

RoboSub 2026 Technical Design Report

University of Alberta – Autonomous Robotic Vehicle Project

Daniel Faltaous (Project Co-Lead), Aryan Aryal (Project Co-Lead),
Jonathan Bonham (Software Co-Lead), Taha Esmahi (Electrical Co-Lead),
Sean Theriault (Electrical Co-Lead), Scott Lambert (Mechanical Co-Lead),
Adam Marshall (Mechanical Co-Lead)

Abstract - ARVP presents Kenai for RoboSub 2026, with a season focused entirely on repeatability: with only three scored runs, a single failure mode determines placement. A weighted task selection framework produced 80% task and 95% course navigation targets, yielding a 72% cumulative probability of a clean full-points run across two finals attempts. Three hardware upgrades addressed root causes of 2025 underperformance: a redesigned BMS eliminated intermittent power failures traced to an open ground loop; the Motus IMU resolved yaw bias instability from $8^\circ/\text{hr}$ to $0.2^\circ/\text{hr}$; and a 650mm 2-DOF claw resolved DVL blanking, camera field-of-view, and thruster turbulence issues that kept Octagon success at 15%. A feature-based pose estimation pipeline using XFeat and LighterGlue replaced bounding box detection, enabling full 3D alignment for precision claw pickup. Testing spanned three progressive levels across 104.5 pool hours and 28 sessions, with pool time protected through a three-month hardware freeze and structured test reports produced for each major upgrade.

I. Competition Strategy

Competition Objective: Repeatability

With Kenai, hardware reliability is no longer the primary performance bottleneck. The cylindrical hull, standardized electrical interfaces, and consolidated boards resolved the failure points that plagued previous vehicles, shifting focus to the next limiting factor: repeatability. In 2025, task capability in testing did not translate to competition execution; development was still ongoing the week before finals. The core insight driving this season is that the competition rewards repeatability over capability: with only two semifinal runs and one final, a single failure mode determines placement.

This produced two distinct objectives: task completion repeatability, where individual task success rates set the score ceiling, and course navigation repeatability, where the probability of reaching each task determines whether that ceiling is hit. A high task success rate is worthless if the vehicle cannot get there.

Task Selection

Task selection was driven by six criteria: pre-existing capability, development cost, development risk, pool testing budget required, cascade risk, and point value. Each reflects a dimension of the same question: given ARVP's pool testing budget, engineering capacity, and competition objective, what is each task's expected contribution to a repeatable score? The full weighted decision matrix is in Appendix A and the system requirements derived from task targets are in Appendix B.

Gate, Dropper, and Torpedo exceeded 95% success rates in 2025 with no changes that would meaningfully impact them. Pool time here is maintenance, not development. Slalom's 30% competition rate in 2025 traced to two identifiable failure modes: yaw drift causing up to 45° heading error, and ambiguous pole detection with successive red poles. Both are directly addressed: the Motus tactical IMU improves yaw bias instability from $8^\circ/\text{hr}$ to $0.2^\circ/\text{hr}$ with full EMI certification [1], and a pole thickness comparison resolves detection ambiguity. Physical entanglement remains the one unmitigated cascade risk, managed via a hard abort timeout. Claw's 15% success rate in 2025 traced to DVL blanking at low operating height and insufficient testing time. Both are addressable via mechanical upgrades and software tuning. With core tasks already exceeding 95%, pool capacity exists to develop Claw without cannibalizing reliable scorers. Pool time will only shift toward Claw once core tasks are revalidated. Intersub communication remains 1-way acoustic; increasing complexity introduces hardware development cost, pool testing burden, and a mid-run failure mode that risks the entire run. Return Home is forfeited: no existing capability, high pool cost, and low point value make it a poor allocation under any scenario.

Course navigation is treated as a distinct performance target. In 2025 it succeeded in 65% of runs, driven by the same yaw drift affecting Slalom. With the Motus IMU resolving that root cause, 95% is targeted for 2026 and underlies the entire run model.

From this exercise, ARVP set an 80% task success rate and 95% course navigation target. Multiplying success rates across all tasks yields a

47% probability of a clean full-points run per attempt ($0.95 \times 0.80 \times 0.95 \times 0.95 \times 0.80 \times 0.90 \times 0.95$). With two attempts in the finals window, the cumulative probability of at least one clean run reaches 72%. Tasks below 80% are not considered reliable scorers regardless of point value. See Figure 1.




Figure 1. 2025 Competition Success Rates and 2026 Target Success Rates by Task

Concept of Operations

ARVP deploys two vehicles: Kenai, the primary competition AUV, and Koda, a secondary AUV whose sole role is triggering Kenai's run start via one-way acoustic communication, satisfying the intersub task at no additional development cost. Once started, Kenai executes tasks in a fixed sequence: Gate, Slalom, Bins, Torpedo, and Octagon. The full step-by-step run procedure is detailed in Appendix C.

Strategic Approach

Most requirements derived from task and success rate targets did not demand more hardware development; they demanded more pool testing time. Every capability added is a liability as much as an asset, introducing integration issues, testing burden, and failure modes that compound under time pressure. With hardware downtime already under 10%, the bottleneck was repetition in the pool, making structured, high-volume testing well in advance of competition the single highest-leverage action available. Task success rate targets are met through iteration and tuning in the water.

Pool testing captures what component-level and simulation testing cannot: real underwater dynamics, sensor behavior, and system interactions that are the foundation of repeatability. Protecting that time meant managing scope deliberately: doing less, earlier and more thoroughly, so the vehicle was stable well before competition. Proven systems were left untouched; the hull and frame have had zero

ingress since first deployment, torpedo launchers and droppers survived thousands of cycles in last year's testing, and the control, sonar, and acoustic modem boards have had no failures. The vehicle design was frozen three months before competition, giving the software team a static hardware platform for the final stretch. Koda development was kept secondary to Kenai at all times, as any additional capability development draws pool time from the vehicle that actually scores.

Hardware upgrades were each independently justified against repeatability objectives, scheduled during gaps in the pool calendar, and verified at component and integration level before installation. The Motus IMU addresses the $8^\circ/\text{hr}$ yaw bias that limited both Slalom and course navigation in 2025, reducing this to $0.2^\circ/\text{hr}$ with full EMI certification [1]; course navigation gates access to every downstream task, making this the highest-priority sensor upgrade. The BMS redesign was accepted because power reliability is a prerequisite for every task and a single failure ends the run regardless of software or mechanical readiness. The claw architecture was accepted despite added mechanical complexity because Octagon carries the highest remaining point ceiling, prior testing confirmed achievability, and core scorers have sufficient pool capacity to absorb the development without cannibalization. Pool testing ramped from bi-weekly sessions in the fall to weekly in the spring and extended sessions approaching competition.

II. Design Strategy

Mechanical Design

System requirements derived from 2026 task and success rate targets drove a clear mechanical mandate: improve task success rates where specific, attributable failure modes existed, and leave proven systems untouched. For the mechanical team, this meant a focused optimization effort rather than a redesign. The highest-impact improvement was the 2-DOF claw, which represented the largest gap between current performance (15%) and the 80% Octagon task target; additional improvements were made across the platform including torpedo redesign

and support infrastructure upgrades. The hull and propulsion system were proven in 2025 and deliberately left untouched; zero ingress since first deployment and no thruster failures confirmed these systems were not limiting factors in competition performance (see Appendix D for vehicle overview).

Project Highlight - Two Degree of Freedom

Claw: The 2025 claw extended 140mm below the vehicle and produced three compounding, attributable failure modes: insufficient camera field of view at operating height, DVL degradation observed empirically during early-season baseline testing at elevations below 0.5m (see Appendix E, Nov 8 session) against a Nortek-specified blanking distance of 0.1m [2], and thruster turbulence displacing lightweight objects on the table. A software workaround that hovered higher then pounced worsened the turbulence problem, producing a 15% success rate.

The redesigned claw extends 650mm below the vehicle, resolving all three failure modes simultaneously: the vehicle operates at a height where the DVL functions correctly, the camera field of view covers the full object, and thruster turbulence dissipates before reaching the table. A spring system offloads gravity from the primary servo, sized via moment balance to contribute 50% of the required holding torque (see Appendix F for full derivation and spring selection). Conforming TPU end effectors maximize contact area across varying object geometries, with one stationary effector and one actuating, improving grip consistency over the previous symmetric-closing design. The claw represents the season's most significant systems integration effort: its mechanical redesign set a new DVL operating envelope, which tightened positioning accuracy requirements beyond what the 2D vision pipeline could satisfy, directly motivating the 3D pose estimation pipeline described in Software Design. See Figure 3.




Figure 2. Claw Assembly

Electrical Design

The repeatability mandate that drove mechanical optimization applied equally to the electrical team. In 2025, intermittent failures including camera disconnects, inconsistent startup behavior, and servo control loss were difficult to reproduce and harder to attribute, consistently obscuring whether faults were electrical or software in origin and consuming pool testing time that could not be recovered. Rather than redesigning proven systems, the 2026 electrical focus was eliminating that ambiguity: consolidating the power architecture to reduce distributed failure points, and introducing improved monitoring and protection against transient events to increase confidence in system behavior under extended operating conditions. Full electrical architecture diagrams are provided in Appendix G.

Project Highlight - Battery Monitoring System

(BMS): The 2025 architecture distributed power across three independent voltage rails via a battery monitoring board, carrier board, and three modular converter stages. While functional, this created multiple points of failure and made intermittent faults difficult to isolate. The primary 2025 failure was ultimately traced to an open ground loop in the carrier board caused by a latent ESD failure, producing the camera disconnects and subsystem failures observed during competition.

The redesigned BMS consolidates power distribution, monitoring, and protection into a single centralized board, distributing a regulated 12V rail to all onboard electronics while ESCs receive battery voltage directly. Over and under-voltage/current protection is handled by the LTC4421 [3], with battery inputs and subsystem outputs monitored via the INA3221 [4], providing real-time visibility into system behavior. The primary architectural change eliminates the converter stages and reduces grounding

connection points, meaning subsystems can only fail at the BMS or their own ground rather than anywhere along a distributed rail. See Figure 4. for PCB layout.

A trade study comparing the distributed and consolidated architectures confirmed the efficiency gain (see Appendix H). At nominal load, the previous architecture drew 10.7A across three rails versus 4A on the consolidated 12V rail. Assuming 50m Ω total distribution resistance accounting for wire and contact resistance, this corresponds to 2.22W of resistive losses on the previous architecture versus 0.80W on the new rail, a 63.9% reduction.

Software Design

The 2025 IMU's 8 $^\circ$ /hr yaw bias cascaded through the entire software stack: localization drifted, task alignment failed, and course navigation succeeded in only 65% of runs. The 2026 software mandate was twofold: resolve the sensor dependencies that caused those failures, then use the pool testing capacity freed by stable hardware to harden mission logic through repetition rather than last-minute development. With the Motus IMU resolving the yaw drift at the hardware level, the software team focused its development effort on the highest-precision alignment requirement on the course: 3D pose estimation for claw pickup, where the 650mm arm extension demands positioning accuracy that the previous 2D pipeline could not provide. The full software architecture, including the relationships between state estimation, machine vision, mapping, motion planning, and actuator control, is detailed in Appendix I.

Project Highlight - Feature-Based 3D Pose Estimation from Monocular Vision: The previous vision pipeline detected objects with YOLO and fed axis-aligned bounding box corners into an Image-Based Visual Servoing (IBVS) controller, driving four corner points toward target pixel coordinates via an interaction matrix. This approach had three structural weaknesses that made it insufficient for claw precision requirements: bounding box corners are unstable under object rotation or confidence fluctuation, producing up to ± 15 px jitter in the control signal; there was no outlier rejection, meaning a single misplaced corner corrupted the entire interaction

matrix; and only 2D information was available, with depth assumed constant and yaw misalignment uncorrectable.

The new pipeline replaces bounding box corners with learned visual features matched against reference images using XFeat for feature extraction and LighterGlue for descriptor matching, running on every frame at reduced resolution. A dedicated PnP (Perspective-n-Point) node estimates the object's full 3D pose relative to the camera. Image points are clustered with HDBSCAN to reject stray matches from background textures, and translation and yaw outputs are smoothed over a 50-frame sliding window. Backward compatibility with the existing IBVS controller is preserved by projecting four reference-image corner points into the camera frame via the estimated homography. An optional DepthAnythingV3 node provides an independent depth estimate. See Figure 5. for a feature matching demonstration.



Figure 3. Feature matching between the reference image (left) and the camera feed, with lines drawn between identical features; 43 matches are shown for object localization.

The result is a pose estimation pipeline that is stable, outlier-resistant, and provides full 3D alignment where the previous approach provided none, directly enabling the positioning precision the 650mm claw arm demands.

III. Testing Strategy

Mechanical Testing

Mechanical testing for the claw followed the same progressive structure applied across all subsystems: component validation before integration, integration before pool deployment. This sequencing was deliberate; every failure caught on the bench or in the tub is a pool session preserved for end-to-end repetition. For the claw specifically, component and integration testing were entirely feasible outside the pool, making it

one of the clearest opportunities to protect scarce pool time.

Project Highlight - Two Degree of Freedom Claw: Component testing covered the two highest-risk failure points: servo torque and end effector grip reliability. Servo torque was validated using a weighted moment arm simulating operational load at 1.5x factor of safety; all 20 trials passed. End effector performance was characterized using an underwater jig across 125 permutations of lateral offset (-10 to 10cm), angle (-45° to 45°), and height (0 to 5cm). A pass threshold of 100/125 was set, with the expectation that software could account for fringe failures; 110 passed, with all failures attributable to object behavior in the simulated environment rather than gripper failure (see Appendix J for Results).

Integration testing was conducted in the University of Alberta civil engineering lab water tank, which provided just enough clearance to deploy the arm fully underwater without consuming pool time. 100 arm deploy, grip, and retract cycles were performed with a 100% pass rate, confirming underwater actuation before pool deployment (see Appendix E, May 17 session).

End-to-end pool testing used vision-guided descent to the object table, with a pickup confirmed if the object was retained through surfacing. Of 16 attempts to date, 11 passed at 68%, with remaining pool sessions specifically allocated to closing the gap to the 80% target through software tuning (see Appendix E, June sessions). Testing also produced one design change: end effector shape and material were revised based on observed failure modes, demonstrating the value of structured progressive testing over direct pool deployment.

Electrical Testing

The BMS replacement touched every onboard system, making validation the highest-stakes testing effort of the season. A failure at any level could damage connected electronics permanently, and the previous failure mode was intermittent and load-dependent, meaning bench testing alone was insufficient. Testing was structured across three progressive levels, with each level designed to catch failures before they reached the next.

Project Highlight - Battery Monitoring System (BMS): Component-level testing validated voltage ripple, protection circuit behavior, and current monitoring on the bench before any integration. The I7C DC-DC buck converter measured 100mV peak-to-peak ripple, outperforming the 200mV datasheet specification [5]. However, component testing also caught two issues that would have been significantly harder to diagnose in the field: LTC4421 charge pump routing errors [3] and incorrect capacitor values required a second PCB revision, and undervoltage protection behaved incorrectly on a bench supply due to output regulation interference, operating correctly only when transitioned to vehicle batteries. This is precisely the failure mode that structured component testing exists to catch.

Integration testing validated the BMS under partial load with all onboard electronics, cameras, and NVIDIA Orin connected simultaneously. Ten power-cycling startup cycles across 15-minute runtimes confirmed 100% startup consistency and zero camera failures, directly targeting the failure mode observed in 2025 (see Appendix K).

End-to-end validation was performed across pool testing sessions under full operating load including servo and thruster current draw as seen in Figure 6. Over 24 startup cycles and approximately 72 hours of cumulative in-water operation produced zero dropouts, unexpected resets, or camera failures, achieving the 100% startup rate targeted by the competition strategy. Each issue caught at component level avoided a pool session lost to debugging, directly protecting the testing budget that underlies the entire competition strategy.




Figure 4. Fully integrated electronics enclosure before pool testing

Software Testing

The software testing approach spanned three levels: automated hardware checkout before every pool session, simulation-based integration

testing on land, and pool validation once the first two levels passed. Pool hours doubled relative to 2025, totalling 104.5 hours across 28 sessions (see Appendix E). Making productive use of that time is the central testing challenge. Each level was designed to catch failures before they consumed pool time: only work that genuinely requires water gets tested in the water. Full testing hours by environment, task performance outcomes, and vision system metrics are documented in Appendix L.

The `at_tester` automated checkout covers every electromechanical subsystem in a single command before each pool session. Battery health is verified by reading all five cell voltages from the CAN bus and flagging any below 15.2V, catching degraded cells before they cause instability under load. All eight thrusters receive a 300 RPM command over CAN bus with pass/fail logged per motor position based on RPM feedback within 0.2 seconds. Actuators are cycled through full range: claw open/close, dropper fire and reload, and torpedo arm and fire with an explicit safety warning before arming. A separate map test node publishes a dummy detection to the bottom camera topic at 2Hz, verifying the vision-to-mapping pipeline without requiring hardware in the loop. Any failure at this level is resolved before the vehicle enters the water.

Mission logic is validated in Gazebo prior to pool deployment, with the explicit exception of the vision pipeline. Testing computer vision in simulation, even with simulated noise, does not reflect real-world image uncertainty and produces no actionable data. Vision validation is handled separately through the land-based protocol described below.

Project Highlight - Feature-Based 3D Pose Estimation from Monocular Vision: The pose estimation pipeline was validated on land before pool deployment, with acceptance criteria defined as depth accuracy within ± 1 inch and yaw accuracy within $\pm 5^\circ$, measured against a physical object at known distances and orientations. Land testing was chosen deliberately: feature correspondence geometry is identical in air and water, while real-world image quality, lighting variation, and camera noise are not reproducible

in simulation. This kept vision validation entirely out of the pool testing budget.

Land testing consistently met both acceptance criteria. Pool validation across 10 trials confirmed that water turbidity and underwater lighting did not degrade feature matching, with 100% perpendicular alignment achieved and depth estimation accurate within 1cm, matching land results and confirming the pipeline is pool-ready.

IV. Acknowledgements

We thank our advisors, alumni, industry supporters, and sponsors; see Appendix M for detailed acknowledgements. ARVP's success is not possible without you.

V. References

- [1] Advanced Navigation GNSS Compass Certification, "GNSS Compass Reference Manual – Certification," Advanced Navigation, accessed May 25, 2026.
- [2] Nortek Group, "Nucleus1000 Product PDF," Nortek Group. [Online]. Available: <https://www.nortekgroup.com/products/nucleus1000/pdf>. [Accessed: May 25, 2026].
- [3] Analog Devices, "LTC4421: High Power Prioritized PowerPath Controller," Analog Devices Datasheet. [Online]. Available: <https://www.analog.com/media/en/technical-documentation/data-sheets/LTC4421.pdf>. [Accessed: May 25, 2026].
- [4] Texas Instruments, "INA3221-Q1 Automotive, Three-Channel, High- or Low-Side Measurement, Shunt and Bus Voltage Monitor With I2C- and SMBUS-Compatible Interface," Texas Instruments Datasheet. [Online]. Available: <https://www.ti.com/lit/ds/symlink/ina3221-q1.pdf>. [Accessed: May 25, 2026].
- [5] TDK-Lambda, "i7C Series," TDK Product Center. [Online]. Available: https://product.tdk.com/system/files/dam/doc/product/power/switching-power/dc-dc-converter/catalog/i7c_e.pdf. [Accessed: May 25, 2026].

VI. Appendix

Appendix A: Task Selection Matrix

Criteria	Existing Capability	Development Cost		Development Risk	Pool Testing Budget Required	Cascade Risk	Point Value	Sum Total Score
Description	Coming out of RoboSub 2025, could we do this task at all? If so, what were our success rates?	How much time would it take to build consistent success on this task? How much would it cost? Do we need to develop new hardware?		How confident are we that we can develop this capability? Is there a high chance we sink time into it and end up with nothing?	How many hours in the pool do we think this will take to reach our target success rate?	Does failure on this task lead to a cascade of failures? Will a slip-up here risk the rest of the run?	How many points are on the table here?	
Weight (/5)	5	4		3	4	3	4	
	Score (/5)	5	5	5	5	3	2	
Gate	Weighted Score (/25)	25	20	15	20	9	8	97
	Score (/5)	3	4	3	3	1	3	
Slalom	Weighted Score (/25)	15	16	9	12	3	12	67
	Score (/5)	5	5	5	4	4	3	
Dropper	Weighted Score (/25)	25	20	15	16	12	12	100
	Score (/5)	5	5	5	3	4	4	
Torpedo	Weighted Score (/25)	25	20	15	12	12	16	100
	Score (/5)	2	3	3	1	4	5	
Octagon/Claw	Weighted Score (/25)	10	12	9	4	12	20	67
	Score (/5)	0	4	3	2	4	1	
Return Home	Weighted Score (/25)	0	16	9	8	12	4	49
	Score (/5)	3	2	4	2	2	1	
Intersub Communication	Weighted Score (/25)	15	8	12	8	6	4	53

Appendix B: System Requirements Matrix

Description	Category
Robot shall be able to detect leak and shut down power when detected	Environmental
Shall operate at temperatures from 10-35°C external	Environmental
Shall maintain vision capabilities in sunny conditions	Environmental
All external parts shall be pressure rated for at least 20m	Environmental
All internal parts shall be rated for 50kPa pressure	Environmental
Kenai should be able to know its yaw and position at any position in the pool based on starting position	Functional
Kenai must be able to locate the pinger in the pool	Functional
Kenai must be able to traverse in all 6 axes of motion	Functional
All external parts shall be replaceable	Functional
Kenai shall one way communicate with Koda using sound (receive)	Functional
Kenai must be able to grasp and move objects	Functional
Kenai must be able to fire torpedoes	Functional
Kenai must have onboard computer error logging (shutdown, journald, dmesg)	Functional
Kenai must be able to surface at the octagon	Functional
Kenai must be capable of attempting all selected RoboSub tasks	Functional
Kenai must be battery powered	Functional
Kenai must use metric fasteners when possible	Functional
All internal fasteners shall be non-conductive when possible	Functional
Must be able to perform a barrel roll	Functional
Must record temperature and pressure internally	Functional
Kenai shall have handles built on to put it in the water	Human
Kenai should be able to be carried by two people with handles from MEC to ETLC	Human
Shall have an external indication of temperature and pressure	Human
All connectors should be keyed and labelled	Human
External components shall be serviceable (fully take apart and put back together) in under 15 minutes	Human

Kenai shall have an obviously immediately accessible kill switch for thrusters	Human
Battery pods shall be able to be swapped in under 5 minutes by an experienced member	Human
Internal electronics (aside from onboard computer) shall be able to be fully swapped in under 15 minutes	Human
Sensitive sensors and parts (cameras, DVL) shall fit within bounds of thruster guards	Human
External cables on Kenai shall not sag or hang enough to get caught or ripped	Human
All connections should be clearly labelled	Interface
All mechanical subsystems should not need software to reload	Interface
There shall be an external power port	Interface
There shall be external battery ports	Interface
Kenai will have a script that tests every essential function at once	Interface
Kenai should be able to complete the claw task within 5 minutes of starting it	Performance
Kenai should be accurate enough to shoot through the target at least 95% of the time	Performance
The dropper task should be accurate 99% of the time	Performance
Kenai shall operate 1 hour on a full battery charge	Performance
Kenai should be able to receive acoustic frequencies between 25 and 50 kHz	Performance
Kenai shall be able to complete gate 99% of the time	Performance
Kenai shall be able to complete the slalom task 80% of the time	Performance
Kenai should not go faster than 0.5 m/s so it can be easily stopped	Performance
Kenai shall be able to hold position $\pm 3\text{cm}$ in any axis and $\pm 1^\circ$ in yaw	Performance
Kenai shall not accumulate yaw drift greater than $\pm 2^\circ$ during movement	Performance
Kenai shall determine the directionality of the pinger with $\pm 10^\circ$ accuracy	Performance
Kenai shall be at least 0.5% positively buoyant	Performance
Kenai shall have carrying points for users	Safety
Kenai shall have no sharp corners or edges	Safety
Kenai's thrusters shall be caged, shall not affect performance, and shall not be able to hurt anyone	Safety

Appendix C: Detailed Concept of Operations - Full Run Procedure

The following describes Kenai's complete competition run sequence, including role assignment logic, task-specific behaviors, and abort conditions.

Pre-Run Setup

Both vehicles are tethered at the surface. Koda's role is confirmed as acoustic trigger only; no additional mission logic is loaded. Kenai's coin-flip role assignment (Survey/Repair or Search and Rescue) is pre-loaded and determines gate side, bin targets, and Octagon surfacing image selection throughout the run.

Step 1: Intersub Communication

Both tethers are unplugged. Koda transmits a start signal to Kenai via its acoustic modem. Successful receipt is confirmed by a light flashing on Kenai underwater. Koda remains stationary for the remainder of the run.

Step 2: Gate

Kenai takes a coin flip to determine starting orientation and a second coin flip to pre-assign role. Kenai passes through the corresponding gate side while performing two barrel rolls for style points.

Step 3: Slalom

Kenai aligns with the centre pole and passes on the same side as it passed through the gate, repeating this two more times while maintaining the depth of the slalom poles. A hard abort timeout is active throughout to mitigate physical entanglement cascade risk. If the timeout triggers, Kenai exits the slalom area and proceeds directly to Bins.

Step 4: Bins

Kenai drops both markers into the two bins corresponding to its pre-assigned role. The magnetic light interaction is forfeited this season: the mechanism was undefined within one month of competition, making it an unreliable point source under the team's complexity-reliability framework. Kenai proceeds to Torpedo immediately after both drops are confirmed.

Step 5: Torpedo

From 0.46m standoff distance, Kenai fires torpedoes through the large opening first, then the small opening, collecting the correct sequence bonus and the farther distance bonus. Perpendicular alignment is achieved via the XFeat feature-matching pipeline prior to firing. If alignment cannot be confirmed within the timeout window, Kenai fires at best available alignment rather than aborting.

Step 6: Octagon

Kenai surfaces immediately upon arrival at the Octagon to collect surfacing points. Kenai then attempts to pick up all objects, confirming each pickup via current spike detection matched to servo angle. After each pickup, Kenai surfaces facing the correct image based on the number of items currently in the basket:


- 1 item: Survey and Repair navigation icon or Search and Rescue life ring icon
- 2+ items: repair tools icon or distress icon


Kenai rotates once per item collected to collect the yaw bonus. Object placement continues until all objects are placed or the run times out.

Abort and Contingency Logic

Each task has an independent timeout. On timeout expiry, Kenai logs the failure, skips to the next task in sequence, and attempts course navigation to the following task start position. Additionally, if a vision detection is missed, there is specialized recovery behavior at every task that dictates how the AUV should physically respond. No individual task failure can prevent downstream task attempts except a power failure, which is mitigated by the BMS redesign described in Section II.

Appendix D: Vehicle Overview Diagrams





- 01 Blue Robotics aluminum tube**
Off-the-shelf pressure vessel.
- 02 End caps (*2)**
Double O-ring sealed.
- 03 3D-printed electronics tray**
Mounted to rear end cap; removable as a unit.
- 04 Sheet-metal frame**
Attaches directly to both sides of the enclosure.
- 05 Thrusters**
Mounted to the side plates.
- 06 Hydroshell**
Trim and balance. Replaces Buoyancy Pods
- 07 Camera**
Forward-facing.
- 08 Acoustic modem**
Transmit-only; sends signals to Kenai.
- 09 External battery Pod**
2x 14.4 V packs.

Appendix E: Pool Testing Budget and Time Allocation

Date	Hours	Time Allocation
Early season — onboarding & baseline		
Sept 13	2.5	New exec and software member training
Sept 27	2.5	New exec and software member training
Oct 11	2.5	New exec and software member training
Oct 25	2.5	New exec and software member training
IMU failure investigation		
Nov 8	2.5	PID tuning and RoboSub 2025 IMU failure troubleshooting — baseline torpedo runs (10×) with DVL+IMU; first attempt with MicroStrain IMU alone
November 22	2.5	Recreating RoboSub 2025 IMU failure — continued torpedo runs to isolate yaw drift root cause; confirming MicroStrain 8°/hr bias instability under load
Localization refactor & Motus integration prep		
Dec 20	2.5	PID model validation and localization refactor verification — first pool session after merging the localization refactor and Advanced Navigation (Motus) package; baseline PID re-tuning following thruster and electrical swaps; initial Koda gate runs filled
Jan 10	2.5	MPC/alternative PID controller testing and claw software bringup — validating Arthur's C++ PID against existing Python controller; first in-water test of claw depth estimation node (Part 1); Koda IMU orientation and localization baseline filled

Motus tactical IMU validation

Feb 7	2.5	Motus tactical IMU testing — initial calibration and yaw bias characterization; comparing course navigation success rate against MicroStrain baseline
-------	-----	-------------------------------------------------------------------------------------------------------------------------------------------------------

Feb 19	3	Motus tactical IMU testing — extended validation under full vehicle load; confirming 0.2°/hr bias stability and EMI certification compliance
--------	---	----------------------------------------------------------------------------------------------------------------------------------------------

Core task revalidation

Feb 21	2.5	Gate and style points — revalidating gate pass with Motus IMU; barrel roll timing and coin-flip orientation logic
--------	-----	-------------------------------------------------------------------------------------------------------------------

Feb 28	2.75	Gate and style points — continued gate repeatability runs targeting 99% threshold; refining vision-relative waypoint approach
--------	------	-------------------------------------------------------------------------------------------------------------------------------

Mar 21	3.5	Torpedo mission — revalidating perpendicular banner alignment with new XFeat feature-matching pipeline; testing 0.46m firing distance for large and small opening sequence
--------	-----	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Mar 22	2.25	Dropper mission — revalidating bin drop accuracy; confirming 99% target success rate post-Motus IMU integration
--------	------	-----------------------------------------------------------------------------------------------------------------

Apr 4	3.5	Slalom mission — first full slalom runs with Motus IMU; testing pole detection ambiguity fix (thickness comparison); validating yaw-stable passes through all three sets
-------	-----	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Apr 18	2.5	Slalom mission — continued repeatability tuning; hardening abort timeout for physical entanglement cascade risk
--------	-----	-----------------------------------------------------------------------------------------------------------------

Community showcase preparation

Apr 26	8	Gate, slalom, dropper, torpedo, and octagon surface stringing for 2026 community showcase
--------	---	-------------------------------------------------------------------------------------------

Apr 27	4	Gate, slalom, dropper, torpedo, and octagon surface stringing — 2026 community showcase
--------	---	-----------------------------------------------------------------------------------------

Apr 30	5	Gate, slalom, dropper, torpedo, and octagon surface stringing — 2026 community showcase
--------	---	-----------------------------------------------------------------------------------------

Claw system development

May 17	2.5	DVL depth troubleshooting — diagnosing blanking distance issues with old claw geometry; validating new 650mm arm clears the 0.1m DVL blanking threshold
--------	-----	---------------------------------------------------------------------------------------------------------------------------------------------------------

May 24	4.5	Claw mission — first full pool tests of new 2-DOF claw; vision-guided descent, pick-and-place cycles; initial end-to-end success rate characterization (16 attempts)
--------	-----	----------------------------------------------------------------------------------------------------------------------------------------------------------------------

Pre-competition full-system validation (Not Completed at Time of Writing)

May 30	2.5	New task data collection — collecting labelled vision data for updated 2026 task objects; Avery segmentation model training inputs
--------	-----	------------------------------------------------------------------------------------------------------------------------------------

Jun 7	9	New claw mission — extended claw iteration session; tuning XY alignment, depth estimation, and grab confirmation via current spike detection; improving end-to-end success rate toward 80% target
-------	---	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------


Jun 13	5	New dropper and torpedo mission — final revalidation of dropper and torpedo against 2026 task configurations; confirming success rates at competition-ready thresholds
--------	---	------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Jun 20	5	New claw mission — continued claw refinement; post-grab surface and rotate sequence for octagon yaw bonus; bin placement confirmation
--------	---	---------------------------------------------------------------------------------------------------------------------------------------

Competition runs (Not Completed at Time of Writing)

June 14	4.5	RoboSub 2026 full run — practice full-course run; gate → slalom → dropper → torpedo → octagon sequence; course navigation rate measurement
Jun 21	7	RoboSub 2026 full run — semifinal and finals runs
Jun 27	5	RoboSub 2026 full run — semifinal and finals runs
Total	104.5	28 sessions

Appendix F: Claw Design Analysis



Taking the moment around point p , including the torque of the servo itself is,

$$\sum_p M = 0$$

$$F_g \cdot l - F_s \cdot \sin(45^\circ) \cdot s - M_{servo} = 0$$

The design decision was that the spring halves the required torque of the servo, so,

$$F_s \cdot \sin(45^\circ) \cdot s = M_{servo}$$

Therefore,

$$F_g \cdot l = 2F_s \cdot \sin(45^\circ) \cdot s$$

Or,

$$mgl = 2F_s \cdot \sin(45^\circ) \cdot s$$

$$\frac{mgl}{2 \cdot \sin(45^\circ) \cdot s} = F_s$$

Where:

$$m = 0.550 \text{ kg},$$

$$g = 9.81 \frac{\text{m}}{\text{s}^2},$$

$$l = 0.320 \text{ m},$$

$$s = 0.055 \text{ m}.$$

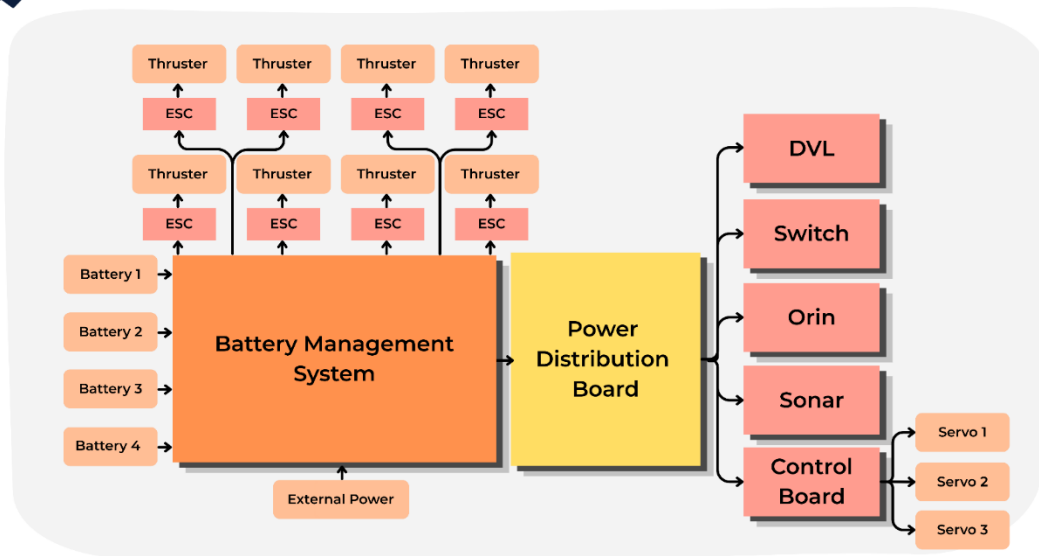
$$F_s = \frac{0.55 \text{ kg} \cdot 9.81 \frac{\text{m}}{\text{s}^2} \cdot 0.320 \text{ m}}{2 \cdot \sin(45^\circ) \cdot 0.055 \text{ m}}$$

$$F_s = 22.197 \text{ N}$$

Appendix G: Electrical Architecture Diagrams

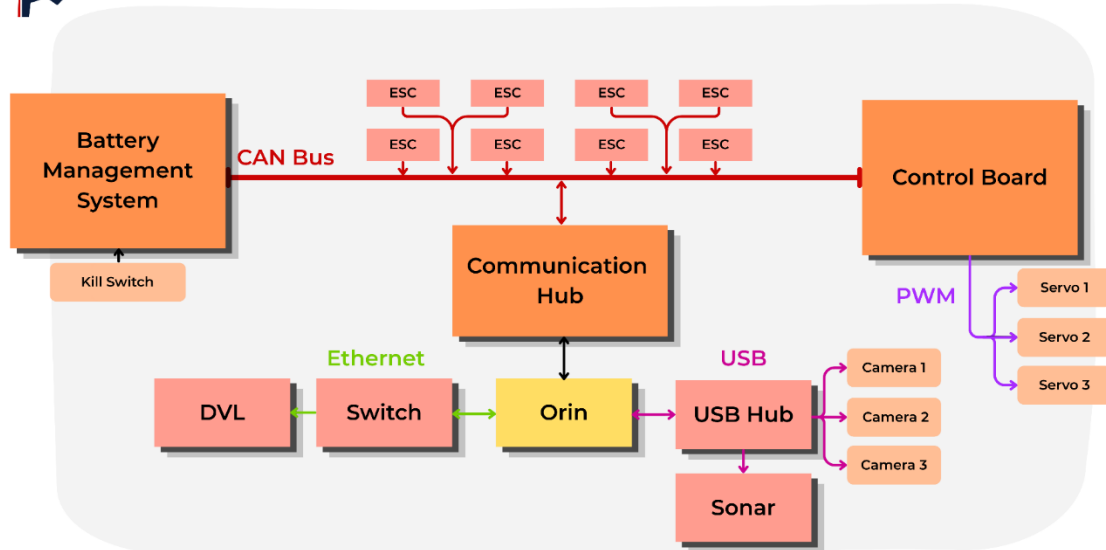


Overview: Power Diagram





Overview: Communication Diagram



Appendix H: BMS Design Analysis

H.1 Conductor Resistance Derivation

Using $R = \rho L/A$ where:

- $\rho = 1.72 \times 10^{-8} \Omega \cdot \text{m}$ (resistivity of copper)
- $A = 0.823 \text{ mm}^2 = 0.823 \times 10^{-6} \text{ m}^2$ (18 AWG cross-sectional area)
- $L = 0.35 \text{ m}$ (average cable length)

$$R = (1.72 \times 10^{-8} \times 0.35) / (0.823 \times 10^{-6}) = 6.02 \times 10^{-9} / 8.23 \times 10^{-7} = 7.3 \text{ m}\Omega \text{ per conductor}$$

Including return path and connector contact resistance (estimated 15 m Ω per connection point):

$$R_{\text{total}} = 2 \times 7.3 + 2 \times 15 = 14.6 + 30 = 44.6 \text{ m}\Omega \approx 50 \text{ m}\Omega \text{ per run}$$

H.2 Previous Architecture — Per Rail Current Draw

Rail Current

12 V I = 2.0 A

7.3 V I = 3.2 A

5 V I = 5.5 A

H.3 Previous Architecture — Resistive Loss Per Rail

$P_{\text{loss}} = I^2 \times R$, where $R = 0.05 \Omega$

- 12 V rail: $(2.0)^2 \times 0.05 = 4.00 \times 0.05 = 0.20 \text{ W}$
- 7.3 V rail: $(3.2)^2 \times 0.05 = 10.24 \times 0.05 = 0.51 \text{ W}$
- 5 V rail: $(5.5)^2 \times 0.05 = 30.25 \times 0.05 = 1.51 \text{ W}$

Total previous architecture losses: $0.20 + 0.51 + 1.51 = 2.22 \text{ W}$

H.4 New Architecture — 12V Consolidated Rail

12 V rail: $I = 4 \text{ A}$

$P_{\text{loss}} = (4)^2 \times 0.05 = 16 \times 0.05 = 0.80 \text{ W}$

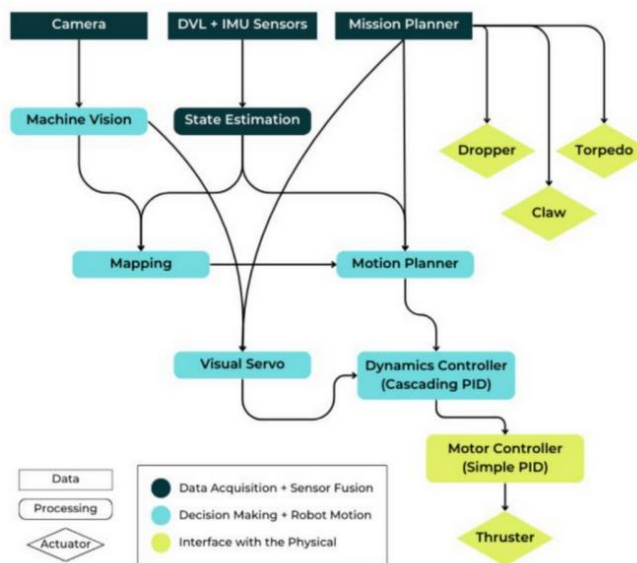
H.5 Reduction

$\Delta P = 2.22 - 0.80 = 1.42 \text{ W}$

% reduction = $(1.42 / 2.22) \times 100 = 63.9\%$

The consolidation of three distributed rails into a single 12 V rail reduces total resistive distribution losses from 2.22 W to 0.80 W, a 63.9% reduction in resistive distribution losses.

Appendix I: Software Architecture Diagram



Appendix J: Claw Testing Report

Summary

Test Phase	Cycles / Duration	Key Metric	Result	Requirements Verified
Component-Level	20 Servo stress tests 125 end effector trials	Sufficient servo torque / 80% gripper success rate	Pass - 88% Success Rate	Validated basic functionality of servo and end effectors.
Integration	100 arm/claw cycles	100% actuation success rate	Pass	Validated system actuation with all components underwater.
End-to-End (Pool)	16 task attempts	Ability to pick up and hold both objects	Pass	Validated full system operation and task achievability during pool testing.

Distance from Center (cm) [x-y plane]	Angle (Deg)	Height (cm) [z-axis]	Pass?
-10	-45	0	Pass
-10	-45	1	Pass
-10	-45	2	Fail
-10	-45	3	Pass
-10	-45	5	Pass
-10	-30	0	Pass

-10	-30	1	Pass
-10	-30	2	Pass
-10	-30	3	Pass
-10	-30	5	Pass
-10	0	0	Pass
-10	0	1	Pass
-10	0	2	Pass
-10	0	3	Fail
-10	0	5	Pass
-10	30	0	Pass
-10	30	1	Fail
-10	30	2	Pass
-10	30	3	Pass
-10	30	5	Fail
-10	-45	0	Pass
-10	-45	1	Pass
-10	-45	2	Pass
-10	-45	3	Pass

-10	-45	5	Pass
-5	-45	0	Pass
-5	-45	1	Pass
-5	-45	2	Fail
-5	-45	3	Pass
-5	-45	5	Pass
-5	-30	0	Pass
-5	-30	1	Pass
-5	-30	2	Fail
-5	-30	3	Pass
-5	-30	5	Pass
-5	0	0	Pass
-5	0	1	Pass
-5	0	2	Pass
-5	0	3	Pass
-5	0	5	Pass
-5	30	0	Pass
-5	30	1	Pass

-5	30	2	Pass
-5	30	3	Pass
-5	30	5	Pass
-5	-45	0	Pass
-5	-45	1	Pass
-5	-45	2	Pass
-5	-45	3	Pass
-5	-45	5	Pass
0	-45	0	Pass
0	-45	1	Pass
0	-45	2	Fail
0	-45	3	Pass
0	-45	5	Pass
0	-30	0	Pass
0	-30	1	Pass
0	-30	2	Pass
0	-30	3	Pass
0	-30	5	Pass

0	0	0	Pass
0	0	1	Pass
0	0	2	Fail
0	0	3	Pass
0	0	5	Pass
0	30	0	Pass
0	30	1	Pass
0	30	2	Pass
0	30	3	Pass
0	30	5	Pass
0	-45	0	Pass
0	-45	1	Pass
0	-45	2	Pass
0	-45	3	Pass
0	-45	5	Pass
5	-45	0	Pass
5	-45	1	Pass
5	-45	2	Pass

5	-45	3	Pass
5	-45	5	Pass
5	-30	0	Pass
5	-30	1	Pass
5	-30	2	Pass
5	-30	3	Pass
5	-30	5	Pass
5	0	0	Pass
5	0	1	Pass
5	0	2	Pass
5	0	3	Fail
5	0	5	Pass
5	30	0	Pass
5	30	1	Pass
5	30	2	Pass
5	30	3	Pass
5	30	5	Pass
5	-45	0	Pass

5	-45	1	Pass
5	-45	2	Fail
5	-45	3	Pass
5	-45	5	Pass
10	-45	0	Pass
10	-45	1	Pass
10	-45	2	Pass
10	-45	3	Pass
10	-45	5	Fail
10	-30	0	Fail
10	-30	1	Pass
10	-30	2	Fail
10	-30	3	Pass
10	-30	5	Pass
10	0	0	Pass
10	0	1	Pass
10	0	2	Pass
10	0	3	Pass

10	0	5	Fail
10	30	0	Pass
10	30	1	Pass
10	30	2	Pass
10	30	3	Pass
10	30	5	Pass
10	-45	0	Pass
10	-45	1	Pass
10	-45	2	Fail
10	-45	3	Pass
10	-45	5	Pass

Appendix K: BMS Testing Report

Summary

Test Phase	Cycles / Duration	Key Metric	Result	Requirements Verified
Component-Level	Multiple bench tests	Voltage ripple, protection circuits	Pass	Validated basic functionality of onboard power components before integration.
Integration	10 startup cycles, 15 min each	100% startup success rate	Pass	Validated system startup and basic functionality with all electronics

connected under static conditions.

End-to-End (Pool)	24+ startups, ~72 hrs in-water	Zero dropouts or camera failures	Pass	Validated full system operation under active load conditions during pool testing.
-------------------	--------------------------------	----------------------------------	------	-----------------------------------------------------------------------------------

1. Component-Level Testing

Validated basic functionality of onboard power components prior to integration.

Test	Observations	Expected	Result	Requirements Verified
DC-DC Buck Converter Ripple (I7C)	Measured 100 mV peak-to-peak ripple; outperforms datasheet spec	<200 mV (datasheet)	Pass	Verified output voltage and ripple voltage within datasheet spec.
LTC4421 — Hotswap	Correctly switched to external power connection as intended	Switch to external power port on connection	Pass	Verified higher voltage source is selected when external connection is present.
LTC4421 — Undervoltage Protection	Triggered near 20 V on bench supply due to output regulation interference; correct at 12 V threshold on vehicle battery	12 V threshold	Pass	Verified output is clipped when undervoltage threshold is triggered.
LTC4421 — Overvoltage Protection	Routing errors and incorrect capacitor values caused incorrect operation; second PCB revision required and validated	Trigger at OV limit	Revised & Pass	Verified output is clipped when overvoltage threshold is triggered.

Note: LTC4421 routing and capacitor errors affecting overvoltage protection were corrected in a second PCB revision before integration. All other tests passed without revision.

2. Integration Testing

BMS validated with cameras, onboard electronics, and NVIDIA Orin connected simultaneously under partial load.

Parameter	Result	Status	Requirements Verified
Startup cycles completed	10 / 10	Pass	The system shall have a 100% successful startup rate.
Startup success rate	100%	Pass	The system shall have a 100% successful startup rate.
Runtime per cycle	15 minutes under partial load	Pass	The system shall operate as intended for 15 minutes under partial load.
Camera failures observed	None (previous season: intermittent failures under same conditions)	Pass	Cameras shall maintain connection during operation.

3. End-to-End Outcomes

Observed totals from pool testing sessions under full operating conditions including servo and thruster current draw. No fixed cycle target was defined; testing continued across sessions until battery depletion.

Metric	Observed Result	Requirements Verified
Session duration	~1 hour per session until battery depletion	Kenai shall operate for 1 hour uninterrupted on full battery charge.
Camera failures & unexpected resets.	Zero	Cameras shall maintain connection with zero dropouts during operation.
Overall startup success rate	100% across full season	Kenai must be capable of completing all RoboSub tasks

Appendix L: Software Testing Report

Testing Hours & Resource Allocation

We prioritized finding logical bugs in simulation to maximize our in-water hardware testing. Pool time was heavily weighted toward localization and revalidating baseline tasks before moving to new logic.

Testing Environment	Total Hours	Focus
Simulation	> 100 hrs	Mission logic, state machine transitions, basic system stability.
Pool Testing (Total)	104.5 hrs (28 sessions)	Hardware-in-the-loop, environmental edge cases, end-to-end runs.
<i>↳ Baseline & T-Shoot</i>	<i>10 hrs</i>	Onboarding, 2025 IMU failure investigation, baseline core tasks.
<i>↳ Localization (Motus)</i>	<i>10.5 hrs</i>	PID/MPC validation, Motus IMU integration (confirmed 0.2°/hr bias stability).
<i>↳ Core Task Reval</i>	<i>14.5 hrs</i>	Pushing Gate, Torpedo, and Dropper logic to 99% success thresholds.
<i>↳ Slalom Integration</i>	<i>6 hrs</i>	Yaw-stable passes, physical entanglement abort timeout logic.

Vision System: 3D Pose Estimation Upgrade

Dropped the bounding box/segmentation models in favor of XFeat (feature matching) + HDBSCAN (clustering). This fixed our precision alignment issues and eliminated the need to label datasets.

Metric	Previous (Segmentation)	Current (XFeat + HDBSCAN)	Impact / Notes
---------------	--------------------------------	----------------------------------	-----------------------

Processing Speed	< 15 fps	> 20 fps	Better reaction time for dynamic alignment.
Yaw Accuracy	± 30°	± 5°	Crucial for Slalom and Octagon approach.
Depth Accuracy	± 4 inches	± 1 inch	Allowed us to tighten Claw drop tolerances.
Outlier Rejection	Inconsistent	High	HDBSCAN isolates only the largest data clusters.
Training Time	~12 hrs / session	0 hrs	Saved ~36 hrs total (bypassed ~3 labeling sessions).

Task Performance Outcomes

Reliability was the main objective this year. Integrating the Motus IMU and the new vision pipeline drastically improved consistency across the board.

Task	Previous Success	2026 Success	Notes	TDR System Requirement Target
Course Navigation	65%	Targeting 95%	Underlies the entire run model.	95% navigation rate
Core (Gate, Drop, Torpedo)	>95%	~90-99%	Base reliability held at ~90% during upgrades; pushed to ~99% post-Motus integration.	Minimum 80% task rate

Slalom	30%	Targeting 99%	Fixed via Motus yaw stability + new 5° vision yaw accuracy.	Minimum 80% task rate
Octagon (End-to-End)	15%	Targeting 80%	Baseline established May 24 (16 attempts). Still refining post-grab sequencing.	Minimum 80% task rate

Claw & Octagon Integration Parameters

To fix the DVL blanking and thruster wash issues from last year, we extended the claw to 650mm and dialed in the approach sequence over 5-10 pool iterations.

Parameter	Final Value	Validation Method / Notes	TDR System Requirement Target
Min. Safe Hover Height	0.5 meters	Verified physically by diver w/ measuring tape. Prevents thruster wash from blowing objects off the table.	Maintain > 0.5m elevation to prevent DVL faulty data
Pickup Confirmation	~2.5A	Software monitors servo current during close. A 2.5A stall spike confirms a successful grasp vs. an empty close.	Confirm pickups via matching current spikes to servo angles
DVL Blanking	Cleared	New 650mm arm length successfully bypasses the minimum-range restriction during descent.	Clear 0.1m DVL blanking distance

Appendix M: Acknowledgements

ARVP's 2026 season would not have been possible without the support of many individuals and organizations. The University of Alberta remains our single largest backer, providing both the financial foundation and the facilities that make our work possible. The Water Resources Lab granted us access to their testing space, often on short notice, and Engineering Safety and Shipping supported our competition logistics. The Project Group Services team went above and beyond in placing orders, arranging workshops, advocating for us within the university, and

advising members through personal challenges. The Engineering Students Society provided additional financial and marketing support.

Our sponsors amplified the impact of every hour we put in. ISS MineSafety as our Title sponsor, Aramark and Altium at the Partner level, Railshop Services, Marl Technologies, Elegoo, and Teledyne MEMS at the Support level, and Negmax, GitLab, and Anderson Anodizing as Donors. We are grateful for their continued investment in student engineering.

Special thanks to alumni Cole Dewis, Sergii Penner, Brendan Flynn, Kendra Smart, and Nathan Chong, who guided us through a significant team turnover and advised us throughout the year on everything from FMEA to underwater acoustics. More broadly, our 29-year alumni community continues to engage with us technically, attend our events, and open doors for our members professionally. That continuity is one of ARVP's greatest assets.

To everyone who attended our annual community showcase and donated: thank you. And to every team member's friends and family who endured the late nights, long hours, and time away, we could not have done this without your patience and support.