RoboSub 2026 – Technical Design Report
Stanford University

Mechanical: Kai Song*, Félicie Hoffmann, Eden Lim, Caroline Goodrich
Electrical: Ryota Sato*, Eli Silverstein, Samantha Lau, Zhikai Huang, Kayla Lee, John Li
Software: Rohan Bhowmik*, Malti John, Maya Bridgman, Svea Strassburger, Mason Condello
*indicates Sub-Team Lead

May 27th, 2026

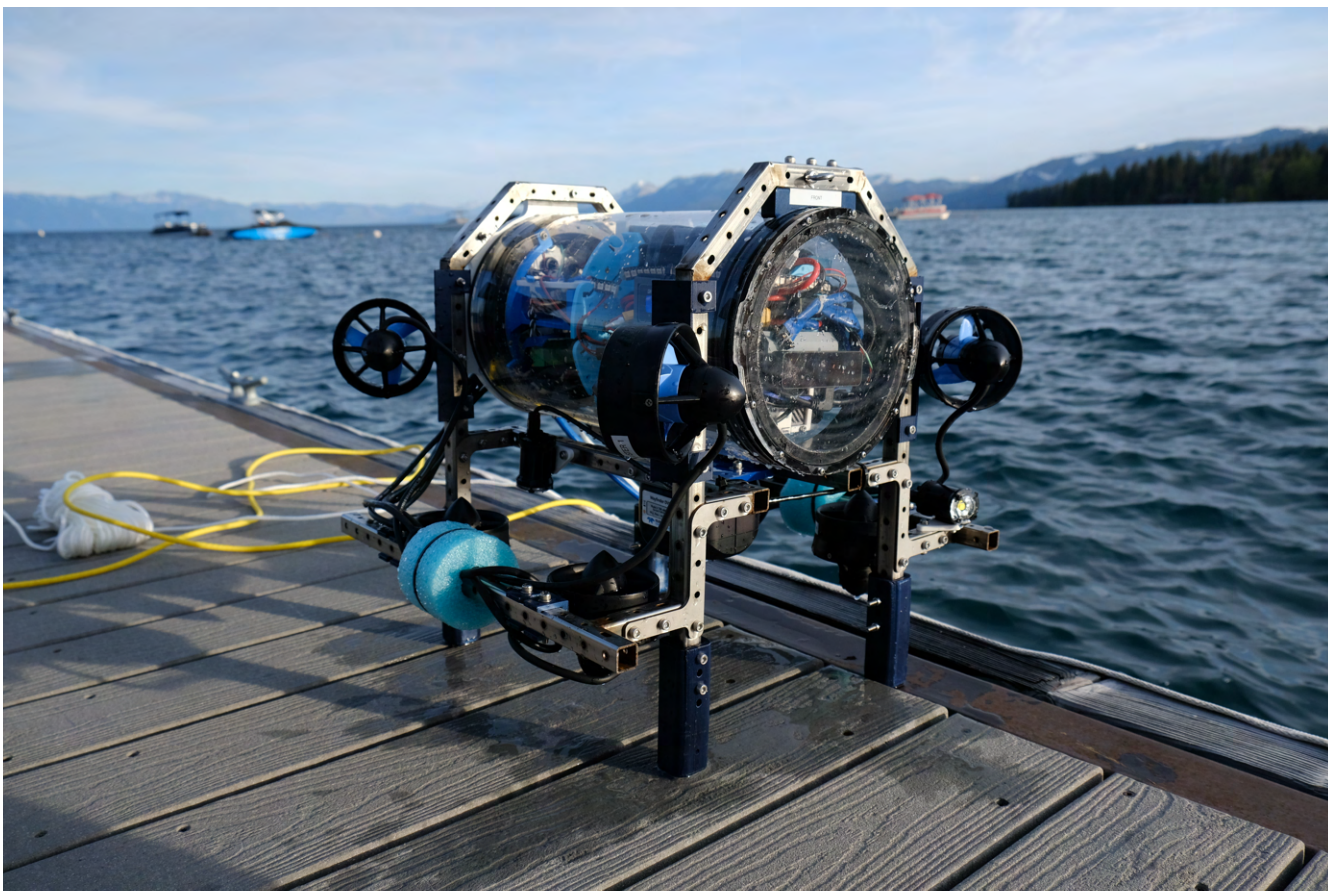

STANFORD ROBOSUB

Abstract

This report presents Stanford's next-generation Autonomous Underwater Vehicle program for RoboSub 2026, built around a modular, rapid-deployment ecosystem that prioritizes maintainability and subsystem decoupling under a compressed development timeline. Mechanically, a sliding internal electronics chassis on carbon-fiber guide rails minimizes in-water maintenance downtime. The electrical subteam contributes three custom PCBs, an intelligent dual-battery rail enabling hot-swappable power transitions, and the full audio pipeline—a dual-microcontroller four-hydrophone array running Time-Difference-of-Arrival (TDOA) processing for the pinger task. The software stack is a hardware-agnostic ROS2 infrastructure built around an edge-optimized YOLOv8 perception pipeline. To overcome underwater optical distortion and limited pool access, a JTA-GAN synthetic-data pipeline trains the vision models largely in simulation. Weekly testing at the Stanford Avery Aquatic Center plus a 34 ft hydrostatic dive at MBARI on our legacy platform validate the pressure-hull and sealing methodology carried into the new build, anchoring competition readiness across the core task sequence.

1 Competition Strategy

Under a compressed engineering lifecycle, our strategy is to attempt every task on the RoboSub course while ranking engineering effort by input modality: **localization** > **vision** > **audio**. Localization is the substrate every other task runs on top of; vision gates point-yielding manipulation; acoustics is the most environment-sensitive input. This hierarchy determines where marginal hours go, not which tasks we attempt.

1.1 Two-Platform Deployment

We run two vehicles concurrently to de-risk both schedule and competition. Gen 3 is the primary scorer, optimized for the competition loop. Gen 2 is a stable testbed during development—keeping the software team in the water while Gen 3 hardware is fabricated—and travels to competition on deck as a backup that can pick up the autonomy stack if Gen 3 is downed, providing graceful degradation rather than a single point of failure.

1.2 Full Task Coverage

The Gen 3 vehicle is integrated to attempt every task in the autonomy challenge: Heading Out (Coin Flip), Begin Assessment (Gate), Avoid Debris (Slalom), Recon (Bins), Deploy (Torpedoes), Resupply (Octagon), Return Home, Pingers, and Inter-vehicle Communications (IVC). The mechanical actuators for the manipulation tasks (three-DOF claw, spring torpedo launcher, pendulum-release dropper) and the four-hydrophone acoustic array are all physically integrated, and the master `.yaml` task block lets any subset be enabled or skipped at deploy time.

1.3 Capability Ranked by Input Modality

Engineering hours are allocated by the dependency stack each task sits on, ordered **localization** > **vision** > **audio**:

Tier 1 – Localization. These tasks include Coin Flip, Gate, Slalom, and Return Home. They are prioritized because every downstream task inherits their failure modes: a vehicle that cannot hold heading at the Gate cannot align torpedoes either. Reliability investments concentrate here—EKF tuning, DVL/IMU/depth fusion, PID step responses, and marker-corrected dead-reckoning bounds.

Tier 2 – Vision. These tasks include Bins, Torpedoes, and the Octagon. Once pose is solved, vision becomes the constraint. These tasks yield the highest per-task points but require detection robust to underwater attenuation, scattering, and shaded profiles. The JTA-GAN synthetic pipeline plus YOLOv8 on a hand-labeled dataset (Section 2.1.1) is the bulk of our software investment after the localization foundation is locked.

Tier 3 – Audio. These tasks include the Pingers and IVC. The electrical-subteam-owned audio pipeline was field-validated at last year's RoboSub competition, where it captured pinger data up to 45 kHz. We are also implementing a novel ML-based intervehicle-communication (IVC) channel on top of the same hydrophone front end. Audio still ranks below vision because pool-by-pool acoustic conditions (multipath, SNR) vary more than vision conditions.

1.4 Task Sequence and Reliability Gating

$$\text{Coin Flip} \longrightarrow \text{Gate} \longrightarrow \text{Slalom} \longrightarrow \text{Bins} \longrightarrow \text{Pingers} \longrightarrow \text{Torpedoes} \longrightarrow \text{Octagon} \longrightarrow \text{Return Home} \longrightarrow \text{IVC} \quad (1)$$

Coin Flip sets initial orientation. Gate and Slalom exercise localization first. Bins follows because downward-facing dropping has wide lateral tolerance, cross-verifying altitude and position-hold before the harder forward-vision targets. Pingers then runs as the ordering decision: from the acoustic localization, the state machine commits to either Torpedoes → Octagon or Octagon → Torpedoes, so the nearer target is hit first and the further one second, minimizing traversal between high-value scoring stations. Return Home re-closes the loop on the same localization stack. IVC runs last as a stretch demo, attempted only once scoring tasks are banked.

1.5 Complexity vs. Reliability Trade-Off

Within the three-tier hierarchy, each task's relative priority is governed by an empirical Task Priority Score P_t:

$$P_t = \frac{R_m \times R_s \times \text{Points}}{\lambda_h + \lambda_s} \quad (2)$$

where R_m, R_s are estimated mechanical and software reliability coefficients ($0 \leq R \leq 1$), Points is the competition valuation, and λ_h, λ_s are hardware and software development costs in developer-hours.

Task	$\mathbf{R_m}$	$\mathbf{R_s}$	$\boldsymbol{\lambda_h}$	$\boldsymbol{\lambda_s}$	**Tier**
Coin Flip	1.00	0.95	0	5	**T1 Loc.**
Gate	0.95	0.95	10	15	**T1 Loc.**
Slalom	0.90	0.85	15	40	**T1 Loc.**
Return Home	0.95	0.90	5	20	**T1 Loc.**
Bins	0.85	0.75	30	60	**T2 Vis.**
Torpedoes	0.70	0.65	55	85	**T2 Vis.**
Octagon	0.65	0.55	30	80	**T2 Vis.**
Pingers	0.85	0.70	15	60	**T3 Aud.**
IVC	0.70	0.30	5	100	**T3 Aud.**

Table 1: Task allocation matrix (λ in hrs). All tasks remain in the autonomy state machine; the tier column reflects the localization > vision > audio priority that drives where each subteam's marginal hours go.

Table 1 captures the allocation. Tier-1 localization absorbs disproportionate hours per point of yield because every downstream task inherits its outputs; Tier-2 vision concentrates the largest absolute hour count because the perception stack is reused across the course; Tier-3 audio is funded at a stretch level.

2 Design Strategy

2.1 Software Architecture

Our stack is structured into decoupled ROS2 modules linked by well-defined message nodes (Fig. 1), with central `.yaml` parameters across all modules to support fast on-site tuning.

The planning module maps the routine as waypoint vectors plus action tasks; a master `.yaml` holds the state machine in task blocks specifying *what* runs, while `path_generator` decides *how* from spatial coordinate files (e.g., adapting to the start-gate coin-flip). The perception stack inherits from a `generic_camera` base so models and orientations change via config—not code—publishing hardware-agnostic depth/RGB topics. The hardware tier normalizes sensor output through

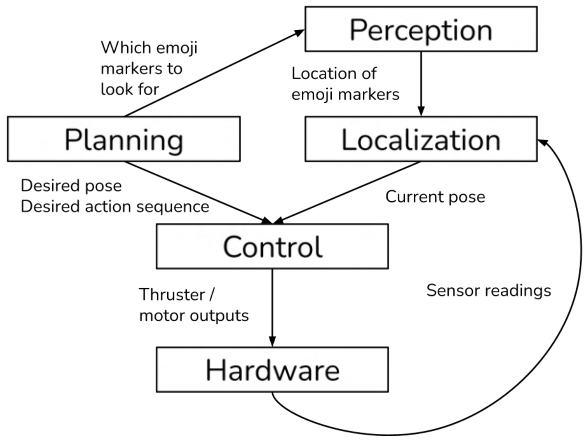

Figure 1: Software System Architecture and Module Interactions.

a `generic_sensor` base into five variables (linear position/velocity, angular orientation/velocity, linear acceleration), filtered per axis before EKF fusion. Thruster mixing is a standalone text file.

2.1.1 Perception Pipeline

Underwater tracking is impeded by light absorption, particulate scattering, and wavelength-dependent attenuation (red desaturates first). For low-latency edge inference we deployed YOLOv8 [1] for its light single-stage architecture, planned for extension with additional convolutional layers for feature retention [2]. Because our physical imagery is restricted to the Stanford pool, we built a 1,000+ instance hand-labeled dataset (Fig. 4, Appendix A) split into "dry," "wet," and a shaded validation set.

Models trained on dry samples generalized poorly underwater (Fig. 5, Appendix A), and mixed-domain training lost precision on shaded profiles. We therefore apply physics-based underwater style transfer to dry imagery using the JTA-GAN framework [3].

2.1.2 Control: from PID toward Model-Predictive

We currently run ROS `robot_localization` EKF fusion with decoupled PID control. Because underwater forces are non-linear and coupled, PID requires exhaustive tuning and introduces cross-axis error. We are deriving analytical hydrodynamics to inject a state-dependent Jacobian into our estimator and model-predictive trackers [4]. From the rigid-body acceleration $\dot{\boldsymbol{\nu}} = M^{-1}(\boldsymbol{\tau} - \boldsymbol{C}(\boldsymbol{\nu})\boldsymbol{\nu} - \boldsymbol{D}(\boldsymbol{\nu})\boldsymbol{\nu} - \boldsymbol{g}(\boldsymbol{\eta}))$ we define $\dot{\boldsymbol{x}} = f(\boldsymbol{x}, \boldsymbol{\tau})$ and compute $\mathbf{F}_k = \partial f / \partial \boldsymbol{x}$, which we expect to cut dead-reckoning drift versus stock ROS2 localization.

2.2 Electrical Architecture

The RoboSub electrical system was fully revamped with three custom-redesigned PCBs and a unified power and communication rail, enabling a streamlined PCB slide design. Fig. 2 shows the full system diagram; components inside the dark outline are housed within the submarine, while those outside face outward in the bulkhead.

2.2.1 Battery Management Board

Battery replacement was a major time sink at competition: swapping a battery required unpowering the submarine and fully rebooting the main compute device. To speed this up, we designed a PCB (Fig. 6, Appendix B) that connects two batteries to the submarine simultaneously at high current (up to 100A), letting us replace one battery without unpowering the onboard compute module. Solid-state devices were ruled out by the thermal constraints of a sealed underwater tube with no thermal coupling, and relays by size constraints. Instead, we connect a low-current, low-voltage secondary supply to the power bus through a diode; this supply keeps the compute module alive while the main battery, connected directly to the bus, is replaced.

2.2.2 Microcontroller Breakout Board

This year we redesigned the microcontroller breakout board (Fig. 7, Appendix B) with every relevant sensor and I/O predefined: the BlueRobotics Bar02 10m depth sensor, BlueRobotics switch (kill switch), green and red LEDs, four headlights, a leak sensor, and up to seven PWM servos (for the torpedo, dropper, and claw systems). A 10-pin interconnect ribbon cable drives the eight thrusters

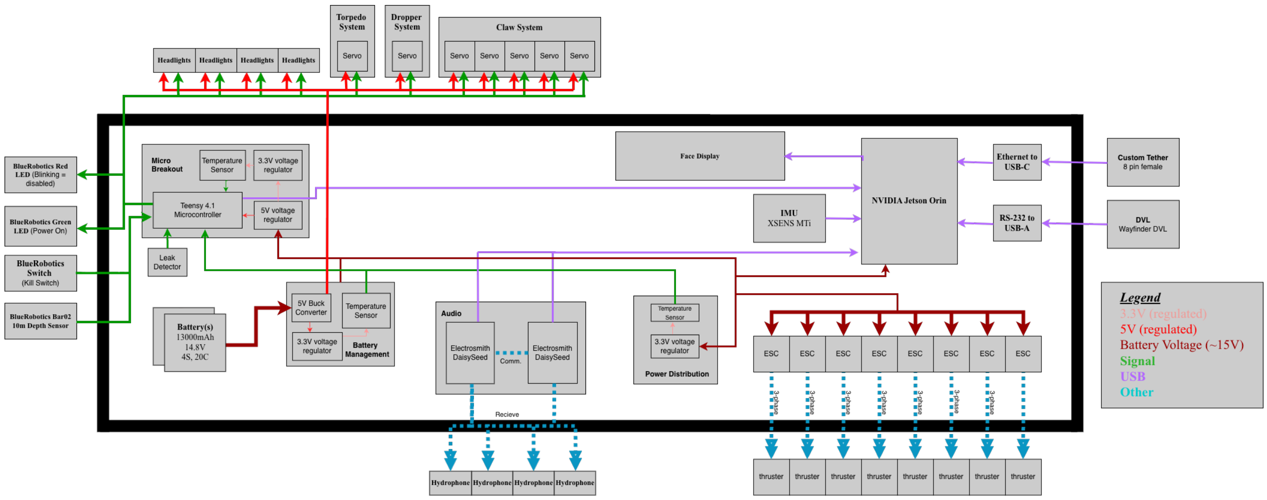

Figure 2: Gen 3 RoboSub electrical subsystem.

via PWM on the power distribution board, and we added I2C, UART, and SPI connectivity. All connections use easy-to-attach Molex connectors. The board also carries an MCP9808 temperature sensor, an MS5607-02BA03 pressure sensor, and an SHT31 humidity sensor. Firmware runs on a Teensy 4.1, which communicates over serial with the main NVIDIA Jetson Orin host. The latest revision is currently in progress; the figure shows the most up-to-date board layout as of this TDR submission.

2.2.3 Power Distribution Board

The power distribution board (Fig. 8, Appendix B) powers the eight thrusters, connecting them to battery voltage through terminal blocks and to PWM control from the microcontroller through Molex connectors. An ACS713 Hall-effect sensor measures the current to each thruster; an ADS7830 analog-to-digital converter serializes the eight sensor outputs into a single I2C stream to the microcontroller. An on-board MCP9808 temperature sensor also reports over I2C.

2.2.4 Connectivity

The connectivity system uses an Ethernet port for live data monitoring and tethered operation. To raise bandwidth from roughly 1 Mbps to at least 1 Gbps, we selected gel-filled Cat6 cable for durability in submerged environments, entering the submarine through Blue Robotics WetLink penetrators; underwater splices use the BlueRobotics WetLink cable splice kit, with all connections epoxied for water resistance.

2.2.5 Audio

The audio subsystem supports two threads: the competition pinger localization task and an in-development intervehicle-communication capability. Both share a four-hydrophone front end of custom-fabricated piezoelectric elements, built with guidance from the Stanford CCRMA and mounted at the four corners of the AUV for maximum baseline.

The pinger pipeline runs on two Electrosmith Daisy Seed boards (STM32, C++), each sampling two channels at 96 kHz and synchronized master/slave over a GPIO clock line; we developed custom serial and FFT libraries for the platform. A 64-sample FFT gives 0.66 ms time resolution ($\sim$1 m localization error at 1500 m/s) and 1500 Hz bins (96000/64), separating the 2500 Hz-spaced pinger frequencies cleanly. The full processing chain (Fig. 9, Appendix B)—A/D, filtering, envelope detection, first-arrival decision—runs entirely on the Daisy Seed and was field-validated at last year's RoboSub competition at pinger excitations up to 45 kHz (Fig. 10, Appendix B).

The intervehicle-communication thread (Fig. 11, Appendix B), currently being integrated, uses natural spoken language as the medium. A WaveNetstyle decoder [5]—a topology common in neural amplifier modeling—trained on real underwater audio undoes water-induced distortion and feeds OpenAI Whisper [6] for transcription. An off-the-shelf speaker drives the channel; our hydrophone array receives.

2.2.6 Auxiliary Components

The torpedo launcher, dropper, and three-DOF claw are actuated by PWM servos from the Teensy, powered from an external regulated supply since servo current exceeds MCU drive. Waterproofing was validated by submerged sweep tests on multiple servo models with epoxy-sealed connections. Additional components: PWM-controlled lumen headlights, IMU, DVL, and an OLED display serving as the vehicle's face.

2.2.7 Safety & Failsafes

The primary failsafe is a BlueRobotics kill switch wired into the microcontroller breakout board: twisting the switch asserts a hardware interrupt that drops all eight ESC PWM lines to neutral, deactivating propulsion regardless of software state. In parallel, the onboard leak sensor and humidity/pressure telemetry are polled by the Teensy at every control cycle. Status of the submarine is easily observable via green/red LEDs facing outside the bulkhead.

2.3 Mechanical Architecture

Our 2026 strategy trades the legacy deep-water platform's mass-heavy, COTS-heavy design—which restricted electrical-bay access and hurt agility—for a Gen 3 chassis optimized purely for the competition loop, modeled in OnShape (Fig. 3).

Sliding electronics bay. Carbon-fiber tubes with marine-epoxied linear bearings form a low-friction sliding tray; flexible cable chains keep umbilicals managed so the full circuit stack is serviceable without disconnecting structure (Fig. 12, Appendix C).

Material & modularity. Freshwater arenas let us laser-cut hull plates and bulkheads from low-cost aluminum instead of titanium/acrylic, cutting prototyping turnaround. Thruster brackets clip onto a magnetic rail between the carbon rods for rapid in-pool rebalancing.

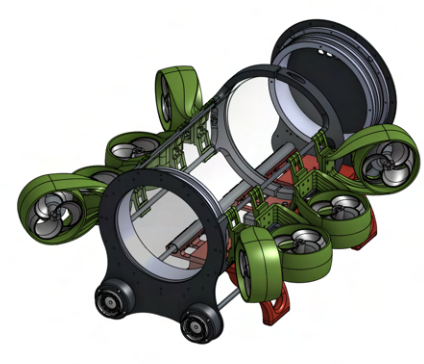

Figure 3: Gen 3 full-system CAD assembly.

Custom propulsion & end-effectors. 5010 360 kV BLDC motors with fluid-saturated anti-corrosion potting and resin-printed props cut per-motor cost $\sim$90% while letting us iterate blade pitch. Phase lines exit through MT60 bulkheads. Three lower-frame actuators—a three-DOF retrieval claw, a spring torpedo launcher (Fig. 13, Appendix C), and a pendulum-release dropper—cover the manipulation tasks.

2.4 Systems Engineering & Cadence

The program runs on a quarter-gated cadence mirroring the testing tiers of Section 3. **Autumn:** Preliminary Design Review locks scope and interfaces per subteam. **Winter:** subteam-owned execution against those interfaces. **Spring:** full-stack integration and competition prep. Within each quarter, weekly subteam-lead syncs surface cross-team dependencies, full-team meetings stage integration, and dedicated subteam workdays focus development. We are currently in prequalification prep with all three subteams converging on Gen 3.

3 Testing Strategy

Every capability we score must be demonstrated repeatedly before competition. We run a tiered regimen with explicit entry/exit criteria per tier: lower tiers isolate faults cheaply on the bench, higher tiers confirm integrated behavior in water. The full plan—schedule, environments, risk register, and logged results—lives in Appendix D; this section summarizes methodology and how it gates the task sequence of Section 1.

Tier	Validates	Environment
Component	PCB load, firmware, sensor I/O	Lab bench
Subsystem	power swap, comms, actuators	Lab bench
Integrated	full-stack power, noise, thermal	Assembled, dry
In-water	sensors, control loops, sealing	AOERC pool
Mission	autonomy task sequence	AOERC / MBARI

Table 2: Tiered test plan; each tier gates the next.

3.1 Component, Sensor, and Control-System Testing

Hardware and software develop in parallel so testing decouples from fabrication. Every custom PCB—power distribution, battery management, microcontroller breakout—is bench load-tested and firmware-validated before hull integration, then re-tested under operational load for noise isolation and thermal stability. Subsystem checks then prove the engineered reliability features: dual-battery hot-swap without compute reboot, 1 Gbps end-to-end through the WetLink path, and leak/pressure telemetry tripping before any seal failure.

Sensors are validated against known references before entering the estimator: depth against a measured column, IMU bias on a level fixture, DVL over a known traverse, camera intrinsics via calibration target, and per-channel hydrophone response. Control-system testing then proceeds bottom-up: thrust allocation against per-thruster ACS713 telemetry, axis-wise PID tuning from pool step responses, and EKF validation via mid-run marker corrections that bound dead-reckoning drift. This ordering gates the mission sequence—Bins precedes Torpedoes only once position/altitude estimates are stable.

3.2 In-Water Validation and Operational Lifecycle

Physical validation runs weekly at the Stanford Avery Aquatic Center (AOERC), augmented by weekend field ops. Waterproofing and buoyancy are first vetted in campus fountains. The Gen 2 platform has been deployed to 34 ft at MBARI, confirming the pressure-hull and bulkhead-sealing methodology carried into Gen 3, which is now being qualified in pool and pressure-tank fixtures before its own deep dive. Every deployment is bracketed by written objectives and full telemetry logging, so pool time—our scarcest resource—is never spent rediscovering known state.

3.3 Simulation and Perception Testing

Pool access is limited, so iteration moves off-deck where possible. JTA-GAN [3] applies physics-based underwater style transfer to a dry marker set—synthesizing turbidity, scattering, color attenuation—so YOLOv8 [1] can be retrained between water sessions, not only during them. Planned trajectories are simulated in parallel: the path generator's animated overlays (Appendix F) verify continuity before a route is sent to the vehicle. In-water time is then reserved for failure modes only a real deployment reveals.

Acknowledgments

We thank Professor Oussama Khatib for academic guidance and curriculum alignment; Stanford Student Robotics for organizational framework, leadership, financial administration, and lab space; Jeff Stribling and the Lab64 Course Assistants for test equipment, benchtop troubleshooting, and a home lab; the Stanford CCRMA for high-fidelity acoustic testing access; and the Avery Aquatic Center staff for accommodating our weekly validation cadence. We are grateful to Teledyne Marine for sensor subsidizations, and to Reveille VC for the financial support that let us scale prototyping across multiple vehicles from simulation to deep-water deployment.

References

[1] G. Jocher, A. Chaurasia, and J. Qiu, "Ultralytics yolov8," 2023.

[2] M. Tan, R. Pang, and Q. V. Le, "Efficientdet: Scalable and efficient object detection," 2020.

[3] Y.-H. Chen, L.-Y. Yu, and Y.-Y. Chen, "Jtagan: A physics-informed framework for realistic underwater image generation and improved object detection," *Mathematics*, vol. 14, no. 4, 2026.

[4] T. I. Fossen, *Handbook of Marine Craft Hydrodynamics and Motion Control*. Chichester, UK: John Wiley & Sons, 2011.

[5] A. Wright, E.-P. Damskägg, L. Juvela, and V. Välimäki, "Real-time guitar amplifier emulation with deep learning," *Applied Sciences*, vol. 10, no. 3, 2020.

[6] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," 2022.

[7] G. Maiello, "Metaprobe-2.1 github." `https://github.com/GiuliaMaiello/Metaprobe-2.1`, 2024. Accessed: 2025-06-30.

4 Appendix

A Software Figures

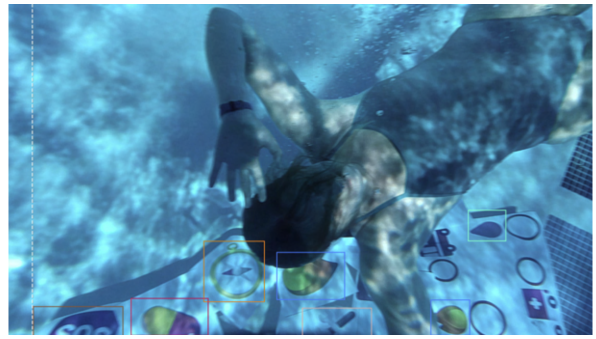

Figure 4: Physical collection of the training target dataset.

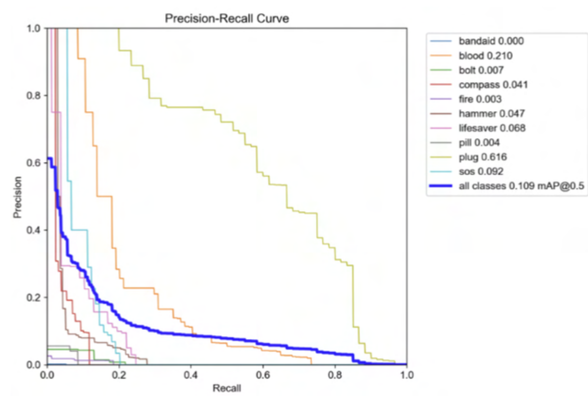

Figure 5: Precision–Recall: dry-trained baseline on wet targets.

B Electrical Figures

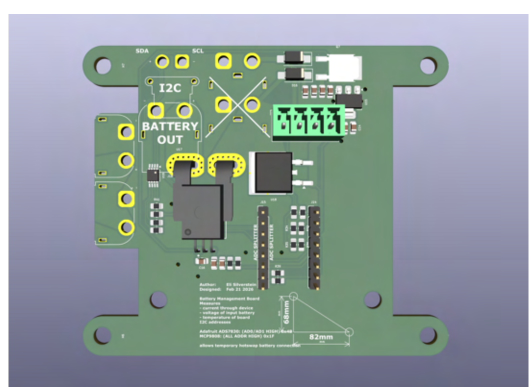

Figure 6: 3D rendering of the battery management board.

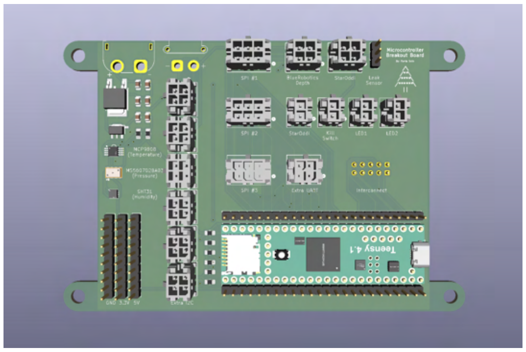

Figure 7: 3D rendering of the microcontroller breakout board.

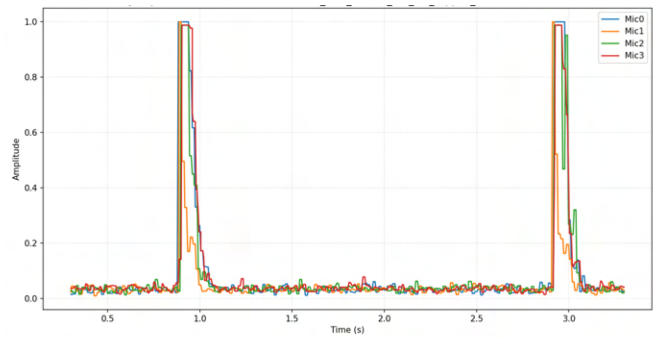

Figure 10: Captured pinger data from last year's semifinals run.

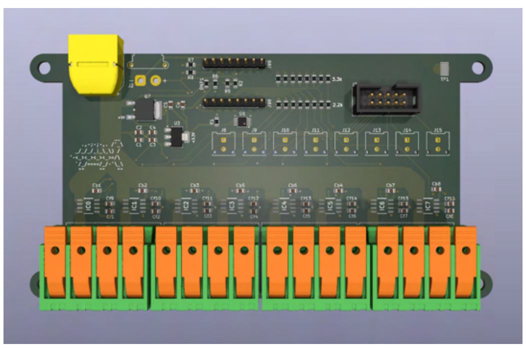

Figure 8: 3D rendering of the power distribution board.

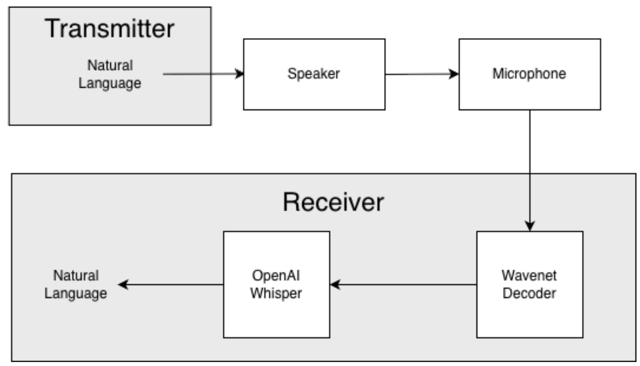

Figure 11: Intervehicle communication flowchart.

C Mechanical Figures

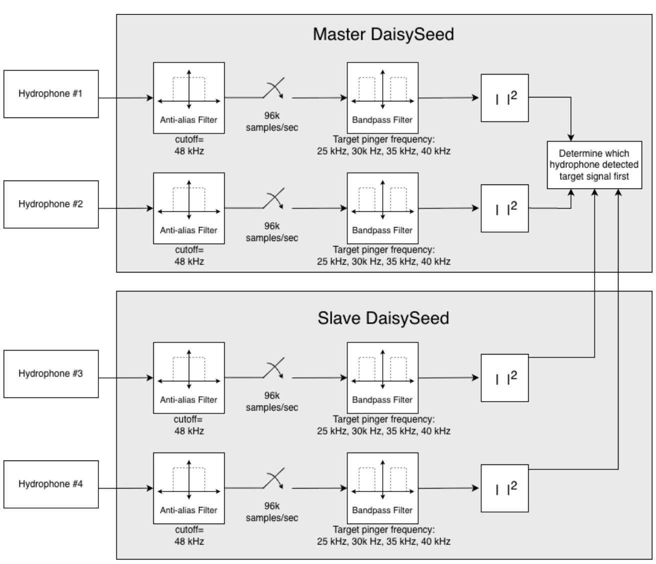

Figure 9: TDOA algorithm flowchart.

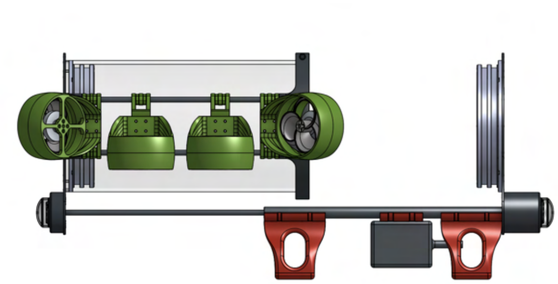

Figure 12: Gen 3 internal sliding chassis array.

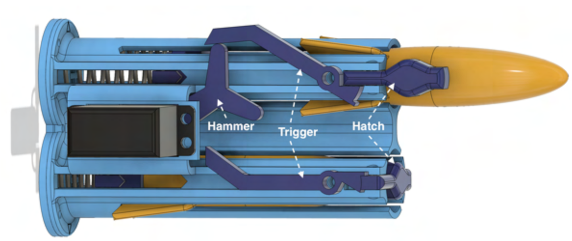

Figure 13: Labeled spring torpedo launcher.

D Test Plan & Results

This appendix expands the tiered methodology of Section 3 with concrete entry/exit criteria, the cumulative results to date, and the active risk register that drives our remaining pool sessions.

Entry / Exit Criteria

Component tier. Entry: PCB fabricated and visually inspected. Exit: each rail holds nominal voltage under rated load for $\geq$ 30 min with $<$ 2 °C rise above ambient; all firmware unit tests pass on the Teensy 4.1.

Subsystem tier. Entry: component tier passed. Exit: dual-battery swap demonstrated three consecutive times without compute reboot; 1 Gbps sustained over the WetLink path for $\geq$ 10 min; every PWM servo (claw, dropper, torpedo) cycles 100 times submerged with no leak alarm.

Integrated tier. Entry: all custom PCBs mounted in the sliding bay. Exit: full stack draws within $\pm 5\%$ of predicted current per axis under representative mission load; no thermal hotspot above 60 °C after a 20-minute dry soak.

In-water tier. Entry: integrated tier passed plus pressure/leak telemetry green. Exit: depth-hold within ± 10 cm for 60 s; heading-hold within $\pm 3°$; EKF residuals bounded under marker corrections.

Mission tier. Entry: in-water tier passed for the involved actuators and sensors. Exit: each task completed end-to-end in three out of five runs at AOERC under the same configuration that will be deployed at competition.

Logged Results

MBARI 34 ft dive (Gen 2). Pressure hull and bulkhead seals held under sustained hydrostatic load with no leak telemetry triggered; the same sealing approach is carried into Gen 3.

Dual-battery hot-swap. Demonstrated on the bench: main-pack removal does not interrupt Jetson Orin uptime.

Gigabit link. iperf3 reports 1 Gbps sustained end-to-end through the WetLink penetrator path.

Pinger localization. First-arrival detection demonstrated on data captured at last year's RoboSub semifinals run (Fig. 10), with per-board synchronization holding within one FFT block.

Active Risk Register

R1 – pool-access cadence. Mitigation: Gen 2 testbed keeps software in the water during Gen 3 fabrication.

R2 – perception domain gap. Mitigation: JTA-GAN synthetic pipeline plus shaded validation set.

R3 – acoustic SNR under competition pool acoustics. Mitigation: CCRMA-guided hydrophone build; first-arrival algorithm tolerant of multipath.

R4 – single-source sensor failure. Mitigation: EKF residual monitoring with marker-based position correction.

E Electrical Component Selection

This section records the rationale behind the vehicle-level component choices that complement the custom PCBs described in the body.

Battery. A Lumenier 13000 mAh 4S 20C LiPo. The 4S configuration spans 12.8–16.8 V, within limits for both thrusters and compute, which removes the need for high-current buck/boost conversion. The 20C rating supplies burst current for all eight thrusters simultaneously, and 13 Ah gives roughly 40 min of runtime under a continuous 20 A load.

Thrusters & ESCs. Eight Blue Robotics T200 thrusters, each driven by its own BlueRobotics ESC. Because T200 thrust decays with input voltage, each PWM command is compensated against the real-time battery voltage for consistent thrust.

Low-Level Control. A Teensy 4.1 handles time-critical, hardware-facing work—ESC PWM generation, non-critical sensor polling, bulkhead interfacing, and the bay display—and talks to the high-level system over a high-speed serial link, cleanly separating real-time control from mission logic.

High-Level Control. A Jetson Orin AGX runs perception, navigation, and planning with GPU-accelerated vision and sensor fusion, freed from low-level constraints by the Teensy split.

DVL. Teledyne's Wayfinder DVL met our weight, cost, and performance constraints; its 60 m bottom-track range covers the anticipated depth profiles.

IMU. A Movella Xsens MTi-610, fused with DVL velocity and Bar02 depth to limit drift.

Vision. Two Luxonis Oak-D S2 stereo-depth cameras supply RGB frames to the YOLOv8 detection pipeline described in Section 2.1.1.

Connectors. XT60/XT30 for power, Molex Micro-Fit for sensors, and ribbon cable for board-to-board. Harness reliability proved critical—bad connectors were a recurring failure source.

F Control & Planning Formulation

PID control. The control module reduces 6-DOF state error $\boldsymbol{e}(t)$ (position and orientation) to a control wrench:

$$\boldsymbol{u}(t) = K_P\,\boldsymbol{e}(t) + K_I \int_0^t \boldsymbol{e}(\tau)\,d\tau + K_D\,\frac{d}{dt}\boldsymbol{e}(t).$$

Orientation is represented as a unit quaternion; the attitude error is the axis–angle residual between current $\mathbf{q}_{\text{cur}}$ and desired $\mathbf{q}_{\text{ref}}$,

$$\boldsymbol{e}_\theta = 2\,\mathrm{vec}\big(\mathbf{q}_{\text{ref}} \cdot \mathbf{q}_{\text{cur}}^{-1}\big)\,,$$

where $\mathrm{vec}(\cdot)$ is the vector part of the quaternion. This is fed into the PID loop alongside positional error.

Path generation from waypoints. To turn sparse waypoints into dynamically feasible trajectories, we fit a C^5-continuous spline through both position and orientation so downstream controllers get smooth velocity and acceleration references. For position we fit a quintic B-spline through $\{\boldsymbol{p}_i \in \mathbb{R}^3\}_{i=1}^N$ with zero first/second derivatives at the endpoints ($\boldsymbol{p}'(0) = \boldsymbol{p}''(0) = \boldsymbol{p}'(T) = \boldsymbol{p}''(T) = \mathbf{0}$), guaranteeing continuity through jerk. Orientation is interpolated on $\mathrm{SO}(3)$ via the logarithmic map $\boldsymbol{\theta}_i = \log(\mathbf{q}_i) \in \mathbb{R}^3$, fitting independent quintic splines per component and recovering $\mathbf{q}(t) = \exp(\boldsymbol{\theta}(t))$, with $\boldsymbol{\omega}(t) = \dot{\boldsymbol{\theta}}(t)$ and $\boldsymbol{\alpha}(t) = \ddot{\boldsymbol{\theta}}(t)$. The method supports waypoint closure and animated overlays for visual confirmation before deployment. The full open-source stack is available at our repository.

G Mechanical Design Notes & Lessons Learned

Design philosophies.

Iterability. Standardized mounting-hole patterns let subsystems move and the frame carry over between vehicles. The shared pattern already paid off when the DVL shifted to make room for the camera tube and arm, and it will again when next year's sub builds on the same frame.

Cross-team, smart design. A MECH is **M**echanical **E**lectrical **C**omputational **H**appiness—nothing is designed inside one subteam. Concretely: add wire cutouts and standoffs to mounts before they reach electrical, not after.

Good CAD. Fully constrain and dimension sketches and design in millimeters; with many parts, small CAD errors compound into assembly headaches.

Heat-set inserts. Prefer heat-set inserts (plastic) or tapped threads (metal) over nuts; see gm0's custom-manufacturing guide for sizing.

H Research Directions

Instant-NGP / NeRFs. Building on Neural Radiance Fields, we fuse posed 2D images into a continuous 3D map. A NeRF represents a scene as a 5D function

$$F_\Theta : (\mathbf{x}, \mathbf{d}) \to (\mathbf{c}, \sigma),$$

mapping a 3D location $\mathbf{x}$ and viewing direction $\mathbf{d}$ to emitted color $\mathbf{c}$ and volume density σ (network parameters Θ).

eDNA collection. Metaprobes—small ($\sim$4 in) hollow spheres housing medical-gauze filter rolls—capture environmental DNA in aquatic settings, printed from the open-source design of Maiello et al. [7].

Environmental sensing. We log temperature and salinity gradients during deployments to contextualize sensor and vehicle behavior.

Underwater ocean-floor mapping. We are scoping a follow-on capability that uses the Blue Robotics Ping360 mechanically-scanning imaging sonar to build 2D occupancy / bathymetric maps of the surrounding seabed during operations. The Ping360 sweeps a narrow acoustic beam over a full 360° horizontal field at configurable range ($\sim$50 m) and returns intensity-vs-range echoes per bearing; fusing successive sweeps with the EKF pose estimate yields a vehicle-relative ground-plane map suitable for obstacle avoidance and longer-horizon exploration tasks beyond the RoboSub course. The same data pipeline feeds naturally into our existing NeRF / 3D-mapping research thread, supplying acoustic depth where optical visibility falls off.