

DREAM: Domain-aware Reasoning for Efficient Autonomous Underwater Monitoring

Zhenqi Wu^{1*}, Abhinav Modi^{2*}, Angelos Mavrogiannis², Kaustubh Joshi²,
Nikhil Chopra², Yiannis Aloimonos², Nare Karapetyan³, Ioannis Rekleitis⁴, Xiaomin Lin^{1,†},

Abstract—The ocean is warming and acidifying, increasing the risk of mass mortality events for temperature-sensitive shellfish such as oysters. This motivates the development of long-term monitoring systems. However, human labor is costly and long-duration underwater work is highly hazardous, thus favoring robotic solutions as a safer and more efficient option. Yet deploying such robots for persistent, wide-area benthic monitoring demands real-time, environment-aware decision-making without human intervention, a capability that existing systems still lack. To this end, we present *DREAM*, a Vision Language Model (VLM)-guided autonomy framework for long-term underwater exploration and monitoring. It autonomously explores the seafloor, detects and localizes objects of interest such as oyster clusters, and builds a spatial coverage map of their distribution. *DREAM* couples (i) a reasoning-augmented prompt that guides VLM planning with (ii) an occupancy map providing memory and overview, and (iii) a low-level controller to realize actions. Our framework outperforms all baselines across both tasks: in oyster monitoring, it uses 23.0% fewer steps and covers 8.9% more oysters than the vanilla VLM, and completes the task 31.5% faster than UIVNAV. In shipwreck exploration, it achieves 100% coverage versus 60.2% for the vanilla model with 27.5% fewer steps. All code and prompts can be found at <https://github.com/zhenqi72/DREAM>.

I. INTRODUCTION

Oceans are the lifeline of our planet, nurturing entire ecosystems [1] that are essential to human survival and controlling the global climate of the earth [2]. Despite their vast size, a significant portion of the oceans has not yet been explored or mapped, hindering our ability to safeguard vital ecosystems and anticipate their responses to global environmental shifts [3]. Direct human exploration can often be too risky. For example, monitoring a coral reef might require maneuvering through confined areas or environments that often lie beyond the reach of human divers due to additional adverse conditions such as extreme pressure, limited visibility, and the physical risks involved.

Deploying Remotely Operated Vehicles (ROVs) offers a safer and more practical solution but can be challenging

This work was supported by USDA NIFA sustainable agriculture system program under award number 20206801231805.

* Equal Contributors, † Corresponding Author

¹Electrical Engineering, University of South Florida, Tampa, FL, 33620, USA. Emails: {zhenqi, xlin2}@usf.edu

²Maryland Robotics Center (MRC), University of Maryland, College Park, MD 20742, USA. Emails: {abhil625, angelosm, kjoshi, nchopra, jyaloimo}@umd.edu.

³Woods Hole Oceanographic Institution (WHOI), Woods Hole, MA, 02543, USA. Emails: nare@whoi.edu

⁴Mechanical Engineering, University of Delaware, Newark, DE 19716, USA. Emails: yiannisr@udel.edu




Fig. 1: Example of *DREAM* deployed in the real world on a BlueROV surveying an oyster reef in a pool. The top left image shows a sample observation from the robot's camera.

in uncharted underwater environments and inefficient for long-term monitoring [4]. ROV operators often need to collaborate closely with ocean scientists, but differences in background and expertise can complicate communication. Furthermore, the survival of vital ecosystems requires long-term continuous monitoring, which can make the entire operation prohibitively expensive due to the costs associated with logistics, specialized personnel, and equipment [5]. Vision-based monitoring methods have shown promise in specific domains, with applications in surveying coral reefs [6], shipwrecks [7], and oysters [8]. However, these methods are inherently imitation learning-based and have no feedback from the outer world. Recent systems such as OceanPlan [9], OceanChat [10], and AquaChat++ [11] advance natural language piloting, and Xu et al. [12]–[14] contribute multiple works on multi-AUV decision making, coordination, and adaptive control, but still lack robust reasoning, adaptability, and persistent spatial memory. To address these limitations, we propose a new VLM-based underwater monitoring System, *Domain-aware Reasoning for Efficient Autonomous Underwater Monitoring (DREAM)*. The system leverages human diver knowledge, oceanic domain knowledge, together with real-time visual input from an onboard camera to generate adaptive monitoring policies that underwater robots can execute. The main contributions of this work are:

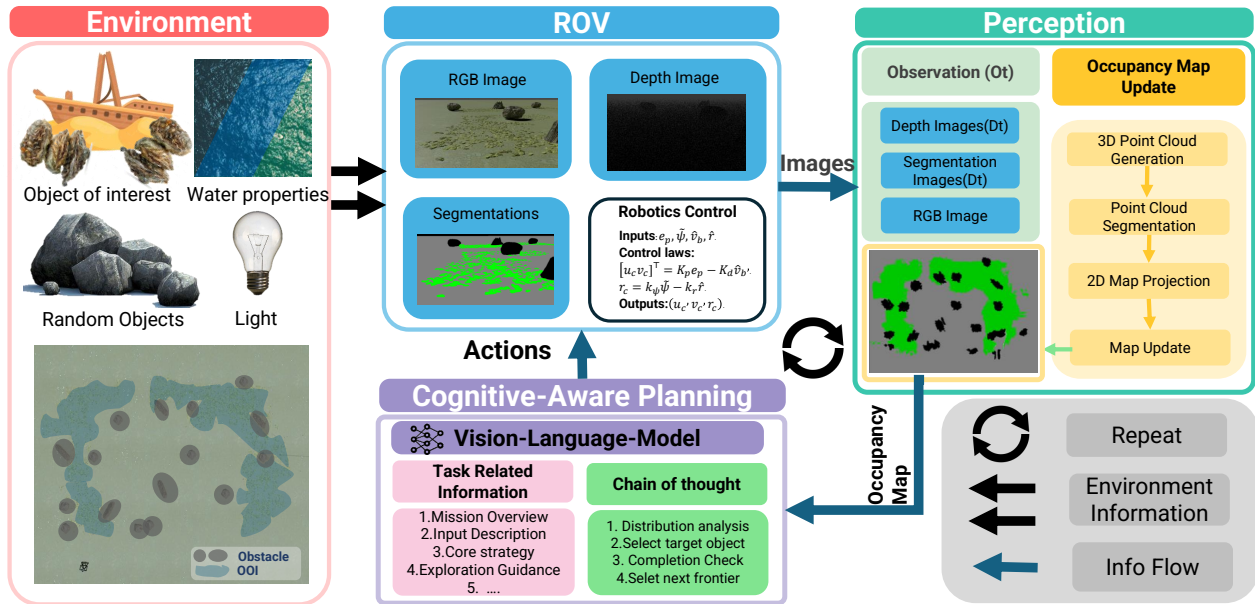


Fig. 2: An overview of the DREAM framework. The environment provides multimodal inputs (RGB, depth, and segmentation) captured by the ROV. These observations are fused in the Perception module to build and update occupancy maps. Cognitive-Aware Planning leverages a vision–language model with chain-of-thought reasoning to guide frontier selection, mission planning, and persistent monitoring of underwater objects of interest. Robotic control then executes movement commands, closing the loop for continual exploration and monitoring.

- 1) A 3-layer architecture of perception, cognitive-aware planning, and control, that enables end-to-end monitoring, from perception to high-level reasoning and finally to low-level adaptive execution.
- 2) A VLM-based framework that integrates domain knowledge with sequential reasoning (via Chain-of-Thought [15]) to produce persistent monitoring policies without relying solely on localization.
- 3) Real-world deployment (Fig. 1) of our framework on an underwater robot, demonstrating feasibility of efficient exploration in marine environments for various object of interests.
- 4) We will open-source our synthetic environments, real-world dataset, and our code to facilitate further underwater robotics monitoring capabilities.

The rest of this paper is organized in the following manner. First, in Section II, we discuss the recent advances in the field. The problem formulation and proposed approach are presented in Section III. Experimental validation in simulation and the lab environment are discussed in Sections IV, V, respectively. Finally, a conclusion with discussion on future work is presented in Section VI.

II. RELATED WORK

Large Language Models have revolutionized the field of robotics as a source of commonsense reasoning that bridges multimodal input from the environment, such as textual instructions or visual observations, to high-level actions [16]–[20]. These actions are often coupled with low-level controllers [21] or robot-specific APIs [22] to

yield explicit robot control commands and hence entire actionable motion plans. In our work, we leverage visual observations to generate high-level robot actions and bridge simulation to reality by generating control commands through an integrated PD controller. While our system yields high-level actions in simulation, our approach is modular and can be combined with complementary work that performs obstacle avoidance and local navigation in underwater environments [23], [24].

Autonomous robot navigation has been studied extensively across both ground environments [25]–[28] and underwater environments [29]–[31]. While ground navigation benefits from reliable sensing and localization [32], underwater environments introduce unique challenges [8], [33]–[36]. Many approaches have been developed for underwater monitoring and exploration, with applications ranging from cave exploration [37], [38] to oyster detection [39]–[42]. Coverage path planning methods, from boustrophedon decomposition [43] to frontier-based exploration [44], address area coverage through geometric criteria but lack semantic understanding of the environment. Furthermore, these approaches lack the adaptive feedback loop that our approach introduces, where semantic and ecological context, such as the distribution of oyster clusters, continuously guides replanning.

Recent advances include adapting LLMs and VLMs to underwater and marine robots, and researchers have even finetuned such models [45]–[47] on ocean science tasks as they enable online decision making powered by commonsense reasoning, making them suitable for

long-term autonomous monitoring without the constant need of a human in the loop. Xie et al. [48] present an LLM-enhanced reinforcement learning controller that jointly tunes rewards and control parameters, improving maneuvering under extreme sea conditions. SeafloorAI [49] expands data resources for geology-focused perception, strengthening visual grounding and instruction following for seafloor surveys. Word2Wave [50] offers a language-driven mission interface that maps high-level goals to efficient subsea deployments. Furthermore, OceanChat [10] enables natural language piloting with dynamic replanning, AquaChat++ [11] coordinates multi-ROV aquaculture inspection with fault-aware task planning, and Autonomous Vehicle Maneuvering [51] integrates visual grounding with LLM-guided path planning for surface vehicles. While OceanPlan [9] translates natural language instructions into AUV trajectories for long-horizon tasks, it does not construct a spatial memory of the environment. In contrast, our approach integrates a dynamically built occupancy map with structured Chain-of-Thought reasoning, enabling persistent and efficient long-term monitoring for benthic environments. These works collectively progress from instruction understanding to robust low-level control utilizing data and tooling that will allow long-horizon planning and safety-aware autonomy in the ocean. Our approach advances this line of work by coupling visual observations with planning and integrated control.

III. APPROACH

As depicted in Fig. 2, our framework has three modules: perception, cognitive-aware planning, and control. In the perception module, the ROV’s front-facing camera provides depth, semantic-segmentation, and RGB images. We fuse the depth and segmentation images to build an occupancy map, which is incrementally updated throughout the exploration process. In the planning module, the VLM receives these inputs together with a well-crafted prompt and chain-of-thought instructions; it focuses on the objects of interest and outputs high-level actions. Because VLM inference is relatively slow, we only use it for high-level navigation rather than direct low-level control. Finally, the control module executes those actions, translating discrete speed, direction, and turn-angle commands into motion to achieve efficient, long-term monitoring.

A. Perception phase

The perception phase encompasses how the robot senses, processes, and interprets information from its environment to build a structured representation that supports downstream decision-making and planning. At each step, the ROV collects front-facing RGB, depth, and semantic segmentation images. These multimodal inputs are fused through a multi-stage pipeline to generate spatial representations in the form of occupancy maps. Depth pixels are first transformed into 3D point clouds and projected into global coordinates using camera-to-world transformations. Given a pixel (u, v)

with depth $z(u, v)$ and intrinsics (f_x, f_y, c_x, c_y) , its 3D position in the world frame is obtained by

$$\mathbf{p}_w(u, v) = \Pi \left(\mathbf{T}_c^w \begin{bmatrix} \frac{(u-c_x)}{f_x} z(u, v) \\ \frac{(v-c_y)}{f_y} z(u, v) \\ z(u, v) \\ 1 \end{bmatrix} \right), \quad (1)$$

where $\Pi([x \ y \ z \ w]^\top) = [x/w, y/w, z/w]^\top$, subject to the depth validity constraint

$$z_{\min} \leq z(u, v) \leq z_{\max}. \quad (2)$$

where

- z_{\min} is the minimum reliable depth (e.g., 0.1 m) used to suppress sensor noise or self-reflections, and
- z_{\max} is the maximum reliable depth (e.g., 20 m) beyond which measurements are treated as invalid due to noise and range limitations.

Height thresholds classify the resulting points: elevated regions are marked as obstacles, while target objects (e.g., oysters) are filtered using semantic labels. Let $m(u, v) \in \{0, 1\}$ denote the semantic mask for the object of interest (e.g., oysters), with $m = 1$ indicating a positive detection. After applying the depth-validity constraint and the camera-to-world transform (Equation 1), we define the set of valid world points

$$\mathcal{P}_{\text{valid}} = \{\mathbf{p}_w = (x_w, y_w, z_w)^\top : z_{\min} \leq z(u, v) \leq z_{\max}\}. \quad (3)$$

We partition $\mathcal{P}_{\text{valid}}$ into three disjoint regions:

$$\mathcal{P}_{\text{obj}} = \{\mathbf{p}_w \in \mathcal{P}_{\text{valid}} : m(u, v) = 1\}, \quad (4)$$

$$\mathcal{P}_{\text{obs}} = \{\mathbf{p}_w \in \mathcal{P}_{\text{valid}} : h_{\min} \leq z_w \leq h_{\max}\} \setminus \mathcal{P}_{\text{obj}}, \quad (5)$$

$$\mathcal{P}_\emptyset = \mathcal{P}_{\text{valid}} \setminus (\mathcal{P}_{\text{obj}} \cup \mathcal{P}_{\text{obs}}), \quad (6)$$

so that $\mathcal{P}_{\text{valid}} = \mathcal{P}_{\text{obj}} \dot{\cup} \mathcal{P}_{\text{obs}} \dot{\cup} \mathcal{P}_\emptyset$. Here, \mathcal{P}_{obj} denotes target objects on (or near) the seafloor using semantics, \mathcal{P}_{obs} represents the elevated structures considered as obstacles, and \mathcal{P}_\emptyset defines empty (free) space.

Given the robot’s grid pose at time t , denoted by \mathbf{c}_t with yaw ψ_t , the occupancy map is updated using a ray-casting process within the camera’s field-of-view sector (ψ_t, θ) and sensing range d_{\max} over the free space N . This yields the visible set V_t , which incrementally updates the explored mask, E_t :

$$E_t = E_{t-1} \cup V_t, \quad (7)$$

$$V_t = \text{RayCast}(\mathbf{c}_t, \psi_t, \theta, d_{\max}, N). \quad (8)$$

As illustrated in Fig. 3, the blue cone represents the instantaneous field of view (FOV) of the camera and the occupancy map is generated by tracing rays through this FOV. Once a ray intersects an occupied cell, all subsequent cells along that line of sight are marked as *unknown*, since the sensor cannot see beyond obstacles. These occluded regions appear as white gaps in the occupancy map and can only be

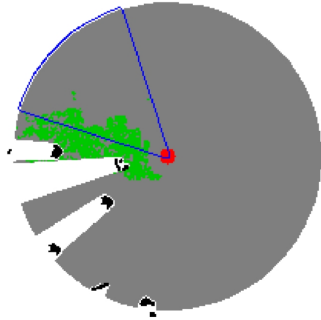


Fig. 3: FOV demonstration on the occupancy map

revealed through exploration from different viewpoints. This distinction between navigable space, occupied space, target objects, and unknown space is critical for safe navigation and for guiding next-best-view exploration. The occupancy grid also incorporates safety margins to account for the robot’s physical dimensions. In real-world trials, the same RGB, depth, and segmentation modalities are captured by the ROV’s onboard sensors and the occupancy map is incrementally updated using depth, semantic segmentation, and pose estimates. This map provides the VLM with both a concise representation of the surrounding structure and a persistent memory of explored regions.

The exploration process proceeds in two stages. First, the robot performs an in-place rotation to acquire a panoramic view, constructing a 360° occupancy map of its immediate surroundings. This initialization phase enables the robot to identify safe entry directions and plan its first movements. In the second stage, the robot executes exploration actions guided by next-best-view predictions, incrementally expanding its map and knowledge of the environment.

B. Cognitive-Aware Planning phase

After the cognition phase, the VLM receives the front-camera RGB, segmentation, and depth images, along with the occupancy map. These inputs provide the ROV with situational context, the locations of target objects, and their spatial distribution. To ensure the model understands the task, we supply a carefully crafted prompt that enumerates the inputs and units, defines the mission, specifies clear standards for abstract instructions, outlines exploration guidance and decision priorities, and lists the available actions. Since we want to have long-term monitoring on our target objects, we define our task: “*Your mission is to efficiently and comprehensively discover and map all oyster clusters (target object) on the seafloor.*” Upon experimentation, we find that VLMs can struggle with vague directives (e.g., “*finish exploring the nearby oyster area*”), therefore we define completion precisely: green (oyster) regions on the occupancy map must be fully enclosed by gray (explored) cells, with no white (missed) patches inside. As with human operators, explicit guidance is crucial; without it, the robot tends to move randomly. We therefore instruct the agent to follow the local continuity of oyster distributions,

expanding coverage around nearby clusters rather than jumping to distant ones, to improve efficiency. Nevertheless, a well-crafted prompt alone does not always yield coherent inference; the model can still lack a clear reasoning path.

Thus, we leverage Chain of Thought (CoT) [15], a sequence of intermediate reasoning steps that can substantially improve complex reasoning. We design a hand-crafted CoT prompt that encodes an efficient, human-like search logic and yields a robust, transparent reasoning path that consists of the following steps:

- 1) Distribution analysis: infer local oyster layout and density from segmentation/occupancy.
- 2) Select current target area: choose the nearest/high-value frontier on the oyster boundary.
- 3) Completion check: verify whether the current region satisfies the completion criterion.
- 4) Select next target/frontier: prioritize adjacent, denser oyster areas; avoid long, low-yield jumps.
- 5) Safety & feasibility: confirm standoff distance and collision-free motion using depth/occupancy.
- 6) Action selection: output direction (left/right/forward), turn angle, and step length from the discrete action set.

C. Control

We use a planar PD controller for positioning and yaw regulation, while vertical motion is stabilized by the vehicle’s depth/altitude hold. Let $p = [x, y, z]^T$ be the position of the robot and ψ be its yaw angle. Like most ROVs, the robot is dynamically stable in its roll and pitch. For a desired waypoint p^* with desired heading ψ^* , the body-frame pose error by rotating the world frame error into the vehicle frame $R(\hat{\psi})$, e_p is defined as

$$e_p = R(\hat{\psi})^T \tilde{p}, \quad R(\hat{\psi}) = \begin{bmatrix} \cos \hat{\psi} & -\sin \hat{\psi} \\ \sin \hat{\psi} & \cos \hat{\psi} \end{bmatrix} \quad (9)$$

where $\tilde{(\cdot)}$ denotes the error $(\cdot)^* - \hat{(\cdot)}$. The velocity command $[u_c, v_c]^T$ and angular yaw rate r_c sent to the robot is driven by the PD controller

$$\begin{bmatrix} u_c \\ v_c \end{bmatrix} = K_p e_p - K_d \hat{v}_b \quad (10)$$

$$r_c = k_\psi \tilde{\psi} - k_r \hat{r} \quad (11)$$

where $\hat{v}_b = [\hat{u}, \hat{v}]^T$ and \hat{r} are the estimated surge/sway velocities and yaw rates, respectively.

An IMU–DVL estimator supplies this estimated state feedback for the PD law. We use an invariant extended Kalman filter (InEKF) [31], [52], [53] as well as a Cascade Iteratively Preconditioned Gradient Descent (C-IPG) [34] observer to propagate the discrete-time inertial dynamics and update the state with body-frame water-relative velocities from the DVL (and pressure depth) when available. These methods give us robust attitude/velocity estimates even during aggressive yawing and intermittent DVL bottom-lock. The update treats DVL as a direct velocity observation in the vehicle frame and applies zero-velocity or near-hover

pseudo-measurements whenever the policy asks the robot to “Stop,” which helps bound drift.

Together, these layers form a simple stack: language actions → short setpoint segments → PD tracking, with state supplied by an observer in the form of InEKF or C-IPG (IMU+DVL) and gently regularized by continuous-time alignment when the scene supports it. We explicitly decoupled perception (oyster segmentation) and control stability, thereby maintaining safety.

IV. EXPERIMENTAL SETUP

A. Environment

We evaluate our approach in Oystersim [54] across 15 environments with obstacles, of which 10 are oyster reefs and 5 are shipwrecks (Fig. 4 top row). Our environment’s length range and width range are both from 40 meters to 50 meters. To mimic natural oyster morphology and show our framework’s adaptability, we put fringing, string (linear), and patch reefs, appearing either individually or in combination. For the shipwreck set, we select five distinct wreck models and place them on the seafloor.

B. Metrics

For all environments, we report the total exploration time, coverage percentage (for oysters we consider the number of oysters and for shipwrecks the area). The total exploration time scores the efficiency of an agent by measuring the total time it takes to finish the entire exploration. An agent’s efficiency is inversely related to its exploration time; greater time denotes lower efficiency and vice versa. We set max steps to 200. If in 200 steps exploration has not concluded, we set the total exploration time to 200 steps. The percentage of coverage shows the quality of exploration. If equal to 100, it shows that the agent can cover all target objects in the desired area.

C. Baseline

We compare our framework with UIVNAV [8], a previous AI-based method for zero-shot object navigation. UIVNAV is an end-to-end underwater navigation approach that leverages intermediate representations to enable zero-shot transfer to new target objects without retraining. Beyond method comparisons, we evaluate our crafted chain-of-thought (CoT) against a naïve prompt (no CoT) to isolate the benefit of structured reasoning. Both prompting strategies use the same VLM; in our experiments, we adopt GPT-5 [55] as the reasoning model.

V. RESULTS

A. Comparison on oysterbed

As shown in Table I and Fig. 4(e)(f), our framework outperforms the baseline and the vanilla model. In this figure, guided by our framework (yellow trajectory), the ROV follows the oyster distribution, remains focused on the target regions, and exhibits minimal redundant exploration. In contrast, the baseline and vanilla model do not track the oyster pattern and often diverge to unrelated areas. More

TABLE I: Simulation Results — Comparison with UIVNAV (Oyster Environment)

Method	Prompt Type	Avg. Exploration	Avg. Coverage
		Time (steps)	Rate (%)
UIVNAV		200	55.5%
Our Framework		137.3	98.3%

TABLE II: Simulation Results — Comparison with Vanilla VLM Baseline

Environment	Framework	Our framework	GPT-5
	Exploration Time-Oyster (steps)		137.3
Exploration Time-Shipwreck (steps)		145	200
Average Coverage Rate-Oyster		98.3%	90.6%
Average Coverage Rate-Shipwreck		100%	60.2%

specifically, in oyster scenarios, our framework uses 31.5% less time than UIVNAV and achieves 98.3% coverage within the same time budget. In contrast, the baseline only has a 55.5% average coverage within the limit, and the vanilla model has an average coverage of 90.6%. Therefore, we observe that no baseline can cover more oysters than our approach. On the other hand, our framework efficiently explores and maps the entire wreck without collisions, attaining an average coverage of 100% with an average exploration time of 145 steps.

B. Comparison on shipwreck

Using vanilla mode, the ROV exhibits inconsistent behavior. As shown in Table II and Fig. 4(g)(h), our path (yellow trajectory) circumnavigates the wreck and achieves full coverage without collisions, whereas the vanilla model fails to cover the entire structure and at times drifts outside the boundary. Finally, the vanilla model has 200 steps with 60.2% coverage in shipwreck environments. Coupled with our framework, the ROV uses 27.5% fewer steps and achieves 66.0% higher coverage in shipwreck scenarios. Note that UIVNAV is omitted from the shipwreck evaluation, as it was designed for overhead survey tasks where the ROV flies directly above the target, a fundamentally different navigation regime from circumnavigating a vertical structure.

C. Real-world deployment

To demonstrate feasibility in real-world settings, the framework was tested in a scenario where oyster shells and a pipe (to simulate a shipwreck) were laid out in a tank, and a BlueROV2 robot platform was used to survey the environment. The robot platform used is a BlueROV2 equipped with an onboard camera. Other sensors on the BlueROV2 include an onboard IMU and a retrofitted Waterlinked DVL-A50 to localize and control the ROV’s position.

We evaluated the system in a circular tank that is 12ft in diameter and 5ft deep. 200 oyster shells were

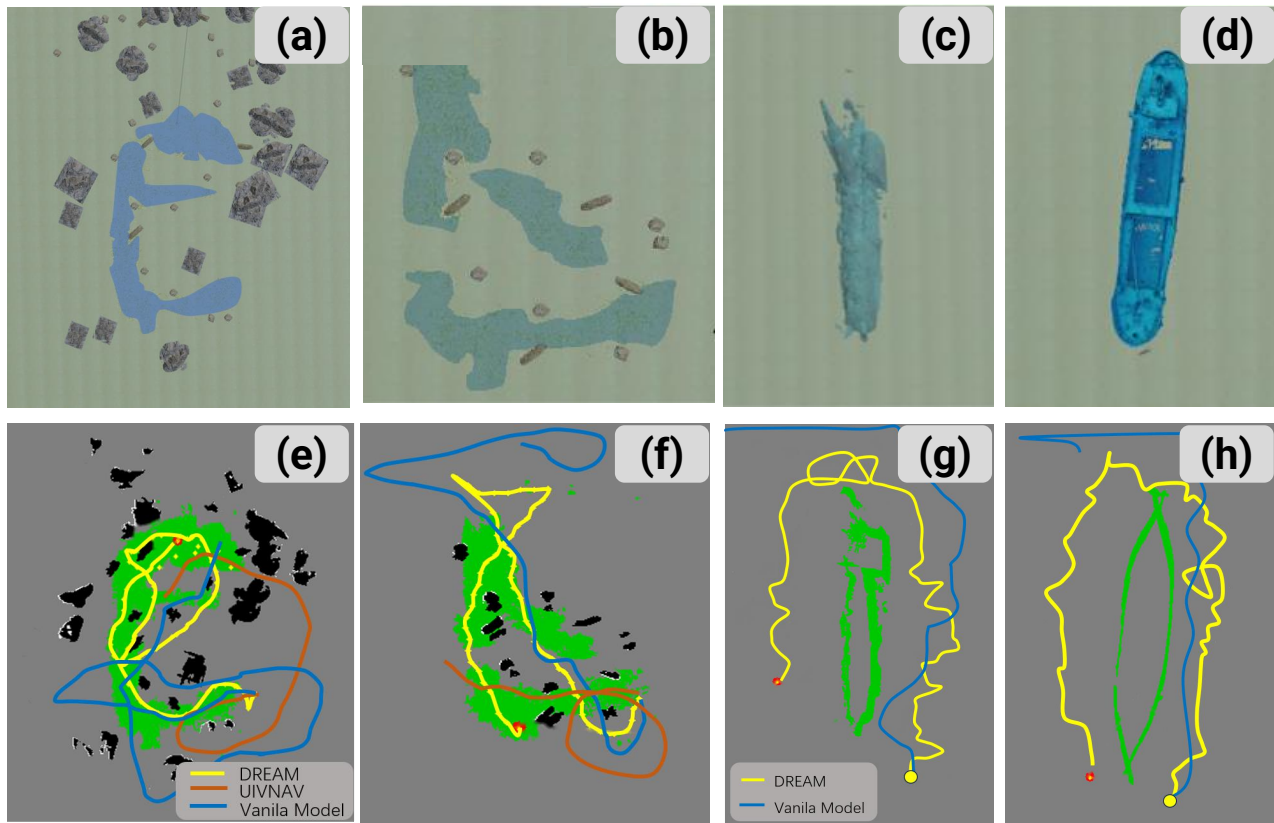


Fig. 4: Testing environments and algorithmic comparisons. Top row: simulated oyster reef (left) and shipwreck (right) environments. Bottom row: comparison of our framework against baseline methods, showing improved efficiency and coverage in both oyster and shipwreck monitoring tasks.

arranged in a circular arc pattern and two pipes linearly as shown in Fig. 5. During trials, the vehicle faced realistic pool-scale disturbances: slow cross-currents, tether drag, small turn-radius asymmetries due to hydrodynamic coupling, and intermittent visual degradation from surface reflections. Despite these effects, the pipeline executed end-to-end in real time, with the front camera producing frames for on-board oyster segmentation, and the PD controller ensuring stability and positions by estimating odometry from C-IPG, as well as InEKF.

A representative run is shown in Fig. 6. Starting without a prior map, the robot first stabilized yaw and altitude, then advanced with gentle, $0.1m$ steps, favoring directions where segmented oysters densified. The step size of $0.1m$ was selected based on the limited dimensions of the tank of $15ft$. Qualitatively, this experiment exhibits near-complete traversal of the oyster patch with only a few missed interior pockets.

VI. CONCLUSION AND FUTURE WORK

In this paper, we propose a VLM-based framework for underwater monitoring. The system constructs an occupancy map to provide the VLM with a compact overview of the surroundings and a persistent memory of previously visited locations. Conditioned on target object, the VLM outputs high-level actions. By supplying a well-crafted prompt




Fig. 5: Pool setup with oyster shells arranged in a circular arc and two pipes positioned to emulate a shipwreck.

and a chain-of-thought (CoT) scaffold, we endow the model with stronger situational understanding and efficient, human-like navigation capabilities. In comprehensive simulation and real-world experiments, our framework demonstrates superior performance, adaptability, efficiency, and complete coverage capability. Although our experiments are conducted in relatively small-scale environments due to time and hardware limitations, they serve as an initial step

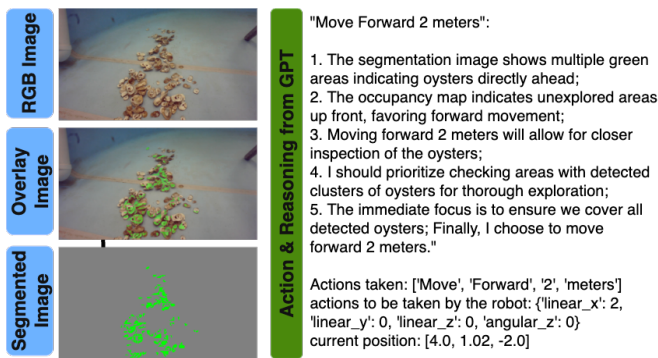


Fig. 6: Example of a real-world experiment and the low-level actions produced by our framework.

towards persistent long-term monitoring.

For future work, we plan to leverage domain-adapted visual encoders and underwater foundation models such as AquaticCLIP [47] to improve generalization to open-water and high-turbidity conditions [56], [57]. We also intend to integrate disturbance-aware control strategies to handle ocean currents and tether drag in open-water deployments, and incorporate formal safety constraints into the planning framework to provide stronger guarantees in cluttered environments.

REFERENCES

- [1] D. Ward, J. Melbourne-Thomas, G. T. Pecl, K. Evans, M. Green, P. C. McCormack, C. Novaglio, R. Trebilco, N. Bax, M. J. Brasier, *et al.*, "Safeguarding marine life: conservation of biodiversity and ecosystems," *Reviews in fish biology and fisheries*, vol. 32, no. 1, pp. 65–100, 2022.
- [2] G. R. Bigg, T. D. Jickells, P. S. Liss, and T. J. Osborn, "The role of the oceans in climate," *International Journal of Climatology: A Journal of the Royal Meteorological Society*, vol. 23, no. 10, pp. 1127–1159, 2003.
- [3] L. J. McKenzie, L. M. Nordlund, B. L. Jones, L. C. Cullen-Unsworth, C. Roelfsema, and R. K. Unsworth, "The global distribution of seagrass meadows," *Environmental Research Letters*, vol. 15, no. 7, p. 074041, 2020.
- [4] G. Li, T.-W. Wong, B. Shih, C. Guo, L. Wang, J. Liu, T. Wang, X. Liu, J. Yan, B. Wu, *et al.*, "Bioinspired soft robots for deep-sea exploration," *Nature Communications*, vol. 14, no. 1, p. 7097, 2023.
- [5] D. L. McLean, M. J. Parsons, A. R. Gates, M. C. Benfield, T. Bond, D. J. Booth, M. Bunce, A. M. Fowler, E. S. Harvey, P. I. Macreadie, *et al.*, "Enhancing the scientific value of industry remotely operated vehicles (rovs) in our oceans," *Frontiers in Marine Science*, vol. 7, p. 220, 2020.
- [6] T. Manderson, J. C. G. Higuera, R. Cheng, and G. Dudek, "Vision-based autonomous underwater swimming in dense coral for combined collision avoidance and target selection," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1885–1891.
- [7] N. Karapetyan, J. V. Johnson, and I. Rekleitis, "Human diver-inspired visual navigation: Towards coverage path planning of shipwrecks," *Marine Technology Society Journal*, vol. 55, no. 4, pp. 24–32, 2021.
- [8] X. Lin, N. Karapetyan, K. Joshi, T. Liu, N. Chopra, M. Yu, P. Tokekar, and Y. Aloimonos, "Uivnav: Underwater information-driven vision-based navigation via imitation learning," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 5250–5256.
- [9] R. Yang, F. Zhang, and M. Hou, "Oceanplan: Hierarchical planning and replanning for natural language auv piloting in large-scale unexplored ocean environments," in *Proceedings of the 18th International Conference on Underwater Networks & Systems*, 2024, pp. 1–5.
- [10] R. Yang, M. Hou, J. Wang, and F. Zhang, "Oceanchat: Piloting autonomous underwater vehicles in natural language," *arXiv preprint arXiv:2309.16052*, 2023.
- [11] A. Saad, W. Akram, and I. Hussain, "Aquachat++: Llm-assisted multi-rov inspection for aquaculture net pens with integrated battery management and thruster fault tolerance," *arXiv preprint arXiv:2508.06554*, 2025.
- [12] J. Xu, Z. Zhang, J. Wang, Z. Han, and Y. Ren, "Multi-auv pursuit-evasion game in the internet of underwater things: An efficient training framework via offline reinforcement learning," *IEEE Internet of Things Journal*, vol. 11, no. 19, pp. 31 273–31 286, 2024.
- [13] G. Xie, J. Xu, Y. Ding, Z. Zhang, S. Zhang, and Y. Li, "Never too prim to swim: An llm-enhanced rl-based adaptive s-surface controller for auvs under extreme sea conditions," in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025, pp. 8990–8997.
- [14] J. Xu, G. Xie, J. Tang, Y. Ding, W. Liu, J. Huang, S. Zhang, and Y. Li, "Never too cocky to cooperate: An fim and rl-based usv-auv collaborative system for underwater tasks in extreme sea conditions," 2025. [Online]. Available: <https://arxiv.org/abs/2504.14894>
- [15] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou, *et al.*, "Chain-of-thought prompting elicits reasoning in large language models," *Advances in neural information processing systems*, vol. 35, pp. 24 824–24 837, 2022.
- [16] B. Ichter, A. Brohan, Y. Chebotar, C. Finn, K. Hausman, A. Herzog, D. Ho, J. Ibarz, A. Irpan, E. Jang, R. Julian, D. Kalashnikov, S. Levine, Y. Lu, C. Parada, K. Rao, P. Sermanet, A. T. Toshev, V. Vanhoucke, F. Xia, T. Xiao, P. Xu, M. Yan, N. Brown, M. Ahn, O. Cortes, N. Sievers, C. Tan, S. Xu, D. Reyes, J. Rettinghouse, J. Quiambao, P. Pastor, L. Luu, K.-H. Lee, Y. Kuang, S. Jesmonth, N. J. Joshi, K. Jeffrey, R. J. Ruano, J. Hsu, K. Gopalakrishnan, B. David, A. Zeng, and C. K. Fu, "Do as i can, not as i say: Grounding language in robotic affordances," in *Proceedings of the Conference on Robot Learning (CoRL)*, vol. 205, 2023, pp. 287–318.
- [17] A. Mavrogiannis, C. Mavrogiannis, and Y. Aloimonos, "Cook2l1l: Translating cooking recipes to l1l formulae using large language models," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 17 679–17 686.
- [18] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, X. Chen, K. Choromanski, T. Ding, D. Driess, A. Dubey, C. Finn, *et al.*, "Rt-2: Vision-language-action models transfer web knowledge to robotic control," *arXiv preprint arXiv:2307.15818*, 2023.
- [19] Y. Mu, Q. Zhang, M. Hu, W. Wang, M. Ding, J. Jin, B. Wang, J. Dai, Y. Qiao, and P. Luo, "Embodiedgpt: Vision-language pre-training via embodied chain of thought," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [20] M. Walczak, U. Kallakuri, and T. Mohsenin, "Atlasv2: Llm-guided adaptive landmark acquisition and navigation on the edge," in *Proceedings of the AAAI Symposium Series*, vol. 5, no. 1, 2025, pp. 196–203.
- [21] J. Liang, W. Huang, F. Xia, P. Xu, K. Hausman, B. Ichter, P. Florence, and A. Zeng, "Code as policies: Language model programs for embodied control," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 9493–9500.
- [22] A. Mavrogiannis, D. Yuan, and Y. Aloimonos, "Discovering object attributes by prompting large language models with perception-action apis," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 9003–9010.
- [23] P. Mane, A. J. George, R. Makam, R. Majumder, and S. Sundaram, "Eroas: 3d efficient reactive obstacle avoidance system for autonomous underwater vehicles using 2.5 d forward-looking sonar," *arXiv preprint arXiv:2411.05516*, 2024.
- [24] J. Yang, M.-Q. Le, M. Gong, and Y. Pu, "Duvin: Diffusion-based underwater visual navigation via knowledge-transferred depth features," *arXiv preprint arXiv:2509.02983*, 2025.
- [25] Y. Lu and E. Plaku, "Leveraging single-goal predictions to improve the efficiency of multi-goal motion planning with dynamics," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 850–857.
- [26] Y. Lu, T. Xu, L. Wang, N. Hawes, and X. Xiao, "Decremental dynamics planning for robot navigation," in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2025, pp. 4559–4565.
- [27] T. Xu, C. Pan, M. B. Rao, A. Datar, A. Pokhrel, Y. Lu, and X. Xiao,

- “Verti-bench: A general and scalable off-road mobility benchmark for vertically challenging terrain,” *arXiv preprint arXiv:2502.11426*, 2025.
- [28] L. Wang, T. Xu, Y. Lu, and X. Xiao, “Reward training wheels: Adaptive auxiliary rewards for robotics reinforcement learning,” in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2025, pp. 15 262–15 267.
- [29] L. Christensen, J. de Gea Fernández, M. Hildebrandt, C. E. S. Koch, and B. Wehbe, “Recent advances in ai for navigation and control of underwater robots,” *Current Robotics Reports*, pp. 1–11, 2022.
- [30] Y. Wu, X. Ta, R. Xiao, Y. Wei, D. An, and D. Li, “Survey of underwater robot positioning navigation,” *Applied Ocean Research*, vol. 90, p. 101845, 2019.
- [31] K. Joshi, T. Liu, A. Williams, M. Gray, X. Lin, and N. Chopra, “3d water quality mapping using invariant extended kalman filtering for underwater robot localization,” *arXiv preprint arXiv:2409.11578*, 2024.
- [32] X. Xiao, Z. Xu, S. A. Ghani, A. Datar, D. Song, P. Stone, A. Mazen, K. Yazdipaz, I. Mateyaunga, M. Faied, *et al.*, “Autonomous ground navigation in highly constrained spaces: Lessons learned from the fourth barn challenge at icra 2025.”
- [33] J. Xu, G. Xie, J. Tang, Y. Ding, W. Liu, S. Zhang, and Y. Li, “Never too cocky to cooperate: An fm and rl-based usv-auv collaborative system for underwater tasks in extreme sea conditions,” *arXiv preprint arXiv:2504.14894*, 2025.
- [34] K. Joshi, T. Liu, and N. Chopra, “Cascade ipg observer for underwater robot state estimation,” *arXiv preprint arXiv:2504.15235*, 2025.
- [35] M. Xanthidis, M. Kalaitzakis, N. Karapetyan, J. Johnson, N. Vitzilaios, J. O’Kane, and I. Rekleitis, “Aquavis: A perception-aware autonomous navigation framework for underwater vehicles,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Prague, Czech Republic, 2021, pp. 5387–5394.
- [36] M. Xanthidis, N. Karapetyan, H. Damron, S. Rahman, J. Johnson, A. O’Connell, J. M. O’Kane, and I. Rekleitis, “Navigation in the presence of obstacles for an agile autonomous underwater vehicle,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 892–899.
- [37] A. Gupta, A. Abdullah, X. Li, V. Ramesh, I. Rekleitis, and M. J. Islam, “Demonstrating cavepi: Autonomous exploration of underwater caves by semantic guidance,” *arXiv preprint arXiv:2502.05384*, 2025.
- [38] M. Chatzisprou, L. Horgan, H. Hwang, H. Sathishchandra, M. Roznere, A. Q. Li, P. Mordohai, and I. Rekleitis, “Mapping the catacombs: An underwater cave segment of the devil’s eye system,” *arXiv preprint arXiv:2507.06397*, 2025.
- [39] X. Lin, V. Mange, A. Suresh, B. Neuberger, A. Palnitkar, B. Campbell, A. Williams, K. Baxevani, J. Mallette, A. Vera, *et al.*, “Odyssey: Oyster detection yielded by sensor systems on edge electronics,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 5290–5297.
- [40] X. Lin, N. J. Sanket, N. Karapetyan, and Y. Aloimonos, “Oysternet: Enhanced oyster detection using simulation,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 5170–5176.
- [41] N. Karapetyan, M. Xu, Y. Zhang, K. Joshi, N. Anil, S. Mayenkar, X. Lin, Y. Aloimonos, N. Chopra, P. Malone, and M. Yu, “Oysterbot: Autonomous surface vehicle based system for aquaculture monitoring operations,” in *OCEANS 2025, Great Lakes*. IEEE, 2025, pp. 1–6.
- [42] C.-Y. Wang, A. G. Nandhan, Y.-T. Shen, W.-Y. Chen, S. S. S. Kumar, A. Long, A. Williams, G. Magnusson, A. Pattillo, D. Webster, *et al.*, “Shellcollect: a framework for smart precision shellfish harvesting using data collection path planning,” *IEEE Access*, 2024.
- [43] E. Galceran and M. Carreras, “A survey on coverage path planning for robotics,” *Robotics and Autonomous systems*, 2013.
- [44] B. Yamauchi, “A frontier-based approach for autonomous exploration,” in *Proceedings 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA’97)*. IEEE, 1997, pp. 146–151.
- [45] Z. Bi, N. Zhang, Y. Xue, Y. Ou, D. Ji, G. Zheng, and H. Chen, “Oceangpt: A large language model for ocean science tasks,” in *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2024, pp. 3357–3372.
- [46] W. Nguyen, A. Phan, K. Kimura, H. Maeno, M. Tanaka, Q. Le, W. Poucher, and C. Nguyen, “Llamarine: Open-source maritime industry-specific large language model,” *arXiv preprint arXiv:2503.00203*, 2025.
- [47] B. Alawode, I. I. Ganapathi, S. Javed, N. Werghi, M. Bennamoun, and A. Mahmood, “Aquaticclip: A vision-language foundation model for underwater scene analysis,” *arXiv preprint arXiv:2502.01785*, 2025.
- [48] G. Xie, J. Xu, Y. Ding, Z. Zhang, S. Zhang, and Y. Li, “Never too prim to swim: An llm-enhanced rl-based adaptive s-surface controller for auvs under extreme sea conditions,” *arXiv preprint arXiv:2503.00527*, 2025.
- [49] K. X. Nguyen, F. Qiao, A. Trembanis, and X. Peng, “Seafloorai: a large-scale vision-language dataset for seafloor geological survey,” *Advances in Neural Information Processing Systems*, vol. 37, pp. 22 107–22 123, 2024.
- [50] R. Chen, D. Blow, A. Abdullah, and M. J. Islam, “Word2wave: Language driven mission programming for efficient subsea deployments of marine robots,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 4107–4114.
- [51] T.-Y. Kim and W.-S. Choi, “Autonomous vehicle maneuvering using vision-llm models for marine surface vehicles,” *Journal of Marine Science and Engineering*, vol. 13, no. 8, p. 1553, 2025.
- [52] E. R. Potokar, K. Norman, and J. G. Mangelson, “Invariant extended kalman filtering for underwater navigation,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5792–5799, 2021.
- [53] A. Barrau and S. Bonnabel, “The invariant extended kalman filter as a stable observer,” *IEEE Transactions on Automatic Control*, vol. 62, no. 4, pp. 1797–1812, 2016.
- [54] X. Lin, N. Jha, M. Joshi, N. Karapetyan, Y. Aloimonos, and M. Yu, “Oystersim: Underwater simulation for enhancing oyster reef monitoring,” in *OCEANS 2022, Hampton Roads*. IEEE, 2022, pp. 1–6.
- [55] OpenAI, “Chatgpt-5: Large language model,” <https://chat.openai.com/>, 2025, accessed: 2025-09-14.
- [56] J. Li, K. A. Skinner, R. M. Eustice, and M. Johnson-Roberson, “WaterGAN: Unsupervised generative network to enable real-time color correction of monocular underwater images,” *IEEE Robotics and Automation Letters*, vol. 3, no. 1, pp. 387–394, 2018.
- [57] M. J. Islam, Y. Xia, and J. Sattar, “Fast underwater image enhancement for improved visual perception,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3227–3234, 2020.