

USIM and U0: A Vision-Language-Action Dataset and Model for General Underwater Robots

Junwen Gu^{1,3†}, Zhiheng Wu^{2†}, Pengxuan Si³, Shuang Qiu^{1,3}, Yukai Feng^{1,3},
Luoyang Sun^{1,3}, Laien Luo^{1,3}, Lianyi Yu^{1,3}, Jian Wang^{1,3*}, and Zhengxing Wu^{1,3}

Abstract—Underwater environments present unique challenges for robotic operation, including complex hydrodynamics, limited visibility, and constrained communication. Although data-driven approaches have advanced embodied intelligence in terrestrial robots and enabled task-specific autonomous underwater robots, developing underwater intelligence capable of autonomously performing multiple tasks remains highly challenging, as large-scale, high-quality underwater datasets are still scarce. To address these limitations, we introduce USIM, a simulation-based multi-task Vision-Language-Action (VLA) dataset for underwater robots. USIM comprises over 561K frames from 1,852 trajectories, totaling approximately 15.6 hours of BlueROV2 interactions across 20 tasks in 9 diverse scenarios, ranging from visual navigation to mobile manipulation. Building upon this dataset, we propose U0, a VLA model for general underwater robots, which integrates binocular vision and other sensor modalities through multimodal fusion, and further incorporates a convolution-attention-based perception focus enhancement module (CAP) to improve spatial understanding and mobile manipulation. Across tasks such as inspection, obstacle avoidance, scanning, and dynamic tracking, the framework achieves a success rate of 80%, while in challenging mobile manipulation tasks, it reduces the distance to the target by 21.2% compared with baseline methods, demonstrating its effectiveness. USIM and U0 show that VLA models can be effectively applied to underwater robotic applications, providing a foundation for scalable dataset construction, improved task autonomy, and the practical realization of intelligent general underwater robots.

I. INTRODUCTION

The underwater environment poses growing demands, encompassing diverse applications such as marine ecological surveys, resource exploitation, inspection of pipelines, and underwater salvage [1], [2]. However, the unique characteristics of the underwater environment make human operations far more challenging and hazardous compared to terrestrial or aerial scenarios. Consequently, underwater robots have gradually emerged as a key solution for executing underwater tasks. The development of underwater robots faces distinctive challenges, including complex hydrodynamics, limited visibility, and restricted communication. These challenges highlight the pressing need to enhance the autonomy and

intelligence of underwater robots. Autonomous systems capable of executing complex underwater missions will greatly extend humanity’s ability to operate in the oceans, which cover 71% of the Earth’s surface, thereby advancing our understanding and exploration of the planet.

In recent years, the construction of embodied intelligence with physical-world manipulation capabilities has become a trend in artificial intelligence and robotics. Data-driven approaches are reshaping the field, enabling remarkable advances in autonomy, particularly in humanoid [3] and quadruped robots [4] powered by large-scale datasets and advanced algorithms. In contrast, many complex underwater tasks still heavily rely on human teleoperation. For example, in underwater grasping tasks, even experienced operators require extensive practice to improve grasping performance [5]. As a result, collecting large-scale, high-quality data in real underwater environments remains highly costly. Although existing studies have developed frameworks for teleoperation [6] and autonomous underwater intervention tasks [7], significant challenges remain in integrating multiple tasks. Taken together, these limitations have led to fragmented “data islands”, where the lack of unified frameworks and accessible datasets stands in stark contrast to the urgent demand for higher autonomy and intelligence in underwater robots.

To address these challenges, we introduce USIM, an underwater simulation-based multi-task Vision-Language-Action (VLA) dataset, which comprises both perception and control traces of a BlueROV2 underwater robot platform. Collected within simulation environments constructed using the Stonefish simulator [8], [9], USIM consists of 561K frames spanning 1,852 trajectories and approximately 15.6 hours of robot–environment interactions. It covers 20 tasks across 9 scenarios, ranging from visual navigation to mobile manipulation. Building upon this dataset, we further propose U0, a VLA model for general underwater robots. Built on Isaac-GR00T N1.5 [10], U0 combines multimodal fusion of multi-source sensor data with convolution-attention-based perception focus enhancement to better perform underwater tasks. It integrates binocular vision with additional sensors to enhance spatial understanding, addressing the specific embodiment of underwater robots that combines thruster locomotion and manipulator control. Furthermore, the convolution-attention-based perception focus enhancement module (CAP), guided by features of the Vision-Language Model (VLM), improves orientation estimation and facilitates mobile manipulation. We conduct extensive

Junwen Gu and Zhiheng Wu contributed equally to this work. Corresponding author: Jian Wang.

[†]The Key Laboratory of Cognition and Decision Intelligence for Complex Systems, Institute of Automation, Chinese Academy of Sciences, Beijing 100190, China (e-mail: gujunwen2022@ia.ac.cn jianwang@ia.ac.cn; zhengxing.wu@ia.ac.cn).

²Baidu Inc., Beijing 100085, China (e-mail: wzh404.ai@gmail.com).

³The School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing 100049, China.

<https://vincentgu2000.github.io/u0project>

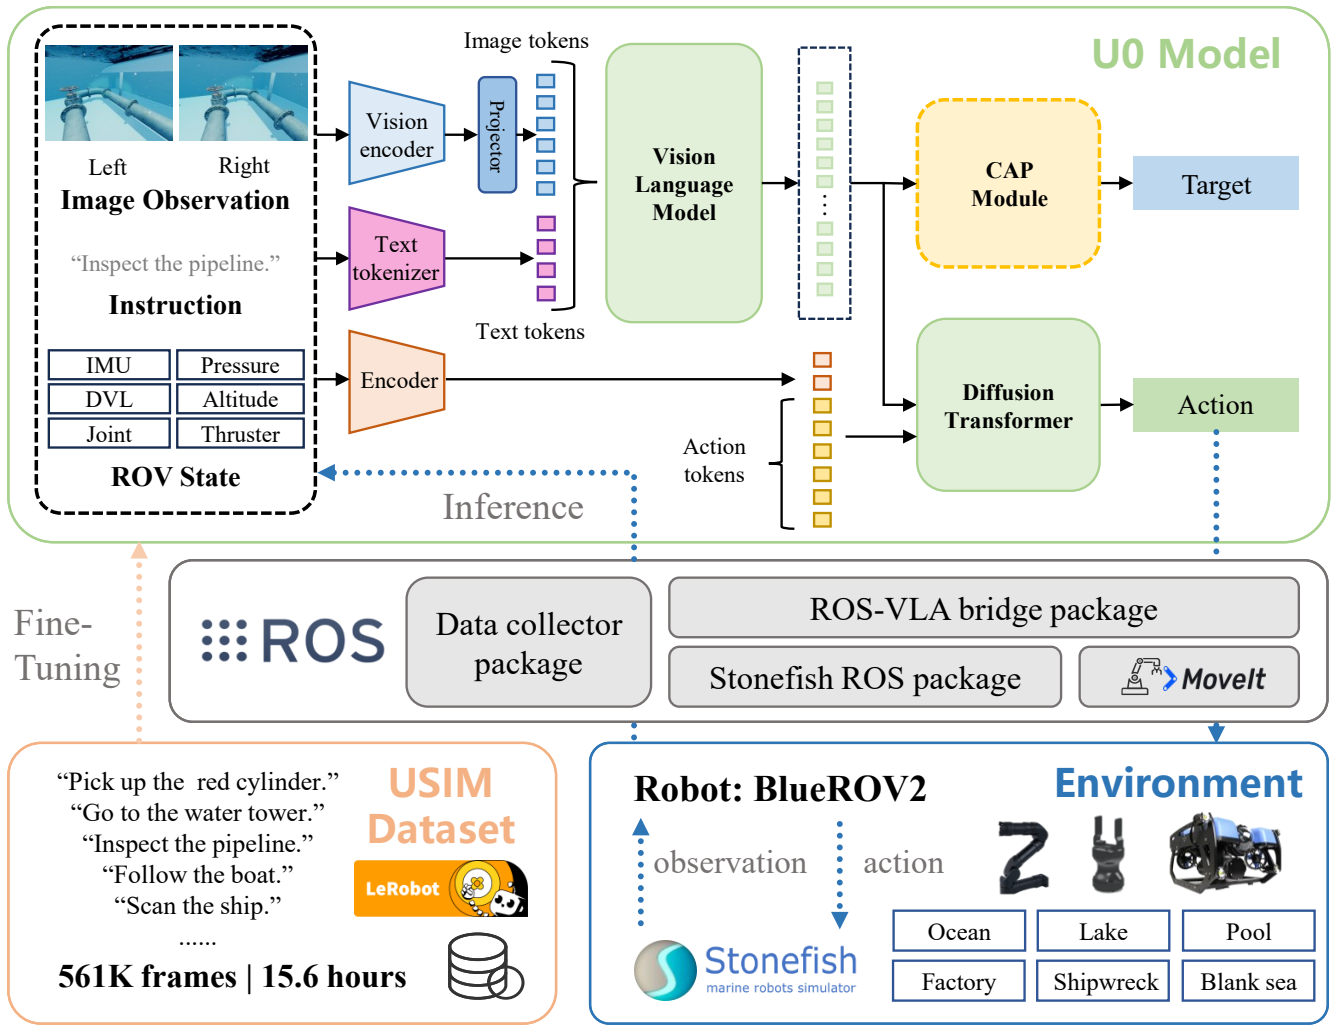


Fig. 1: **Overall Framework.** Diverse underwater scenarios and a BlueROV2 robot equipped with a manipulator and gripper are first constructed using the Stonefish simulator. Data collection and control are implemented via ROS, resulting in the USIM dataset of 561K frames (approximately 15.6 hours) covering 20 tasks. Based on USIM, U0 is developed with a dual-system architecture, incorporating multimodal sensor fusion and convolution-attention-based perception focus enhancement, while producing both target perception and robot action outputs.

validation in the testing dataset and simulation environments. With a model size of 3B parameters, U0 is lightweight enough for deployment on embedded AI computing platforms such as NVIDIA Jetson, facilitating practical realization of embodied intelligence in underwater scenarios.

Overall, we present a scalable data-to-task framework that integrates a self-constructed VLA dataset and a VLA model suitable for diverse underwater tasks. To the best of our knowledge, this is the first effort to jointly address multi-task underwater perception and action guided by language instructions. The main contributions of this paper are summarized as follows:

- A high-quality underwater VLA dataset (USIM) is constructed, comprising more than 561K frames from 1,852 trajectories, totaling around 15.6 hours of robot-environment interactions, and covering 20 tasks

across 9 diverse scenarios.

- A VLA model for general underwater robots (U0) is proposed, which achieves enhanced spatial understanding and improved performance in mobile manipulation tasks through multimodal fusion and convolution-attention-based perception focus enhancement.
- A scalable data-to-task framework is established for underwater robots. Across tasks such as inspection, navigation, and dynamic tracking, the framework achieves a success rate of 80%. In challenging mobile manipulation tasks, it reduces the distance between the robot and the target by 21.2% compared with baseline model, demonstrating its potential for handling such tasks effectively.

U0 demonstrates that VLA models can be effectively trained for underwater robotic applications, providing new

perspectives for dataset construction and utilization toward developing more intelligent underwater robots.

II. RELATED WORK

A. Towards Intelligent Underwater Robots

Underwater robots have evolved from remotely operated vehicles (ROVs) to autonomous underwater vehicles (AUVs). Traditional ROVs rely on tethers and human teleoperation, which not only restrict their operational range but also require operators with specialized knowledge and proficient control skills [6]. Moreover, the per-joint manual control mode used in some tasks proves inefficient. To address these limitations, the development of AUVs that integrate perception, planning, and control capabilities, enabling semi-autonomous navigation and intervention, has become a major direction in underwater robotics. Recently, data-driven approaches such as reinforcement learning (RL) and imitation learning (IL) have been explored for various underwater tasks, including motion control [11], path planning [12], formation control and obstacle avoidance [13], target tracking [14], [15], autonomous navigation [16], and manipulation [5], [17]. In addition, large language models and vision-based deep learning methods have been employed to enhance high-level cognition [18] and human-robot interaction [19] in autonomous underwater systems. Nevertheless, developing general-purpose models capable of handling multiple underwater tasks remains challenging and warrants further research.

B. Vision-Language-Action Datasets and Models

The robotics community has witnessed rapid progress in indoor embodied datasets and VLA models. Based on datasets such as DROID [20], Open X-Embodiment [21], and AgiBot World [22], pioneering VLA models including RT-2 [23], OpenVLA [24], π_0 [25], and GR00T N1.5 [10], demonstrate that combining multimodal perception with large-scale pretraining enables robots to generalize across tasks and environments. These advances have significantly propelled the development of domestic and indoor service robots, facilitating zero-shot generalization in navigation and manipulation. However, to our knowledge, research on constructing VLA models specifically for underwater robots remains scarce.

C. Underwater Datasets and Simulation Platforms

The performance of VLA models critically depends on large-scale training data, which can originate from video demonstrations, synthetic data, and real-world multimodal datasets [10]. Existing underwater datasets, including VAROS, UIEB, AQUALOC, and Underwater Caves Sonar Data Set [26]–[29], provide images, sonar data, and other sensory modalities, while various synthetic data generation methods have also been proposed [30]. These datasets primarily aim to enhance the perception capabilities of underwater robots. Given the high cost and risk of real-world underwater experiments, collecting data in simulation offers

advantages in efficiency and safety. Popular underwater simulators include HoloOcean, Dave, FishGym, and Stonefish [8], [9], [31], [32], which simulate fluid dynamics, visual degradation, and common underwater sensors. Overall, existing underwater datasets are largely task-specific and provide insufficient diversity required for developing unified frameworks, underscoring the benefits of leveraging underwater simulators to generate more comprehensive datasets.

III. METHOD

The overall framework of this work is illustrated in Fig. 1, consisting of simulation environment, dataset, and the VLA model. Details are described below.

A. Simulation Environment Construction

To generate sufficient high-quality data for training and evaluating the VLA model, we constructed nine distinct underwater scenarios using the Stonefish simulator. These scenarios include seabed environment, subsea pipeline, industrial pool, solar charging station, lake environment, open sea surface, underwater factory, modern shipwreck, and ancient shipwreck, as shown in Fig. 2. Subsequently, we built a BlueROV2 underwater robot equipped with a manipulator and parallel gripper within Stonefish, allowing real-time simulation of hydrodynamics, collisions, and grasping behavior. We further integrated the simulation with ROS using Stonefish’s ROS package, facilitating automated data collection. To enhance scene diversity, we implemented multiple map randomization modules, generating environments with variable object placement. Leveraging Stonefish’s rendering capabilities, we randomized sunlight conditions and water clarity as shown in Fig. 3. By incorporating scenario design, map randomization, and varying lighting conditions, our simulation environment generates diverse and realistic visual scenes, resulting in a dataset that exhibits both high diversity and realism. Such characteristics are expected to improve the generalization of models trained on this dataset.

B. Dataset Generation


The dataset comprises 20 tasks across 9 scenarios, totaling 561K frames and 1,852 trajectories (recorded at 10 Hz, approximately 15.6 hours) of high-quality robotic data. Among these, 526K frames with 1,752 trajectories are allocated for training, and 35K frames with 100 trajectories are reserved for testing, each provided in separate datasets formatted according to the LeRobot specification. The 20 tasks include 12 grasping tasks, 2 pipeline inspection tasks, 2 underwater shipwreck scanning tasks, 2 obstacle-avoidance navigation tasks, 1 dynamic tracking task, and 1 transport task, with example task trajectories illustrated in Fig. 4. The number of episodes, average duration, and total frames for each task are presented in Table I. As a VLA dataset, USIM provides language instructions, robot sensory inputs, and action data. The distribution of frames across different language instructions is shown in Fig. 5. Sensor modalities include binocular camera images, a pressure sensor, an inertial measurement unit (IMU), and a Doppler velocity log (DVL), while action




(a) Seabed

(b) Subsea pipeline


(c) Industrial pool




(d) Solar charging station



(e) Lake



(f) Open sea surface



(a) Elevation and azimuth



(g) Underwater factory



(h) Modern shipwreck



(i) Ancient shipwreck




(b) Jerlov

Fig. 2: Detailed visualization of simulation scenarios, comprising a total of nine distinct environments.

Fig. 3: Views under varying sunlight and water clarity.


Instruction: "Inspect the pipeline."

Duration: 110 s



Instruction: "Go to the charge station."

Duration: 20 s



Instruction: "Scan the ship."

Duration: 68 s



Instruction: "Pick up the red cylinder and transfer it to the box."

Duration: 27 s



Instruction: "Pick up the red cylinder."

Duration: 20 s



Fig. 4: Example task trajectories of USIM dataset, including pipeline inspection, obstacle avoidance navigation, shipwreck scanning, transportation, and grasping tasks.

TABLE I: Details of the USIM Dataset

Task ID	Episodes	Avg. duration (s)	Total frames
Pick red factory	105	23	23,862
Pick red shallow	105	24	25,148
Pick redx factory	105	22	23,595
Pick redx shallow	105	25	25,817
Pick blue factory	105	23	24,013
Pick blue shallow	105	26	27,262
Pick bluex factory	105	22	23,111
Pick bluex shallow	105	25	25,860
Pick pipe0 factory	105	22	23,591
Pick pipe0 shallow	105	23	23,838
Pick pipe1 factory	105	22	23,121
Pick pipe1 shallow	105	23	24,639
Transfer red shallow	105	29	30,106
Goto charge station	105	15	15,228
Goto water tower	107	28	30,478
Follow boat	55	36	19,564
Scan ship modern	55	67	36,794
Scan ship ancient	55	72	39,339
Inspect pipeline sea	55	65	35,670
Inspect pipeline pool	55	110	60,261




Fig. 5: Task distribution of the USIM dataset.

signals consist of thruster pulse-width modulation (PWM) signals and manipulator joint angles. To enable large-scale data acquisition, we developed an automated parallel data collection pipeline with task-specific execution logic. At the control level, a PID controller ensures accurate ROV pose tracking, whereas grasping tasks leverage MoveIt for manipulator planning and control.

C. Underwater VLA Model

The scarcity of large-scale underwater robotic datasets makes training a VLA model from scratch impractical. Therefore, we propose U0, a VLA model for underwater tasks built upon the pre-trained Isaac-GR00T N1.5 backbone. The upper part of Fig. 1 illustrates the overall architecture. The model accepts multi-modal sensor data as input. Visual images and language instructions are first processed through respective encoders and then fed into the VLM, while additional sensor modalities and robot action data are provided to a diffusion transformer. The transformer integrates these inputs with the VLM features via cross-attention to generate actions. Concurrently, the VLM features are passed through the CAP module, which leverages

convolution-attention mechanisms to enhance visual feature representations, thereby improving the model’s ability to capture the relationships between underwater motion states and task objectives.

Multimodal Fusion of Multi-Source Sensor Data. Unlike terrestrial robots, underwater robots operate predominantly in a floating state and must coordinate movement with manipulator actions. They also frequently change depth, in contrast to most ground robots that are limited to horizontal movement. In addition, underwater visual perception is strongly affected by environmental disturbances such as turbidity, lighting variations, and water currents. To address these challenges, we incorporate underwater-specific sensors into the state space, including pressure sensor, IMU, and DVL, enabling the model to develop accurate proprioception and localization. Furthermore, binocular camera images from the robot are also provided as input to enhance visual perception capabilities. The action space of underwater robots also differs from that of terrestrial robots. Beyond manipulator control, locomotion is primarily achieved through multiple thrusters, whose control signals occupy a distinct force space. We normalize thruster PWM signals and combine them with manipulator joint angles to construct the action space, while including both in the state space to improve the model’s understanding of current robot actions.

To enable the model to reason effectively about task goals relative to the robot’s current state, we represent target positions and orientations in a robot-centric coordinate system. Specifically, let the six-degree-of-freedom poses of the target and the robot in the world coordinate system be

$$p_t = (R_t, \mathbf{t}_t), \quad p_r = (R_r, \mathbf{t}_r), \quad (1)$$

where $R \in SO(3)$ denotes the rotation matrix and $\mathbf{t} \in \mathbb{R}^3$ the translation vector. Then, the target pose expressed in the robot-centric coordinate system is given by

$$p_{t2r} = (R_r^\top R_t, R_r^\top (\mathbf{t}_t - \mathbf{t}_r)). \quad (2)$$

We use p_{t2r} to construct ground-truth labels for training. Compared with the global coordinate representation, this relative formulation better captures the dynamic motion characteristics of underwater scenarios and mitigates the dependency on a fixed world reference frame. Furthermore, it aligns with the egocentric reasoning principle observed in biological decision-making processes, thereby enhancing the model’s generalization ability across diverse tasks.

Convolution-Attention-based Perception Focus Enhancement Module. Due to severe visual degradation in underwater environments, relying solely on the vision capability of GR00T N1.5 is insufficient for accurately capturing target objects. To address this limitation, U0 extends Isaac-GR00T N1.5 with a convolution-attention-based perception focus enhancement module guided by VLM features. This module strengthens the model’s ability to detect and localize target objects. Fig. 6 illustrates the architecture of the CAP module, which is composed of convolutional layers, attention layers, pooling operations, and a multilayer perceptron. The




Fig. 6: CAP module architecture.

computation process is formulated as:

$$\text{Token} = \text{VLM}(\text{Img}_{\text{left}}, \text{Img}_{\text{right}}), \quad (3)$$

$$F = \text{Conv}(\text{Token}, \text{MASK}), \quad (4)$$

$$\text{Att} = \text{Conv}(F), \quad (5)$$

$$T = \text{MLP}(\text{pool}(F \cdot \text{Att})), \quad (6)$$

where Img_{left} and $\text{Img}_{\text{right}}$ denote stereo images from the binocular vision sensor, $\text{VLM}(\cdot)$ represents the Vision-Language model, and Token denotes the extracted features. $\text{Conv}(\cdot)$ refers to convolutional operations, while MASK avoids computation over padding features. Att represents channel-wise attention, $\text{pool}(\cdot)$ is a pooling operation, $\text{MLP}(\cdot)$ denotes the multilayer perceptron, F is the intermediate feature, and T is the predicted target.

The training objective of the CAP module is formulated as a mean squared error (MSE) loss between the predicted target T and the ground truth T_{gt} :

$$L_{\text{CAP}} = \text{MSE}(T, T_{\text{gt}}). \quad (7)$$

For action module, we follow the original GR00T N1.5 training objectives. The sum of the losses of the two modules serves as the overall loss L .

$$L = L_{\text{action}} + \alpha L_{\text{CAP}}, \quad (8)$$

where, L_{action} is the loss of the action module and α represents the weight factor.

Training Details. Our training was conducted using a total batch size of 1024 for 5000 steps. During training, the CAP

module functions as an auxiliary task branch to enhance perception of target-related visual features. This branch can be disabled during inference, ensuring that deployment does not incur any increase in model size or computational latency.

IV. EXPERIMENTS

In this section, we establish a set of challenging benchmarks to evaluate the effectiveness of the proposed U0 model and the USIM dataset. The evaluation is primarily conducted through open-loop offline evaluation and closed-loop online testing.

A. Open-Loop Offline Evaluation

In the offline evaluation, we constructed a simulation-based test dataset, comprising 20 tasks with 5 trajectories each, totaling 35K frames (approximately 1 hour). We compared the performance of three models on this test dataset: the original GR00T N1.5 without fine-tuning, GR00T N1.5 fine-tuned on the proposed USIM dataset (denoted GR00T FT), and U0, which incorporates the CAP module. To further assess visual perception in underwater settings, we conducted experiments with both monocular and binocular image inputs. For GR00T FT and U0, fine-tuning was performed separately on the monocular and binocular variants of the training dataset. The results are summarized in Table II, where tasks sharing the same instruction are aggregated. Here, e_{action} denotes the error of the action module, and e_{target} represents the error of the CAP module.

The results indicate that directly applying GR00T N1.5 to underwater tasks yields substantially higher errors—its average e_{action} is an order of magnitude larger than that of the fine-tuned models—highlighting the significant domain gap between the humanoid robot datasets used during pretraining and the underwater domain. Fine-tuning on the USIM dataset substantially improves performance, validating the effectiveness of our proposed dataset. Interestingly, under the pretrained weights, GR00T N1.5 performs better with monocular images than with binocular inputs, suggesting that without domain adaptation, the model favors monocular data more consistent with its pretraining distribution. However, after training on the USIM dataset, binocular data exhibit a clear advantage, with models achieving higher action accuracy and overall better perceptual precision compared to monocular inputs.

Compared with GR00T FT, U0 achieves a reduction of 7.7% and 4.2% in average e_{action} under monocular and binocular inputs, respectively. This demonstrates that incorporating the convolution-attention-based perception focus enhancement module significantly improves the model's ability to perceive target objects and, consequently, enhances action accuracy. Notably, in monocular grasping tasks, U0 achieves substantial performance gains over GR00T FT, indicating that the CAP module effectively mitigates the inherent limitations of monocular 3D spatial perception. This advantage, however, becomes less pronounced when binocular vision is available.

TABLE II: Comparison of e_{action} and e_{target} Across Tasks

Task	Number	e_{action} (Mono)			e_{action} (Bino)			e_{target} (U0)	
		GR00T N1.5	GR00T FT	U0	GR00T N1.5	GR00T FT	U0	Mono	Bino
Inspect the pipeline.	10	0.3186	0.1036	0.1065	0.3228	0.1183	0.1104	0.1434	0.1488
Scan the ship.	10	0.3469	0.1012	0.1004	0.3461	0.1054	0.1020	1.2566	1.0339
Go to the water tower.	5	0.3936	0.0724	0.0740	0.3828	0.0857	0.0936	1.5797	1.7723
Go to the charge station.	5	0.3425	0.0705	0.1053	0.3511	0.1112	0.0920	0.6363	0.8541
Follow the boat.	5	0.3403	0.3586	0.3430	0.3407	0.0397	0.0310	0.2438	0.1237
Pick up the red cylinder.	20	0.1850	0.0512	0.0448	0.1835	0.0510	0.0523	0.0253	0.0330
Pick up the blue cylinder.	20	0.1861	0.0447	0.0393	0.1915	0.0412	0.0412	0.0250	0.0303
Pick up the pipe.	20	0.1878	0.0509	0.0330	0.1888	0.0397	0.0359	0.0231	0.0330
Pick up the red cylinder and transfer it to the box.	5	0.2264	0.0842	0.0563	0.2356	0.0265	0.0271	0.7309	0.0557
ALL	100	0.2435	0.0791	0.0730	0.2451	0.0619	0.0593	0.3142	0.2778




Fig. 7: Comparison of online success rates across multiple tasks.

TABLE III: Mean Robot–Target Distance in Mobile Grasping Tasks

Task ID	Mean distance		
	GR00T FT	U0 Mono	U0 Bino
Pick red shallow	0.3580	0.3100	0.2515
Pick blue shallow	0.3801	0.2651	0.2884
Pick pipe0 shallow	0.2932	0.2850	0.2810
Pick redx shallow	0.3312	0.3721	0.2346
Pick red factory	0.3837	0.3496	0.3204
ALL	0.3492	0.3164	0.2752

Overall, the offline evaluation results confirm the feasibility of using the USIM dataset for training underwater VLA models, while also demonstrating the effectiveness of U0’s multimodal fusion of multi-source sensory information and convolution-attention-based perception focus enhancement.

B. Closed-loop Online Testing

To further validate the effectiveness of U0, we integrated GR00T N1.5, GR00T FT, and U0 into the simulation environment and evaluated their performance on actual task executions, covering five grasping tasks and seven non-grasping tasks. Since the untrained GR00T N1.5 failed to complete any of the tasks (with a success rate of 0%), its results

were excluded from the statistics. For U0, we examined its performance under both monocular and binocular visual inputs.

In the seven non-grasping tasks, each model was executed 10 times per task, and the corresponding success rates were recorded, as shown in Fig. 7. Across these tasks, U0 achieved an average success rate of 80%, demonstrating its capability to perform general-purpose task execution. The average success rate also indicates that U0 equipped with binocular vision consistently outperforms its monocular counterpart and surpasses GR00T FT across tasks, corroborating the results observed in the open-loop offline evaluation. At the task level, the performance patterns of GR00T FT and U0 also closely align with those reported in the offline evaluation.

For the five challenging mobile grasping tasks, each model was executed five times per task, and the distances between the ROV and the target objects were measured, as summarized in Table III. U0 with binocular input reduced the average distance to the target by 21.2% compared with GR00T FT. These results demonstrate that U0 more accurately identifies target locations and initiates grasp attempts, further validating the effectiveness of the CAP module. In addition, U0 achieves higher localization accuracy when using binocular vision compared to monocular vision.

Overall, the online evaluation results are highly consistent with those of the offline evaluation. They not only highlight U0's capabilities in underwater navigation, obstacle avoidance, inspection, and scanning, but also reveal its potential in handling challenging mobile grasping tasks that require accounting for fluid dynamics and the complex interactions among the vehicle body, manipulator, and target objects. These findings further confirm the feasibility of leveraging simulation environments to generate data and train underwater VLA models with generalizable task-execution capabilities.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we presented USIM, a large-scale simulation-based multi-task dataset, and U0, a Vision-Language-Action model for general underwater robots. By leveraging binocular vision, multimodal sensor fusion, and the convolution-attention-based perception focus enhancement module, U0 achieves an 80% average success rate on non-grasping tasks and reduces the distance to targets in challenging mobile grasping tasks by 21.2% compared with the baseline, demonstrating the feasibility of training underwater VLA models with simulation-generated data.

Future work includes improving U0's performance in mobile grasping through richer simulation scenarios and more diverse data collection strategies. Incorporating additional modalities such as sonar could enhance perception in deep-water or low-visibility environments. Finally, real-world deployment and field validation are necessary steps toward fully autonomous, multi-task underwater robots.

REFERENCES

- [1] T. Liu *et al.*, "A bioinspired multimotion modality underwater micro-robot," *Sci. Adv.*, vol. 11, no. 19, May 2025, Art. no. eadu2527.
- [2] R. K. Katzschmann, J. DelPreto, R. MacCurdy, and D. Rus, "Exploration of underwater life with an acoustically controlled soft robotic fish," *Sci. Robot.*, vol. 3, no. 16, Mar. 2018, Art. no. eaar3449.
- [3] Z. Su *et al.*, "Hitter: A humanoid table tennis robot via hierarchical planning and learning," *arXiv preprint arXiv:2508.21043*, Sep. 2025.
- [4] X. Tong *et al.*, "Quart-online: Latency-free multimodal large language model for quadruped robot learning," in *Proc. 2025 IEEE Int. Conf. Robot. Autom. (ICRA)*, Atlanta, GA, USA, May 2025, pp. 9533–9539.
- [5] R. Liu, H. Ha, M. Hou, S. Song, and C. Vondrick, "Self-improving autonomous underwater manipulation," in *Proc. 2025 IEEE Int. Conf. Robot. Autom. (ICRA)*, Atlanta, GA, USA, May 2025, pp. 16915–16922.
- [6] A. Phung, G. Billings, A. F. Daniele, M. R. Walter, and R. Camilli, "A shared autonomy system for precise and efficient remote underwater manipulation," *IEEE Trans. Robot.*, vol. 40, pp. 4147–4159, 2024.
- [7] E. Palmer, C. Holm, and G. Hollinger, "Angler: An autonomy framework for intervention tasks with lightweight underwater vehicle manipulator systems," in *Proc. 2024 IEEE Int. Conf. Robot. Autom. (ICRA)*, Yokohama, Japan, May 2024, pp. 6126–6132.
- [8] P. Cieslak, "Stonefish: An advanced open-source simulation tool designed for marine robotics, with a ros interface," in *Proc. OCEANS 2019 – Marseille*, Marseille, France, Jun. 2019, pp. 1–6.
- [9] M. Grimaldi *et al.*, "Stonefish: Supporting machine learning research in marine robotics," in *Proc. 2025 IEEE Int. Conf. Robot. Autom. (ICRA)*, Atlanta, GA, USA, May 2025, pp. 13 605–13 611.
- [10] NVIDIA *et al.*, "Gr00t N1: An open foundation model for generalist humanoid robots," *arXiv preprint arXiv:2503.14734*, Mar. 2025.
- [11] J. Gu, J. Wang, Z. Liu, M. Tan, J. Yu, and Z. Wu, "Deformation control and thrust analysis of a flexible fishtail with muscle-like actuation," *IEEE Trans. Robot.*, vol. 41, pp. 159–179, 2025.
- [12] J. Yang, J. Ni, M. Xi, J. Wen, and Y. Li, "Intelligent path planning of underwater robot based on reinforcement learning," *IEEE Trans. Automat. Sci. Eng.*, vol. 20, no. 3, pp. 1983–1996, Jul. 2023.
- [13] Z. Fang, T. Chen, T. Shen, D. Jiang, Z. Zhang, and G. Li, "Multi-agent generative adversarial interactive self-imitation learning for auv formation control and obstacle avoidance," *IEEE Robot. Autom. Lett.*, vol. 10, no. 5, pp. 4356–4363, May 2025.
- [14] Y. Wang *et al.*, "Target tracking control of a biomimetic underwater vehicle through deep reinforcement learning," *IEEE Trans. Neural Netw. Learning Syst.*, vol. 33, no. 8, pp. 3741–3752, Aug. 2022.
- [15] I. Masmitja *et al.*, "Dynamic robotic tracking of underwater targets using reinforcement learning," *Sci. Robot.*, vol. 8, no. 80, Jul. 2023, Art. no. eade7811.
- [16] X. Lin *et al.*, "Uivnav: Underwater information-driven vision-based navigation via imitation learning," in *Proc. 2024 IEEE Int. Conf. Robot. Autom. (ICRA)*, Yokohama, Japan, May 2024, pp. 5250–5256.
- [17] J. Gao, Y. Li, Y. Chen, Y. He, and J. Guo, "An improved SAC-based deep reinforcement learning framework for collaborative pushing and grasping in underwater environments," *IEEE Trans. Instrum. Meas.*, vol. 73, pp. 1–14, 2024.
- [18] M. Buchholz, I. Carlucho, M. Grimaldi, and Y. R. Petillot, "Distributed AI agents for cognitive underwater robot autonomy," *arXiv preprint arXiv:2507.23735*, Aug. 2025.
- [19] A. Gomez Chavez, A. Ranieri, D. Chiarella, and A. Birk, "Underwater vision-based gesture recognition: A robustness validation for safe human-robot interaction," *IEEE Robot. Automat. Mag.*, vol. 28, no. 3, pp. 67–78, Sep. 2021.
- [20] A. Khazatsky *et al.*, "Droid: A large-scale in-the-wild robot manipulation dataset," *arXiv preprint arXiv:2403.12945*, Apr. 2025.
- [21] A. O'Neill *et al.*, "Open X-embodiment: Robotic learning datasets and RT-X models : Open X-embodiment collaboration0," in *Proc. 2024 IEEE Int. Conf. Robot. Autom. (ICRA)*, Yokohama, Japan, 2024, pp. 6892–6903.
- [22] AgiBot-World-Contributors *et al.*, "Agibot world colosseum: A large-scale manipulation platform for scalable and intelligent embodied systems," *arXiv preprint arXiv:2503.06669*, Aug. 2025.
- [23] B. Zitkovich *et al.*, "Rt-2: Vision-language-action models transfer web knowledge to robotic control," in *Proc. 7th Conf. Robot Learning (CoRL)*, vol. 229, Nov. 2023, pp. 2165–2183.
- [24] M. J. Kim *et al.*, "Openvla: An open-source vision-language-action model," *arXiv preprint arXiv:2406.09246*, Sep. 2024.
- [25] K. Black *et al.*, "π0: A vision-language-action flow model for general robot control," *arXiv preprint arXiv:2410.24164*, Nov. 2024.
- [26] P. Georg Olofsson Zwilmeyer, M. Yip, A. Langeland Teigen, R. Mester, and A. Stahl, "The varos synthetic underwater data set: Towards realistic multi-sensor underwater data with ground truth," in *Proc. 2021 IEEE/CVF Int. Conf. Comput. Vis. Workshops (ICCVW)*, Montreal, BC, Canada, Oct. 2021, pp. 3715–3723.
- [27] C. Li *et al.*, "An underwater image enhancement benchmark dataset and beyond," *IEEE Trans. on Image Process.*, vol. 29, pp. 4376–4389, 2020.
- [28] M. Ferrera, V. Creuze, J. Moras, and P. Trouvé-Peloux, "Aqualoc: An underwater dataset for visual-inertial-pressure localization," *Int. J. Robot. Res.*, vol. 38, no. 14, pp. 1549–1559, Dec. 2019.
- [29] A. Mallios, E. Vidal, R. Campos, and M. Carreras, "Underwater caves sonar data set," *Int. J. Robot. Res.*, vol. 36, no. 12, pp. 1247–1251, Oct. 2017.
- [30] Z. Wu, Z. Wu, X. Chen, Y. Lu, and J. Yu, "Self-supervised underwater image generation for underwater domain pre-training," *IEEE Trans. Instrum. Meas.*, vol. 73, pp. 1–14, 2024.
- [31] Z. Huang, M. Buchholz, M. Grimaldi, H. Yu, I. Carlucho, and Y. R. Petillot, "Urobench: Comparative analyses of underwater robotics simulators from reinforcement learning perspective," in *Proc. OCEANS 2024 – Singapore*, Singapore, Singapore, Apr. 2024, pp. 1–8.
- [32] W. Liu, K. Bai, X. He, S. Song, C. Zheng, and X. Liu, "Fishgym: A high-performance physics-based simulation framework for underwater robot learning," in *Proc. 2022 IEEE Int. Conf. Robot. Autom. (ICRA)*, Philadelphia, PA, USA, May 2022, pp. 6268–6275.