Zhenqi Wu

Ph.D. Researcher | Multimodal AI, Embodied Intelligence & Robot Autonomy

Tampa, FL | zhenqi@usf.edu | +1 352-214-8099 | Google Scholar

RESEARCH PROFILE

Ph.D. researcher working at the intersection of multimodal machine learning, vision-language reasoning, learning-based planning, and embodied AI. Developed VLM/VLA methods for autonomous underwater navigation and exploration, with first-author research at ICRA, IROS, and OCEANS. Experience spans LoRA fine-tuning, reinforcement learning, closed-loop evaluation, real-robot deployment, and sim-to-real data pipelines.

EDUCATION

University of South Florida — Ph.D. in Electrical Engineering (in progress); GPA 4.0/4.0 | *Aug 2025–Present*

Expected graduation: 05/29

Arizona State University — Ph.D. studies in Computer Science; GPA 3.92/4.0 | *Aug 2023–May 2025*

University of Florida — M.S. in Electrical and Computer Engineering | *Aug 2021–May 2023*

Dalian Maritime University — B.S. in Electrical Engineering and Automation | *Sep 2016–Jul 2020*

SELECTED PUBLICATIONS

- Z. Wu et al., "DREAM: Domain-Aware Reasoning for Efficient Autonomous Underwater Monitoring," IEEE ICRA, 2026. [First author]
- Z. Wu, Y. Lu, X. Xiao, X. Lin, "CORAL: Contextual Reasoning and Local Planning in a Hierarchical VLM Framework for Underwater Monitoring," IEEE/RSJ IROS, 2026. [First author]
- Z. Wu et al., "CAVE-NAV: VLM-Based Autonomous 3D Navigation in Underwater Cave Environments," IEEE OCEANS, 2026. [First author]
- Y. Lu, B. Wang, Z. Wu et al., "APPLV: Adaptive Planner Parameter Learning from Vision-Language-Action Model," IEEE/RSJ IROS, 2026.
- I. Duporge*, Z. Wu*, M. E. Black* et al., "Enhancing Satellite-Based Detection of Marine Fauna Using Synthetic Data," Ecological Informatics 95:103767, 2026. [*Equal contribution]

RESEARCH EXPERIENCE — SELECTED PROJECTS

APPLV — Learning-based Planner Adaptation with VLA Models | *2026* | *IROS 2026*

- Fine-tuned Qwen2.5-VL-3B using LoRA; developed temporal-context and DPT-style regression components to predict safety-critical planner parameters rather than direct control actions.
- Implemented supervised learning and TD3 reinforcement learning fine-tuning across DWA, TEB, MPPI, and DDP; evaluated generalization on 300 unseen BARN environments (up to 94.34% success) and physical Clearpath Jackal trials.

CORAL — Hierarchical VLM Reasoning for Autonomous Underwater Exploration | *2026* | *IROS 2026* | *First author*

- Proposed a hierarchical perception-to-action architecture combining semantic waypoint reasoning, geometric generate–verify–correct validation, and dynamics-aware MPPI/DDP control.
- Achieved 94.28% coverage with zero collisions across ten simulated environments and 57% fewer VLM queries; real-world BlueROV2 experiments found all objects of interest with zero collisions.

RESEARCH EXPERIENCE — CONTINUED

DREAM — VLM-Guided Autonomous Underwater Navigation & Exploration | *2026* | *IEEE ICRA 2026* | *First author*

- Designed a vision-language autonomy framework integrating RGB, depth, segmentation, occupancy mapping, environment-aware reasoning, and adaptive exploration in previously unknown underwater scenes.
- Demonstrated collision-free exploration and strong coverage in simulated environments; in real BlueROV2 experiments, located all objects of interest without collisions.

AQUAACT — Underwater Robot Demonstration Dataset & Benchmark | *Ongoing* | *VLA policy training in progress*

- Built a language-annotated teleoperation dataset with approximately 2.0 hours of real and 17.9 hours of simulated demonstrations, synchronized forward-view video, IMU, pressure, and human joystick commands.
- Developed a firmware-in-the-loop ArduSub simulation/data collection pipeline calibrated to measured ROV velocity and turn response; defined closed-loop tasks and offline diagnostics for policy benchmarking.
- Currently training VLA/imitation-learning policies; closed-loop policy evaluation is planned.

CAVE-NAV — VLM-Guided 3D Underwater Cave Navigation | *2026* | *IEEE OCEANS 2026* | *First author*

- Developed multimodal scene reasoning over RGB, depth, and sonar-derived vertical clearance for safe 3D exploration through constrained cave passages; validated collision-free traversal in simulation.

Synthetic Data for Marine Fauna Detection | *2026* | *Ecological Informatics* | *Equal contribution*

- Co-developed a score-based diffusion synthetic imagery pipeline and YOLO-based evaluation, demonstrating how synthetic data improves multimodal visual perception under limited labeled data.

RESEARCH METHODS & TECHNICAL SKILLS

Machine Learning: Python, PyTorch, Hugging Face Transformers, LoRA fine-tuning, reinforcement learning (TD3), supervised learning, computer vision, diffusion-based data synthesis, scikit-learn.

Robotics & Embodied AI: ROS, NVIDIA Isaac Sim, MuJoCo, Gazebo, OpenCV, ArduSub/MAVROS, BlueROV2, simulation-to-real calibration, VLM/VLA systems, multimodal sensing.

Planning & Control: Model Predictive Path Integral (MPPI), Differential Dynamic Programming (DDP), Dynamic Window Approach (DWA), Timed Elastic Band (TEB), autonomous exploration, collision avoidance.

Additional programming: C, MATLAB; simulation/visualization: Blender.

INDUSTRY EXPERIENCE & ACADEMIC SERVICE

Engineer — Shandong Hi-speed Co., Ltd., China (Oct 2020–Jan 2021): Collaborated on ROS-based autonomous driving simulation, sensor integration, data logging, and testbed validation.

Reviewer: IEEE Robotics and Automation Letters (RA-L); IEEE International Conference on Robotics and Automation (ICRA); IEEE/RSJ IROS; IEEE Transactions on Vehicular Technology (TVT).