Navigating Dynamic Waters: Learning-Based Control and Potential-Field Planning for ASVs

X. Vicent, R. Velasco, E. Diaz, C. Wilson, and Dr. M. Dhanak

***Abstract*—This paper presents the design, integration, and testing of a Deep Reinforcement Learning (DRL)–enabled Autonomous Surface Vehicle (ASV) developed by Team Owltonomous for the 2026 RoboBoat Competition. The ASV is designed to operate reliably in dynamic maritime environments, including wind, waves, and moving obstacles. Building on prior competition experience, the platform was redesigned with a new asymmetrical catamaran hull and a streamlined electronics architecture to improve stability, serviceability, and reliability. The autonomy stack integrates AI-based perception, an Artificial Potential Field (APF) planner for collision-aware navigation, and a DRL-based control framework trained offline to enhance robustness under uncertain dynamics. Vision-based object detection is combined with LiDAR-based clustering for situational awareness. Simulation, data-driven validation, and on-water testing support the deployment of a reliable and adaptable ASV for RoboBoat 2026.**

***Keywords*— Autonomous Surface Vehicle, Deep Reinforcement Learning, Artificial Potential Fields, AI-Based Perception, Maritime Navigation**

I. Introduction

Integrating Artificial Intelligence (AI) into the perception, modeling, and control systems of an Autonomous Surface Vehicle (ASV) is the primary objective of Team Owltonomous in achieving reliable autonomy for the 2026 RoboBoat Competition. The dynamic environments the ASV faces during the competition, driven by currents and winds, demand adaptive systems capable of learning from data and refining their performance over time. Building upon prior competition experience and data gathered from previous competitions and teams, Team Owltonomous continues to leverage AI-driven methods to enhance navigation, perception, and control, while emphasizing robustness and autonomy for the competition tasks.

For the 2026 competition, the team worked on a comprehensive redesign of the ASV platform, including a new hull to address leaks in past competitions and inadequate weight handling, as well as a redesign of the electrical system due to the previous system's origins in the WAM-V 16 platform used in previous RobotX competitions. The updated platform enables easier integration between the vehicle's hardware and its autonomy stack, with greater consideration of competition demands rather than a general ASV design. Figure 1 shows the new hull tested in the water.

In parallel with the hardware redesign, AI methodologies were further incorporated into the ASV's system identification and control framework. Data-driven modeling techniques are employed to capture vehicle dynamics more accurately, allowing the control system to adapt to changes in mass distribution, hydrodynamic effects, and environmental conditions. Accurate system identification is critical when using nonlinear controllers and training neural networks for vehicle autonomy and navigation.

Complementing the modeling and control efforts, AI-based vision algorithms continue to play a similar role in object detection, localization, and situational awareness as they did in RoboNation competitions. The ASV is designed to execute navigation and task-specific behaviors with increased precision and resilience. This report presents the competition strategy, design decisions, and testing methodology that collectively support the team's goal of deploying a reliable, adaptable, and AI-enabled ASV for RoboBoat 2026.

Fig.1: Team Owltonomous new hull for RoboBoat 2025

II. Competition Strategy

For RoboBoat 2026, the team's strategy focuses on deploying a reliable autonomous surface vehicle optimized for navigation-intensive tasks under realistic environmental conditions, including wind and dynamic obstacles commonly observed in Sarasota. The ASV employs a deep reinforcement learning (DRL)– based control framework [1], trained offline to enhance robustness and adaptability under uncertain dynamics.

For this year, the required Evacuation Route & Return task is treated as the highest priority and is completed first to validate perception, planning, and control performance. Subsequent tasks are selected dynamically based on vehicle state, perception, confidence, and proximity within the course, allowing the ASV to opportunistically attempt additional challenges without compromising overall reliability. Tasks requiring higher uncertainty or rapid override behavior, such as Harbor Alert, are intentionally deprioritized and reserved for the final phase of a run due to lower confidence in robust air acoustic detection under competition noise conditions.

a) Course Approach

Across all competition tasks, the ASV operates under a unified autonomy framework that integrates perception, planning, and control within a ROS 2 architecture. Camera-based object detection helps recognize objects, while LiDAR-based clustering helps locate obstacles and understand their shape and position. These perception outputs are fused with onboard navigation data to estimate vehicle state and inform decision-making. Navigation is handled using an Artificial Potential Field (APF) planner [2], which generates collision-aware trajectories in real time under dynamic environmental conditions. Low-level actuation commands are generated by a deep reinforcement learning–based control policy. This shared autonomy stack allows task-specific behaviors to be implemented at the mission level while reusing validated perception, planning, and control components throughout the course.

b) Task Selection in Semifinals/Finals

Task selection is governed by a risk-aware mission policy that balances point potential with execution reliability. Core navigation tasks are prioritized early, while higher-uncertainty tasks are attempted only when system readiness thresholds are met. Readiness is evaluated based on perception confidence from vision detections, navigation risk inferred from LiDAR-based obstacle density, and the overall vehicle state. If confidence degrades, the ASV aborts the current objective and transitions to a conservative recovery behavior before reattempting tasks. This approach minimizes cascading failures and improves overall run consistency under variable environmental conditions.

c) Situational Awareness and Report

A graphical user interface (GUI) is implemented on the base station computer to provide real-time situational awareness of the ASV. The interface displays system health, vehicle state, and the task currently being executed by the autonomy stack. A heartbeat message is generated from this information using a dedicated C++ reporting node and transmitted to the competition judges in accordance with the defined communications protocol.

d) Task 1 - Evacuation Route & Return

Task 1 is approached as a core navigation validation task, serving as the baseline for subsequent mission execution. The strategy relies on AI-based object detection to identify and classify navigation buoys, combined with an APF planner for real-time obstacle avoidance and smooth trajectory generation. Classification algorithms differentiate buoy types and geometries, enabling robust gate identification under varying lighting and environmental conditions. Successful completion of this task confirms reliable perception, planning, and control performance, which is necessary before advancing to higher-complexity objectives.

e) Task 2 – Debris Clearance

Task 2 extends the navigation strategy from Task 1 into a cluttered, dynamic environment. Camera-based classification is prioritized to distinguish debris indicators and hazards, while LiDAR is used to localize obstacles and estimate their centroids for safe path planning. The APF planner dynamically adjusts the vehicle trajectory to remain within the channel and avoid obstacles without introducing task-specific control logic. Detected objects and task-relevant information are reported through the graphical user interface and heartbeat system.

f) Task 3 - Emergency Response Sprint

Task 3 is approached as a high-speed navigation challenge that emphasizes rapid perception, decision-making, and maneuverability. The ASV leverages camera-based classification to identify gate buoys, the color indicator, and the central yellow buoy, while LiDAR is used for obstacle localization and situational awareness. Once the indicator color is identified, the desired circumnavigation direction is selected at the mission level. The APF planner generates a collision-aware trajectory that prioritizes smooth, continuous motion through the gates and around the marker buoy, without reliance on station-keeping behaviors. Task progress, indicator color, and completion timing are reported through the graphical user interface and heartbeat system in accordance with competition requirements.

g) Task 4 - Supply Drop

The Supply Drop task is executed whenever delivery vessels are detected within a favorable distance of the

ASV's current objective. Camera-based classification distinguishes between yellow and black stationary vessels, while LiDAR supports relative localization and obstacle awareness. Based on prior competition experience, water delivery to yellow vessels is prioritized, as this capability demonstrated higher reliability in previous RoboBoat events. When conditions allow, the ASV aims a steady water stream at the designated target using vision-guided alignment. Ball delivery to black vessels using the racquetball launcher is treated as a secondary objective and is attempted only when system confidence is high..

h) Task 5 - Navigate the Marina

For the Navigate the Marina task, the ASV prioritizes vision-based classification to identify available docking slips, along with their associated color indicators and number signage. Camera-based perception determines slip availability and relative priority, while LiDAR supports obstacle detection and the estimation of the docking bay geometry. Once a valid slip is identified, the ASV computes the centroid of the target bay and generates a collision-aware approach using the APF planner. The vehicle executes a controlled, low-speed trajectory into the slip, emphasizing alignment and clearance within the constrained marina environment. If the initially selected slip is unavailable or obstructed, the ASV continues scanning adjacent bays until a valid docking opportunity is confirmed.

i) Task 6 - Harbor Alert

The Harbor Alert task is approached as an interrupt-driven behavior that overrides the current mission state when a valid acoustic signal is detected. The ASV is equipped with a USB omnidirectional microphone to monitor the acoustic environment and classify audible alerts in real time. Upon detection of a valid signal, the mission controller abandons the active task and transitions to the appropriate response behavior based on the classified alert. Due to uncertainty in reliably detecting signals amid high ambient noise from wind, waves, and nearby vessels, this task is deprioritized and reserved for the final qualification task. If high confidence is encountered during the competition, this task will be attempted during the semifinals/finals.

III. Design Strategy

a) Hull Design

The hull design for the RoboBoat 2026 Autonomous Surface Vehicle was revised in response to operational limitations observed during RoboBoat 2025 [3]. The previous hull was originally designed for a lower payload mass. As the total system weight increased to approximately 110 lb, the vehicle operated with reduced freeboard and an undesirably low waterline. This condition increases susceptibility to splashing and water ingress during maneuvering and in the presence of wakes. Additionally, mechanical failures were observed in the T200 thruster mounting structures, and leaks developed in locations that were difficult to access or repair in the field, resulting in lost testing and competition time. These issues motivated a redesign focused on increased displacement margin, structural robustness, and improved maintainability.

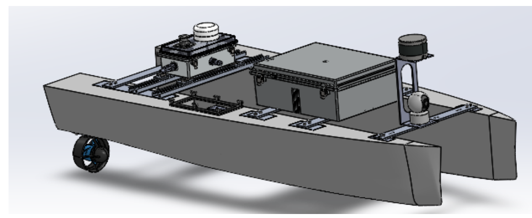

Fig. 2: CAD of the new hull with electronics and vision system

For RoboBoat 2026, the team adopted an outboard asymmetrical catamaran hull to address these limitations. Compared to the previous symmetric configuration, the asymmetrical hull geometry shifts the centers of buoyancy outward, increasing transverse stability and improving low-speed maneuverability for navigation and docking tasks. The slender pontoon design reduces wetted surface area and hydrodynamic interference between hulls, improving propulsion efficiency while providing increased buoyancy and freeboard at full system weight. Structural aluminum crossbeams connect the hulls and provide accessible mounting interfaces for propulsion components and the central electronics enclosure, improving rigidity and serviceability. This hull configuration establishes a more stable and reliable platform for the ASV's autonomy. Figure 3 shows the 2025 and 2026 hulls during the fabrication phase. More details of the design can be found in Appendix A.

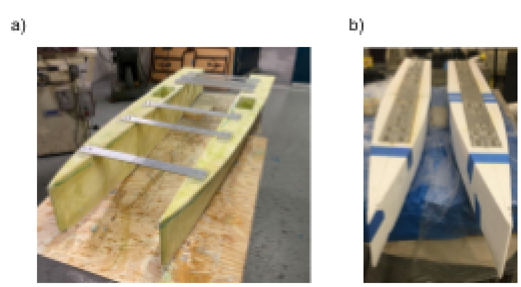

Fig. 3: a) Fabrication process of the new hull, b) fabrication process of the 2025 hull.

b) Propulsion System

The propulsion system for RoboBoat 2026 was redesigned to address thrust limitations observed during previous competition cycles. In prior years, the ASV used Blue Robotics T200 thrusters [4]; however, on-water testing and competition experience in Sarasota demonstrated that the thrusters lacked sufficient thrust authority under sustained wind conditions. To enhance robustness and ensure reliable navigation performance, the propulsion system was upgraded to Blue Robotics T500 thrusters [5], providing increased thrust margin and improved disturbance rejection capabilities.

For the 2026 competition, the ASV employs a simplified differential-drive configuration using two stern-mounted T500 thrusters. Unlike previous designs that incorporated additional bow thrusters for station-keeping, this configuration reflects a strategic decision to simplify the system and focus on reliable navigation tasks. Differential thrust provides sufficient maneuverability for the selected challenges while simplifying control allocation and reducing mechanical and electrical integration risk. The thrusters are mounted using a reinforced aluminum mounting system integrated into the hull structure. Figure 4 shows the CAD of the mounting system for the T500 thruster.

Fig. 4: CAD of the T500 mounting system

c) Guidance, Navigation, and Control System (GNC)

The Guidance, Navigation, and Control (GNC) system is built around a centralized electronics architecture for maximum reliability and ease of use. For RoboBoat 2026, a completely new electronics enclosure and internal layout were developed. The previous system was inherited from a WAM-V 16 platform and was both oversized and unnecessarily complex for a smaller ASV. The redesigned architecture better matches the RoboBoat platform's scale, power requirements, and maintenance needs, while significantly reducing integration and debugging overhead.

The electronics architecture is divided into two primary subsystems. The first subsystem is dedicated to high-power and safety-critical components, including the ASV's battery pack, safety interlock system, and motor controllers. Physically and electrically isolating these components enhances safety, facilitates fault diagnosis, and minimizes electromagnetic interference with sensitive navigation electronics.

The second subsystem houses the ASV's primary computing and sensing infrastructure. This includes the main on-board processor, the Jetson Orin AGX [6], networking hardware, battery-monitoring circuitry, power regulation, system power control, a custom-designed printed circuit board (PCB) for sensor data aggregation and power distribution, the RC receiver, and all vision-system power and data interfaces. Centralizing these components within a dedicated enclosure improves cable management, accessibility, and overall system robustness.

Design decisions for the 2026 system were heavily informed by lessons learned from previous competitions, with reliability as the primary driver of change. Historically, Team Owltonomous encountered recurring issues with two critical components: Pixhawk flight controllers used for GPS and IMU data, and USB-based vision cameras. Both systems were replaced in the current design.

Flight-controller functionality was migrated to a custom PCB integrating a dedicated IMU chipset, paired with an external marine-grade Garmin GPS sensor optimized for surface vessels. All navigation sensor data is processed directly on the custom PCB, reducing software complexity and eliminating the need for general-purpose flight-control hardware.

Overall, the redesigned GNC electronics architecture prioritizes reliability, modularity, and ease of maintenance while providing a scalable foundation for future development.

d) Subsystems (Water Pump and Ball Shooter)

The ball shooter was originally developed for the Maritime RobotX 2024 competition and utilizes a NEO 550 brushless motor [7] controlled by a SPARK MAX motor controller [8] to propel racquetballs. A smart servo governs the feeding mechanism, ensuring that only one ball is loaded into the shooter at a time. The transparent acrylic housing holds up to 4 racquetballs, providing ample capacity for competition. This system was successfully deployed during both the Maritime RobotX 2024 and RoboBoat 2025 competitions. Figure 5 shows the racquetball shooter.

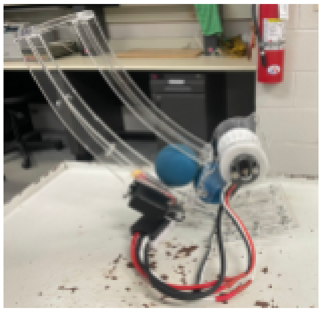

Fig. 5: Racquetball shooter

e) Electrical Integrations and Safety

The ASV safety system is designed to ensure reliable shutdown and fault handling during both autonomous and manual operation. Safety functionality is enforced through a combination of hardware-level power protection and software-based supervision, minimizing risk to personnel, the vehicle, and surrounding assets during competition and testing.

1) *System Overview:* Safety-critical electronics are housed within a newly designed electronics enclosure and supported by a PCB that provides power distribution, voltage regulation, and protection for onboard systems. While the overall safety philosophy remains consistent with previous platforms, the updated enclosure and PCB improve integration, accessibility, and reliability within the RoboBoat ASV form factor.

2) Battery Management and Emergency Stop: Battery health and power integrity are monitored through a Battery Management System (BMS) that provides protection against undervoltage, overcurrent, and fault conditions. The propulsion system, powered by TATTU 6S lithium-polymer batteries [9], can be electrically isolated through an emergency stop mechanism that enables immediate shutdown of thrust while preserving low-power operation of essential electronics when possible.

3) *Software Supervision:* Software-based safety monitors operate alongside hardware protections to detect conditions such as loss of communication or low battery levels. When triggered, these conditions initiate a controlled transition to a safe state. Hardware-level protections remain authoritative, ensuring safe operation even in the event of software failure.

f) Vision and Audio System

The ASV perception system integrates camera, LiDAR, and audio sensing to support object detection, classification, and localization during competition tasks. For RoboBoat 2026, USB cameras were replaced with Power-over-Ethernet (PoE) Amcrest 4K outdoor security cameras [9]. The ASV is equipped with three such cameras, each supporting low-light operation, automatic exposure adjustment, and integrated one-way audio capability. These features provide robust perception performance under the highly variable lighting and environmental conditions typical of RoboBoat competitions. The use of RJ45 connections and sealed, high-quality passthrough connectors significantly improves data integrity and long-term reliability compared to USB-based solutions.

Geometric perception and obstacle localization are facilitated by a Velodyne VLP-16 Hi-Res LiDAR [10], which provides 360° point-cloud data for buoy detection, obstacle avoidance, and environmental awareness. Camera-based semantic detections are fused with LiDAR clustering to improve robustness under varying lighting and weather conditions.

Audio perception is handled by an external omnidirectional USB microphone. Audible alerts are detected using a Goertzel-based frequency detector and published to the mission controller as ROS 2 messages, with interrupts driven by the frequency detector. Figure 6 shows the ASV vision and audio system flowchart.

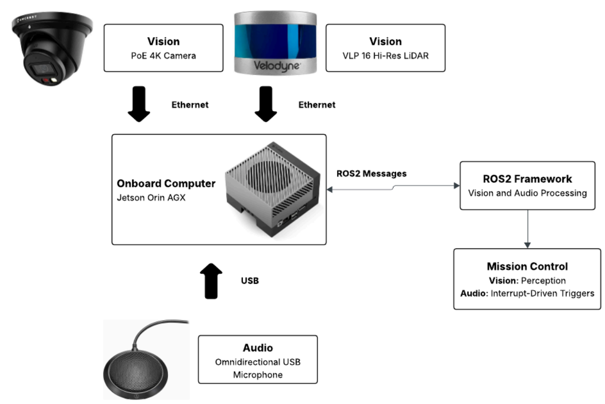

Fig. 6: Flowchart of the vision and audio system

g) Software Stack

The autonomy software for RoboBoat 2026 is implemented in ROS 2 and structured to clearly separate mission decision-making, perception, planning, and control. At the highest level, a Mission Control node selects and manages competition tasks based on vehicle state and strategy. Perception is handled by a dedicated AI vision system that uses a YOLO26-based detector [11] for object classification and localization, complemented by K-means clustering [12] applied to LiDAR point clouds for obstacle detection. Perception outputs are published to the autonomy stack as

ROS 2 messages. Details of the algorithms are in Appendices B and C.

Path planning is performed using an APF–based planner that generates collision-aware reference trajectories toward mission objectives. Low-level control is executed using a DRL controller trained with Proximal Policy Optimization (PPO)[13], which maps vehicle state and reference information into actuation commands. Final control outputs are transmitted via serial communication to a Teensy 4.1 microcontroller [14], which generates PWM signals for the electronic speed controllers driving the T500 thrusters. Detailed formulations of the AI-based controls are provided in Appendices D and E. Figure 7 shows the high-level flowchart of the software stack.

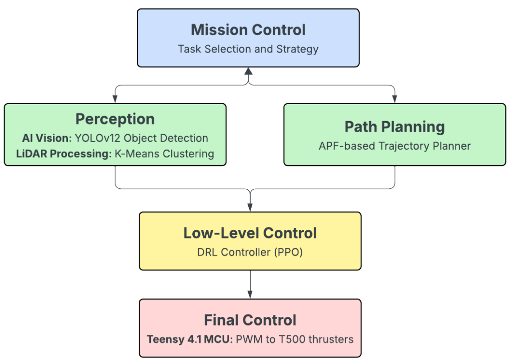

Fig. 7. Flowchart of the software stack

IV. Testing Strategy

Based on experience from previous RoboBoat and RobotX competitions, the team's testing strategy emphasizes early software validation, progressive system integration, and frequent iteration without constraining hardware development. To achieve this, testing is conducted across three complementary layers: simulation, offline data-driven validation, and on-water testing. This layered approach enables rapid development of autonomy algorithms while progressively increasing environmental realism and system complexity.

A. *Simulation*

Simulation is the primary environment for early development and regression testing of the autonomy stack. The team uses the Virtual RobotX (VRX) Gazebo simulator to develop and validate mission logic, perception pipelines, and APF-based path-planning algorithms under controlled conditions. Custom simulation worlds are created to approximate RoboBoat tasks, including navigation channels, debris fields, and docking scenarios. Simulation enables the rapid iteration of mission logic and control behaviors, independent of hardware availability, while providing a consistent framework for testing failure cases and edge conditions.

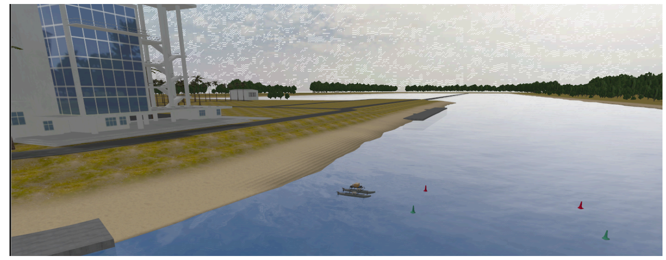

Fig. 8. ASV attempting Task 1 in the VRX Simulator

B. *Data Sharing Testing*

Offline testing using shared datasets is critical for validating perception and planning algorithms. The team leverages the RoboNation Data Sharing platform and archived datasets from previous RoboBoat and RobotX competitions to train and evaluate YOLO-based object detection models and LiDAR clustering algorithms. Recorded ROS bag files containing camera and point cloud data are used to verify detection robustness, classification accuracy, and planner behavior prior to deployment on the vehicle. This approach allows perception performance to be evaluated across a wide range of environmental conditions without requiring repeated field testing.

C. *On-Water Testing*

On-water testing is conducted to validate full system integration and assess ASV performance under real-world conditions, including wind, current, and waves. Testing progresses from basic navigation and obstacle avoidance to task-level execution, evaluating perception robustness, APF-based planning, and the deep reinforcement learning control policy under unmodeled dynamics. In addition to RoboBoat-specific testing, the team's research group has previously exercised core autonomy components on a WAM-V 16 platform, providing added confidence in the transferability of the underlying algorithms.

V. Acknowledgments

Team Owltonomous gratefully acknowledges the support of RoboNation, FAU's College of Mechanical and Ocean Engineering, FAU Student Government, and The Link Foundation for their financial support of this project. We also thank our faculty advisors, Dr. Manhar Dhanak and Dr. Pierre-Philippe Beaujean, for their continued technical guidance and mentorship, which have been instrumental in the development, integration, and testing of the Autonomous Surface Vehicle. Their expertise and support have played a critical role in advancing the team's research and competition efforts.

References

[1] Y. Li, “Deep reinforcement learning: An overview,” *arXiv preprint* arXiv:1701.07274, 2017. doi: 10.48550/arXiv.1701.07274.

[2] S. Xie *et al*., "The obstacle avoidance planning of USV based on improved artificial potential field," *2014 IEEE International Conference on Information and Automation (ICIA)*, Hailar, China, 2014, pp. 746-751, doi: 10.1109/ICInfA.2014.6932751.

[3] X. Vicent, T. Davies, C. Wilson, C. Workman, and M. Dhanak, “Optimizing Maritime Autonomy: AI-Enhanced Evolutionary Strategies for Intelligent Navigation,” *in Proc. RoboBoat 2025*, Sarasota, FL, USA, Mar. 2025.

[4] Blue Robotics, "T200 Thruster," [Online]. Available:https://bluerobotics.com/store/thrusters/t100-t200-thrusters/t500-thruster/. [Accessed: 27-Jan-2025]

[5] Blue Robotics, "T500 Thruster," [Online]. Available:https://bluerobotics.com/store/thrusters/t100-t200-thrusters/t500-thruster/. [Accessed: 14-Jan-2026]

[6] NVIDIA, "NVIDIA Jetson AGX Orin," [Online]. Available: https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-orin/. [Accessed: Oct. 06, 2024].

[7] REV Robotics, "REV Robotics UltraPlanetary Gearbox," [Online]. Available: https://www.revrobotics.com/rev-21-1651/. [Accessed: 27-Jan-2025].

[8] REV Robotics, "REV Robotics SPARK MAX Motor Controller," [Online]. Available: https://www.revrobotics.com/rev-11-2158/. [Accessed: 27-Jan-2025].

[9] Amazon, "Tattu 30000mAh 6S 25C LiPo Battery," [Online]. Available:https://www.getfpv.com/tattu-plus-10000mah-22-2v-25c-6s-lipo-smart-battery-pack.html .[Accessed: 14-Jan-2026.

[10] Mapix Technologies, "Velodyne VLP-16 Hi-Res | Lidar scanner with increased field of view," [Online]. Available: https://www.mapixtech.com. [Accessed: Oct. 06, 2024].

[11]G. Jocher, A. Chaurasia, and J. Qiu, “Ultralytics YOLO,” *GitHub repository*,2023.[Online]https://github.com/ultralytics/ultralytics[Accessed: Jan-18-2026]

[12] F. Samadzadegan, M. Maboodi, S. Saeedi, and A. Javaheri, "Evaluating the potential of clustering techniques for 3D object extraction from LIDAR data," in *Proc. First Int. Conf. Comput. Vis. Theory Appl. (VISAPP)*, vol. 2, 2006, pp. 149-154

[13] L. Engstrom, A. Ilyas, S. Santurkar, D. Tsipras, F. Janoos, L. Rudolph, and A. Madry, “Implementation matters in deep policy gradients: A case study on PPO and TRPO,” *arXiv preprint arXiv:2005.12729*, 2020, doi: 10.48550/arXiv.2005.12729.

[14] PJRC, "Teensy® 4.1," [Online]. Available: https://www.pjrc.com/store/teensy41.html. [Accessed: Oct. 06, 2024].

Appendix A: Hull Design

The RoboBoat 2026 Autonomous Surface Vehicle utilizes an outboard asymmetrical catamaran hull to address buoyancy, stability, and serviceability limitations identified in the previous platform. The prior hull operated with insufficient freeboard at an overall vehicle mass of approximately 110 lb, increasing susceptibility to water ingress and reducing reliability during competition. The redesigned hull increases displacement margin while remaining compliant with RoboBoat dimensional constraints, providing improved freeboard at full system weight.

The asymmetrical catamaran geometry features a curved outboard hull surface and a flat vertical inboard face, which shifts the centers of buoyancy outward and increases transverse stability (Fig. A.1). This configuration reduces hydrodynamic interference between hulls and decreases wetted surface area, improving propulsion efficiency and low-speed maneuverability. The hull was designed with a waterline length of 1.83 m; assuming a conservative Froude number of 1.1, the corresponding displacement velocity is 4.66 m/s, which is attainable given the available propulsion thrust.

Hull geometry was developed using a set of cross-sectional stations lofted in SolidWorks along guide curves aligned with the deck, keel, and design waterline. Structural aluminum cross-beams connect the hulls and provide rigid, accessible mounting interfaces for the deck and central electronics enclosure, improving structural integrity and maintainability.

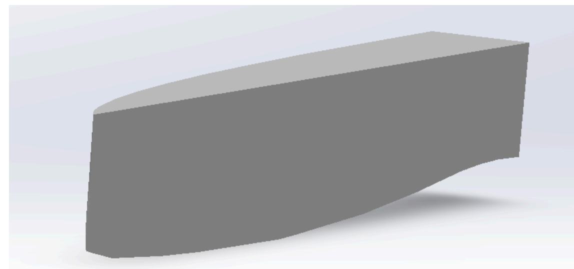

Fig. A1. CAD of the new hull

Table A1: Hull Geometry and Design Ratios

Length	5.92 ft
Length/Beam	9.4
Beam (hull)	0.63 ft
Beam (total)	1.77 ft
Separation Ratio	0.2
Height	0.86 ft
Waterline	0.33ft
Volume (1 hull)	1.63 cubic feet

Appendix B: ASV Hydrodynamic Modeling and Parameter Optimization

This appendix describes the dynamic modeling and system identification process used to obtain a high-fidelity three-degree-of-freedom (3-DOF) planar model of the ASV. Vehicle dynamics are formulated in a body-fixed reference frame, capturing surge, sway, and yaw motion.

The nonlinear equations of motion are expressed as:

$$\mathbf{M}_b\dot{\boldsymbol{\nu}} + \mathbf{C}_b(\boldsymbol{\nu})\boldsymbol{\nu} + \mathbf{D}_b(\boldsymbol{\nu})\boldsymbol{\nu} = \boldsymbol{\tau}_b$$

where

$$\boldsymbol{\nu} = [u\ v\ r]^T$$

denotes the body-frame velocities, and

$$\boldsymbol{\tau}_b = [F_x \; F_y \; M_z]^T$$

represents the control forces and yaw moment. The inertia, Coriolis, and damping matrices incorporate both rigid-body and hydrodynamic added-mass effects, in accordance with standard formulations for marine vehicle modeling.

Initial hydrodynamic coefficients are obtained from benchmark parameterizations for WAM-V–class vessels and refined using experimental data collected from acceleration, circle, and zig-zag maneuvers performed on the physical ASV. These tests capture surge drag, coupled sway–yaw dynamics, and turning behavior.

Parameter refinement is performed using a Covariance Matrix Adaptation Evolution Strategy (CMA-ES). For each candidate parameter set, the simulated ASV response is compared against measured data using identical thrust inputs. The optimization objective minimizes the weighted root mean square error (RMSE):

$$J = \mathrm{RMSE}(u) + \mathrm{RMSE}(v) + \lambda\,\mathrm{RMSE}(r)$$

where an increased weighting λ is applied to the yaw-rate error to prioritize accurate turning dynamics.

The resulting optimized model closely matches the physical ASV response and serves as the simulation environment for training and evaluating deep reinforcement learning–based controllers.

Table B1: Performance results of system identification between baseline and CMA-ES

Metrics	Baseline	CMA-ES
Acceleration RMSE (u) m/s	0.232	0.072
Acceleration RMSE (v) m/s	0.068	0.066
Acceleration RMSE (r) rad/s	0.013	0.014
Circle RMSE (u) m/s	0.317	0.099
Circle RMSE (v) m/s	0.097	0.097
Circle RMSE (r) rad/s	0.153	0.019

The CMA-ES optimization significantly improves the accuracy of the ASV model compared to the benchmark parameters. In the acceleration tests, surge RMSE is reduced from 0.232 m/s to 0.072 m/s, indicating a much better representation of longitudinal dynamics. In the circle tests, yaw-rate RMSE is reduced from 0.153 rad/s to 0.019 rad/s (≈88% reduction), demonstrating a substantial improvement in turning behavior. These results confirm that the optimized model more accurately captures both straight-line and maneuvering dynamics, making it suitable for controller design and DRL training.

Appendix C: Deep Reinforcement Learning Controller for ASVs

This appendix details the deep reinforcement learning (DRL) controller used for low-level ASV navigation and trajectory tracking. The controller is trained using Proximal Policy Optimization (PPO) on the CMA-ES–optimized hydrodynamic model described in Appendix B.

The observation space is defined using line-of-sight navigation variables, cross-track error, heading alignment, and body-frame velocities, enabling the policy to generalize across varying initial conditions without reliance on predefined paths. The action space consists of normalized differential thrust commands applied to the port and starboard thrusters.

The reward function is designed to promote efficient waypoint convergence, heading alignment, and stable motion while penalizing lateral drift and excessive yaw rate:

$$R_t = k_d(d_{t-1} - d_t) + k_\psi \cos(\psi_{\mathrm{LOS}}) - k_v|v_t| - k_\omega|\omega_t| + k_{\mathrm{time}}$$

where d_t is the distance to the waypoint, ψ_{LOS} is the line-of-sight heading error, v_t is lateral velocity, and w_t is yaw rate. A terminal bonus is applied upon reaching the waypoint within a specified tolerance.

Training episodes consist of waypoint navigation tasks with randomized initial positions and headings, and waypoint distances ranging from 3 to 10 meters.

Table C1: Performance results between PID controller and DRL

Controller	Mean Speed (ms)	Heading Rate RMS (rad/s)	Mean WP error (m)	CRE RMS (m)
PID HS	0.99	0.062	1.00	1.05
DRL	1.03	0.068	0.76	0.97

Controller performance is evaluated in simulation and compared against a baseline PID heading–speed controller. As shown in Table X, the DRL controller achieves lower mean waypoint error (0.76 m vs. 1.00 m) and reduced cross-track error RMS (0.97 m vs. 1.05 m), while maintaining comparable speed and heading-rate stability. These results demonstrate improved path-following accuracy and robustness, validating the effectiveness of combining learning-based control with a high-fidelity, data-driven vehicle model.

Appendix D: Perception Algorithms: LiDAR Clustering and Vision-Based Object Detection

K-means clustering is employed to segment LiDAR point cloud data into distinct object candidates for obstacle detection and localization. The algorithm partitions a dataset into k clusters by minimizing the intra-cluster variance. The set of cluster centroids is defined as

$$C = \{c_1, c_2, ..., c_k\}$$

Each data point xjx_jxj is assigned to the nearest centroid based on Euclidean distance, forming the cluster set

$$S_i = \{ x_j : \| x_j - c_i \| \leq \| x_j - c_m \|, \forall m, 1 \leq m \leq k\}$$

Centroids are updated by averaging all points assigned to each cluster:

$$c_i = \frac{1}{|S_i|} \sum_{x_j \in S_i} x_j$$

The clustering process minimizes the objective function

$$J = \sum_{i=1}^{k} \sum_{x_j \in S_i} \| x_j - c_i \|^2$$

and iterates until convergence, defined by

$$\| c_i(t + 1) - c_i(t) \| < \epsilon, \forall i \in \{1, ..., k\}$$

After convergence, clusters are filtered using geometric constraints such as size, height, and spatial consistency to identify competition-relevant objects, including buoys and obstacles.

To complement LiDAR-based perception, the ASV employs a deep-learning vision pipeline based on the *You Only Look Once* (YOLO) object detection framework. YOLO performs single-stage, real-time object detection by processing the entire image in one forward pass, enabling high inference speed while maintaining accurate localization and classification.

The model is fine-tuned using a custom maritime dataset containing competition-specific objects, including navigation buoys, docking structures, and vessels, captured under varying lighting conditions and water surface reflections. This training improves robustness in low-contrast and low-visibility environments commonly encountered during RoboBoat operations.

The optimized YOLO model is deployed on NVIDIA Jetson-class embedded processors, enabling efficient, real-time inference within the ASV's perception pipeline. Detected objects are published as ROS 2 messages and fused with LiDAR-based clustering outputs to provide semantic classification and geometric localization. This combined perception approach enhances detection reliability and situational awareness across navigation and task execution scenarios.

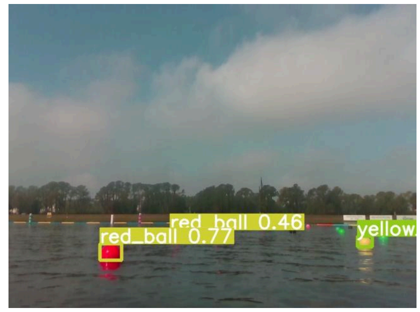

Fig. D1. YOLO26 Custom Model Detector Output

Appendix E: List of major components

Component	Vendor	Model/Type	Specs	Custom/Purchase	Cost	Year of Purchase	Reasoning
ASV Computer	NVIDIA	Jetson Orin AGX	32 GB RAM 2 Tb SSD	Loan from FAU Engineering Department	1999.99	N/A	ASV Vision Processing
ASV Custom PCB	FAU E-Lab	ASV Main Control Board	Teensy 4.1 Microcontroller, 8 PWM out, embedded safety, fused outputs, 6 subsystem communications, computer power control	Custom	520.00	2025	ASV motor Control, Radio communication, subsystem communication and power distribution, RC monitoring, light stack control, safety system monitoring
ASV LiDAR	Velodyne	VLP-16 HighRes	16 Beams, 100 m radius	Loan from FAU Engineering Department	15,000.00	N/A	Sensor for autonomous navigation, short and medium ranged performance
ASV Camera	Amcrest	4k PoE Camera	15 fps 129 DEG FOV	Purchased	399.97	2025	Object classification and outdoor water resistant sensor
ASV GPS	Garmin	P19x HVS	10Hz position, velocity and time data NMEA	Purchased	169.99	2024	ASV Sensors
ASV IMU	Adafruit	BNO055 Absolute Orientation Sensor	100Hz Orientation, Angular Velocity, Acceleration	Purchased	34.95	2025	ASV Sensors

ASV Network Bridge Module	Ubiquiti	RP-5AC-GEN2	5 GHz	Loan from FAU Engineering Department	230.00	N/A	ASV Communications
ASV Light Tower	Banner Engineering Corporation	TL50BLGYRQ	Stack Light Complete Unit Green, Red, Yellow LED Wire Leads	Loan from FAU Engineering Department	263.00	N/A	ASV Safety and Status
ASV Network Edgerouter	Ubiquiti	ER-X-SFP	5-PORT GIGABIT Router w/POE	Purchased	99.00	N/A	ASV Communications
ASV PoE Switch	NETGEAR	5-Port PoE Gigabit Unmanaged Switch	4x PoE 5x Gigbabit Ethernet Ports 63W	Purchased	69.99	2025	Networking and power for PoE cameras
ASV Thrusters	Blue Robotics	T-500	16 Kg f thrust	Purchased	952.00	2025	ASV Propulsion
ASV Batteries	Amazon	6S LiPO TATTU 30000 mAh	6S LiPO TATTU 30000 mAh	Purchased	999.99	2025	ASV Propulsion
Ball Shooter Brushless Motor	Rev Robotics	NEO 550 Brushless Motor	Brushless Motor	Purchased	28.00	2025	Ball Shooter
Ball Shooter Motor Controller	Rev Robotics	SPARK MAX Motor Controller	Motor controller for brushless motor	Purchased	93.50	2025	Ball Shooter
Ball Shooter Servo	Rev Robotics	Smart Robot Servo	Servo for Balls	Purchasd	30.00	2025	Ball Shooter

VI. Appendix F: Testing Plan

I. Scope

Team Owltonomous employs a layered testing strategy to validate the Autonomous Surface Vehicle (ASV) across perception, planning, and control subsystems while minimizing dependency on hardware availability. Testing is structured to progress from simulation and offline data-driven validation to full system integration and on-water deployment. The Virtual RobotX (VRX) simulator and RoboNation Data Sharing resources form the foundation of early development, enabling rapid iteration prior to physical testing.

II. Schedule

The testing schedule is organized around three primary phases: simulation-based development, offline validation using recorded datasets, and on-water system integration. Table I summarizes the testing timeline, objectives, and environments.

Table F1. Team Owltonomous' testing schedule and results

Start/End Date	Documentation	Target	Result	Environment	Member Presence
06/01/25 - 07-31-25	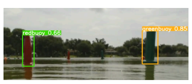	YOLO vision model training using shared datasets	≥ 0.60 confidence in competition objects	RoboNation Data Sharing, RoboFlow, MRC Server	AI Team
08/20/25 - 09/15/25	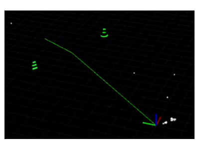	LiDAR clustering and object detection validation	Reliable detection of buoys and docking bays	Shared ROS bags, RViz	Software & Controls Team
09/01/25 - 09/30/25		Mission logic and APF path planning validation	Task completion without vision dependency	VRX Simulator	Software & Controls Team
10/01/25 - 11/15/25		DRL controller training and evaluation	Stable waypoint tracking (3–10 m)	Simulation	Software & Controls Team

01/15/26 – 01/31/26	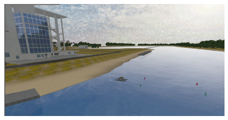	Integrated mission testing	End-to-end autonomy validation	VRX Simulator	Full Team
02/01/26 - 02/17/26		On-water system testing	Full system integration under disturbances	On-Water tests at the US-Intracoastal Waterway in Dania Beach, FL	Full Team

III. Resource & Tools

Testing activities leverage a combination of software and physical resources. Simulation and offline testing are conducted using the VRX Gazebo simulator on dedicated desktop workstations. Perception development relies on datasets from the RoboNation Data Sharing platform and archived ROS bag files. Hardware development and integration are supported by FAU's electronics and machine shop facilities.

For on-water testing, the ASV is deployed using a davit system, and anchored buoys are used to replicate elements of the competition task. These resources allow progressive validation of perception, planning, and control subsystems in increasingly realistic environments.

IV. Environment

Software testing is initially conducted in simulation and using shared datasets due to limited access to competition-grade physical objects. This approach enables validation of perception algorithms and mission logic prior to physical deployment.

On-water testing is conducted at FAU's SeaTech facility along the US Intracoastal Waterway in Dania Beach, Florida. Testing is performed within a protected marina environment to minimize external traffic and environmental risks while enabling the controlled evaluation of full system autonomy.

V. Risk Management

Safety considerations are integrated into all testing activities. Battery handling and charging procedures follow established protocols to mitigate risks associated with lithium-polymer batteries. All work conducted in the electronics and machine shops adheres strictly to FAU laboratory safety guidelines.

On-water testing introduces additional risks due to wildlife, vessel traffic, and deployment operations. To mitigate these risks, testing sessions are coordinated with FAU SeaTech staff, required documentation is completed in advance, and all team members receive a safety briefing prior to deployment. Personal protective equipment, including hard hats, is mandatory during davit operations.

VI. Results

Between Fall 2025 and early Spring 2026, the team validated key components of the autonomy stack through simulation, shared datasets, and limited physical testing. Notable outcomes include stable waypoint tracking with the DRL controller, successful task execution without vision-based navigation in simulation, and verification of the electronics enclosure's waterproofing. The use of shared RoboNation datasets significantly accelerated perception development and reduced reliance on early field testing.

VII. Appendix G: Acronyms

Acronym	Definition
AI	Artificial Intelligence
ASV	Autonomous Surface Vehicle
APF	Artificial Potential Field
DRL	Deep Reinforcement Learning
PPO	Proximal Policy Optimization
CMA-ES	Covariance Matrix Adaptation Evolution Strategy
LiDAR	Light Detection and Ranging
YOLO	You Only Look Once
PWM	Pulse Width Modulation
GNC	Guidance, Navigation, and Control
PCB	Printed Circuit Board
CMA-ES	Covariance Matrix Adaptation Evolution Strategy
PID	Proportional-Integral-Derivative
ROS2	Robot Operating System 2
VRX	Virtual RobotX
FAU	Florida Atlantic University
USB	Universal Serial Bus
PoE	Power over Ethernet
USB	Universal Serial Bus
BMS	Battery Management System
GUI	Graphical User Interface