

A Two-Loop Finite-State-Machine Controller with Bounded-Retry Verification for Resource-Constrained Autonomous Marine Debris Investigation

Advaith Rajagopalan
BASIS Independent Silicon Valley
San Jose, California, USA
rajagopalan.advaith@gmail.com

Abstract—Persistent marine debris monitoring requires autonomous surface vehicles (ASVs) that detect, investigate, and validate floating targets autonomously. ASV autonomy often takes it for granted that perception operates reliably; however, a resource-constrained ASV that stages a thermal proposal before an expensive vision-based confirmation must deal with the possibility of ambiguous detections. We introduce a two-loop finite-state-machine (FSM) controller—an exploration loop to cover an area and an investigation loop to validate a detection—where the key idea is to bound the number of rotational retries in verification. Such a verification stage allows for distinguishing between confirmation (DONE) and verification rejection (ABORT) while giving an ambiguous detection a chance of retry. Rather than fixing design parameters arbitrarily, we explore their effect through an experimental voting threshold sweep, an analysis of the failure mode of ABORT outcomes from 46 trials of the FSM controller in a dry run, and a coverage-efficiency framework. The dry run yields 35 DONE (76%) and 11 ABORT (24%) outcomes, demonstrating simulation-level correctness of the FSM logic; we show that every one of the 11 observed ABORTs stems from a verification rejection rather than a navigation or state-transition failure. On-water mission validation is the immediate next step of the work.

Index Terms—autonomous surface vehicle, finite-state machine, marine debris monitoring, perception uncertainty, multi-view voting, coverage path planning, resource-constrained robotics

I. INTRODUCTION

Exposure to ultraviolet light and wave action causes floating plastic to degrade into microplastics, rendering their extraction from the water column impractical. Hence, interception of macro-debris prior to fragmentation motivates monitoring efforts at an intermediate spatial scale, between satellite remote sensing and static collectors [1], [2]. Autonomous surface vehicles (ASVs) are promising in this scenario owing to their mobility, low cost, and ability to revisit dispersed targets.

However, much ASV autonomy research assumes reliable perception, treating it as the clean input to a controller [3], [4]. A resource-constrained ASV, however, that stages an inexpensive thermal anomaly proposal ahead of a costly vision-based confirmation cannot assume that. There exists the possibility that an anomaly proposed by thermal detection fails to confirm

at one viewpoint and confirms at a slightly different angle, or that an anomaly represents a non-target such as a reflection or a biological mat. In other words, an ambiguous result is obtained, and the controller—not only the perception stack—has to deal with that. In particular, the latter implies that the controller must explicitly determine when to re-observe the anomaly, when to confirm, and when to resume the search. The present paper makes the following contributions:

- 1) a two-loop FSM controller that couples an exploration loop with an investigation loop and distinguishes between verification *rejection* (ABORT) and task completion (DONE);
- 2) a bounded-retry verification mechanism that gives an ambiguous anomaly a small, fixed, capped rotational retry;
- 3) an experimental characterization of controller design parameters: a voting threshold sensitivity sweep, an ABORT failure-mode breakdown, and an analytical coverage-efficiency framework.

The emphasis is made on explicitly capturing the discrete decision layer atop the continuous-time motion control and the perception stack, which we treat as established components [5], [15]. The paper continues with a literature review (Sec. II), the problem description and system context (Sec. III), the two-loop FSM controller design (Sec. IV), the experimental and analytical characterization of design parameters (Sec. V), the baseline of operational performance (Sec. VI), a discussion of limitations (Sec. VII), and conclusions (Sec. VIII).

II. RELATED WORK

ASV autonomy and FSM control. The design of autonomous marine mission FSMs is a widely explored topic, and many approaches combine a low-level controller with a finite-state decision layer. Li *et al.* [3] use an FSM to guide ASV motion in a vision-based docking procedure, and low-cost catamaran USVs with layered control frameworks have been demonstrated for shallow-water environmental monitoring [4]. The continuous-time motion control problem itself is

well studied for marine craft [5]; our focus is the discrete decision layer that processes perception information and sequences task execution, rather than the trajectory tracking beneath it. Unlike reactive behavior-based or behavior-tree controllers [9], an explicit FSM keeps a small and auditable set of mission modes and transitions—which matters here because a safety-relevant outcome (ABORT) must be traceable to a specific cause, as our failure-mode analysis exploits.

Multi-view voting in perception. Aggregating detections across multiple viewpoints to suppress single-view errors is a long-standing idea in robotic perception. Consensus voting across multiple camera-equipped agents has been shown to substantially improve classification accuracy over single-view baselines [6], and multi-view fusion has been used to remove spurious or intersecting detections in robot scene understanding [7]. We adopt the fixed multi-view voting rule as a given and, unlike previous studies, analyze the *threshold* of that rule as a controller design parameter.

Coverage path planning for surface vehicles. Systematic area coverage by marine craft has a long history [8], with recent treatments specific to unmanned surface vehicles for environmental monitoring [10] and analyses of how sweep geometry affects efficiency [11]. Our coverage contribution is intentionally light: a closed-form efficiency estimate meant to guide exploration-loop parameter choices at design time, complementary to the full planners above.

ASV debris detection systems. Recent systems pair surface vehicles with deep detectors for floating litter retrieval and surveillance [12], [13]. These works concentrate on the detector and the collection task; this work, however, addresses the orthogonal issue of handling the ambiguity that staged perception delivers to the controller, so that an unreliable detection does not corrupt the mission.

III. SYSTEM OVERVIEW

The controller operates on a low-cost catamaran ASV platform that serves only as the substrate for this work: differential-thrust motion control, three RGB cameras at the bow for multi-view imaging, a thermal camera mounted on a servo for panoramic scanning, and a Raspberry Pi 5 for computation. The hardware design and bench validation of the platform are reported elsewhere; we take its availability as a prerequisite for this work. Perception operates in a staged fashion [16]: a panoramic thermal sweep proposes bearings of possible anomalies at low cost, and a vision-based investigation stage is triggered only when necessary, where a separately evaluated YOLOv8n detector [14] [15] produces per-camera classifications that are combined via a multi-view vote. The controller coordinates these stages and navigates the platform: it performs grid coverage with periodic scans for thermal anomalies, initiates an investigation for a candidate found at a given bearing, resolves ambiguous cases during verification, and manages coverage of the search region. Fig. 1 illustrates this relationship.

71 pt
0.986 in
25 mm




Fig. 1. Functional view of the system. The two-loop FSM controller orchestrates the staged perception pipeline and commands the platform; this paper concerns the controller layer, treating perception and motion control as components.

TABLE I
FSM TRANSITIONS. τ IS THE VOTE THRESHOLD, $R_{\max}$ THE RETRY CAP, AND $\sum \Delta\theta$ THE CUMULATIVE RETRY ANGLE.

State	Trigger / condition	Next state
EXPLORE	anomaly at bearing β	COARSE_ALIGN
EXPLORE	no anomaly; stations remain	EXPLORE (next stn.)
EXPLORE	no anomaly; grid covered	HALT
COARSE_ALIGN	aligned to β	FINE_ALIGN
FINE_ALIGN	heading refined	APPROACH
APPROACH	candidate range reached	VERIFY
VERIFY	votes $\geq \tau$	DONE $\rightarrow$ EXPLORE
VERIFY	votes $< \tau$, retry allowed	VERIFY (after retry)
VERIFY	votes $< \tau$, retry capped	ABORT $\rightarrow$ EXPLORE

IV. CONTROLLER ARCHITECTURE

A. Two-loop FSM design

The controller operates in two nested loops (Fig. 2). The outer *exploration loop* performs coverage of a predefined grid; at each grid station the vehicle scans for thermal anomalies; if none are found it moves to the next station (or halts once the grid is covered), and if any is found at bearing β control passes to the inner loop. The inner *investigation loop* completes the investigation of the detected anomaly at β , returning either DONE, indicating a confirmed investigation, or ABORT, indicating a verification rejection, after which the exploration loop resumes. The advantage of two separate loops is the independence of exploration logic from confirmation logic, together with a clear division where perception uncertainty enters the system: in the verification phase of the inner loop. The complete set of transitions is given in Table I.

B. States of the investigation loop

In the investigation loop, four states must be completed before verification: COARSE_ALIGN roughly rotates the vehicle toward the candidate bearing β , FINE_ALIGN fine-tunes the orientation, APPROACH moves closer, and finally VERIFY captures RGB images and classifies. Each alignment state is given a fixed timing budget—the investigation loop targets a three-second turn window per alignment action—to make

57 pt
0.792 in
20.1 mm

57 pt
0.792 in
20.1 mm

71 pt
0.986 in
25 mm




Fig. 2. Two-loop FSM. The exploration loop (top) advances between grid stations until a thermal anomaly triggers the investigation loop (COARSE_ALIGN → FINE_ALIGN → APPROACH → VERIFY). VERIFY yields DONE on a passing vote, a bounded rotational RETRY on an ambiguous vote, or ABORT once retries are exhausted; both terminal outcomes return control to exploration. When all grid stations have been visited without a confirmation, the exploration loop enters a terminal HALT state (Algorithm 1; omitted from the figure for clarity).

sure that the mission cannot get stuck indefinitely due to a particularly recalcitrant anomaly candidate. The states are kept coarse: they encode the decision logic rather than the actual motion, which is taken care of by the platform's continuous-time controller.

C. Bounded-retry mechanism

The core of the inner loop is its ability to deal with ambiguous verification. Instead of aborting on the very first failed verification, we issue a small but fixed rotational retry, because an ambiguity may resolve from a slightly different viewpoint. Retries are bounded in both their cumulative angle (total rotation $\leq 45^\circ$) and their count ($R_{\max}$), guaranteeing termination. Once a retry yields a passing verdict, the outcome becomes DONE; if the cap is reached without a passing vote, the verdict remains ABORT. A single-shot verification call is thus converted into a bounded re-observation policy, making verification rejection a deliberate, first-class controller outcome rather than an uncontrolled failure.

D. Voting verification logic

The staged thermal-to-vision perception architecture on which this controller runs is described in a companion paper [16], which owns the perception pipeline and its multi-view voting rationale; here we treat that pipeline as given and study the voting *threshold* as a controller parameter. The controller consumes the verification evidence at the frame level: the three RGB cameras and multiple frames per camera yield five verification outputs per investigation, and the candidate

is accepted when at least τ of the five outputs are positive (threshold τ). This is the same consensus mechanism described at camera level in [16] (a candidate confirmed by two of the three cameras), expressed here at the granularity of the five frame/camera outputs that the FSM actually consumes. The threshold determines the sensitivity/specificity trade-off: a low τ accepts weak evidence but also admits marginal detections, while a high τ rejects plausible detections and has good specificity. Algorithm 1 outlines the main FSM loop; the choice of τ is the subject of the empirical sensitivity sweep in Sec. V.

E. Termination and well-formedness

Two properties follow directly from the construction. First, every reachable state has at least one outgoing transition for every input it can receive (Table I), so the FSM cannot deadlock. Second, each investigation terminates in a finite number of cycles: VERIFY is left as soon as a vote passes, and otherwise after at most $R_{\max}$ retries (and never beyond the 45° cumulative-angle cap, which the guard $\sum \Delta\theta + \Delta\theta \leq 45^\circ$ enforces *before* each rotation), so the loop reaches DONE or ABORT and returns to EXPLORE. Because the exploration grid is finite and each station is marked visited, a coverage pass that confirms no target visits every station once and then enters the terminal HALT state (Algorithm 1); persistent monitoring is obtained by relaunching passes at a higher level. The bounded-retry cap and the grid-exhaustion exit are therefore not merely heuristics but the properties that make the controller provably halting and progress-preserving.

F. Implementation

The controller and its perception interfaces are implemented in Python as composable modules—sensor stubs, state-
transition handlers, the voting rule, and structured logging—so the complete FSM can run in a hardware-free dry-run harness (Sec. VI) by swapping physical drivers for software proxies. All randomized experiments use fixed seeds, and per-trial logs (visited states, vote counts, and confidences) are retained, so the analyses of Sec. V are reproducible from the recorded outputs. This modular separation is precisely what allows the controller logic to be validated independently of the platform and the detector, which is the methodological stance of the present paper.

V. EXPERIMENTAL ANALYSIS

We characterize three controller design parameters: a voting threshold for verification, an ABORT failure mode, and exploration geometry efficiency. The objective is to demonstrate that each design choice is thoroughly explored across the design space rather than fixed arbitrarily.

A. Voting-threshold sensitivity sweep

To decouple verification logic from the remainder of the controller, we simulated only the verification stage. Each of the five votes was sampled independently with a synthetic positive-output probability of 0.6, calibrated to the DONE rate

Algorithm 1 Two-loop FSM main loop (simplified)

```

1:  $state \leftarrow \text{EXPLORE}$ ; mark all grid stations unvisited
2: while  $state \neq \text{HALT}$  do
3:   if  $state = \text{EXPLORE}$  then
4:     perform thermal sweep at current station; mark visited
5:     if anomaly at bearing  $\beta$  then
6:        $state \leftarrow \text{COARSE\_ALIGN}$ ;  $retries \leftarrow 0$ 
7:     else if an unvisited station remains then
8:       move to next grid station
9:     else
10:       $state \leftarrow \text{HALT}$  {grid covered}
11:    end if
12:  else if  $state$  is an alignment or approach state then
13:    execute alignment/approach toward  $\beta$ ; advance  $state$ 
14:  else if  $state = \text{VERIFY}$  then
15:     $v \leftarrow$  positive votes among the 5 RGB outputs
16:    if  $v \geq \tau$  then
17:      transmit coordinates;  $state \leftarrow \text{EXPLORE}$  {DONE}
18:    else if  $retries < R_{\max}$  and  $\sum \Delta\theta + \Delta\theta \leq 45^\circ$  then
19:      rotate by  $\Delta\theta$ ;  $retries++$  {RETRY}
20:    else
21:       $state \leftarrow \text{EXPLORE}$  {ABORT}
22:    end if
23:  end if
24: end while

```

TABLE II
VOTING-THRESHOLD SENSITIVITY SWEEP (1000 TRIALS PER THRESHOLD).

Threshold	DONE rate	ABORT rate	Pos. votes	Conf.
2/5	0.911	0.089	3.00	0.558
3/5	0.675	0.325	2.97	0.552
4/5	0.342	0.658	3.02	0.560

of the 46-trial baseline (Sec. VI). The threshold τ was swept over 2/5, 3/5, and 4/5, with 1000 trials per setting and fixed random seeds for reproducibility. Table II and Fig. 3 report the result.

The positive-output probability was selected so that the modeled 3/5 DONE rate (0.675) would approximate the baseline dry-run DONE rate; the small gap from the observed 0.76 stems from the fact that the baseline mixes 3/5, 4/5, and 5/5 vote outcomes whose empirical frequencies an independent-vote model only approximates. The sweep is thus a sensitivity study *around* the baseline operating point rather than an exact replica of it.

As the positive-output probability stays constant (0.6) throughout the sweep, the mean number of positive votes stays near three across all settings; thus the decrease in DONE rate is caused only by stricter thresholds. The permissive 2/5 threshold maximizes the number of completions, the strict 4/5 threshold maximizes specificity, and the 3/5 threshold sits




Fig. 3. DONE rate versus voting threshold. The permissive 2/5 rule maximizes completion, the conservative 4/5 rule maximizes specificity, and 3/5 is the middle operating point used in the baseline trials.

TABLE III
POSITIVE-VOTE DISTRIBUTION AMONG THE 11 BASELINE ABORTS. ALL REACHED VERIFY AND FELL BELOW THE 3/5 THRESHOLD.

Positive votes (of 5)	ABORT trials	Mean conf.
1	5	0.335
2	6	0.474
≥ 3	0	—

between them, matching the threshold used in the baseline dry run—which supports it as a balanced operating point rather than an arbitrary default.

B. ABORT failure-mode breakdown

The 46-trial FSM dry run resulted in 35 DONE outcomes and 11 ABORT outcomes. To investigate the ABORT failures, we examined the states each visited and the voting results. All eleven reached VERIFY after visiting the full chain $\text{COARSE_ALIGN} \rightarrow \text{FINE_ALIGN} \rightarrow \text{APPROACH} \rightarrow \text{VERIFY}$ and returned a complete set of five verification outputs, yet failed to reach the 3/5 acceptance threshold (Table III). None reached ABORT as a consequence of a navigation timeout, an earlier-stage failure, or a state-transition error.

Hence, ABORT occurred only in cases when the vote fell below the acceptance threshold. The upstream FSM transitioned correctly in every one of the 46 trials—alignment, approach, and verification sequencing worked as designed—and ABORT was attributable *solely* to the verification rule. This agrees with the bounded-retry design, which makes ambiguous evidence the cause of ABORT but not the upstream state transitions: ambiguous evidence ends an investigation cleanly and returns the vehicle to searching, without contaminating navigation or coverage.

C. Coverage-efficiency design framework

We conclude this section by relating exploration-loop geometry to expected throughput analytically. For a square grid with spacing g and thermal sweep radius r , the farthest point

TABLE IV

ANALYTICAL COVERAGE ESTIMATE ($r=5$ m, $v=0.5$ m/s, $T_{\text{SCAN}}=10$ s).
DESIGN-TIME VALUES, NOT MEASURED PERFORMANCE.

g/r	g (m)	Gap-free	Stat./hr	Rate (m ² /hr)	100 m sq.
0.50	2.50	Yes	240	1500	15.0%
1.00	5.00	Yes	180	4500	45.0%
1.41	7.07	Yes	149	7456	74.6%
1.50	7.50	No	144	8100	81.0%
2.00	10.0	No	120	9425	94.2%




Fig. 4. Analytical coverage rate versus normalized grid spacing g/r . The dashed line marks the gap-free limit $g = \sqrt{2}r$; beyond it apparent coverage keeps rising but gaps open between scan footprints.

in a cell lies at distance $g/\sqrt{2}$ from the nearest scan center, giving the conservative gap-free condition

$$g \leq \sqrt{2}r. \quad (1)$$

Assuming a station scan time T_{scan} and a transit speed v , the time per station is $T_{\text{station}} = T_{\text{scan}} + g/v$, so $N_{\text{stations}} = 3600/T_{\text{station}}$ stations are scanned per hour. With effective unique area per station $A_{\text{eff}} \approx \min(g^2, \pi r^2)$, the approximate hourly coverage is $N_{\text{stations}} A_{\text{eff}}$. Under the illustrative assumptions $r = 5$ m, $v = 0.5$ m/s, $T_{\text{scan}} = 10$ s, and a square 100×100 m region, Table IV and Fig. 4 plot efficiency as a function of normalized spacing.

As we see, increasing the spacing reduces the number of stations while increasing the unique area each station contributes. Around $g = \sqrt{2}r$ the grid remains gap-free while avoiding excessive overlap; beyond that point, additional apparent coverage is achieved at the expense of gaps between scan footprints. We stress that this is a design-time tool for selecting exploration parameters, with real-world values substituted from field measurement, rather than a claim of measured operational performance.

The estimate is intentionally simple and omits second-order effects: it charges only straight-line transit (g/v) per station and neglects turning overhead, acceleration, and inter-row transitions, and it assumes still water. Environmental forcing such as currents can shift the energy-optimal sweep direction away from the geometric optimum [11], and a deployed planner would fold these terms in. For the design-comparison

71 pt
0.986 in
25 mm

purpose here, the closed form is nonetheless sufficient to rank spacing choices and to locate the gap-free knee at $g = \sqrt{2}r$.

VI. OPERATIONAL BASELINE: 46-TRIAL DRY RUN

The full FSM design was verified in a software dry run that exercises all states of the FSM, from exploration through investigation up to verification, with software substitutes for the physical sensors and actuators. For each candidate, the verification proxy outputs a binary “debris / not debris” decision and a confidence score, calibrated to match expected detector behavior, so that the voting logic executes on every cycle. Across 46 unique runs with the threshold set at $3/5$, the controller reached DONE in 35 cases (76%) and ABORT in 11 cases (24%). As established in Sec. V, all 11 ABORTs were verification rejections, and the upstream FSM operated as expected: no navigation nor state-transition failure was observed. In other words, the 46-trial dry run demonstrates correctness of the FSM pipeline—state transitions, bounded retry, and safe termination—in simulation. Its drawback is clear: the dry run is an instrument for validating control logic and cannot replace the actual sensing and ASV maneuvers, which alone require on-water trials.

VII. DISCUSSION AND LIMITATIONS

The core strength of the controller is that it handles uncertain perception results as a controller responsibility and not as a detector one: ambiguous verification outcomes are absorbed by a bounded re-observation policy, and the ABORT/DONE distinction allows the result to be clearly communicated to the rest of the mission. The design is parameterized rather than hard-coded, and each parameter is characterized—the voting threshold by a sensitivity sweep, the failure modes by an ABORT breakdown, and the exploration geometry by an efficiency estimate.

Why bounded retry? Two simpler policies bound the design. Immediately aborting on the first failed vote discards potentially valid viewpoints and lowers the DONE rate, while accepting an unlimited number of unsuccessful retries risks stalling the mission indefinitely on a persistent non-target. The bounded-retry policy strikes a balance between these extremes—recovering recoverable cases while guaranteeing progress—and the voting threshold τ additionally gives an operator a single, interpretable knob for the sensitivity/specificity trade-off. The same bounded-retry-and-vote pattern transfers to other ASV tasks with staged perception, where a cheap proposal precedes an expensive confirmation.

The design, however, is accompanied by the following limitations. First, validation to date is in simulation, and the voting sweep is calibrated with a synthetic verifier rather than measured per-frame detector outputs, which will become available from experiments on water. Second, the coverage framework rests on design assumptions (r , v , T_{scan}) that will have to be estimated in the field later, so the reported rates are characterizations of the *design* rather than of deployed performance. Finally, the assumption of multi-view independence in voting is not valid in reality: the three cameras observe

57 pt
0.792 in
20.1 mm

57 pt
0.792 in
20.1 mm

71 pt
0.986 in
25 mm

correlated viewpoints, so their errors are not independent—a deceptive appearance that survives at one bow camera may also survive at an adjacent one. Independence therefore likely *overstates* the benefit of voting, and the true operating point may be less favorable than the synthetic sweep suggests.

Planned future work will address these shortcomings: on-water mission trials (planned for summer 2026) will exercise sensing and dynamics end to end; measured coverage will be compared against the analytical prediction; and measured detector statistics, together with an estimate of the inter-camera correlation, will be incorporated into the voting model so that the threshold can be tuned to the true sensitivity/specificity operating point.

VIII. CONCLUSION

We presented a two-loop FSM controller that detects, investigates, and verifies marine debris on a resource-constrained ASV, where the central design principle is a bounded-retry verification stage that distinguishes a confirmed detection (DONE) from a verification rejection (ABORT). The design was characterized through a voting-threshold sweep, an analysis of the ABORT failure mode, and a coverage-efficiency framework; its full pipeline was validated in a 46-trial dry run that yielded 35 DONE and 11 ABORT outcomes, where every ABORT was caused by verification rejection rather than a navigation or transition failure. By capturing the perception uncertainty introduced by a staged detection process within the discrete decision-making layer, the architecture provides a validated basis for on-water testing, which is the immediate next phase of this work.

REFERENCES

- [1] J. R. Jambeck *et al.*, “Plastic waste inputs from land into the ocean,” *Science*, vol. 347, no. 6223, pp. 768–771, 2015.
- [2] A. L. Andrady, “Microplastics in the marine environment,” *Mar. Pollut. Bull.*, vol. 62, no. 8, pp. 1596–1605, 2011.
- [3] J. Li, J. Chavez-Galaviz, and N. Mahmoudian, “Vision driven trailer loading for autonomous surface vehicles in dynamic environments,” *Front. Robot. AI*, vol. 12, art. 1607676, 2025.
- [4] A. Temilolorun and Y. Singh, “Towards design and development of a low-cost unmanned surface vehicle for aquaculture water quality monitoring in shallow water environments,” arXiv:2410.09513, 2024.
- [5] T. I. Fossen, *Handbook of Marine Craft Hydrodynamics and Motion Control*. Chichester, U.K.: Wiley, 2011.
- [6] E. Richards, B. Thapa, and L. Mashayekhy, “Edge-enabled collaborative object detection for real-time multi-vehicle perception,” arXiv:2506.06474, 2025.
- [7] L. Wang *et al.*, “Multi-view fusion-based 3D object detection for robot indoor scene perception,” *Sensors*, vol. 19, no. 19, art. 4092, 2019.
- [8] E. Galceran and M. Carreras, “A survey on coverage path planning for robotics,” *Robot. Auton. Syst.*, vol. 61, no. 12, pp. 1258–1276, 2013.
- [9] M. Colledanchise and P. Ögren, *Behavior Trees in Robotics and AI: An Introduction*. Boca Raton, FL, USA: CRC Press, 2018.
- [10] S. K. R. Sudha, D. Mishra, and I. Hameed, “A coverage path planning approach for environmental monitoring using an unmanned surface vehicle,” *Ocean Eng.*, vol. 310, art. 118645, 2024.
- [11] L. Zhao and Y. Bai, “Energy efficient coverage path planning for USV-assisted inland bathymetry under current effects: an analysis on sweep direction,” *Ocean Eng.*, vol. 305, art. 117910, 2024.
- [12] J. Abanes, H. Jang, B. Erkinov, J. Awadalla, and A. Tzes, “ATRON: Autonomous trash retrieval for oceanic neatness,” *Front. Robot. AI*, vol. 12, art. 1718177, 2026.

- [13] J. P. Q. Tomas, J. E. E. Tupas, M. T. Soniel, C. H. M. E. Caruz, and D. B. Babar, “Real-time detection of floating debris in waterways using YOLOv8,” in *Proc. 14th Int. Workshop Comput. Sci. Eng. (WCSE)*, Phuket, Thailand, 2024.
- [14] G. Jocher, A. Chaurasia, and J. Qiu, “Ultralytics YOLO (version 8.x),” 2023. [Online]. Available: <https://github.com/ultralytics/ultralytics>
- [15] A. Rajagopalan, “Deployment-based hard-negative stress test of marine debris detection using water-dependent negative examples for resource-constrained ASVs,” in *Proc. 2026 IEEE Canadian Atlantic Ocean Symp. (CAOS)*, Halifax, NS, Canada, 2026, to be published.
- [16] A. Rajagopalan, “Ocean Voyager: A staged thermal-to-vision perception architecture for autonomous marine debris monitoring,” in *Proc. OCEANS 2026 MTS/IEEE Monterey*, Monterey, CA, USA, 2026, to be published.