular shapes in Spring riverine water, and (c) low-light underwater cave scenarios (note that the tether is only for safety and easy recovery). Images are best viewed digitally at $2\times$ zoom; full video demonstration is submitted in the supplementary files; also available online at: <https://youtu.be/9BPpB1nu98E>.

perpendicular to the current (see Fig. 17b), as there was no active lateral control to counter crossflow. Additionally, the vehicle’s slightly back-heavy design induced a pitch-up motion during upstream travel, as shown in Fig. 17c, further compromising stability and navigation accuracy.

Low-light tests: nighttime cave environments. These trials utilized an earlier (ablation) iteration of CavePI equipped with a three-thruster configuration, in contrast to the newly proposed four-thruster design. The three-thruster setup required precise orientation of the downward thruster to achieve proper heave motion; otherwise, thrust forces could inadvertently induce *roll*. Additionally, the absence of a dedicated roll-control mechanism posed challenges for stabilization. In contrast, the four-thruster configuration incorporates two thrusters for heave motion, enabling dedicated roll control and significantly improving overall stability during navigation.

Despite being equipped with two onboard lights, CavePI faced considerable challenges in low-light environments due to glare and backscattering effects, which impeded accurate semantic perception. Nighttime trials conducted in an

underwater cave/grotto system revealed that the lightweight segmentation models struggled to detect the caveline from camera images. During two one-hour dives, support divers reported instances where submerged tree roots, mosses, and other thin structures were mistakenly identified as the caveline. However, once manually repositioned onto the caveline by the divers, CavePI was able to maintain tracking for up to one minute. These findings underscore the need for integrating a more powerful computational platform capable of running robust segmentation models to enhance CavePI’s perception capabilities in challenging environments.

VII. LIMITATIONS & CHALLENGES

A. Design Aspects

CavePI’s onboard components are packed within a single pressure-sealed tube, leaving minimal space for additional hardware. Notably, major housing volume is occupied by the LiPo battery, which can be replaced with a more compact alternative to create space for future add-ons. We also plan to upgrade the perception SBC from the existing Jetson Nano to a Jetson Orin Nano that offers $4\times$ memory with more




Fig. 16: A few perception failure modes are shown. The down-facing camera falsely detects (a) the edge of the laboratory tank; (b) tree roots as the caveline, resulting in incorrect tracking.

advanced GPU resources. The redesigned system will also incorporate a magnetic switching mechanism for seamless power control. Outside of the enclosure, the 4 thrusters are arranged to allow control over surge, heave, roll, and yaw motions, however, their adjacent placement induces mutual turbulence and reduces thrust efficiency. To address this, we are investigating alternative thruster placements that will enhance motion dynamics without affecting the robot's buoyancy properties and dynamic stability.

B. Perception Challenges

The current onboard sensor suite offers a limited understanding of the surrounding 3D environment. For instance, the Ping2 sonar provides only 1D depth measurements, which we will replace with an advanced 360° scanning sonar system for enhanced spatial awareness. The scanning sonar, combined with other state estimation sensors, will map the environment and effectively avoid obstacles during navigation. Moreover, the two cameras currently operate independently for different purposes without synchronization. In the next iteration, we propose incorporating a 45° slanted camera and combining all the visual feeds into a mosaic vision. This advanced vision system will offer wider FOV with more peripheral information and improve visual servoing performance. Additionally, the current caveline detection model occasionally misidentifies objects such as submerged tree roots as part of the caveline as shown in Fig. 16b, and it sometimes fails to detect the line under low-light conditions. These issues lead to significant tracking inaccuracies. Although more computationally intensive models might address these shortcomings, they are currently infeasible due to hardware limitations. With the proposed design modifications for the next iteration, we plan to adopt a more robust detection model to improve CavePI's tracking performance. Furthermore, hand gesture recognition [64] will be integrated to enable seamless cooperation between divers and CavePI during underwater cave operations. With this advanced sensor setup, we will deploy advanced SLAM algorithms, such as SVIn2 [44] for robust navigation in GPS-denied underwater cave environments.

C. Smooth 6-DOF Control

Although CavePI is a 6-DOF AUV capable of maneuvering in a 3D environment, it currently offers active control over




Fig. 17: A few tracking failure modes are shown: (a) CavePI is correctly following the line; (b) lateral drift under strong currents; (c) pitching upward due to currents, attributed to its back-heavy design; (d) overshooting its intended trajectory while moving downstream.

four DOFs: surge, heave, roll, and yaw. In future iterations, we intend to reposition its thrusters to enable control over the remaining two DOFs – pitch and sway – thereby enhancing the robot's maneuverability in complex underwater environments. Furthermore, CavePI's autonomous control primarily utilizes a proportional-derivative (PD) controller, which performs effectively under conditions with minimal environmental disturbances. However, this straightforward control strategy becomes unstable in more complex scenarios, such as when water currents are present. In such conditions, lower proportional gains are insufficient for adjusting CavePI's yaw to align with the cave line, while higher proportional gains cause overshooting in the robot's trajectory during sharp turns (see Fig. 17d), attributed to both perception latency and the limitations of the current control method. Additionally, nonlinear (primarily quadratic) drag forces significantly impact the robot's stability and must be incorporated into the control system design. To overcome these challenges, we are developing a more robust nonlinear adaptive control system to enhance stability. Concurrently, we are optimizing the communication between the perception and control modules to reduce latency and improve overall responsiveness.

VIII. DISCUSSION

Our investigation identifies four primary challenges impeding effective visual servoing and safe navigation of AUVs within underwater cave environments: (1) accurate scene parsing under noisy sensing conditions, with very little ambient light; (2) robust state estimation in feature-scarce environments lacking GPS or external localization aids; (3) reliable tracking and control mechanisms to maintain visibility of navigation markers during maneuvering; and (4) the inherent logistical complexity and risks involved in field deployments for in-situ experimental validation. These challenges are further exacerbated in our experiments by the compact form factor, low inertia, and energy-efficient design of the CavePI platform, which collectively reduce motion stability in the presence of unpredictable hydrodynamic disturbances, particularly near

cave entrance and spring outlets. Furthermore, the performance of our semantic-guided navigation framework hinges on the accuracy of the underlying semantic segmentation model, which is challenged by variations in waterbody types, turbidity, and backscatter commonly encountered in confined and cluttered cave systems. These findings highlight the pressing need for more robust perception and control strategies tailored to the unique operational demands of underwater caves.

To this end, the broader research community stands to benefit from our open documentation of system performance under real-world conditions for both successful outcomes and failure cases. These findings emphasize the critical need to balance perception accuracy with computational efficiency to enhance mission reliability in natural underwater environments – an aspect often underrepresented in controlled laboratory settings. Another key insight from our field deployments is the influence of external disturbances on caveline-following performance. Our results indicate that effective *sway control* is essential for enabling compact marine robotic systems to maintain stability and trajectory in the presence of strong currents. This observation underscores the need for further research into adaptive control strategies that can dynamically respond to environmental disturbances. Moreover, our integration of visual tags as a mechanism for human-robot interaction offers a scalable and practical alternative to tethered communication for short-range underwater tasks, advancing the potential for more flexible and intuitive operations.

Building upon this foundational system, our ongoing work focuses on advancing toward fully autonomous underwater cave exploration with greater resilience and mission safety. Recent upgrades include the integration of more powerful compute (Nvidia™ Jetson AGX Orin instead of Orin Nano), an adaptive control framework to dynamically compensate for external disturbances, and multi-modal sensor fusion [57] to enhance perception robustness and state estimation accuracy. Ongoing work is also integrating a more dense segmentation pipeline (using CaveSeg [1]) for detecting overhead structures and companion divers from sparse pointcloud (by SVIn2 [44]) to improve autonomy. Furthermore, the ability to operate in partially mapped or entirely unknown caves introduces opportunities for innovation in semantic SLAM, risk-aware path planning, and safety-critical control. These enhancements will be particularly valuable in human-robot collaborative missions, where the robot must support shared situational awareness and dynamic decision-making in real-time.

IX. CONCLUSION & FUTURE WORK

Underwater cave exploration is important in hydrology for freshwater resource management, as well as in geological and archaeological studies. Since cave surveys are labor-intensive and hazardous, fully autonomous platforms can help significantly. This work demonstrates CavePI, a portable and cost-effective AUV tailored for semantic-guided navigation in GPS-denied underwater environments. The platform’s key strength lies in extracting sparse semantic cues from noisy, low-light conditions through a fast and precise segmentation pipeline,

enabling intelligent and adaptive navigation. The planning and control systems have been meticulously designed, rigorously tested, and fine-tuned within controlled laboratory settings to ensure robust and reliable performance. Comprehensive field deployments in diverse, challenging underwater environments, along with extensive digital twin simulations, have demonstrated CavePI’s reliability, adaptability, and operational efficiency. These results underscore its potential as a versatile and dependable tool for safe, long-term autonomous exploration of underwater caves. Future iterations will optimize system performance and capabilities to support a wider range of applications for exploring any overhead structures (ship hulls, pipelines) without predefined markers. To achieve this, future work will focus on enabling zero-shot semantic understanding and active planning capabilities.

ACKNOWLEDGMENT

This research is supported in part by the U.S. National Science Foundation (NSF) grants #2330416, #2024741, and #1943205; and the University of Florida (UF) Research grant #132763. We are thankful to Dr. Nare Karapetyan, Ruo Chen, and David Blow for facilitating our field trials at Ginnie open-water springs and Blue Grotto. Furthermore, we 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 Halcyon Dive Systems, Teledyne FLIR LLC, and KELDAN GmbH lights.

REFERENCES

- [1] Adnan Abdullah, Titon Barua, Reagan Tibbetts, Zijie Chen, Md Jahidul Islam, and Ioannis Rekleitis. CaveSeg: Deep Semantic Segmentation and Scene Parsing for Autonomous Underwater Cave Exploration. In *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024.
- [2] Dimitrios Alexiou, Georgios Zampokas, Evangelos Skartados, Kosmas Tsiakas, Ioannis Kostavelis, Dimitrios Giakoumis, Antonios Gasteratos, and Dimitrios Tzovaras. Visual Navigation Based on Deep Semantic Cues for Real-Time Autonomous Power Line Inspection. In *International Conference on Unmanned Aircraft Systems (ICUAS)*, pages 1262–1269, 2023.
- [3] Thor Morales Bieze, Frederick Largilliere, Alexandre Kruszewski, Zhongkai Zhang, Rochdi Merzouki, and Christian Duriez. Finite Element Method-Based Kinematics and Closed-Loop Control of Soft, Continuum Manipulators. *Soft robotics*, 5(3):348–364, 2018.
- [4] Diego Renan Bruno and Fernando Santos Osorio. Image Classification System Based on Deep Learning Applied to The Recognition of Traffic Signs for Intelligent Robotic Vehicle Navigation Purposes. In *Latin American Robotics Symposium (LARS) and Brazilian Symposium on Robotics (SBR)*, 2017.
- [5] Peter L Buzzacott, Erin Zeigler, Petar Denoble, and Richard Vann. American Cave Diving Fatalities 1969-2007. *International Journal of Aquatic Research and Education*, 3(2):7, 2009.
- [6] Massimo Caccia, Gabriele Bruzzone, and Gianmarco Veruggio. Sonar-Based Guidance of Unmanned Underwater Vehicles. *Advanced robotics*, 15(5):551–573, 2001.

- [7] Alexander Cerón, Iván Mondragón, and Flavio Prieto. Onboard Visual-Based Navigation System for Power Line Following with UAV. *Int. Journal of Advanced Robotic Systems*, 15(2), 2018.
- [8] Liang-Chieh Chen. Rethinking Atrous Convolution for Semantic Image Segmentation. *arXiv preprint arXiv:1706.05587*, 2017.
- [9] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4):834–848, 2017.
- [10] Ling Ding, Huyin Zhang, Jinsheng Xiao, Cheng Shu, and Shejie Lu. A Lane Detection Method Based on Semantic Segmentation. *Computer Modeling in Engineering & Sciences*, 122(3):1039–1053, 2020.
- [11] Samsung Energy Business Division. Introduction of INR18650-30Q. <https://bluerobotics.com/wp-content/uploads/2018/10/INR18650-30Q-Data-Sheet.pdf?x70095>, 2014. Accessed: 04-18-2025.
- [12] Vishnu Sashank Dorbala, Gunnar Sigurdsson, Robinson Piramuthu, Jesse Thomason, and Gaurav S Sukhatme. CLIP-Nav: Using CLIP for Zero-Shot Vision-and-Language Navigation. *arXiv preprint arXiv:2211.16649*, 2022.
- [13] Heming Du, Xin Yu, and Liang Zheng. VTNet: Visual Transformer Network for Object Goal Navigation. *arXiv preprint arXiv:2105.09447*, 2021.
- [14] Sheck Exley. *Basic Cave Diving: A Blueprint for Survival*. Cave Diving Section of the National Speleological Society, 1986.
- [15] Derek Ford and Paul Williams. *Introduction to Karst*. John Wiley & Sons, Ltd, 2007.
- [16] Evgeny Gershikov, Tzvika Libe, and Samuel Kosolapov. Horizon Line Detection in Marine Images: Which Method to Choose? *International Journal on Advances in Intelligent Systems*, 6(1), 2013.
- [17] Yogesh Girdhar, Nathan McGuire, Levi Cai, Stewart Jamieson, Seth McCammon, Brian Claus, John E San Soucie, Jessica E Todd, and T Aran Mooney. Curree: A curious underwater robot for ecosystem exploration. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11411–11417. IEEE, 2023.
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [19] Khadija Himri, Pere Ridao, Nuno Gracias, Albert Palomer, Narcís Palomeras, and Roger Pi. Semantic SLAM for an AUV Using Object Recognition from Point Clouds. *IFAC-PapersOnLine*, 51(29):360–365, 2018.
- [20] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. Searching for MobileNetV3. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1314–1324, 2019.
- [21] Chenguang Huang, Oier Mees, Andy Zeng, and Wolfram Burgard. Visual Language Maps for Robot Navigation. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 10608–10615, 2023.
- [22] Md Jahidul Islam, Youya Xia, and Junaed Sattar. Fast Underwater Image Enhancement for Improved Visual Perception. *IEEE Robotics and Automation Letters (RA-L)*, 5(2):3227–3234, 2020.
- [23] Md Jahidul Islam, Alberto Quattrini Li, Yogesh A Girdhar, and Ioannis Rekleitis. Computer Vision Applications in Underwater Robotics and Oceanography. *Computer Vision: Challenges, Trends, and Opportunities*, 2024.
- [24] Bharat Joshi, Sharmin Rahman, Michail Kalaitzakis, Brennan Cain, James Johnson, Marios Xanthidis, Nare Karapetyan, Alan Hernandez, Alberto Quattrini Li, Nikolaos Vitzilaos, and Ioannis Rekleitis. Experimental Comparison of Open Source Visual-Inertial-Based State Estimation Algorithms in the Underwater Domain. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7221–7227, 2019.
- [25] Masoud R Kazemi. Reliability Based Analysis and Design of Anchor Retrofitted Concrete Gravity Dams. In *World Conference on Earthquake Engineering*, 2004.
- [26] Michael J Lace and John E Mylroie. The Biological and Archaeological Significance of Coastal Caves and Karst Features. In *Coastal Karst Landforms*, pages 111–126. Springer, 2013.
- [27] Jin Han Lee, Sehyung Lee, Guoxuan Zhang, Jongwoo Lim, Wan Kyun Chung, and Il Hong Suh. Outdoor Place Recognition in Urban Environments Using Straight Lines. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 5550–5557, 2014.
- [28] Yi Li, Andres Villada, Shao-Hao Lu, He Sun, Jianliang Xiao, and Xueju Wang. Soft, Flexible Pressure Sensors for Pressure Monitoring Under Large Hydrostatic Pressure and Harsh Ocean Environments. *Soft Matter*, 19(30):5772–5780, 2023.
- [29] Marianela Machuca Macías, José Hermenegildo García-Ortiz, Taygoara Felamingo Oliveira, and Antonio Cesar Pinho Brasil Junior. Numerical Investigation of Dimensionless Parameters in Carangiform Fish Swimming Hydrodynamics. *Biomimetics*, 9(1), 2024.
- [30] Angelos Mallios, Pere Ridao, David Ribas, Marc Carreras, and Richard Camilli. Toward Autonomous Exploration in Confined Underwater Environments. *Journal of Field Robotics*, 33(7): 994–1012, 2016.
- [31] Travis Manderson, Juan Camilo Gamboa Higuera, Stefan Wapnick, Jean-François Tremblay, Florian Shkurti, David Meger, and Gregory Dudek. Vision-Based Goal-Conditioned Policies for Underwater Navigation in the Presence of Obstacles, 2020.
- [32] Alfredo Martins, José Almeida, Carlos Almeida, André Dias, Nuno Dias, Jussi Aaltonen, Arttu Heininen, Kari T Koskinen, Claudio Rossi, Sergio Dominguez, et al. Ux 1 system design—a robotic system for underwater mining exploration. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1494–1500. IEEE, 2018.
- [33] Quentin Massone, Sebastien Druon, Yohan Breux, and Jean Triboulet. Contour-based Approach for 3D Mapping of Underwater Galleries. In *Global Oceans: Singapore–US Gulf Coast*, 2020.
- [34] Mohammadreza Mohammadi, Adnan Abdullah, Aishneet Juneja, Ioannis Rekleitis, Jahidul Islam, and Ramtin Zand. Edge-Centric Real-Time Segmentation for Autonomous Underwater Cave Exploration. In *IEEE International Conference on Machine Learning and Applications (ICMLA)*, 2024.
- [35] Nvidia. Nvidia TensorRT. <https://github.com/NVIDIA/TensorRT>, 2016. Accessed: 01-22-2025.
- [36] ONNX. Open Neural Network Exchange. <https://onnx.ai/>, 2017. Accessed: 01-22-2025.
- [37] Shivam K Panda, Yongkyu Lee, and M Khalid Jawed. Agronav: Autonomous Navigation Framework for Agricultural Robots and Vehicles using Semantic Segmentation and Semantic Line Detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6272–6281, 2023.
- [38] Pedro Patrón, Emilio Miguelanez, Joel Cartwright, and Yvan R Petillot. Semantic Knowledge-Based Representation for Improving Situation Awareness in Service Oriented Agents of Autonomous Underwater Vehicles. In *OCEANS*, 2008.
- [39] Pedro Patron et al. *Semantic-based Adaptive Mission Planning for Unmanned Underwater Vehicles*. PhD thesis, Citeseer, 2010.
- [40] Leah Potts, Peter Buzzacott, and Petar Denoble. Thirty Years of American Cave Diving Fatalities. *Diving Hyperbaric Medicine*, 46:150–154, 2016.
- [41] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning Transferable Visual

- Models from Natural Language Supervision. In *International conference on machine learning*, pages 8748–8763, 2021.
- [42] Sharmin Rahman, Alberto Quattrini Li, and Ioannis Rekleitis. Sonar Visual Inertial SLAM of Underwater Structures. In *IEEE International Conference on Robotics and Automation*, pages 5190–5196, 2018.
- [43] Sharmin Rahman, Alberto Quattrini Li, and Ioannis 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)*, pages 8048–8053, 2019.
- [44] Sharmin Rahman, Alberto Quattrini Li, and Ioannis Rekleitis. SVIn2: A Multi-sensor Fusion-based Underwater SLAM System. *International Journal of Robotics Research*, 41(11-12): 1022–1042, July 2022.
- [45] Arturo L Rankin, Carl D Crane III, and David G Armstrong II. Evaluating a pid, pure pursuit, and weighted steering controller for an autonomous land vehicle. In *Mobile Robots XII*, volume 3210, pages 1–12. SPIE, 1998.
- [46] Kristof Richmond, Chris Flesher, Laura Lindzey, Neal Tanner, and William C Stone. Sunfish®: A human-portable exploration auv for complex 3d environments. In *OCEANS 2018 MTS/IEEE Charleston*, pages 1–9. IEEE, 2018.
- [47] Kristof Richmond, Chris Flesher, Neal Tanner, Vickie Siegel, and William C Stone. Autonomous Exploration and 3-D Mapping of Underwater Caves with the Human-portable SUN-FISH® AUV. In *Oceans: Singapore–US Gulf Coast*, 2020.
- [48] Avilash Sahoo, Santosha K Dwivedy, and PS Robi. Advancements in The Field of Autonomous Underwater Vehicle. *Ocean Engineering*, 181:145–160, 2019.
- [49] Moveh Samuel, Mohamed Hussein, and Maziah Binti Mohamad. A review of some pure-pursuit based path tracking techniques for control of autonomous vehicle. *International Journal of Computer Applications*, 135(1):35–38, 2016.
- [50] Louis L Scharf, William P Harthill, and Paul H Moose. A comparison of expected flight times for intercept and pure pursuit missiles. *IEEE Transactions on Aerospace and Electronic Systems*, pages 672–673, 1969.
- [51] Dhruv Shah, Błażej Osiński, Sergey Levine, et al. Lm-Nav: Robotic Navigation with Large Pre-trained Models of Language, Vision, and Action. In *Conference on robot learning*, pages 492–504, 2023.
- [52] Florian Shkurti, Wei-Di Chang, Peter Henderson, Md Jahidul Islam, Juan Camilo Gamboa Higuera, Jimmy Li, Travis Mander-son, Anqi Xu, Gregory Dudek, and Junaed Sattar. Underwater Multi-Robot Convoying using Visual Tracking by Detection. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4189–4196, 2017.
- [53] Thomas Jeongho Song. Experimental Evaluation of Underwater Semantic SLAM. Master’s thesis, Massachusetts Institute of Technology, 2024.
- [54] Sumedh Sontakke, Jesse Zhang, Séb Arnold, Karl Pertsch, Erdem Bıyık, Dorsa Sadigh, Chelsea Finn, and Laurent Itti. Roboclip: One Demonstration is Enough to Learn Robot Policies. *Adv. in Neural Information Processing Systems*, 36, 2024.
- [55] Lorenzo Steccanella, Domenico Bloisi, Jason Blum, and Alessandro Farinelli. Deep Learning Waterline Detection for Low-Cost Autonomous Boats. In *Proceedings of the 15th International Intelligent Autonomous Systems Conference (IAS-15)*, pages 613–625. Springer, 2019.
- [56] Satoshi Suzuki and Keiichi Abe. Topological Structural Analysis of Digitized Binary Images by Border Following. *Computer Vision, Graphics, and Image Processing*, 30(1):32–46, 1985.
- [57] Mehron Talebi, Sultan Mahmud, Adam Khalifa, and Md Jahidul Islam. BlueME: Robust Underwater Robot-to-Robot Communication Using Compact Magnetolectric Antennas. *ArXiv preprint arXiv:2411.09241*, 2025.
- [58] SafetyCulture Content Team. Factor of Safety: Ratio for Safety in Design and Use. <https://safetyculture.com/topics/factor-of-safety/>, 2024. Accessed: 01-29-2025.
- [59] Pedro V Teixeira, Dehann Fourie, Michael Kaess, and John J Leonard. Dense, Sonar-based Reconstruction of Underwater Scenes. In *International Conference on Intelligent Robots and Systems (IROS)*, pages 8060–8066, 2019.
- [60] Guillem Vallicrosa, Khadija Himri, Pere Ridao, and Nuno Gracias. Semantic Mapping for Autonomous Subsea Intervention. *Sensors*, 21(20):6740, 2021.
- [61] Weihan Wang, Bharat Joshi, Nathaniel Burgdorfer, Konstantinos Batsos, Alberto Quattrini Li, Philippos Mordohai, and Ioannis Rekleitis. Real-Time Dense 3D Mapping of Underwater Environments. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2023.
- [62] Nicholas Weidner. Underwater Cave Mapping and Reconstruction Using Stereo Vision. Master’s thesis, Computer Science and Engineering Department, University of South Carolina, 2017.
- [63] Nicholas Weidner, Sharmin Rahman, Alberto Quattrini Li, and Ioannis Rekleitis. Underwater Cave Mapping using Stereo Vision. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 5709 – 5715, 2017.
- [64] Anqi Xu, Gregory Dudek, and Junaed Sattar. A Natural Gesture Interface for Operating Robotic Systems. In *IEEE Int. Conf. on Robotics and Automation*, pages 3557–3563, 2008.
- [65] Ping Yang, Changhui Song, Linke Chen, and Weicheng Cui. Image Based River Navigation System of Catamaran USV with Image Semantic Segmentation. In *WRC Symposium on Advanced Robotics and Automation*, pages 147–151, 2022.
- [66] Boxiao Yu, Reagan Tibbetts, Titon Barna, Ailani Morales, Ioannis Rekleitis, and Md Jahidul Islam. Weakly Supervised Caveline Detection For AUV Navigation Inside Underwater Caves. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 9933–9940. IEEE, 2023.
- [67] Fei Yu, Yuemei Zhu, Qi Wang, Kaige Li, Meihan Wu, Guangliang Li, Tianhong Yan, and Bo He. Segmentation of Side Scan Sonar Images on AUV. In *IEEE Underwater Technology (UT)*, 2019.
- [68] Wenqiang Zhan, Changshi Xiao, Yuanqiao Wen, Chunhui Zhou, Haiwen Yuan, Supu Xiu, Xiong Zou, Cheng Xie, and Qiliang Li. Adaptive Semantic Segmentation for Unmanned Surface Vehicle Navigation. *Electronics*, 9(2):213, 2020.