EasyUUV: An LLM-Enhanced Universal and Lightweight Sim-to-Real Reinforcement Learning Framework for UUV Attitude Control

Jingzehua Xu, Xiaofan Li

Department of Mechanical Engineering

1: introduction

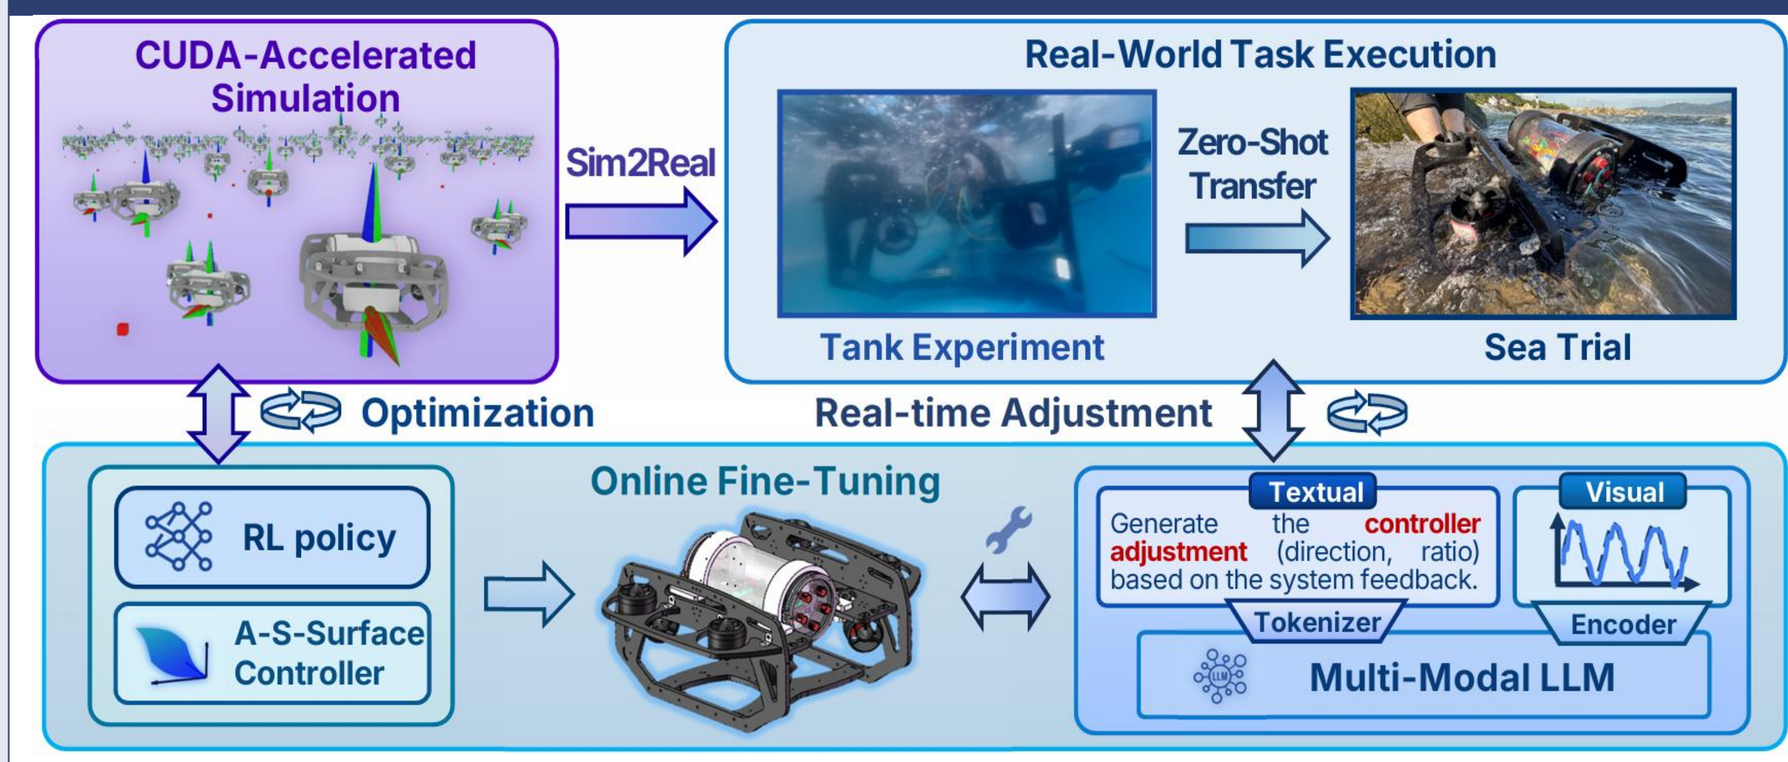

Fig. 1: Illustration of our developed EasyUUV framework. EasyUUV is an LLM-enhanced universal and lightweight Sim2Real RL framework for UUV attitude control, which trains the expert policy via RL in parallelized simulation, while transferring it to a real UUV platform. A multimodal LLM agent further adapts controller parameters using dynamics and sensor feedback for robust performance.

- **Real-world challenge**：AUVs struggle to maintain stable control in nonlinear, disturbed, and partially observed ocean environments.
- **Limitations of traditional control**：Model-based controllers rely on accurate dynamics and adapt poorly to changing conditions.
- **Limitations of RL**：Reinforcement learning suffers from the Sim-to-Real gap and requires heavy manual tuning during deployment.
- **Proposed solution**：The EasyUUV framework combines GPU-trained RL policies, an adaptive A-S-Surface controller, and multimodal LLM-based tuning to achieve robust, transferable control without retraining.

2: Methodology

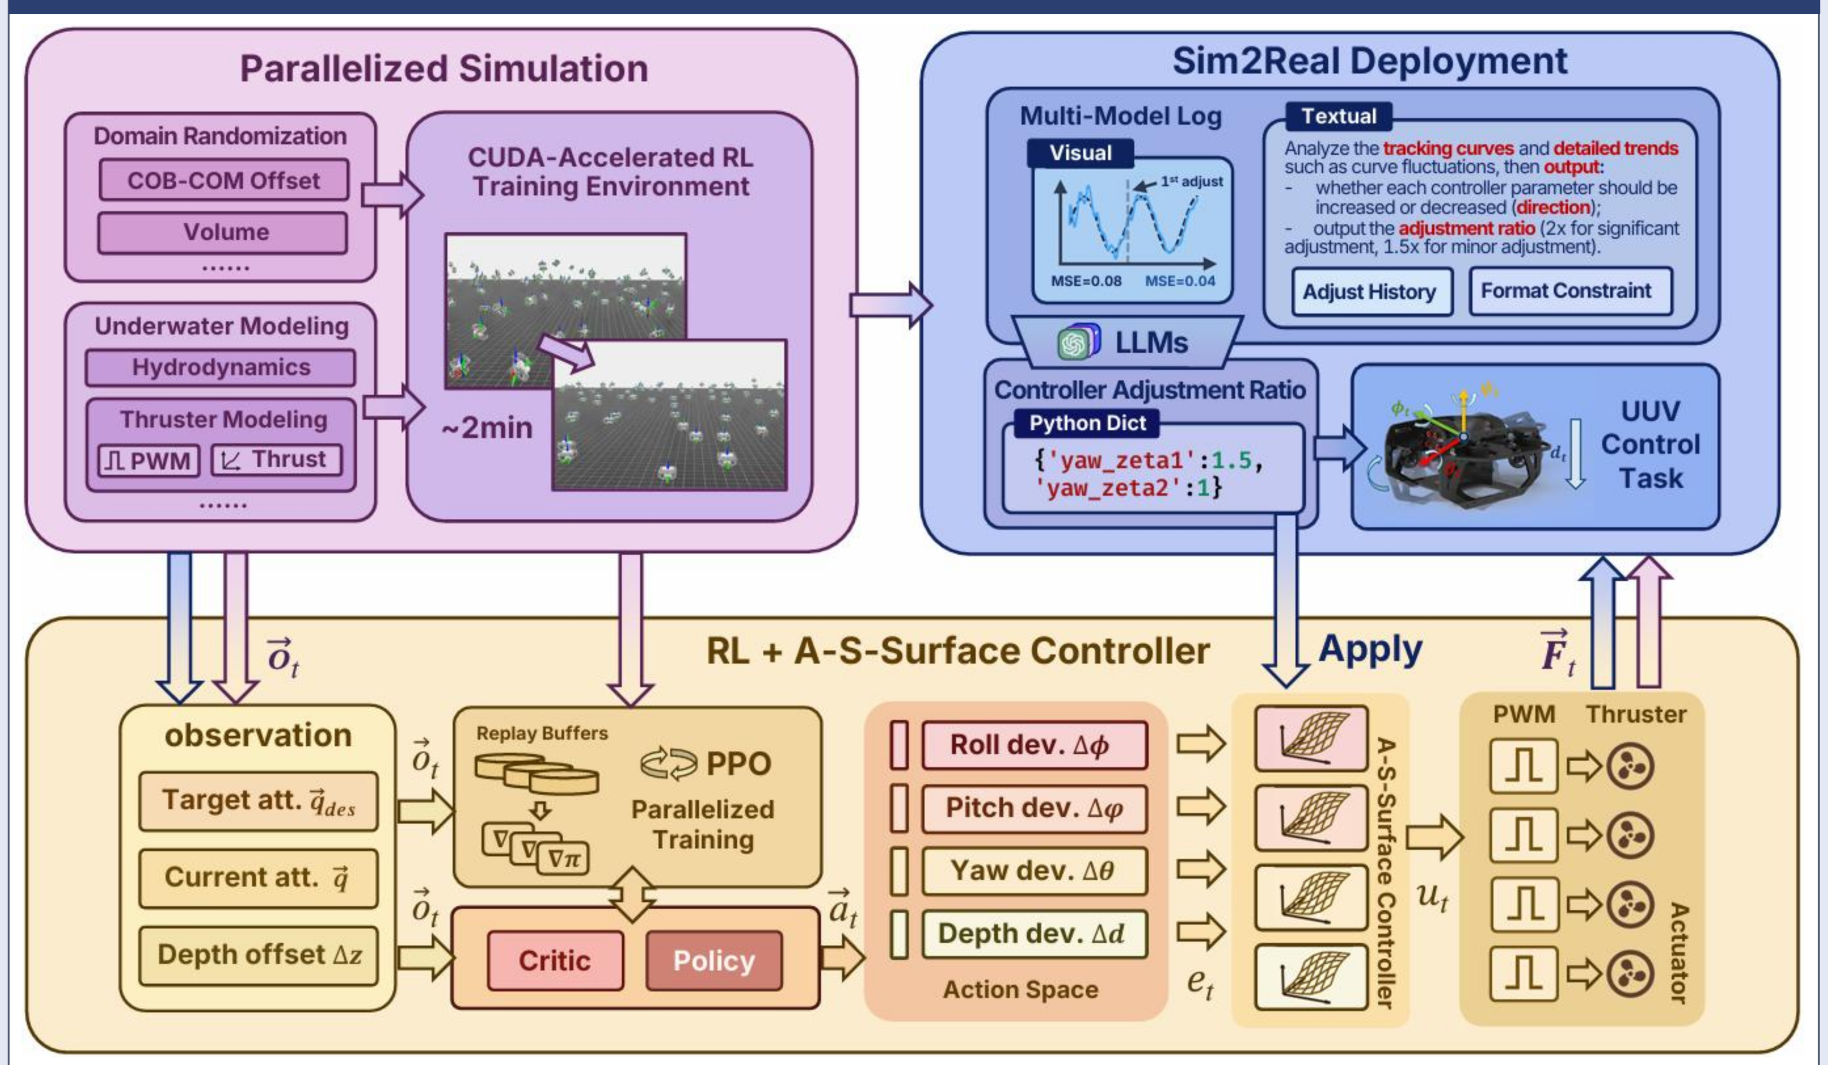

Fig. 2: Architecture of the EasyUUV framework, which comprise three parts: (a) RL and A-S-Surface-based composite controller module; (b)Parallelized RL training simulation environment developed on Isaac Lab; and (c) Sim2Real module for real-world adaption.

- **GPU-Parallel RL Policy Learning**：Train high-level reinforcement learning policies in massively parallel GPU simulation to efficiently acquire disturbance-aware control strategies.
- **Adaptive A-S-Surface Low-Level Control**：Use a nonlinear adaptive A-S-Surface controller to execute stable, real-time actuation under uncertain and highly dynamic underwater conditions.
- **Multimodal LLM-Guided Online Tuning**：Incorporate a multimodal large language model to interpret sensory cues and autonomously adjust control parameters without retraining.

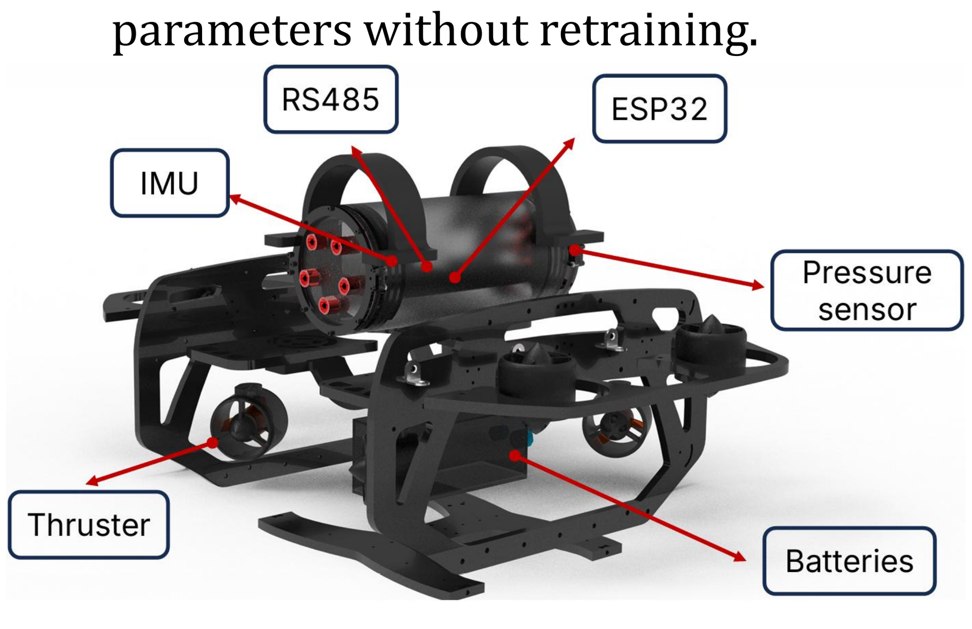

Fig. 3: Exploded view of our EasyUUV hardware platform.

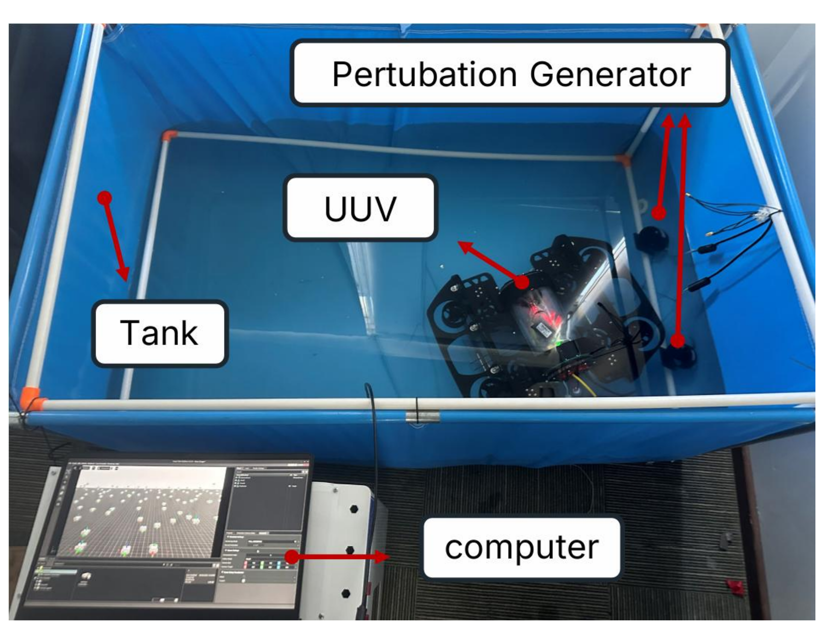

Fig. 4: Experimental testbed for real-world validation of EasyUUV.

3: Research accomplishments

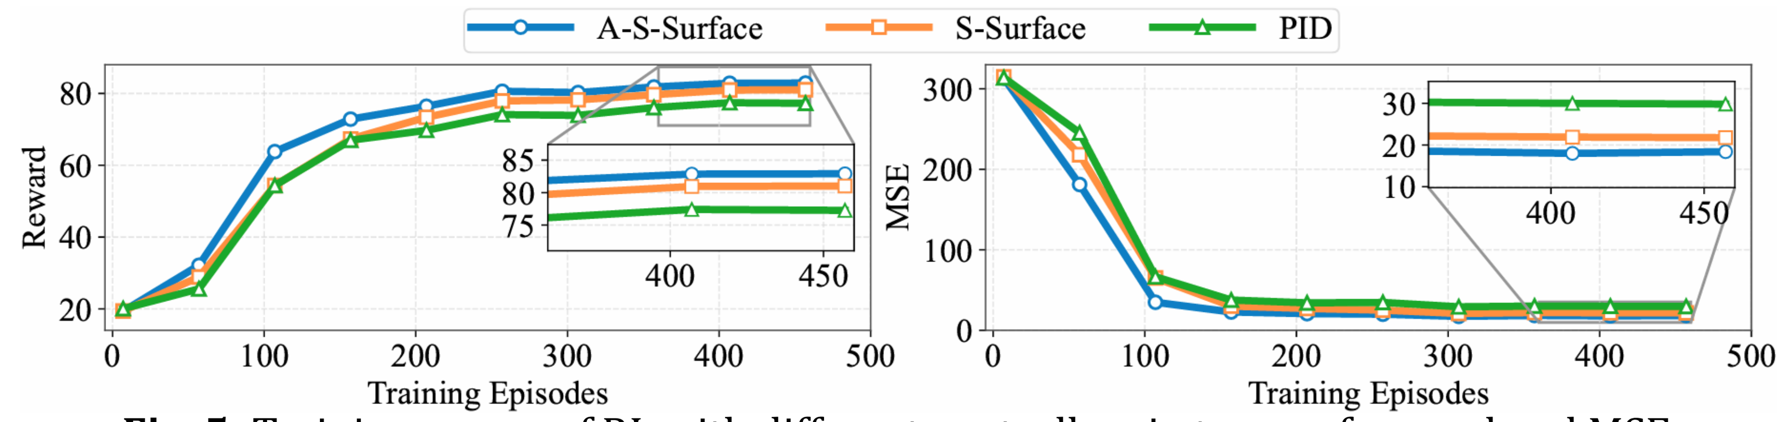

Fig. 5: Training curves of RL with different controllers in terms of reward and MSE. (Left): Reward curves. (Right): MSE curves.

- **RL + Controllers Training Evaluation**：Compare reward and MSE curves across A-S-Surface, S-Surface, and PID to show that the adaptive A-S-Surface controller enables faster learning and lower tracking error.
- **Limitations of RL**：Reinforcement learning suffers from the Sim-to-Real gap and requires heavy manual tuning during deployment.
- **Real-Ocean Deployment Trials**：Validate the framework in open-sea conditions, showing robust control despite stochastic waves and currents.

Fig. 6: Snapshots of EasyUUV operating in an indoor tank at t = 0.0s, 12.5s, and 25.0s in real-world experiments.

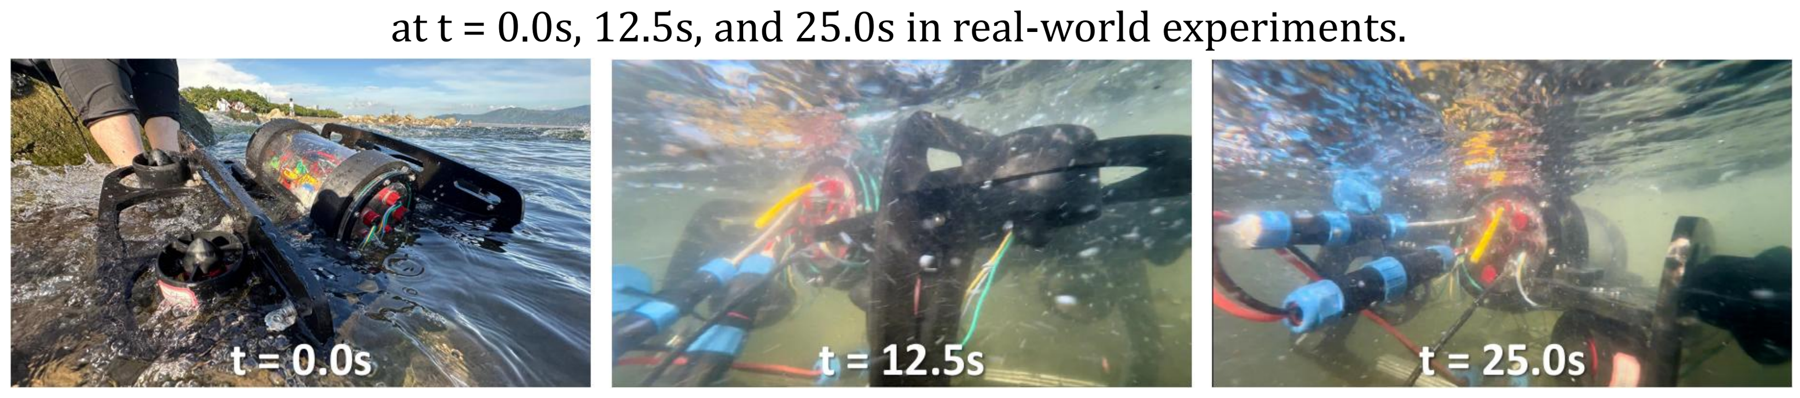

Fig. 7: Snapshots of EasyUUV operating in real ocean conditions at t = 0.0s, 12.5s, and 25.0s in sea trials

- **RL vs Non-RL Attitude Tracking Comparison**：Show that RL-guided control significantly reduces compound attitude errors in real-world experiments.
- **LLM-Guided Online Tuning Under Turbulence**：Demonstrate that multimodal LLM-based parameter adjustment improves tracking accuracy during strong perturbations without retraining.
- **Cross-Environment Robustness Assessment**：Verify that the combined framework generalizes across simulation, indoor tank, and real-ocean settings.

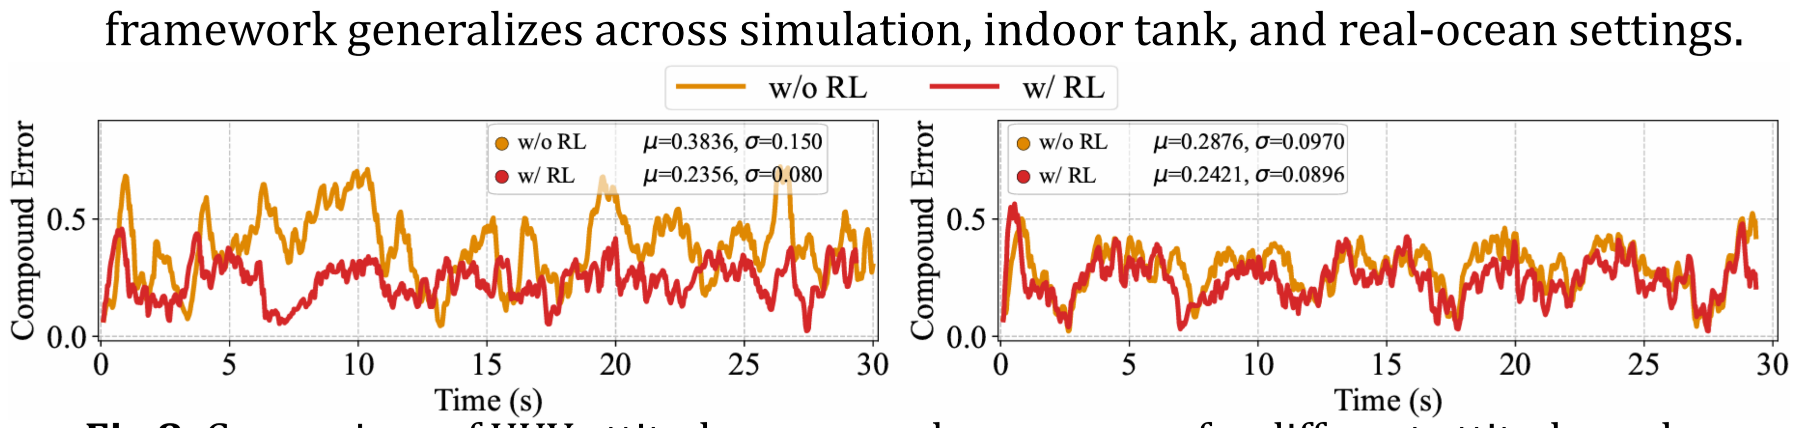

Fig.8: Comparison of UUV attitude compound error curves for different attitude angles combination, under both RL and non-RL setting in real-world experiments. (Left): Yaw and Roll. (Right): Yaw and Pitch.

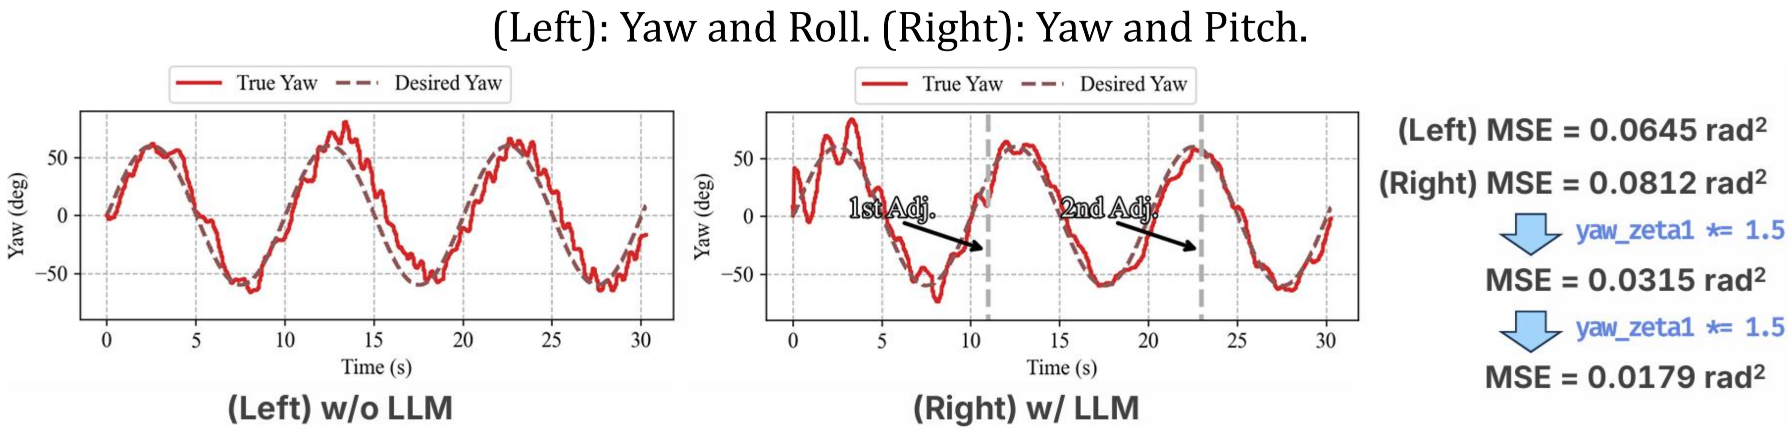

Fig. 9: Tracking response curves under turbulent perturbations with and without LLM-based online fine-tuning of controller parameters.

4: Conclusions

Conclusions: This paper presents EasyUUV, an LLM-enhanced, lightweight Sim2Real RL framework for robust UUV attitude control. The system combines domain-randomized RL training with a hybrid architecture that integrates an A-S-Surface controller, while a multimodal LLM agent enables runtime parameter tuning without retraining. Implemented on a cost-effective 6-DoF UUV platform, EasyUUV supports efficient simulation-based policy learning and achieves zero-shot transfer in real deployments. Extensive simulation and field trials show stable, generalizable control with strong robustness and consistent Sim2Real performance across diverse underwater conditions.

Future Plan: We plan to extend EasyUUV with visual–language models and image enhancement techniques to support higher-level goal interpretation and resilient navigation in unstructured, visually degraded environments.