MM Nautronics RoboSub 2026 Technical Design Report

A. N. Türkeş, B. Eldemir, B. Gökgöz, D. C. İpekçi, D. Engüdar, D. K. Deran,
D. Vuran, E. Demirtaş, E. Erbaltacı, E. İ. Güçtaş, E. Kurt, G. Yıldız,
H. Altaş, H. Ö. Polat, İ. Turmak, M. A. İlhan, M. Akkan, M. B. Gül,
M. Dağlar, M. Gezer, M. İ. Gündüz, M. Yılmaz, Ö. Çıtakoğlu, Ö. Kaplan,
S. Enginar, S. Ercan, S. Y. Eskin, V. B. Dönmez, Y. Akşen, Y. Başaran, Y. Tosun

Supervisor: Lecturer Dr. Barış Yazıcı
Email: mmnautro@metu.edu.tr

***Abstract*—This paper outlines the systems engineering and strategic design of the Arat AUV, developed by MM Nautronics for the 2026 RoboSub. Prioritizing reliability and failure isolation over unvalidated novelty, we target "Advanced" tiers for four visual tasks while deferring the acoustic Octagon task to maximize scoring efficiency. Mechanically, a conventional sealed tube hull eliminates water ingress, completely decoupling external manipulation systems—like the thermoplastic gripper and marker dropper—from the pressure boundary. Electrically, a single-battery architecture with a custom Power Distribution Board ensures stability. Software-wise, an efficient geometry-based vision pipeline utilizing YOLO11-Pose and SolvePnP replaces heavy depth networks for 6D pose estimation. Integrated with a Hybrid Visual Servoing controller and ROS 2 Behavior Trees, this ensures dynamic task transitions. Supported by a rigorous five-phase testing protocol, this report details the trade-offs bridging our strategy with Arat's final architecture.**

***Index Terms*—AUV, autonomy, competition strategy, system design, control systems, testing**

I. Competition Strategy

Our competition strategy is built upon a clear strategic vision, deliberate trade-off analysis, and a task execution plan that directly shaped every design decision described in this report. Entering the competition for our second year, we experienced firsthand that the ambition for design novelty can compromise system reliability under strict time constraints. This section defines our scoring objectives, documents the engineering trade-offs that bound our design space, and describes our course execution approach—including the vehicle design decisions each step depends on.

A. *Strategic Vision*

By analyzing past RoboSub technical design reports, reviewing final run videos, and reflecting on our own on-site observations from last year's event, our team identified the critical system capabilities required for success and the common failure points to avoid. This analysis shaped our strategic priorities: developing an advanced control system, a stable and highly accurate computer vision pipeline, and a simple, modular software architecture capable of seamless task transitions. Our core objective is to execute all selected tasks at the "Advanced" level, as this offers higher scoring potential. Table I maps each selected task to its target capability level and the vehicle subsystem that enables it.

TABLE I: Task–Capability–Vehicle Design Matrix

Task	Target Level	Enabling Subsystem
Gate	Advanced	YOLO11-Pose + PnP
Slalom	Advanced	Wide-angle cam + PnP
Bin	Advanced	IBVS[1]+ marker system
Torpedo	Advanced + range bonus	IBVS[1]+ spring actuator
Octagon	Deferred / Stretch Goal	—

[1] IBVS control algorithms and integration strategies are detailed in Section II-C1.

The Octagon task was designated as optional. Recognizing our strengths in computer vision over acoustics, we determined hydrophone integration would unnecessarily consume nearly 35-40% of our engineering hours. These hours were redirected to refine our four Advanced-level tasks, maximizing our primary scoring potential. While the acoustic infrastructure is deferred to 2027, our mechanically risk-free, fully external gripper was retained as a "stretch goal". This allows us to attempt the Octagon as an opportunity task if the primary objectives are secured early, without jeopardizing system reliability.

B. *Trade-Offs Between Complexity and Reliability*

Every design decision was evaluated against three criteria: failure isolation, integration time, and field repairability. This philosophy was directly shaped by last year's experience, in which a 3D-printed integrated hull—though volumetrically efficient—could not be made watertight, rendering the vehicle inoperable at competition. The lesson was clear: a component that fails under field conditions eliminates not just its own function but every task that depended on the vehicle operating at all.

This year, each subsystem trade-off was framed as a question of where limited engineering hours produce the greatest reliability gain. To systematically address these challenges, we developed a risk mitigation framework, detailed in Table II, alongside specific architectural decisions:

- **Hull architecture:** Reverting to a conventional sealed tube design requires sacrificing internal volume, but it almost entirely eliminates the water leakage risk that previously jeopardized the vehicle. Because the end caps can be easily opened and closed, making the O-rings and penetrator interfaces the only potential points of failure for sealing, and allowing for rapid poolside intervention in case of an issue, this approach scored significantly higher than the previous design.
- **Actuator design:** The magnetic marker dropper eliminates almost all hull penetrations, removing an entire category of sealing failures. The spring-based torpedo actuator eliminates pneumatic complexity and its associated sealing risks. Although both decisions compromise on mechanical elegance, they significantly reduce failures. An actuator that fires reliably scores higher than an actuator that achieves a higher theoretical velocity but fails at the competition.
- **Custom hardware scope:** Custom hardware was largely deprioritised since our manufacturing tolerances could not meet the required accuracy within available resources. The primary hardware exception—a custom Power Distribution Board (PDB)—was justified for two reasons: no off-the-shelf alternative fit our spatial constraints, and embedding per-ESC current sensing directly supports our 2027 thrust-based odometry, turning a one-season cost into a long-term capability.

TABLE II: Risk Management & Fallback Strategies

Category	Primary Risk	Mitigation & Fallback
Hardware	Hull breach; actuator failure	**Mit.:** Min. hull penetrations via magnetic dropper and external gripper. **Fall.:** Leak sensor triggers immediate thruster shutdown.
Software	Task failure; GPU bottleneck	**Mit.:** Vision-based control system eliminates dependency on heavy depth networks. **Fall.:** Behavior Tree skips failed task; navigation points secured first.
Integration Delays	Insufficient pool testing time	**Mit.:** Dockerized ROS2 workspace with Gazebo end-to-end simulation. **Fall.:** Previous vehicle preserved as parallel pool testbed.

C. *Course Execution Strategy*

The vehicle enters the water with Gate traversal, utilizing YOLO11-Pose [1] keypoint detection accelerated via TensorRT [2] on the NVIDIA Jetson Orin NX [3]. During traversal, the visual identifier on the gate is detected to determine the vehicle's assigned mission role. This specific role is then stored in the Behavior Tree's shared memory. Because each role corresponds to a distinct set of target images for the rest of the course, this stored state dynamically dictates which specific visual identifiers the vehicle must target at the Bin and Torpedo stations.

The Slalom task follows immediately, exploiting the same forward-facing wide-angle camera and Perspective-n-Point (PnP) [4] pipeline already active during gate traversal; no sensor reconfiguration is required. The vehicle positions itself to maintain the correct pipe orientation through three passes, targeting Advanced-level scoring.

The Bin and Torpedo tasks are sequenced last. Both require the precise station-keeping capabilities provided by the IBVS controller, and both rely on the mission role memorized at the Gate to identify the correct targets for that specific role. Deferring actuator-dependent tasks ensures that complex maneuvers are attempted only after foundational points have been secured. For the Torpedo task, the vehicle targets the correct firing sequence and attempts the long-range "quest" bonus enabled by the redesigned fin-stabilized torpedo.

In the event of any task failure, the Behavior Tree advances the vehicle to the next objective without freezing the system. In the worst-case scenario (a complete actuator failure), Gate and Slalom points are still secured through stable navigation alone, guaranteeing a meaningful scoring floor before any high-complexity maneuver is attempted.

II. DESIGN STRATEGY

A. *Mechanical Design*

Our mechanical design strategy prioritizes absolute seal integrity, modularity, and operational reliability. By systematically decoupling all manipulation and payload systems from the primary pressure vessel, we eliminated critical hull penetrations, mitigating water ingress risks while optimizing underwater performance.

*1) **Improved Gripper**:* Last year's gripper presented two compounding failure modes. First, its segmented rubber pin contact surfaces failed to conform reliably to irregular object geometries, resulting in grip-loss during manipulation tasks. Second, the gripper was integrated directly into the vehicle hull, introducing a persistent water ingress risk that threatened overall seal integrity.

To eliminate both failure modes simultaneously, the new gripper is designed as a fully external, modular unit — mechanically decoupled from the main hull and therefore removing any penetration-based leakage risk entirely. Contact surfaces were redesigned from segmented pins to monolithic thermoplastic faces, which conform passively to a wider range of object geometries while offering greater mechanical durability [5]. Kinematics were deliberately constrained to a single degree of freedom (radial movement), reducing the total number of moving parts and minimizing underwater failure points without sacrificing adaptability. Power is transmitted via a 35 kg.cm servo motor through a rigid 8 mm shaft, selected to meet estimated payload requirements while keeping the mechanism fully external to the pressure hull. (See Figure 1 for the integrated model and Appendix M for the full exploded assembly view.)

Servo Selection: The 35 kg·cm servo was selected through static grip-force analysis. To achieve negative buoyancy, the target object is ballasted to 1.2 kg, resulting in a gravitational load of $F_w = mg = 1.2 \times 9.81 \approx 11.77\,\text{N}$. The opposing buoyancy force, derived from the object's displaced volume ($V \approx 8.96 \times 10^{-4}\,\text{m}^3$) and water density ($d = 1000\,\text{kg/m}^3$), is calculated as $F_b = Vdg = 8.96 \times 10^{-4} \times 1000 \times 9.81 \approx 8.79\,\text{N}$. The required minimum grip force, accounting for the net downward force and the thermoplastic contact surface friction coefficient ($\mu = 0.3$), is:

$$F_{grip} = \frac{F_w - F_b}{2\mu} = \frac{11.77 - 8.79}{2 \times 0.3} \approx 4.97\ \text{N} \qquad (1)$$

With a moment arm of $L = 15\,\text{cm}$, the selected servo delivers:

$$F_{servo} = \frac{\tau}{L} = \frac{35\ \text{kg}\cdot\text{cm} \times 9.81\ \text{m/s}^2}{15\ \text{cm}} \approx 22.89\ \text{N} \qquad (2)$$

This yields an actual safety factor of $SF_{actual} = 22.89/4.97 \approx 4.6$, confirming the servo meets grip requirements with sufficient margin across all target object geometries under submerged and ballasted conditions. For a detailed calculations refer to Appendix J.

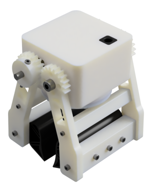

Fig. 1: Grabber.

*2) **Improved Torpedo Launching System**:* RoboSub 2026 awards additional points for long-range hits and sequential firing. Last year's CFD analysis revealed that off-axis angles induced flow separation and a destabilizing pitching moment, while a blunt rear created significant base drag. This year's redesign addressed these aerodynamic and mechanical challenges through deliberate systems-engineering trade-offs. To accommodate the flat pusher surface required by the new solenoid-triggered spring launcher, the torpedo's rear was strictly constrained to a flat profile. Knowing this would induce base drag, the forward section was redesigned to be highly slender, significantly reducing frontal hydrodynamic resistance to compensate [6]. Furthermore, to counteract off-axis tumbling, four rear fins were integrated to pull the center of pressure aft of the center of gravity. Confirmed via CFD at a 5° angle of attack, this provides dart-like stability and passive trajectory self-correction. While the forward fins in our eight-fin symmetric layout introduce localized flow separation, they act as essential mechanical guides centering the torpedo within the tube. Finally, replacing the stepper-motor with a solenoid system eliminated mechanical complexity and sealing risks. Constrained to 51×238×78 mm, the launcher supports simultaneous and sequential firing. Full CFD results are in Appendix G. (See Figure 2 for the external geometry, see Appendix N for the internal actuator assembly and Appendix I for the detailed kinematic calculations and analytical spring ballistic verification)

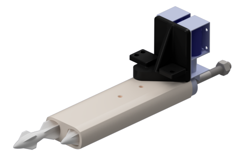

Fig. 2: Torpedo.

*3) **Improved Marker Dropper System**:* The RoboSub marker drop task requires precise placement onto a defined target. Last year's design used a servo-driven rotating mechanism that released markers sequentially via 90-degree rotations. Although functional, the servo actuation required physical penetration into the vehicle body, creating a persistent sealing risk that demanded O-rings and bearings to manage — adding mechanical complexity and potential failure points in the pressure boundary.

This year's design eliminates almost all hull penetrations by replacing servo actuation with an electromagnetic release mechanism. Two markers containing neodymium magnets are passively retained on internal ferromagnetic rods during transit. To release a marker, current is applied to coils wrapped around the rods, generating a repulsive magnetic force that overcomes the passive attraction and initiates the drop. Since the entire mechanism is solid-state with no moving parts, mechanical failure points are significantly reduced. To ensure accurate placement onto the target, markers are teardrop-shaped with a deliberately lowered centre of gravity, minimising drag, vortex shedding, and flutter during free fall [7] — validated by CFD analysis in Appendix H. (See Figure 3 for the module and Appendix L for the solid-state structural breakdown.)

Fig. 3: Marker Dropper.

*B. **Electrical Design***

This year, building on last year's off-the-shelf architecture, electrical design focused on eliminating instabilities introduced by the previous system while laying groundwork for future autonomy features. Custom hardware was pursued only where off-the-shelf solutions were physically inadequate or where accuracy requirements were achievable within manufacturing constraints.

*1) **Improved Power Distribution**:* Last year's triple-battery architecture caused a persistent centre-of-gravity offset that manifested as trim instability during pool testing, requiring repeated manual ballast corrections. Power capacity analysis confirmed that a single 22.2V Li-Po battery meets the vehicle's energy demands across the target mission window — reducing total battery mass by 1475 g without compromising operational endurance.

The transition to a shared power source eliminated last year's hardware-level electromagnetic interference(EMI) isolation. To address this, a custom Power Distribution Board (PDB) (see Figure 4) was designed and integrated directly into the vehicle's electronics organiser, as no off-the-shelf alternative fit within the available space. The PDB provides dedicated signal filtering, voltage regulation, and hardware protection on each power rail. Bench testing confirmed that per-rail noise levels remained within sensor operating limits. Additionally, independent per-ESC current sensing was embedded to support thrust-based odometry planned for next year. A comparison is presented in Table III.

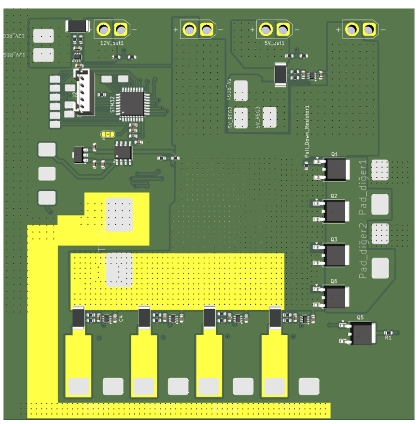

Fig. 4: Custom PDB.

TABLE III: Power Architecture Trade-off: Triple vs. Single Battery

Parameter	Last Year (3×)	This Year (1×)
Thruster supply	2× 14.8V Li-Po	1× 22.2V Li-Po
Electronics supply	1× 7.4V Li-Po	Onboard regulator
Total capacity	333 Wh	115.4 Wh
Total battery mass	2230 g	755 g
EMI isolation	Hardware (separate)	PDB filtering

*2) **Auxiliary Sensor Design**:* Custom sensor development was evaluated for several subsystems but deprioritised after manufacturing tolerance analysis indicated required accuracy could not be met within available resources. The sole exception regarding sensory equipment is the leakage sensor, where low accuracy requirements make a custom design viable. The sensor monitors hull integrity continuously and triggers an immediate failsafe shutdown upon detecting ingress — cutting thruster power before damage occurs.

C. Software Design

Our software architecture is based on the ROS 2 [8], which enables a modular design by dividing tasks into separate packages and establishing robust communication between modules. While most modules are written in Python, performance-critical components that require real-time data access are implemented in C++.

*1) **Vision-Based Control Pipeline**:* Last year's architecture combined SolvePnP-based pose estimation [9] with a Visual-Inertial Odometry (VIO) pipeline intended to provide global localisation. In practice, the VIO system could not be reliably validated even in simulation. This year, the architecture was redesigned around one guiding principle: *eliminate any dependency on global localisation entirely*. Every subsystem must be independently testable, and the control loop must be closed using simple sensors. The new pipeline is structured around two components: a geometry-based pose estimator and a hybrid visual servoing controller.

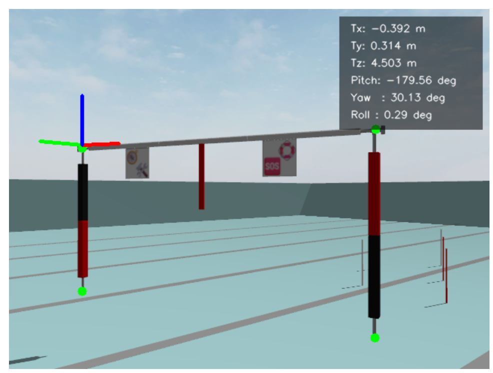

Fig. 5: Pose estimation.

a) Pose Estimation.: A custom-trained YOLO11-Pose model detects 2D keypoints of task elements in the image plane. These keypoints are passed to SolvePnP, which resolves the full 6D pose of the object relative to the camera — providing metric object distance Z and yaw angle θ from a monocular camera alone, without any additional depth sensor.

This design choice was made after evaluating learning-based monocular depth estimation as an alternative (e.g. Depth Anything V3 [10], MiDaS [11]). Three constraints ruled it out. First, self-supervised monocular depth estimators produce scale-ambiguous outputs; metric variants partially resolve this, but generalise poorly to out-of-distribution domains [12] — and the underwater environment, with its chromatic attenuation and backscatter, represents exactly such a domain [13]. Second, running a depth network alongside YOLO11-Pose introduces direct GPU contention, degrading overall pipeline throughput. Third, a dense depth map is unnecessary: the controller requires only the depth of a single detected object, which SolvePnP extracts analytically from keypoints that YOLO already produces, at negligible CPU cost. This geometry-first approach has been independently validated in real-world underwater environments, where RANSAC-based PnP over learned keypoints demonstrated robust metric accuracy in both pool and ocean conditions [13].

b) Hybrid Visual Servoing Controller.: With metric pose available, the control architecture required a strategy for converting pose estimates into velocity commands. Classical IBVS drives the vehicle purely from pixel-plane feature error, which preserves target visibility throughout the approach but produces erratic 3D trajectories when the initial pose error is large [14] — an unavoidable condition when the vehicle first acquires a task target at range.

To address this, a *hybrid visual servoing* controller [15] was adopted that explicitly decouples rotational and translational control. Yaw is corrected via pose-based feedback, exploiting the metric accuracy of SolvePnP for global stability. Translation, however, is driven by image-plane error to guarantee target visibility is maintained throughout the approach. The translational velocity command is given by:

$$\mathbf{v}_c = -\mathbf{L}_v^{+}(\lambda\, \mathbf{e}_t + \mathbf{L}_\omega\, \omega_c) \tag{3}$$

where $\mathbf{e}_t$ is the pixel-domain feature error from YOLO keypoints, $\mathbf{L}_v$ and $\mathbf{L}_\omega$ are partitions of the interaction matrix evaluated at depth Z from SolvePnP, and $\omega_c = -\lambda\,\theta\,\mathbf{u}$ is the yaw correction derived from the estimated yaw angle. The $\mathbf{L}_\omega\,\omega_c$ term explicitly compensates for the translational coupling introduced by yaw correction, ensuring the two channels remain independent.

Velocity commands are passed to a thrust allocation layer that distributes outputs across the vehicle's eight thrusters via a mixing matrix, then forwarded to the flight controller over MAVLink [16] — keeping high-level control logic fully decoupled from low-level actuation. (See Figure 6.)

*2) **Behavior Tree Architecture**:* Last year's behavior tree implementation suffered from a steep learning curve for new team members, slowing mission development and increasing onboarding overhead each competition cycle. While the BT architecture itself remained the right choice — unlike Finite State Machines, BTs support dynamic task prioritisation, conditional execution, and fallback behaviours without the state-explosion problem that arises as mission complexity grows [17] — the underlying library needed to change.

This year the team migrated to `py_trees_ros` [18], selected for two concrete reasons: its Python-based ecosystem significantly reduces the time for new members to contribute to mission logic; and its real-time tree introspection tools allow mission state to be monitored and debugged live during pool testing.

III. Testing Strategy

To prevent the late-stage integration failures and inefficient pool-time utilization experienced in previous cycles, we implemented a dual-track validation architecture. This structure

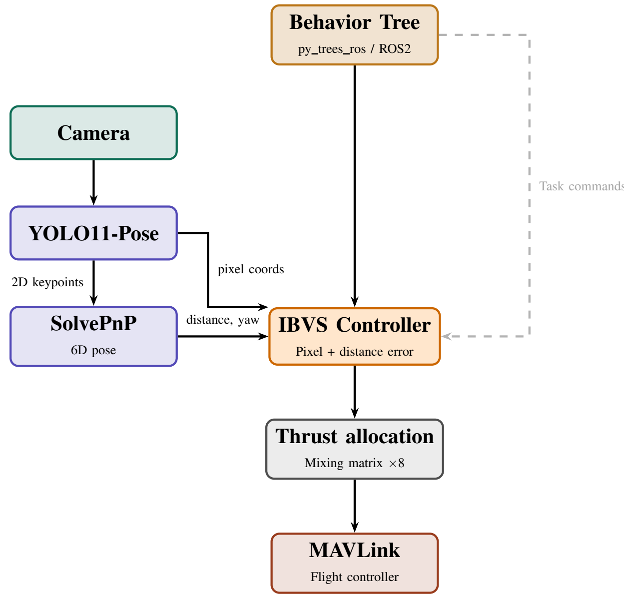

Fig. 6: Vision-based control pipeline.

divides our testing into two parallel workflows: the hardware track (actuator validation and sensor calibration) and the software track (software validation and simulation-based testing). While the software track utilized our legacy AUV for early in-water iteration, both tracks converged at a strict integration gate for our new vehicle, Arat. This ensured all subsystems were fully validated before proceeding to the final phase: full-system pool testing. This decoupled approach allowed teams to operate independently without sacrificing system reliability. The complete testing schedule is detailed in Appendix A, Table IV.

A. *Hardware and Actuator Validation*

All electronic boards and mechanical actuators are tested in isolation prior to system integration. The custom Power Distribution Board (PDB) is validated under load by driving ESC channels at nominal current while monitoring voltage stability and hardware protection behaviors. Servo-driven actuators—the grabber, marker dropper, and torpedo launcher—are cycled repeatedly in dry conditions to verify mechanical consistency, kinematic clearances, and proper trigger timing before any contact with water. This phase isolates hardware faults at the bench, optimizing our limited pool time. Detailed testing protocols and quantitative results for this phase are provided in Appendix B.

B. *Sensor Calibration*

Since our vehicle does not use a DVL to navigate and is fundamentally driven by camera-based feedback, the precision and consistency of the image processing pipeline became a critical system dependency. Both front and bottom cameras undergo meticulous intrinsic calibration. Additionally, the extrinsic rigid transform between each camera and the vehicle's center of mass is validated by comparing SolvePnP pose estimates against known ground-truth configurations. Calibration data is stored and subjected to version control alongside the software stack; any physical adjustment to the camera mounts triggers a mandatory recalibration before the next operational session.

In addition to the vision pipeline, auxiliary sensors are subjected to an equally rigorous pre-dive validation. To ensure precise Z-axis depth control, the external pressure sensor is calibrated against to local atmospheric pressure before every session. The internal IMU is steadily monitored for drift and noise characteristics during dry bench operations. Lastly, the hardware interrupt of the leakage sensor is physically validated by applying moisture to selected parts before the hull endcaps are sealed. The complete subsystem testing matrix and calibration results are documented in Appendix C.

C. *Software Validation*

Our entire ROS2 workspace is containerized using Docker, minimizing environmental inconsistencies between developer systems and the onboard Jetson Orin NX, thereby ensuring the software team maintains a consistent workflow independent of hardware availability. A closed-loop software validation strategy ensures predictable control responses before the vehicle ever enters the water. To accomplish this, every code commit is checked by a continuous integration (CI) pipeline that executes automated unit and integration tests in advance of merging. Furthermore, for mission logic, new object detection model weights and Behavior Tree branches are strictly validated against recorded `rosbag` data ahead of deployment. Detailed software validation protocols are provided in Appendix D.

D. *Simulation-Based Testing*

A comprehensive Gazebo simulation environment, featuring competition-accurate task element models and underwater physics, allows the entire perception-to-control pipeline to be tested end-to-end. Any architectural change related to task execution is subjected to regression testing within this simulation. Decoupling software validation from physical hardware readiness was a critical systems engineering decision. Instead of waiting for a fully integrated vehicle, the software sub-team continuously iterated on task strategies and Behavior Tree navigation while mechanical and electrical integration proceeded asynchronously. The comprehensive simulation testing protocol is available in Appendix D.

E. *In-Water Pool Testing*

In-water testing is strictly reserved for systems that have successfully passed all prior validation stages, and this rigid rule is applied uncompromisingly to our new vehicle, Arat. Before each pool session, the team establishes clear testing objectives: specific capabilities to be validated and measurable acceptance criteria to be met. Following the session, the results are reviewed against these benchmarks, and all environmental parameters are logged to ensure the reproducibility of the observed behaviors.

It was a strategic necessity to ensure that the software team could continue their testing uninterrupted while the isolated validation of the new hardware stack was ongoing. Therefore, our previous year's AUV, which operates on an already validated baseline, was exempted from this strict gating process and utilized as a parallel testing platform. Through this parallelization strategy, the software development process continued without being blocked by hardware schedules

The structured gating process significantly reduced repeated troubleshooting during pool sessions and improved overall testing efficiency. The cumulative risk mitigation outcomes from our entire testing architecture are mapped in Appendix F.

ACKNOWLEDGEMENTS

We extend our deepest gratitude to those who made this project possible. We thank METU and its Mechanical Engineering Department for providing the essential facilities and technical support necessary for our research. We are especially grateful to our mentors: Dean Dr. Taner Zorbay for his continuous guidance and support, our project supervisor Dr. Yazıcı for his dedicated mentorship and leadership throughout the project, Assoc. Prof. Turgut for his invaluable technical guidance, and Assoc. Prof. Koku for his expert insights into design.

Our project heavily relied on the generous resources from our sponsors (further detailed in Appendix O). We thank our Diamond sponsor, Maren Robotics, for crucial material support and underwater electronics. We also acknowledge the substantial financial and institutional backing from Ünides & Gençlik ve Spor Bakanlığı, Aselsan, Mitaş Group, Vodafone, and Adım ODTÜ. For supplying vital mechanical, electronic, and software components, we thank Keçiören Belediyesi Teknomer, Fenac, Fotoniks, Tekyaz (SolidWorks), and Arma. Finally, we deeply appreciate Çetinkaya CNC Ostim and Baca Engineering for their specialized manufacturing support.

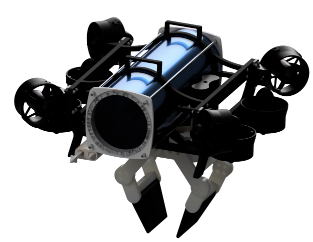

Fig. 7: ARAT AUV.

REFERENCES

[1] Ultralytics, "YOLO11: The latest Ultralytics YOLO model," [Online]. Available: https://docs.ultralytics.com/models/yolo11/. (Accessed: May 18, 2026).

[2] NVIDIA Corporation, "TensorRT SDK: High-Performance Deep Learning Inference," [Online]. Available: https://developer.nvidia.com/tensorrt. (Accessed: May 14, 2026).

[3] NVIDIA Corporation, "NVIDIA Jetson Orin NX System-on-Module," [Online]. Available: https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-orin/. (Accessed: May 15, 2026).

[4] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 2004.

[5] X. Zheng, N. Hou, P. J. D. Dinjens, R. Wang, C. Dong, and G. Xie, "A Thermoplastic Elastomer Belt Based Robotic Gripper," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Las Vegas, NV, USA, 2020, pp. 9275–9280, doi: 10.1109/IROS45743.2020.9341152.

[6] T. Truong, "Drag reduction technologies V," Defence Science and Technology Organisation, Report No. DSTO-GD-0290, 2001. [Online]. Available: https://apps.dtic.mil/sti/html/tr/ADA397790/

[7] J. Tay and T. T. Lim, "Drag reduction with teardrop-shaped dimples," in *2018 Flow Control Conference*, Atlanta, GA, USA, 2018, doi: 10.2514/6.2018-3528.

[8] ROS 2 Developers, "ROS 2 Project Repository," [Online]. Available: https://github.com/ros2. (Accessed: May 5, 2026).

[9] G. Bradski and The OpenCV Team, "cv::solvePnP - OpenCV Documentation," [Online]. Available: https://docs.opencv.org/4.x/d5/d1f/calib3d_solvePnP.html. (Accessed: May 20, 2026).

[10] Depth Anything Authors, "Depth Anything V3," *arXiv*, preprint arXiv:2511.10647, 2025. [Online]. Available: https://arxiv.org/abs/2511.10647.

[11] Intelligent Systems Lab Org, "MiDaS: Monocular Depth Estimation," [Online]. Available: https://github.com/isl-org/MIDAS. (Accessed: May 22, 2026).

[12] S. Atar, Y. Zhi, F. Richter, and M. Yip, "KineDepth: Utilizing Robot Kinematics for Online Metric Depth Estimation," *arXiv*, preprint arXiv:2409.19490, 2024. [Online]. Available: https://arxiv.org/abs/2409.19490.

[13] B. Joshi, S. Rahman, M. Marayati, H. Damron, and Y. Rekleitis, "DeepURL: Deep Pose Estimation Framework for Underwater Relative Localization," *arXiv*, preprint arXiv:2003.05523, 2020. [Online]. Available: https://arxiv.org/abs/2003.05523.

[14] F. Chaumette, "Potential problems of stability and convergence in image-based and position-based visual servoing," in *The Confluence of Vision and Control*, D. Kriegman, G. Hager, and A. S. Morse, Eds., vol. 237, Lecture Notes in Control and Information Sciences. Berlin: Springer-Verlag, 1998, pp. 66–78.

[15] F. Chaumette and S. Hutchinson, "Visual servo control, Part II: Advanced approaches," *IEEE Robot. Autom. Mag.*, vol. 14, no. 1, pp. 109–118, Mar. 2007. [Online]. Available: https://inria.hal.science/inria-00350638v1/document

[16] MAVLink Developers, "MAVLink Micro Air Vehicle Message Marshalling Library," [Online]. Available: https://github.com/mavlink/mavlink. (Accessed: May 8, 2026).

[17] M. Iovino, J. Förster, P. Falco, J. J. Chung, R. Siegwart, and C. Smith, "Comparison between behavior trees and finite state machines," *arXiv*, preprint arXiv:2405.16137, 2024. [Online]. Available: https://arxiv.org/abs/2405.16137.

[18] py_trees_ros Developers, "py_trees_ros: ROS 2 Behavior Trees," [Online]. Available: https://github.com/splintered-reality/py_trees_ros. (Accessed: May 10, 2026).

[19] R. C. Hibbeler and K. B. Yap, *Mechanics for Engineers: Statics and Dynamics*, 13th ed. Singapore: Pearson Prentice Hall, 2013.

[20] Y. A. Çengel and J. M. Cimbala, *Fluid Mechanics: Fundamentals and Applications*, 4th ed. New York, NY, USA: McGraw-Hill Education, 2018.

[21] G. J. Monkman, S. Hesse, R. Steinmann, and H. Schunk, *Robot Grippers*. Weinheim, Germany: Wiley-VCH, 2007.

[22] J. J. Craig, *Introduction to Robotics: Mechanics and Control*, 4th ed. Boston, MA, USA: Pearson, 2018.

[23] S. W. Moore, H. Bohm, and V. Jensen, *Underwater Robotics: Science, Design & Fabrication*. Monterey, CA, USA: Marine Advanced Technology Education (MATE) Center, 2010.

Appendix A
Testing Schedule

Since pool access is always hard to come by, we designed our testing schedule around parallel workflows. We focused heavily on bench testing and simulation for the first two-thirds of the season, making sure we only brought fully validated systems to the water. This setup allowed us to log around 320 hours of simulation and 180 hours of actual in-water testing before the competition. The testing timeline consists of five sequential phases: hardware and actuator checks, sensor calibration, software validation, simulation, and finally, full pool testing. We treated each phase as a strict checkpoint, meaning no subsystem could move forward without passing its tests. More importantly, we didn't want the software team waiting around while the new vehicle (Arat) was being built. To solve this, we kept our legacy AUV operational and used it for early in-water software validation. This kept both the software and hardware teams moving at full speed.

TABLE IV: Testing Timeline and Milestones

Test Phase	**Start Date**	**End Date**	**Prerequisites / Dependencies**
Phase A: Hardware & Actuator Checks	Jan 2026	Mar 2026	Parts procured, PCBs manufactured, CAD frozen
Phase B: Sensor Calibration	Mar 2026	Apr 2026	Machining finished & Phase A cleared
Phase C: Software Validation	Nov 2025	May 2026	Docker & ROS2 workspace set up
Phase D: Simulation-Based Testing	Jan 2026	Jun 2026	Gazebo milestones met & Phase C cleared
Early In-Water Software Validation (Legacy AUV)	Nov 2025	Apr 2026	Gazebo milestones
Phase E: Full System Pool Testing (Arat)	May 2026	Jul 2026	Phases A, B, C cleared & Core Phase D milestones met

Note: Phase D was initiated in parallel once the core modules of Phase C were validated. Full Phase C completion remained a prerequisite for hardware deployment, not for simulation-based iteration.

Appendix B
Testing Protocol and Results for Phase A

TABLE V: Detailed Hardware, Structural, and Actuator Testing Protocol (Phase A)

Test ID	Environment	Test Scenario / Acceptance Criteria	Quantitative Data / Equipment	Result & Design Modification	Status
TS-HW-01	Lab / Bench	**PDB Stability** Verify electrical stability of the custom 4-layer PDB under extreme loads.	**Tools:** Variable power supplies, oscilloscopes, multimeters. **Load:** Nominal continuous.	**Passed.** Line protection circuits and current shunts operated within safe limits.	✓
TS-HW-02	Sim / Bench	**Transient Thermal Analysis** Electronic systems must maintain temperatures below the 80°C critical shutdown threshold during a 10-minute full-load operational regime.	**Sim Setup:** 20°C start. **Load:** Jetson 8W, Camera 2W, ESC 1.2W, Pixhawk 1W, Batt 0.2W.	**Passed.** Peak temperature safely plateaued at 53°C at the ESC units. Validated Aluminum 6061 end-caps as passive heat sinks. No active cooling needed.	✓
TS-HW-03	Lab / Bench	**Dynamic Stress Testing** Thruster mounts must withstand the torsional stress of sudden extreme maneuvers without structural failure.	**Material Specs:** Original HDPE (0.95 g/cm^3) vs. PETG. **Load:** Simulated extreme maneuver torque.	**Modified.** Original HDPE mounts flexed and fractured. Mounts were redesigned and additively manufactured using PETG for superior tensile strength.	✓
TS-HW-04	Dry Bench	**Actuator Mechanical Consistency** Actuators (grabber, marker dropper, torpedo launcher) must execute repeated physical actuations without jamming.	**Cycles:** 50+ dry rapid fire operations per actuator.	**Modified.** Initial torpedo launcher showed friction-induced jamming. 3D print tolerances were increased and reprinted. Mechanical consistency achieved.	✓
TS-HW-05	Dry Bench	**Actuator Triggering Timing** Servo PWM signals must translate to precise, repeatable release timings to ensure synchronization with software commands.	**Delay Target:** $<$ 0.1 sec latency between ROS2 command and physical actuation.	**Passed.** Actuation delays mapped perfectly to command timestamps, guaranteeing predictable task execution before system integration.	✓

Strategic Impact of Thermal, Structural, and Bench Isolation Testing: The hardware validation tests conducted in Phase A served as the first critical line of defense against catastrophic integration failures. The transient thermal analysis (TS-HW-02), conducted under a rigorous 10-minute full-load simulation (assuming peak wattages for the Jetson Orin NX, ESCs, Camera, Pixhawk, and Battery), verified that peak temperatures would safely plateau at 53°C at the ESC units. This finding directly validated our mechanical design choice to rely solely on the Aluminum 6061 end-caps for passive heat dissipation, thereby saving significant internal volume and eliminating the mechanical risk of installing active cooling fans inside a sealed enclosure. Structurally, dynamic stress testing (TS-HW-03) revealed that our original HDPE thruster mounts were prone to flex and fracture under torsional loads. Transitioning these specific mounts to PETG via additive manufacturing ensured structural integrity during extreme maneuvers, while allowing us to preserve the HDPE main chassis for its inherent positive buoyancy. Furthermore, by isolating the custom 4-layer Power Distribution Board (PDB) on the lab bench (TS-HW-01), we verified its electrical stability under continuous load, confirming that line protection circuits and current shunts operated within safe limits. Catching friction-induced jamming in the torpedo launcher early on the dry bench (TS-HW-04), rather than underwater, allowed us to increase 3D printing tolerances and reprint instantly for mechanical consistency. Additionally, dry bench testing of actuator triggering timing (TS-HW-05) confirmed that servo PWM signals translated to precise physical actuations with less than 0.1 sec latency, guaranteeing synchronized task execution. Ultimately, this gating phase ensured that we never exposed expensive processing units to thermal risks or structural collapses, preserving both our hardware budget and our strictly limited pool time.

APPENDIX C
TESTING PROTOCOL AND RESULTS FOR PHASE B

TABLE VI: Detailed Subsystem Testing Protocol and Results Matrix (Phase B)

Test ID	Environment	Test Scenario / Acceptance Criteria	Result & Design Modification	Status
TS-SEN-01	Lab / Bench	**Camera Intrinsic Calibration** Front-facing and downward-facing cameras must be calibrated to eliminate lens distortion. Intrinsic reprojection errors must remain below acceptable thresholds..	**Passed.** Intrinsic matrices generated using custom checkerboard targets. Reprojection error stabilized at < 0.5 pixels for both cameras.	✓
TS-SEN-02	Software / Bench	**Rigid Transform Validation (Extrinsics)** The rigid transform between each camera and the vehicle frame must be validated by comparing SolvePnP pose estimates against known ground-truth configurations.	**Passed.** Angular and translational offsets aligned with CAD models. Rigid mounts validated to ensure transform matrix remains constant under operational vibrations.	✓
TS-SEN-03	Pre-Dive Sequence	**Pressure Sensor Z-Axis Zeroing** To guarantee precise depth keeping on the Z-axis, the pressure sensor must be zeroed according to the local atmospheric pressure before every session.	**Passed.** Automated pre-dive script integrated into the launch sequence to sample atmospheric pressure and tare the sensor to exactly 0.0m.	✓
TS-SEN-04	Lab / Bench	**IMU Drift and Noise Characterization** The IMU must be continuously monitored during dry bench tests to ensure that sensor noise and orientation drift remain within acceptable operational tolerances.	**Passed.** Allan Variance analysis conducted; stationary drift rates confirmed to be within acceptable parameters.	✓
TS-SEN-05	Lab / Bench	**Leakage Sensor Hardware Interrupt** Applying localized moisture to the leak sensor prior to closing the vehicle's caps must trigger an immediate hardware interrupt that executes the failsafe shutdown sequence.	**Passed.** Real-world moisture application successfully bypassed all software layers, instantly severing thruster power via the emergency relay.	✓
TS-SEN-06	Lab / Bench	**Calibration Data Version Control** All calibration configuration files must be stored alongside the software stack under version control. Any physical camera change must enforce a mandatory recalibration.	**Modified.** Gating policy established in repo. If a physical camera replacement occurs, the deployment pipeline blocks software launch until a new verified calibration hash is committed.	✓

Strategic Impact and Risk Mitigation of Sensor Calibration: The low reprojection errors (TS-SEN-01) and stable SolvePnP rigid transforms (TS-SEN-02) presented in Table VI directly shaped our competition strategy by proving that our vehicle could overcome the absence of a hardware DVL through rigorous software calibration discipline. This proven accuracy of our image processing pipeline allowed us to confidently include tasks requiring high positional precision, such as the Torpedo and Marker dropper, among our primary mission objectives. Simultaneously, our sensor validation practices optimized time and risk management during the testing phase. The strict versioning of calibration files (TS-SEN-06) and the standardization of the Z-axis zeroing protocol (TS-SEN-03) eliminated poolside configuration errors and the need to retune depth PID coefficients every session, allowing us to dedicate our strictly limited pool hours entirely to algorithmic testing. Finally, the physical verification that the leakage sensor triggers a direct hardware interrupt (TS-SEN-05)—bypassing all software layers—provided the development team with the structural assurance needed to test autonomy algorithms using more aggressive maneuvers with absolute system confidence.

APPENDIX D
TESTING PROTOCOL AND RESULTS FOR PHASE C & D

TABLE VII: Detailed Software Validation and Simulation-Based Testing Protocol (Phases C & D)

Test ID	Phase	Test Scenario / Acceptance Criteria	Quantitative Data / Equipment	Result & Design Modification	Status
TS-SW-01	Software	**ROS2 IPC Efficiency** High-frequency sensor and object detection data publishing must not bottleneck the Jetson Orin NX during complex maneuvers.	**Metric:** CPU usage & IPC latency. **Tool:** Orin NX Profiler.	**Modified.** Architecture optimized to publish data as primitive types rather than multi-arrays. C++/Python hybrid pipeline eliminated IPC lag.	✓
TS-SW-02	Software	**CI/CD Integrity & Containerization** Automated build pipelines must catch node configuration and import failures prior to hardware deployment.	**Metric:** 100% commit inspection. **Tool:** Docker & GitHub Actions.	**Modified.** A module naming inconsistency causing runtime failures in the Behavior Tree setup was caught and resolved via automated CI gating.	✓
TS-SW-03	Software	**Behavior Tree Fallback Logic** AUV must safely abort the Torpedo task and resume navigation if target tracking is lost.	**Data:** Recorded *rosbag* playback. **Tool:** py_trees framework.	**Passed.** New model weights and BT branches validated offline. Fallback logic successfully managed the abort sequence without vehicle stall.	✓
TS-SIM-01	Gazebo	**IBVS Trajectory Stability** AUV must approach a distant target with a large initial pose error without producing erratic 3D trajectories or oscillating.	**Tool:** Gazebo underwater world. **Metric:** Yaw error margins ($< \pm 2°$).	**Failed & Modified.** Pure pixel-error IBVS caused severe translational coupling. Upgraded to a Hybrid IBVS controller, decoupling yaw via SolvePnP.	✓
TS-SIM-02	Gazebo	**End-to-End Control Pipeline** Detection-to-control data flow must seamlessly execute competition tasks from visual acquisition to mechanical actuation.	**Metric:** Execution success rate of task sequences. **Tool:** Simulated task props.	**Modified.** Direct YOLO-to-IBVS caused conflicts. Pipeline was linearized to strictly feed SolvePnP output into the IBVS controller, eliminating jitter.	✓

Strategic Impact of Software Containerization and End-to-End Simulation: Phases C and D played a fundamental role in decoupling software validation from physical hardware readiness, allowing our sub-teams to work in parallel. By containerizing our entire ROS2 workspace via Docker and enforcing strict CI/CD pipelines (TS-SW-02), we guaranteed identical execution environments and caught critical module naming inconsistencies before poolside deployment. From a systems architecture perspective, evaluating the Jetson Orin NX's processing load (TS-SW-01) revealed significant inter-process communication (IPC) bottlenecks. We resolved this by rewriting our core nodes using a C++/Python hybrid pipeline and publishing strictly as primitive types, enabling our autonomy stack to run flawlessly during high-frequency perception tasks. In parallel, simulation-based testing drove major physical and control shifts. On the controls side, Gazebo simulations revealed that our initial pure pixel-error IBVS caused erratic 3D oscillations (TS-SIM-01). Upgrading to a Hybrid IBVS approach that decouples yaw via SolvePnP restored approach stability entirely in the virtual environment. Ultimately, validating behavior trees with offline *rosbags* (TS-SW-03) and testing the end-to-end pipeline in Gazebo ensured our autonomy stack was robust, safe, and competition-ready long before physical integration was completed.

APPENDIX E
TESTING PROTOCOL AND RESULTS FOR PHASE E

TABLE VIII: Detailed In-Water Pool Testing Protocol and Results (Phase E)

Test ID	Phase	Test Scenario / Acceptance Criteria	Quantitative Data / Equipment	Result & Design Modification	Status	
TS-WET-01	Pool	**Vision Pipeline & Environmental Adaptation** YOLO11-Pose keypoint detection must maintain stable target recognition despite underwater light refraction and surface glare.	Tool: YOLO11-Pose model. Environment: University pool under varying turbidity and lighting conditions.	**Modified.** Bench-validated pipeline failed under pool glare and refraction. OpenCV-based pixelation techniques were applied. Target detection stabilized against all observed pool conditions.	✓	
TS-WET-02	Pool	**Sealing & Pressure Diagnostics** Main hull must demonstrate zero water ingress at competition depth. A positive-pressure method must confirm watertight integrity before first dive.	Tool: 150-psi air compressor, custom sealing fixture, soapy water solution. High-risk zones: flange joints, O-ring grooves, penetrator bases.	**Modified.** Initial vacuum test showed a slight pressure drop but could not localize the leak. Switched to a Positive Pressure Bubble Leak Test. Soapy water immediately revealed capillary micro-leaks at penetrator bases; all were patched. Full watertight integrity confirmed before competition-depth operations.	✓	
TS-WET-03	Pool	**In-Water Hydrodynamic Validation** Torpedo must maintain a stable straight trajectory without tumbling. Teardrop markers must descend without lateral fluttering, validating prior CFD predictions in real fluid dynamics.	Environment: University pool. Observed metrics: torpedo flight trajectory, marker free-fall axis deviation.	**Passed.** Torpedo flew straight without tumbling and markers dropped without lateral flutter. Live pool results validated the CFD simulation data compiled earlier in the season, confirming that geometric and hydrodynamic trade-offs translated correctly from simulation to real-world fluid dynamics.	✓	
TS-WET-04	Pool	**Torpedo Ballistics Validation** The torpedo must reach the 30 cm target distance when fired under real submerged conditions, validating the conservative worst-case drag model ($C_d = 1.05$, flat-ended cylinder) used during spring selection.	Target distance: 30 cm minimum. Design model: $C_d = 1.05$ (flat-ended cylinder, worst-case). Projectile: Fully assembled streamlined torpedo, neutrally buoyant.	**Passed.** The streamlined torpedo exceeded the 30 cm target by 10–20 cm in live pool tests. The real aerodynamic drag was significantly lower than the worst-case model, confirming that the conservative spring selection provided adequate safety margin. Mathematical ballistic models validated for competition use.	✓	
TS-WET-05	Pool	**Gripper Payload & Torque Validation** The external gripper must maintain a secure hold on a 1.2 kg ballasted target object under submerged conditions without servo stall or slip, validating static grip-force calculations.	Target object: 1.2 kg ballasted payload. Servo: 35 kg·cm rated torque. Required safety factor: $\geq$ 4.6.	**Passed.** The 35 kg·cm servo provided the calculated 4.6 safety factor, overcoming net negative buoyancy without stalling. Monolithic thermoplastic contact faces held the payload without slipping. Fully external gripper design eliminated the leakage risks associated with the prior hull-penetrating mechanism.	✓	
TS-WET-06	Pool	**Hybrid IBVS Controller Validation** AUV must approach a distant target without oscillation or overcorrection under real pool fluid dynamics, proving that the Gazebo-validated Hybrid IBVS controller translates to real-world performance.	Metric: Approach trajectory stability, yaw error $< \pm 2°$. Tool: Forward-facing camera, SolvePnP pose estimator.	**Passed.** Pool tests confirmed that yaw decoupling via SolvePnP completely eliminated overcorrections observed with pure pixel-error IBVS. Stable approach trajectories toward distant targets mirrored simulation results exactly, validating the transition to Hybrid IBVS.	✓	
TS-WET-07	Pool	**CPU & Architecture Stress Test (Wet)** The C++/Python hybrid pipeline must handle high-frequency underwater perception tasks on the Jetson Orin NX without IPC lag in a live wet environment.	Metric: CPU load and IPC latency during simultaneous perception and control tasks. Tool: Jetson Orin NX onboard profiler.	**Passed.** The rewritten primitive-type ROS2 publishing pipeline completely eliminated IPC bottlenecks under full wet-environment load. Results confirmed that the architectural fix validated on the dry bench (TS-SW-01) held under live competition-representative conditions.	✓	

Strategic Impact of In-Water Pool Testing: Phase E served as the final validation layer, confirming that every simulation and bench-derived result translated accurately to real-world fluid dynamics. The switch to a Positive Pressure Bubble Leak Test (TS-WET-02) revealed micro-leaks that the initial vacuum method had failed to localize—a pre-competition critical fix that protected all onboard electronics. Torpedo ballistics testing (TS-WET-04) validated our conservative $C_d = 1.05$ worst-case drag model: the streamlined torpedo exceeded the minimum 30 cm target by 10–20 cm, confirming the spring selection margin. In-water hydrodynamic tests (TS-WET-03) confirmed that our CFD-driven geometric trade-offs for the torpedo and teardrop markers were correct, as both exhibited the stable trajectories predicted by simulation. Gripper validation (TS-WET-05) confirmed the calculated 4.6 safety factor under real submerged load and proved that relocating the mechanism fully outside the pressure hull eliminated the leakage risk present in last year's design. Pool validation of the Hybrid IBVS controller (TS-WET-06) proved that the yaw-decoupling fix developed in Gazebo (TS-SIM-02) held under real fluid dynamics, and the wet-environment CPU stress test (TS-WET-07) confirmed full pipeline stability under live perceptual load. Finally, the field integration of hardware color filters and OpenCV pixelation (TS-WET-01) resolved a vision failure mode that only manifested under live pool lighting—directly addressing the risk flagged in Appendix F.

Appendix F
Risk Analysis and Mitigation Strategies

TABLE IX: System Risks, Likelihood, Impact, and Actual Outcomes

Risk	Likelihood	Impact	Mitigation Strategy	Actual Outcome
ESC thermal shutdown under continuous load	Medium	High	10-min full-load thermal simulation at 20°C ambient (**TS-HW-02**)	Peak temperature plateaued at 53°C; passive Al 6061 end-cap cooling validated; no active cooling required
Structural mount fracture during sudden maneuvers	Low	High	Dynamic torsional stress testing; mounts redesigned and reprinted in PETG (**TS-HW-03**)	Torsional strength fully restored; no further fractures observed
Torpedo launcher friction-induced jamming	Medium	Medium	50+ dry rapid-fire actuation cycles; 3D-print tolerances increased and reprinted (**TS-HW-04**)	Friction-induced jamming eliminated; mechanical consistency confirmed before wet deployment
Calibration data drift between sessions causing positioning errors	Medium	High	Mandatory recalibration gating enforced in CI/CD pipeline; versioned calibration hashes (**TS-SEN-06**)	Deployment pipeline reliably blocked launch until verified calibration hash committed; poolside tuning eliminated
IBVS oscillation causing loss of target tracking	High	High	Upgraded to Hybrid IBVS decoupling yaw via SolvePnP in Gazebo (**TS-SIM-01**); validated in pool (**TS-WET-06**)	Yaw decoupled; erratic 3D oscillations eliminated; stable approach trajectories confirmed in both simulation and live pool
YOLO-to-IBVS pipeline conflict causing control jitter	Medium	High	End-to-end pipeline linearized to feed SolvePnP output strictly into IBVS (**TS-SIM-02**)	Jitter fully eliminated; task execution success rate confirmed in simulation before hardware deployment
Behavior Tree stall or abort failure during task execution	Medium	High	BT fallback branches validated via offline *rosbag* playback (**TS-SW-03**)	Abort sequence executed cleanly without vehicle stall; new model weights integrated without regression
CPU bottleneck from sensor data IPC overhead	High	Medium	ROS2 primitive-type publishing and C++/Python hybrid pipeline (**TS-SW-01**); confirmed under wet load (**TS-WET-07**)	IPC lag and processor bottleneck entirely eliminated in both dry-bench and live wet-environment conditions
Water ingress causing catastrophic electronics damage	Low	Critical	Hardware interrupt relay physically verified dry (**TS-SEN-05**); positive-pressure bubble leak test pre-dive (**TS-WET-02**)	Emergency failsafe shutdown confirmed operational; capillary micro-leaks at penetrator bases identified and patched before competition-depth operations
Vision pipeline failure under pool glare and turbidity	Medium	High	Hardware color filters and OpenCV pixelation techniques applied during pool testing (**TS-WET-01**)	YOLO11-Pose detection stabilized against underwater refraction and surface glare across all observed pool conditions
Torpedo spring under-powering; insufficient range at target distance	Medium	Medium	Conservative worst-case drag model ($C_d = 1.05$, flat-ended cylinder) used for spring selection (**TS-WET-04**)	Streamlined torpedo exceeded 30 cm minimum target by 10–20 cm; spring selection margin confirmed adequate
Gripper servo stall or payload slip under submerged load	Low	High	Static grip-force calculations with ≥ 4.6 safety factor; fully external gripper design (**TS-WET-05**)	35 kg·cm servo held 1.2 kg ballasted payload without stall or slip; external design eliminated hull-penetration leakage risk
Behavior Tree import failures upon hardware deployment	Medium	Medium	Docker containerization and CI/CD automated gating (**TS-SW-02**)	Module naming inconsistencies reliably caught and resolved pre-deployment

APPENDIX G
TORPEDO HYDRODYNAMIC AND CFD ANALYSIS

As mentioned in the main text, the torpedo redesign process was heavily driven by Computational Fluid Dynamics (CFD) simulations to adapt the geometry to the new launching mechanism while eliminating the instabilities observed in the previous year's design. The following sections detail the flow separation, pressure distribution, and stability characteristics across different iterations.

A. Analysis of the Legacy Design Flaws

The initial analysis focused on identifying the aerodynamic bottlenecks of the previous torpedo geometries.

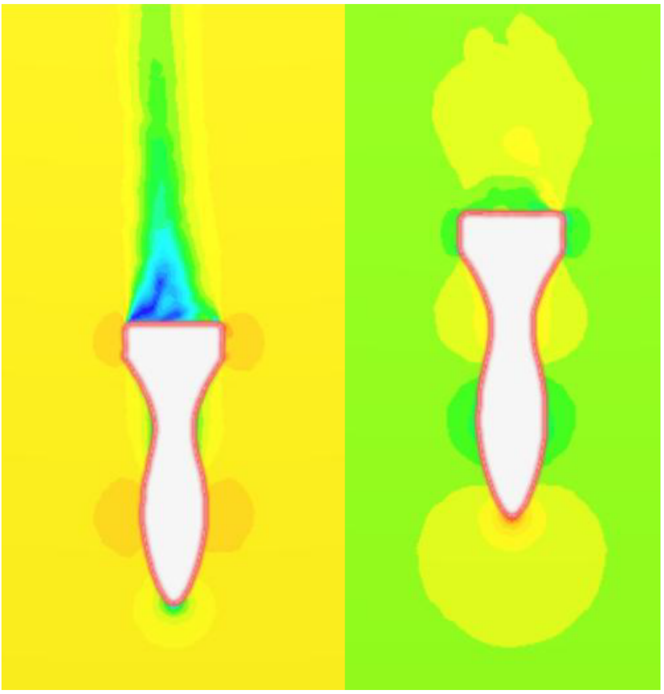

(a) Legacy Design 1: Excessive base drag and wake.

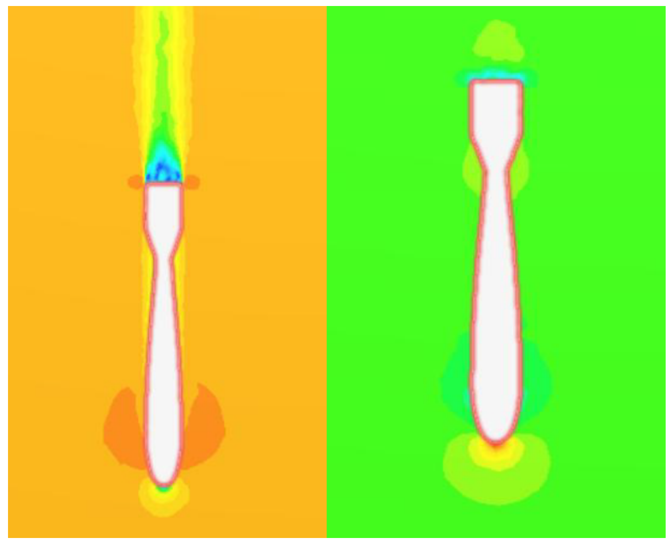

(b) Legacy Design 2: High fineness ratio with tail separation.

Fig. 8: Velocity and pressure contours of legacy torpedo designs.

As seen in Figure 8a, while the pointed nose successfully restricts the stagnation point to a narrow area, the abruptly truncated flat rear creates a massive low-pressure wake (visible as the deep blue region). This geometric suddenness induces severe flow separation and extreme base drag. Furthermore, the narrow "waist" section causes premature flow separation and undesirable pressure fluctuations.

Figure 8b demonstrates an iteration with a higher fineness ratio (long and slender body). While this shape inherently generates less drag than blunt bodies, the sudden contraction and the flat tail still trigger flow separation, resulting in a low-pressure zone at the rear that causes significant form drag.

B. Validation of the Redesigned Torpedo Stability

To solve the tumbling issues and adapt to the strict mechanical constraints of the new solenoid launcher, the torpedo underwent a major geometric overhaul. The concept design was tested under both off-axis (5°) and straight (0°) flow conditions.

Dynamic Stability Validation (5° Angle of Attack): To determine if the projectile would tumble or self-correct, flow was introduced at a 5° angle (Figure 9). The pressure differentials between the upper and lower surfaces generate lift and a pitching moment that tries to push the nose upward. However, by fine-tuning the overall geometry and utilizing the rear fins, the center of pressure (CoP) was firmly positioned aft of the center of gravity (CoG). This creates a massive restoring moment at the tail. The design actively resists tumbling and achieves dart-like passive trajectory self-correction.

C. Final Design Iteration and Mechanical Trade-offs

Following the comprehensive data obtained from the analyses of our legacy and concept designs, the final geometric form was developed (Figure 10). The primary objective of transitioning to this ultimate form was to ensure complete physical

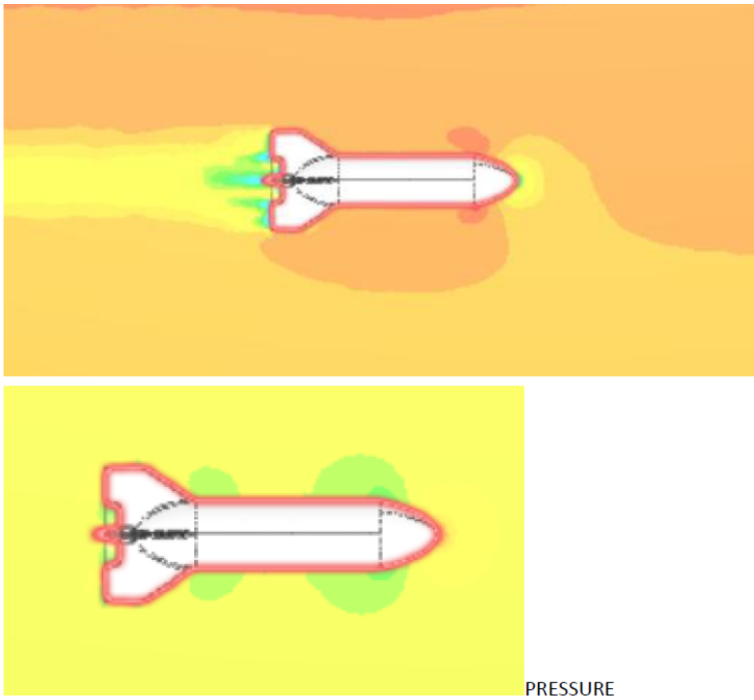

Fig. 9: Concept Design at 5° Angle of Attack: Pressure and velocity contours showing correcting moments.

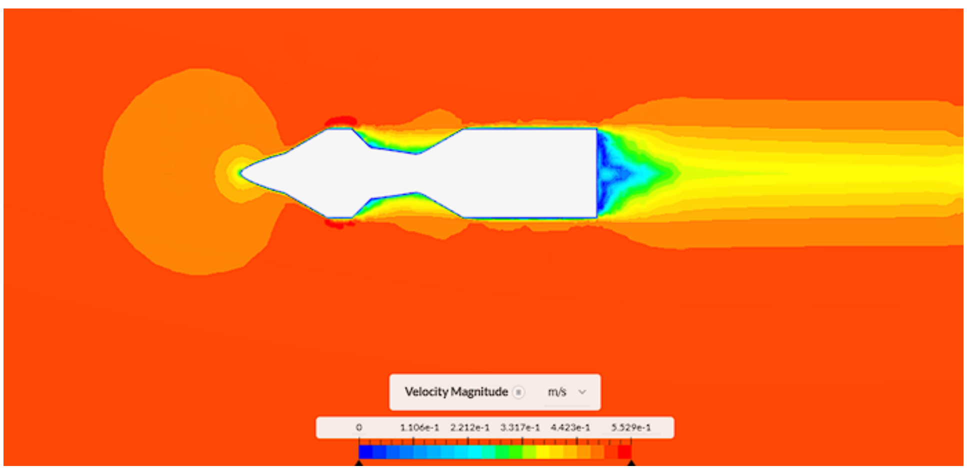

Fig. 10: New Design at 0° Angle of Attack: Slender front profile with integrated guide fins, compensating for the mechanically required flat rear and its localized wake.

compatibility between the torpedo projectile and the launching system. Specifically, to maintain strict linear motion within the launcher mechanism during the firing sequence, four guide fins were integrated into the front section of the torpedo.

Although introducing these forward fins inherently generates additional hydrodynamic drag and energy loss, this design choice was executed as a deliberate and conscious trade-off. Ensuring a perfectly linear, stable launch trajectory holds significantly higher operational value for mission success than a marginal reduction in drag.

Steady-State Performance and Mechanical Trade-offs (0° **Angle of Attack):** As shown in Figure 10, alongside the front fins, the strict requirement for the torpedo to sit flush against the launcher's flat pusher mechanism meant a flat rear geometry was mandatory. This deliberately accepted system-level trade-off results in a prominent low-pressure wake (visible as the blue/green region at the base), inducing unavoidable base drag. However, to effectively offset this penalty and accommodate the new front fins, the overall front section was heavily optimized into a long, slender profile. This forward geometry successfully delays boundary layer separation and significantly reduces frontal hydrodynamic resistance, resulting in a highly balanced overall pressure distribution despite the mechanical constraints imposed by the launcher.

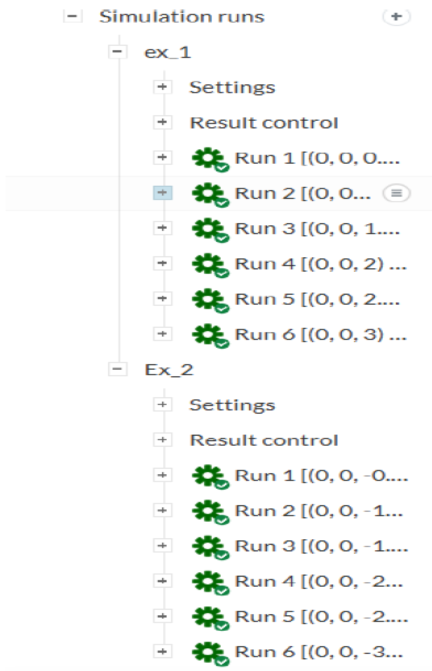

Fig. 11: Parametric simulation sweep across various velocities and conditions.

Simulation Methodology: To guarantee that these stability characteristics hold true across the vehicle's entire operational envelope, an extensive parametric sweep (Figure 11) was conducted. Multiple simulation runs were executed at varying velocities and boundary conditions to pinpoint the optimal launch speed and hydrodynamic constraints.

Appendix H
Marker Dropper Hydrodynamic and Free-Fall Optimization

The marker dropping task demands absolute trajectory precision during the unpowered free-fall phase. As stated in the main text, the marker's geometry was iterated via CFD to minimize drag, prevent flutter, and eliminate vortex shedding that could push the marker off-target.

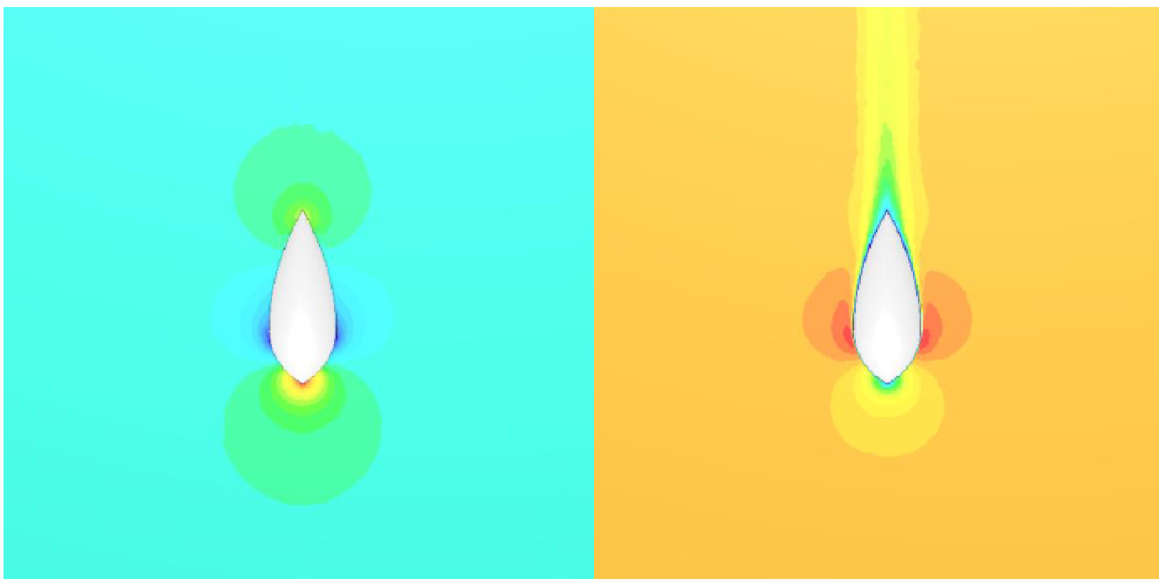

Fig. 12: Initial Marker Design (Velocity and Pressure Contours). Notice the wide low-pressure wake region indicating high form drag and potential flow separation.

Analysis of the Initial Design: Figure 12 illustrates the flow characteristics of our initial marker geometry. While a clear stagnation point forms at the leading edge (nose), the rear section suffers from a significantly wide wake region. This expansive low-pressure zone at the tail indicates early flow separation, resulting in substantial form drag and energy loss. Such a wide wake is highly prone to asymmetric vortex shedding, which would induce lateral fluttering during the marker's free fall, drastically reducing targeting accuracy. Based on these insights, the geometry was refined by tapering the tail by 2° to keep the boundary layer attached for a longer duration.

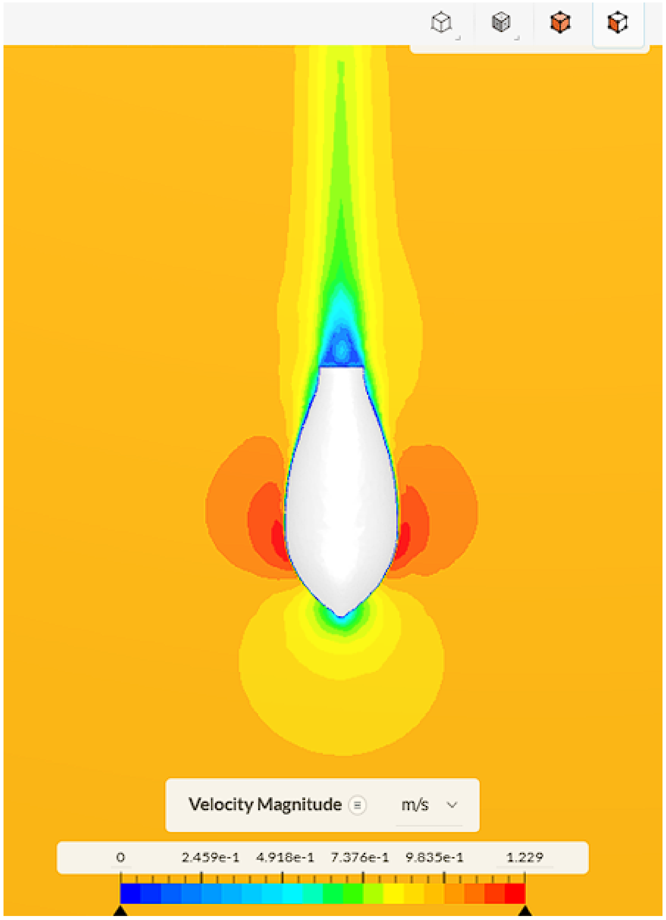

(a) Optimized velocity contour.

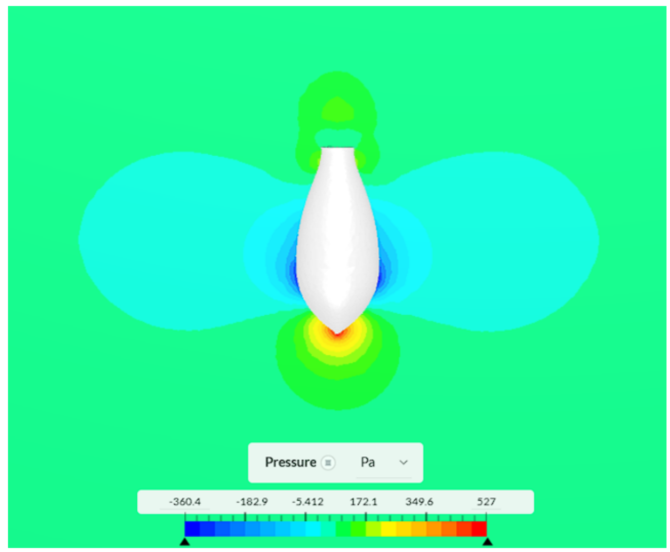

(b) Optimized pressure contour.

Fig. 13: Final teardrop marker geometry demonstrating delayed flow separation and a narrowed wake.

Validation of the Teardrop Geometry: The optimized ”teardrop” design (Figure 13) resolves the issues identified in the initial iteration. The velocity contour (Figure 13a) confirms highly symmetric flow with delayed separation. The wake region at the tail is markedly narrowed, indicating superior pressure recovery. The pressure contour (Figure 13b) validates that the high-pressure stagnation zone is isolated to the very tip, while the rest of the body experiences a smooth, balanced pressure gradient. AI-assisted flow estimations indicate an approximate 60% reduction in drag compared to the initial design.

Center of Gravity Consideration: Even with optimal hydrodynamic form, unpowered projectiles are susceptible to tumbling if perturbed. To lock the marker into a vertical descent path and prevent rotation, the physical manufacturing process dictates a deliberate lowering of the Center of Gravity (CoG), situating it heavily toward the nose. This combination of an optimized teardrop exterior and a bottom-heavy internal mass distribution ensures rapid, flutter-free descent directly to the target.

Design Trade-off for Mechanical Integration: While the optimized teardrop profile significantly minimizes form drag, a purely aerodynamic teardrop converges to a sharp point at the trailing edge. However, the top of our marker is deliberately designed with a flat, truncated surface. Truncating the tail inherently prevents perfect flow reattachment and causes a localized low-pressure wake region (base drag). Despite this minor aerodynamic penalty, the flat top is a strict mechanical necessity. It acts as the mating interface for the electromagnet release mechanism above. Providing a flush, wide contact area ensures stable magnetic coupling and instantaneous release, proving that this geometric deviation is a necessary and justified engineering trade-off.

Appendix I
Torpedo Launcher Spring Selection & Ballistic Analysis

A. Engineering Assumptions and Parameters

This section outlines the conservative analytical model used to verify the torpedo launch mechanism's operational range. The complex geometry of the torpedo (including the ogive nose and stabilization fins) is simplified to a perfect flat-ended solid cylinder. This models a hydrodynamic worst-case scenario with a high drag coefficient ($C_d \approx 1.05$) to ensure a conservative spring power rating.

The torpedo is engineered to achieve a neutrally buoyant state within fresh water ($\rho \approx 1000$ kg/m^3). Consequently, gravity and vertical buoyancy forces balance out perfectly, isolating the analysis to 1D horizontal kinematics. We assume ideal mechanical energy transfer where the potential energy stored in the compressed spring is transferred to the torpedo with 100% efficiency, neglecting mechanical friction and transient fluid drag during the 0.040 m stroke inside the barrel. The launch spring is composed of AISI 304 Stainless Steel (shear modulus $G = 79.3$ GPa) to prevent marine corrosion.

B. Analytical Spring Mechanics and Energy Derivations

Based on a torpedo diameter of 0.0165 m, the frontal cross-sectional area is calculated as $A \approx 2.138 \times 10^{-4}$ m^2. With a length of 0.0380 m, the total volume is $V \approx 8.126 \times 10^{-6}$ m^3. Applying the neutral buoyancy assumption, the effective mass is derived as $m = 0.008126$ kg (8.13 grams).

The selected helical compression spring has a free length of 0.110 m, a wire diameter of $d_w = 0.002$ m, and an outer diameter of $D_{out} = 0.01737$ m. Using 15 active coils ($n_a = 15$) and a mean diameter of $D_m = 0.01537$ m, the spring constant is derived via Hooke's Helical Spring formula [19]:

$$k = \frac{G \cdot d_w^4}{8 \cdot D_m^3 \cdot n_a} \approx 2912 \text{ N/m} \tag{4}$$

The elastic potential energy E_p stored inside the spring at a 0.040 m compression stroke is:

$$E_p = 0.5 \cdot k \cdot x^2 \approx 2.33 \text{ J} \tag{5}$$

Equating the elastic potential energy to the torpedo's translational kinetic energy, the initial muzzle velocity v_0 is extracted [19]:

$$v_0 = \sqrt{\frac{2 \cdot E_p}{m}} \approx 23.90 \text{ m/s} \tag{6}$$

C. Hydrodynamic Ballistic and Range Deceleration Analysis

Once discharged into the fluid medium, the torpedo experiences a velocity-dependent drag force. Setting this equal to inertial resistance yields the space-domain velocity attenuation differential equation [20]:

$$v(x) = v_0 \cdot e^{-\beta x} \tag{7}$$

Where the hydrodynamic damping parameter β is defined as:

$$\beta = \frac{\rho \cdot C_d \cdot A}{2m} \approx 13.80 \text{ m}^{-1} \tag{8}$$

Evaluating the residual velocity at the target travel distance of $x = 0.30$ m confirms sufficient terminal momentum:

$$v(0.30) = 23.90 \cdot e^{-13.80 \cdot 0.30} \approx 0.38 \text{ m/s} \tag{9}$$

Engineering Ballistic Verification Summary: The analysis mathematically verifies that the selected helical spring is fully capable of driving the torpedo beyond the target threshold of 30 cm. Because this mathematical verification is grounded in a conservative, flat-ended cylinder assumption ($C_d = 1.05$), the actual hydrodynamically streamlined torpedo ($C_d \approx 0.3$) will encounter significantly less resistance. Consequently, the actual mechanism will yield an extended operational range of roughly 60 to 80 cm, providing an exceptionally robust safety margin for competitive operations.

Appendix J
Gripper Mechanical Design & Force Analysis Report

A. Engineering Assumptions

The analysis of the underwater vehicle robotic manipulation system is governed by the following engineering assumptions:

- **Static Equilibrium:** The system is assumed to be in static equilibrium within the aquatic environment. The hydrodynamic drag force experienced during the vehicle's low-velocity maneuvers is negligible (∼1-2 N) and is excluded from the primary static analysis.
- **Frictional Characteristics:** The static coefficient of friction between the gripper's contact pads and the cylindrical plastic object is assumed to be uniform at $\mu = 0.3$.
- **Negative Buoyancy (Submerged State Constraints):** The hollow plastic test object has an initial dry mass of 113 grams. To ensure reliable negative buoyancy and complete submersion as per evaluation requirements, it is assumed that internal ballast (e.g., metallic spheres/ball-bearings) is introduced to increase the total mass to 1.2 kg.
- **Power Transmission Efficiency:** Power transmission from the actuator (servo motor) through the 15 cm gripper linkage mechanism is assumed to be ideal, with 100% mechanical efficiency (zero friction and compliance losses at joints).
- **Point of Force Application:** The grasping force is assumed to act as a concentrated normal load applied directly at the extreme terminal tip of the 15 cm gripper arm linkage.

B. System Parameters & Geometric Unit Conversions

The physical properties of the cylindrical object and the electromechanical specifications of the servo actuator are converted from imperial/standard units to SI units for mathematical consistency:

TABLE X: System Parameters for Grip Force Analysis

Parameter Description	**Original Value**	**SI Unit Value**
Object Diameter (Width)	5.9 in	0.14986 m ($r = 0.07493$ m)
Object Height (h)	2.0 in	0.0508 m
Target Ballasted Mass (m)	1200 g	1.20 kg
Fluid Density (ρ_{fluid})	Freshwater Standard	1000 kg/m^3
Gravitational Acceleration (g)	Standard Earth	9.81 m/s^2
Gripper Linkage Length (L)	15 cm	0.15 m
Actuator Stall Torque (τ)	35 kg·cm	3.4335 N·m

C. Mathematical Derivation & Analytical Calculations

Step A: Volumetric and Buoyancy Force Evaluation [20]

The total volume (V) occupied by the cylindrical geometry is determined via:

$$V = \pi \cdot r^2 \cdot h = \pi \cdot (0.07493\ \text{m})^2 \cdot 0.0508\ \text{m} = 8.96 \times 10^{-4}\ \text{m}^3 \tag{10}$$

The upward buoyancy force (F_b) exerted by the displaced fluid medium is:

$$F_b = V \cdot \rho_{fluid} \cdot g = (8.96 \times 10^{-4}\ \text{m}^3) \cdot (1000\ \text{kg/m}^3) \cdot (9.81\ \text{m/s}^2) = 8.79\ \text{N} \tag{11}$$

Step B: Vertical Axis Static Equilibrium ($\Sigma F_y = 0$)

Given that the ballasted weight ($W = m \cdot g = 1.2\ \text{kg} \cdot 9.81\ \text{m/s}^2 = 11.772\ \text{N}$) exceeds the buoyancy force (8.79 N), the object exhibits negative buoyancy and tends to slip downward. To preserve static equilibrium, the dual symmetric friction forces (F_μ) generated by the gripper jaws must act in the upward direction:

$$F_b + 2F_\mu = m \cdot g \implies 2F_\mu = m \cdot g - F_b \tag{12}$$

Substituting the mechanical friction law ($F_\mu = F_{grip} \cdot \mu$):

$$F_{grip} = \frac{m \cdot g - F_b}{2\mu} = \frac{11.772\ \text{N} - 8.79\ \text{N}}{2 \cdot 0.3} = 4.97\ \text{N} \tag{13}$$

Therefore, the minimum required normal force per gripper jaw to prevent slippage is 4.97 N.

Step C: Maximum Available Actuator Grip Force (F_{max}) [19]

The maximum normal clamp force deliverable at the boundary tip of the arm linkage is derived from the nominal actuator stall torque (τ) and the moment arm length (L):

$$F_{max} = \frac{\tau}{L} = \frac{3.4335\ \text{N} \cdot \text{m}}{0.15\ \text{m}} = 22.89\ \text{N} \tag{14}$$

D. Mechanical Safety Factor Verification

The Factor of Safety (FoS) establishes the structural and operational structural margin against unexpected physical disturbances or variations in surface friction coefficients [21], [22]:

$$FoS = \frac{F_{max}}{F_{grip}} = \frac{22.89\ \text{N}}{4.97\ \text{N}} = 4.60 \quad (15)$$

Engineering Evaluation Summary: The calculated Factor of Safety of 4.60 complies exceptionally well with the rigorous standards of underwater robotics design frameworks. In maritime applications, design safety margins between 2.0 and 5.0 are standard practice to absorb unmodeled dynamic variations, transient drag pulses during high acceleration, and surface boundary degradation caused by moisture-induced lubricity loss. The selected 35 kg·cm servo motor coupled with the 15 cm arm linkage geometry is mathematically validated as highly safe and fully capable of executing the manipulation task successfully under the specified ballasted parameters [23].

APPENDIX K
COMPONENT SPECIFICATIONS

TABLE XI: MM Nautronics AUV Components Summary (Updated)

Component	Vendor	Model/Type	Specs	C/P	Cost ($)	Year
Buoyancy Control	MM Nautronics	Inherent Material Buoyancy	3.03% Positive Buoyancy, naturally dampened	Custom	-	2026
Frame	MM Nautronics	HDPE Chassis	Density: 0.95 g/cm³, positive buoyancy	Custom	110	2026
Waterproof Housing	MM Nautronics	PMMA & Al 6061 Tube	PMMA Tube (0.163mm deflection limit), Aluminum 6061 passive heatsink end-caps	Custom	120	2026
Waterproof Connectors	Mucif	Degz-Neo Penetrator	Max Depth: 300m, Sealing: O-Ring and Epoxy, plastic	Purchased	25	2026
Thrusters	CubeMars	W30	Maximum thrust: 7.7kgf, Rated Power: 480W, Rated Current:22A, Weight(in air):420g, Weight(in water): 200g	Purchased	1600	2025
Propellers	CubeMars	Built-in W30 Propellers	-	Purchased	-	2025
Regulator	MM Nautronics	Built-in Converter Regulator	-	Custom	-	2026
Actuators	MM Nautronics	Servo Motor & Solenoid and Mechanisms	Torpedo Launcher & Marker Dropper	Custom	200	2026
Battery	Leopard Power	LiPo	22.2V 6S 5200 mAh 40C configuration	Purchased	200	2026
High Level Control	Holybro	Pixhawk 6C	STM32 based flight controller board	Purchased	175	2026
Battery	Leopard Power	LiPo	22.2V 6S 5200 mAh 40C configuration	Purchased	200	2026

Component	Vendor	Model/Type	Specs	C/P	Cost ($)	Year
Converter	MM Nautronics	Power Distribution Board	Custom-designed PDB for STM32G4	Custom	-	2026
CPU	Wave Share	Jetson Orin Nx	8W Operational Power	Purchased	1100	2025
Compass	MATEK SYSTEMS	Magnetometer RM3100	Professional grade Magnetometer based on RM3100	Purchased	108	2026
Inertial Measurement Unit (IMU)	Holybro(Bosch)	BMI088	Integrated to Pixhawk 6C	Purchased	-	2026
Vision	Intel	RealSense D415	Used strictly for RGB monocular video feed	Purchased	250	2025
Acoustics	-	-	-	-	-	-
Internal Comm Network	MM Nautronics	Ethernet, UART, I2C	ROS2 communication packets	Custom	-	2026
External Comm Interface	Blue Robotics	Fathom-X Tether	Minimum 50-meter slack safety tether	Purchased	-	2026
Manipulator	MM Nautronics	Gripper	Servo-driven mechanical gripper	Custom	-	2026
Algorithms	MM Nautronics	YOLOv11, py_trees_ros, SolvePnP	Behavior Tree task logic, OpenCV pixelation	Custom	-	-
Localization & Mapping	MM Nautronics	2-State Kalman Filter	Error-state Kalman filter with quaternion kinematics	Custom	-	-
Autonomy	MM Nautronics	Behavior Tree	Sequential state memory via py_trees	Custom	-	-
Open Source Software	-	ROS2, Gazebo, OpenCV	-	Free	-	
Inter-Vehicle Communication	-	-	-	-	-	-
Programming Language(s)	-	Python, C++	Hybrid pipeline (C++ core, Python bindings)	-	Free	-

Appendix L
Marker Dropper Exploded View

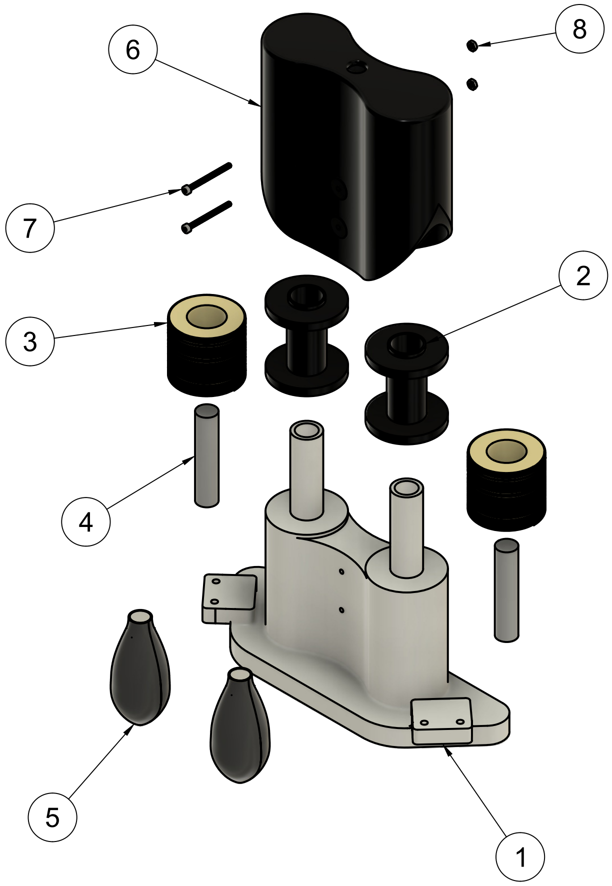

Parts List			
Item	Qty	Part Name	Material
1	1	Dropper	PETG
2	2	Coil	PETG
3	2	Coil Wire	Copper
4	2	Billet	Iron, Cast
5	2	Marker	PETG
6	1	Shell	PETG
7	2	Stainless Steel Socket Head Screw	Steel
8	2	Stainless Steel Hex Nut	Steel

Dept.	Technical reference	Created by MELIH GEZER 5/24/2026	Approved by
		Document type	Document status
		Title Marker-Dropper	DWG No.
			Rev. / Date of issue / Sheet 1/1

Appendix M
Gripper Exploded View

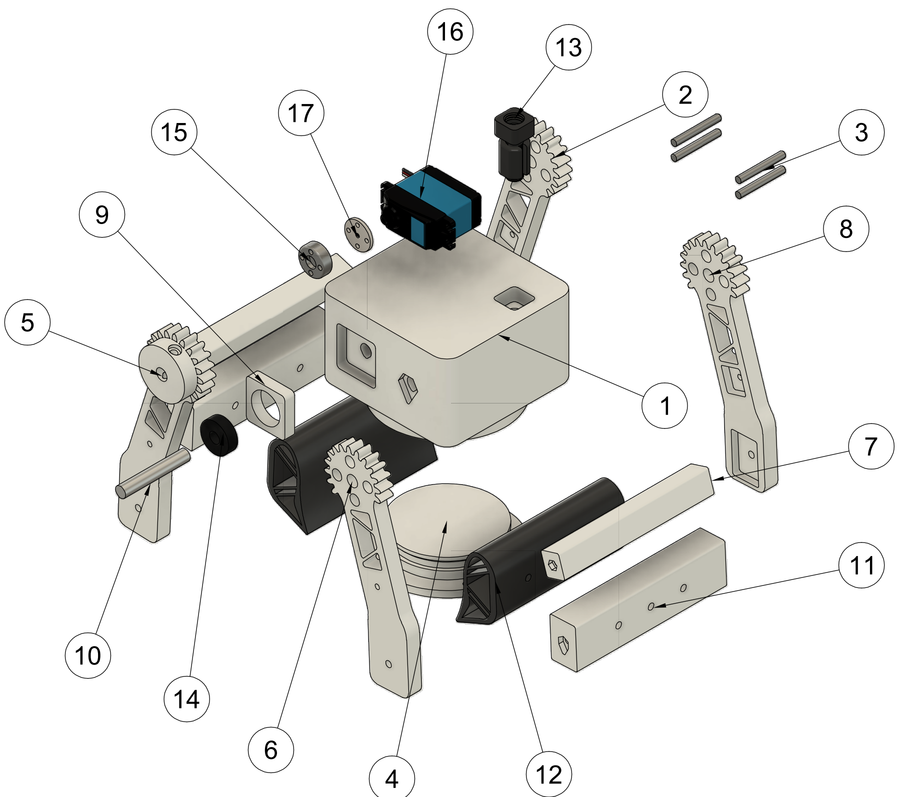

Parts List			
Item	Qty	Part Number	Material
1	1	Sealed Box	ABS Plastic
2	1	Geared-Arm 4	ABS Plastic
3	4	Servo Motor Holder	Steel
4	1	Bottom Cover	ABS Plastic
5	1	Geared-Arm 1	ABS Plastic
6	1	Geared-Arm 2	ABS Plastic
7	2	Upper Link	ABS Plastic
8	1	Geared-Arm 3	ABS Plastic
9	1	Rotary Seal Slot	ABS Plastic
10	1	8 mm Diameter Shaft	Stainless Steel
11	2	Lower Link	ABS Plastic
12	2	Elliptical Contact Surface	Thermoplas tic Resin
13	1	Sealed Connector	----
14	1	Rotary Seal	Rubber
15	1	Servo Coupling	Steel
16	1	35 kg.cm Servo Motor	Generico
17	1	Servo Spline Coupling	Aluminum

Dept.	Technical reference	Created by: MM NAUTRONICS	Approved by
		Document type	Document status
		Title: Grabber System	DWG No.
			Rev. / Date of issue / Sheet: 1/1

APPENDIX N
TORPEDO LAUNCHING SYSTEM EXPLODED VIEW

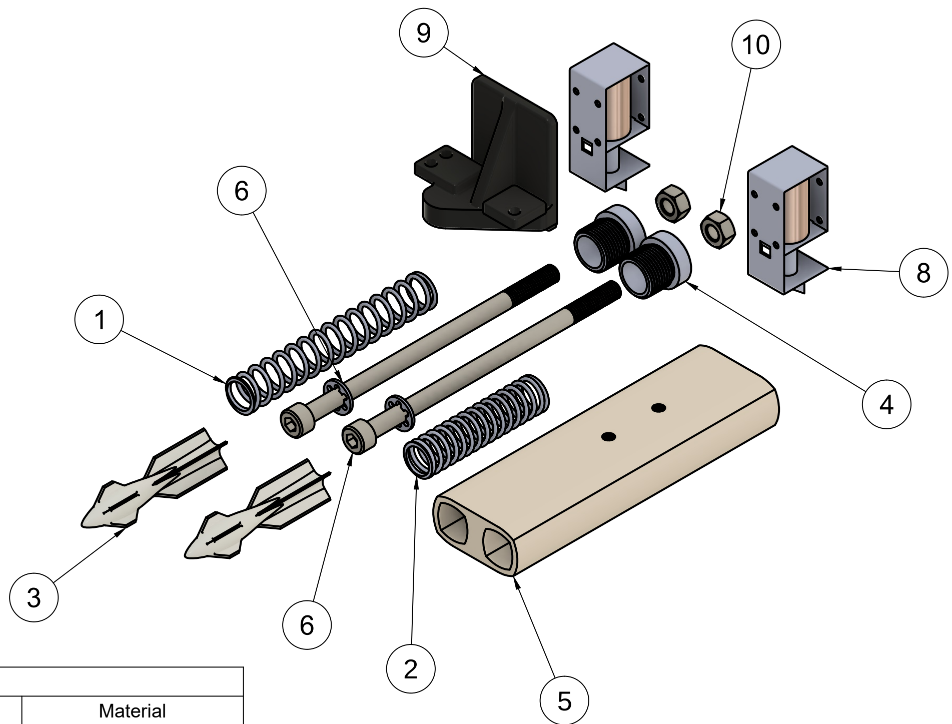

Parts List			
Item	Qty	Part Name	Material
1	1	Spring (Free)	Steel
2	1	Spring (Compressed)	Steel
3	2	Torpedo	ABS
4	2	Blind Plug	Aluminum
5	1	Torpedo Guide Rail	Aluminum
6	2	Socket Head Cap Screw	Steel
7	2	Washer	Steel
8	2	Solenoid Lock	Steel, Copper
9	1	Custom Torpedo Component	PETG
10	2	Hex Nut	Steel

Dept.	Technical reference	Created by MELIH GEZER 5/25/2026	Approved by
		Document type	Document status
		Title Torpedo	DWG No.
			Rev. / Date of issue / Sheet 1/1

Appendix O
Sponsors

Our team is incredibly grateful for the generous support of our sponsors. Their contributions—whether through direct financial funding aligning with our sponsorship tiers, advanced manufacturing, or specialized hardware supply—transformed our engineering designs into reality.

A. Diamond Sponsors

Maren Robotics — For their extensive material support equivalent to our highest tier, providing specialized underwater electronics, and crucial contributions to our vehicle's hardware inventory.

B. Gold Sponsors

Ünides & Gençlik ve Spor Bakanlığı — For their generous financial backing and institutional support empowering our youth-driven engineering initiatives.

C. Silver Sponsors

Aselsan — For their significant corporate financial backing at the Silver tier and valuable contributions to our robotics engineering efforts.
Mitaş Group — For their financial support and general contributions to our project's development.
Keçiören Belediyesi Teknomer — For supplying critical mechanical and electronic components equivalent to our Silver tier funding, required for the vehicle's assembly and rapid prototyping.
Adım ODTÜ — For their institutional backing and continued support of our university-based student project.

D. Bronze Sponsors

Fenac Automation Control — For their support and contributions regarding our automation and control systems.
Fotoniks — For their valuable support in fulfilling our technical, optical, and electronic requirements at the Bronze tier.
Vodafone — For their financial contributions supporting our team's operational activities and communication needs.
Tekyaz — For providing SolidWorks software capabilities and computer-aided design resources equivalent to this tier.
Arma — For their contributions of hardware and mechanical materials to our technical inventory.

E. Manufacturing & Production Support

Çetinkaya CNC Ostim — For their essential machining and production support in manufacturing custom mechanical components.
Baca Engineering — For their specialized manufacturing assistance, specifically the custom production of our vehicle's endcaps.

APPENDIX P
FULLY ASSEMBLED AUV

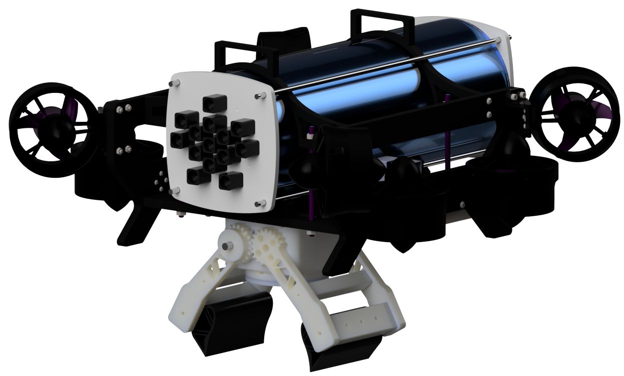

(a) Rear isometric view displaying the main end-cap, cable penetrators, and external gripper integration.

(b) Front isometric view detailing the forward-facing camera interface, thruster configuration, and fully external manipulation system.

Fig. 14: Final integrated system architecture of the Arat AUV prior to deployment.