Plastic Finder Automated Underwater plastic detection system using machine learning and computer vision

Imteyaz shahzad[1], Atufa Afreen[2], Punam Rathod[3], Sabiha Ansari[4], Samiksha Bambode[5]
Kamran Hossain[6]
[1,2,3,4,5,6]*Department. of Computer Science and Engineering Anjuman College of Engineering and Technology, Nagpur*

***Abstract*—Plastic pollution in marine ecosystems has become one of the most pressing environmental concerns of our era. Annual estimates suggest that close to eight million metric tons of plastic waste find their way into the world's oceans, and a significant portion—around 70%—eventually settles on the seabed. Traditional detection methods such as diver-led inspections and trawl-based sampling are resource-intensive, spatially limited, and unable to scale with the magnitude of the problem. This study surveys recent progress in applying deep learning to the automated identification of submerged plastic materials. Through a critical review of eleven peer-reviewed studies published between 2023 and 2025, we assess eight object detection architectures—covering multiple generations of YOLO and Faster R-CNN—using benchmark datasets including TrashCAN 1.0 and DeepTrash. Findings indicate that YOLOv9 and YOLOv12l deliver the strongest detection accuracy, while YOLOv10-s strikes the most favorable trade-off between precision and computational cost. Water turbidity is identified as the principal environmental constraint, with detection performance deteriorating significantly beyond 100 NTU. We conclude by outlining a proposed unified research framework to close existing gaps in microplastic detection and turbidity-aware modeling.**

***Index Terms*—Underwater Plastic Detection; Deep Learning; YOLO; CNN; Marine Pollution; Edge Computing; ROV; TrashCAN**

I. INTRODUCTION

Plastic contamination of aquatic systems has rapidly escalated into a global ecological crisis. Evidence consistently shows that the majority of plastics reaching marine environments originate from land-based human activities, transported into coastal and open-ocean zones via river networks. Once submerged, these materials can persist for hundreds of years, breaking apart into progressively smaller fragments that threaten marine life through entanglement, ingestion, and toxic accumulation along food chains. Human populations are not immune to these effects; microplastics have been identified in seafood intended for consumption, raising significant public health concerns.

Despite growing public and scientific awareness, the spatial distribution and cumulative scale of underwater plastic deposits remain poorly quantified. Existing monitoring methods—including manta trawls, grab samplers, and diver surveys—are restricted to accessible, shallow areas and produce data with limited spatial and temporal coverage. While Remotely Operated Vehicles (ROVs) and Autonomous Underwater Vehicles (AUVs) have extended monitoring capabilities to greater depths, their utility is heavily dependent on the sophistication of the perception systems they carry onboard.

Over the past decade, computer vision and machine learning have transformed the field of object detection. The YOLO (You Only Look Once) family of neural architectures, initially introduced in 2015, has undergone twelve successive generations of development, each bringing meaningful enhancements in detection accuracy, processing speed, and operational efficiency. Convolutional approaches such as Faster R-CNN have similarly matured into high-precision tools suited for offline analysis workflows. Applying these advances to underwater plastic detection represents a well-timed intersection of environmental need and machine learning capability.

The underwater imaging environment, however, introduces a distinct set of difficulties that do not arise in land-based vision tasks. Optical properties of water cause light to refract, scatter, and absorb differentially, degrading image contrast, altering perceived color, and limiting effective viewing range. Turbidity—driven by the concentration of particles suspended in the water column—compounds these effects and fluctuates significantly across geographic regions and seasonal conditions. Plastic debris itself presents a particularly challenging detection target due to its wide variation in size, shape, color, translucency, and the tendency for it to blend into surrounding substrate or biological material.

This paper surveys the landscape of deep learning approaches to underwater plastic detection, synthesizing evidence from eleven research contributions published between 2023 and 2025. Four central questions guide this review: Which model architectures perform best against standardized underwater debris datasets? How does performance change as a function of model complexity and computational demand? Which environmental conditions most severely constrain detection reliability? And what gaps remain between what current research has achieved and what real-world operational deployment requires?

II. PROBLEM CONTEXT AND MOTIVATION

2.1 Scale and Distribution of Underwater Plastic Pollution

Global plastic production has grown dramatically—from approximately two million metric tons annually in 1950 to well over four hundred million metric tons in recent years. A substantial share of this output enters aquatic ecosystems through inadequate waste management systems, industrial effluents, and wind-driven surface transport. Conservative estimates of the total plastic mass currently residing in the world's oceans range between 75 and 200 million metric tons, distributed throughout the water column from the ocean surface down to hadal trenches exceeding ten kilometers depth.

This plastic burden is distributed unevenly. Buoyant macro-sized debris tends to accumulate in large oceanic convergence zones such as the North Pacific Subtropical Gyre, colloquially known as the Great Pacific Garbage Patch, while denser materials sink and accumulate on the seafloor. The uppermost 200-meter layer of the ocean—the epipelagic zone—hosts the highest concentrations of floating debris and represents the primary operational environment for ROV-based detection platforms.

2.2 Limitations of Conventional Monitoring

Current operational monitoring relies primarily on three methods: visual shipboard surveys, net trawls, and aerial or satellite remote sensing. Shipboard visual surveys are labor-intensive and susceptible to observer bias and fatigue. Net trawling provides only snapshot data at discrete locations and can cause collateral damage to sensitive benthic habitats. Remote sensing is fundamentally limited to the ocean surface and therefore misses the large fraction of plastic that has sunk below it. ROVs equipped with standard video cameras can reach deep-sea debris, but the subsequent manual review of video footage by trained analysts is both time-consuming and difficult to scale. An automated, real-time detection system integrated into an ROV's onboard processing pipeline would fundamentally improve the utility of such platforms for both environmental monitoring and targeted cleanup operations.

2.3 Why Deep Learning?

Classical image analysis pipelines depend on manually engineered feature descriptors—such as edge detectors, color histograms, and texture-based co-occurrence matrices—that tend to break down under the variable illumination conditions, turbidity fluctuations, color distortions, and complex background textures characteristic of underwater imagery. Deep convolutional neural networks, by contrast, learn feature representations directly from training data, enabling them to naturally adapt to the statistical properties of underwater images in ways that manual feature engineering cannot replicate. Furthermore, transfer learning—where a network pre-trained on a large-scale dataset such as ImageNet or COCO is subsequently fine-tuned on domain-specific data—allows high-quality detectors to be built even when labeled underwater imagery is scarce. These properties make deep learning the most suitable and practical methodological choice for automated underwater plastic detection.

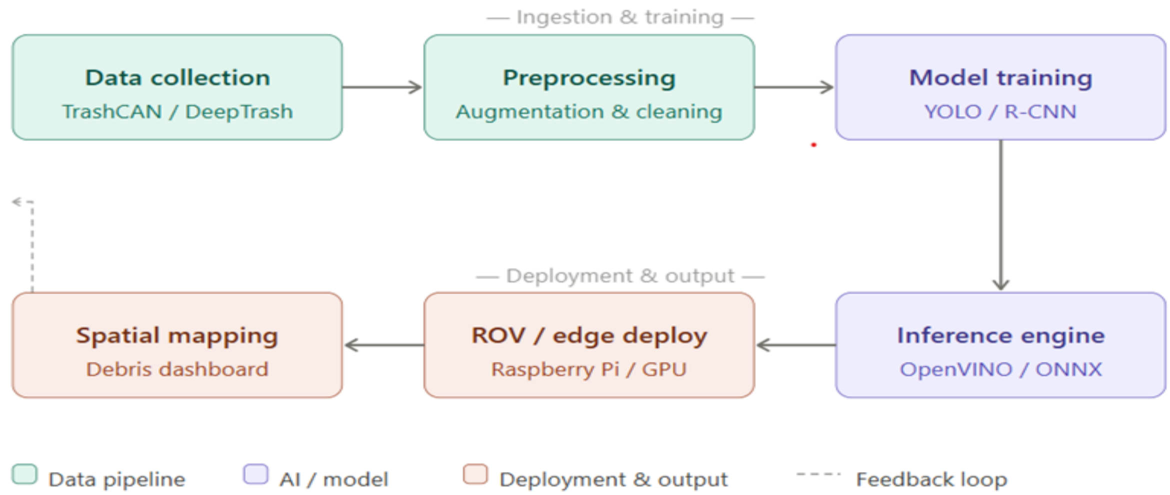

III. DEEP LEARNING ARCHITECTURES FOR OBJECT DETECTION

3.1 The YOLO Family

YOLO architectures reframe object detection as a unified regression task, simultaneously predicting bounding box locations and class probabilities from a full image in a single forward pass through the network. By eliminating the separate region proposal stage required by two-stage detectors, YOLO achieves substantially faster inference while maintaining highly competitive accuracy. The architecture has been refined across twelve generations, with each version introducing targeted improvements to the backbone feature extractor, the neck used for multi-scale feature aggregation, and the detection head.

YOLOv4, released in 2020, paired a CSPDarknet53 backbone with PANet-based feature fusion and incorporated advanced training strategies including mosaic augmentation and self-adversarial training. On epipelagic plastic datasets, YOLOv4-based models typically achieve mAP values in the 80–85% range. YOLOv5 introduced anchor-free prediction support and a modular size-scaling scheme, achieving equivalent or better accuracy at lower computational expense. YOLOv7 added extended efficient layer aggregation networks (E-ELAN) and model reparameterization to maximize accuracy per parameter; one study deploying YOLOv7 via a Flask-based web interface for marine debris monitoring reported a detection accuracy of 0.91 on the TrashCAN benchmark. YOLOv8 advanced the architecture further through a decoupled detection head and fully anchor-free regression, yielding 87% mAP on underwater plastic datasets.

YOLOv9 introduced two key architectural innovations: Programmable Gradient Information (PGI) and Generalized Efficient Layer Aggregation Networks (GELAN), both designed to mitigate information loss during deep network backpropagation. These mechanisms produced notably stable training dynamics and top-tier performance across several debris detection benchmarks. The most recent generation evaluated in this review, YOLOv12l, integrates large-kernel convolutions with attention mechanisms in a lightweight architecture engineered for edge-based deployment. It achieves a 90.2% mAP with real-time inference at ten frames per second on Raspberry Pi hardware.

3.2 Region-Based Convolutional Neural Networks

Faster R-CNN represents the primary two-stage detection alternative examined in the reviewed literature. Unlike YOLO's single-pass design, Faster R-CNN first uses a Region Proposal Network (RPN) to generate candidate object locations, which are then passed to a separate classification and refinement head that operates on features extracted by the backbone. This two-stage pipeline introduces greater inference latency but can offer improved localization precision for small or closely spaced objects. On TrashCAN evaluations, Faster R-CNN achieves mAP values between 81% and 89%, placing it competitively against YOLOv8 but short of YOLOv9. At 13 FPS on a GPU, Faster R-CNN does not meet the throughput requirements for real-time onboard ROV operation, though it remains a useful tool for offline post-processing scenarios where processing speed is less critical.

3.3 Single Shot Detector and Other Architectures

The Single Shot MultiBox Detector (SSD) performs detection at multiple feature map scales within a single unified network, avoiding the need for explicit feature pyramid construction. On underwater debris benchmarks, SSD achieves mAP values of approximately 84–87%, positioning it between the performance bands of YOLOv8 and Faster R-CNN. SSD training can be less stable than YOLO-family models, often characterized by erratic loss behavior that complicates hyperparameter selection. Edge-optimized YOLOv8 implementations using OpenVINO have been shown to exceed SSD-level accuracy on Raspberry Pi hardware, demonstrating that lightweight YOLO configurations offer a more practical path to embedded deployment.

IV. DATASETS AND EVALUATION METHODOLOGY

4.1 Primary Benchmark Datasets

The TrashCAN 1.0 dataset, developed by Hong, Fulton, and Sattar, contains 7,212 annotated images of underwater debris sourced from the JAMSTEC E-Library of Deep-Sea Images maintained by the Japan Agency for Marine-Earth Science and Technology. Images are labeled with both instance segmentation masks and bounding boxes covering 21 categories that span trash items, marine fauna, vegetation, and ROV hardware. TrashCAN serves as the primary training and evaluation benchmark across most of the reviewed studies.

The DeepTrash dataset was assembled from field video recordings of marine plastic captured across multiple Californian locations, including South Lake Tahoe, Bodega Bay, and San Francisco Bay. It contains 2,041 annotated images organized into three classes—plastic, biological material (animals and plants), and shipwrecks—with deliberate variation in image quality, water depth, and visibility levels to simulate real-world deployment conditions. A third dataset, hosted via Roboflow and derived from the JAMSTEC Deep-Sea Debris Database, comprises 1,200 annotated images developed specifically to enable comparative evaluation of YOLOv8 and YOLOv10 architectures.

4.2 Evaluation Metrics

Performance across reviewed studies is quantified using three primary metrics. Mean Average Precision (mAP) integrates precision and recall across all object categories at specified Intersection over Union (IoU) thresholds; mAP50 uses an IoU cutoff of 0.5, while mAP50-95 averages across thresholds between 0.5 and 0.95. Precision captures the fraction of predicted detections that correspond to genuine objects, and Recall measures the proportion of true objects that the model successfully identifies. F1 score combines both into a single balanced indicator. Computational efficiency is reported as inference speed in frames per second (FPS) and as floating-point operations per inference (GFLOPs), both of which are particularly relevant when assessing suitability for edge deployment.

V. COMPARATIVE PERFORMANCE ANALYSIS

5.1 Architecture-Level Comparison on TrashCAN

Table I below presents a consolidated comparison of major detection architectures evaluated on the TrashCAN dataset, drawn from controlled experiments with consistent configurations (NVIDIA Tesla T4, Adam optimizer, learning rate 0.001, batch size 16, input resolution 640×640)

Architecture	Precision (%)	Recall (%)	mAP50 (%)	mAP50-95 (%)
Faster R-CNN	86.9	84.5	88.9	69.0
SSD	86.6	84.0	87.4	68.4
YOLOv4	83.5	80.1	81.9	—
YOLOv5	85.2	82.3	85.0	—
YOLOv7	70.2	75.8	60.6	—
YOLOv8	87.1	84.6	89.0	70.1
YOLOv9	87.8	85.0	90.1	70.2
YOLOv12l	91.8	90.5	90.2	—

YOLOv9 consistently ranks highest on the TrashCAN benchmark in terms of mAP50, reaching 90.1% versus 89.0% for YOLOv8 and 88.9% for Faster R-CNN.
This performance advantage is most visible on challenging categories, including small items such as snack wrappers and nets, and partially obscured objects. YOLOv9 also demonstrated the most stable training behavior, as evidenced by smoother loss curves, an attribute linked directly to its PGI and GELAN components. YOLOv12l recorded the highest precision value in the entire reviewed literature (91.8%), with an overall mAP of 90.2% across four material categories.

5.2 Model Size vs. Accuracy Trade-off
The YOLOv8/YOLOv10 comparison study yields a counterintuitive observation: larger model variants do not reliably surpass smaller ones on underwater plastic detection tasks. YOLOv10-s achieves the best mAP50 (77.2%) among four tested variants while consuming only 24.4 GFLOPs—less than one-sixth the computational load of YOLOv10-l at 166 GFLOPs. Both YOLOv8-l and YOLOv10-l underperform their smaller counterparts, pointing either to overfitting when training data is limited, or to the possibility that the distinctive characteristics of underwater imagery do not benefit from the broader receptive fields that deeper networks develop. For ROV and embedded deployment contexts where processing resources are constrained, compact, purpose-built architectures may be more effective than simply scaling model capacity.

5.3 Deployment Scenario Performance
Three distinct deployment scenarios have been characterized in the reviewed literature. GPU-accelerated workstation deployment using YOLOv9 achieves 85–90% mAP at 30–60 FPS, representing the upper performance ceiling for offline or near-real-time analytical workflows. Raspberry Pi 4 edge deployment with YOLOv12l and OpenVINO optimization reaches 10 FPS at 90.2% mAP, sufficient for real-time monitoring on a slowly-moving ROV platform. ROUV deployment using YOLOv3 at the 100 NTU turbidity boundary achieves a mean confidence score of 77%, defining the practical operational limit for vision-based detection in turbid water conditions.

Platform	Model	Accuracy	Speed	Cost
GPU Workstation	YOLOv9	85–90%	30–60 FPS	High
Raspberry Pi 4	YOLOv12l	90.2%	10 FPS	Low
ROUV (100 NTU)	YOLOv3	77% conf.	Real-time	Medium

TABLE II. Hardware Deployment Scenarios

VI. HARDWARE PLATFORMS AND SYSTEM INTEGRATION

6.1 Edge Computing on Raspberry Pi

The Raspberry Pi 4—featuring a quad-core Cortex-A72 processor and 4 GB of RAM—has become the most commonly used embedded platform for edge-deployed underwater detection systems in the reviewed literature. Its combination of adequate processing power, small physical footprint, low energy consumption, and broad software ecosystem makes it well-suited for integration into ROV payload packages. Two optimization approaches have been documented: deploying YOLOv5 with OpenVINO for shape-based plastic classification at 85% accuracy, and running YOLOv12l with ONNX export and post-training quantization at 10 FPS using Camera Module 2 input. The OpenVINO toolchain converts PyTorch model weights into a format optimized for ARM CPU inference, while quantization from 32-bit floating point to 8-bit integer precision further reduces memory and compute demands with only minimal accuracy degradation. Together, these techniques enable a 90%-accurate detector to operate in real time on hardware available for under 100 USD.

6.2 Remotely Operated Underwater Vehicle Integration

One reviewed study documents a complete hardware integration pipeline for operational ROV deployment using a YOLOv3-based detection system. The vehicle chassis is constructed from PVC pipe to achieve a favorable buoyancy-to-weight ratio, and carries a downward-facing camera for seafloor plastic detection, vertical and horizontal thruster arrays for propulsion and maneuvering, and a Raspberry Pi connected to the camera via micro-USB. Measured power draw totals 3A across all components (1.8A for thrusters, 1.2A for the Raspberry Pi), meaning a 3000 mAh battery pack yields approximately one hour of continuous operation. Experiments conducted in a controlled pond environment with artificially induced turbidity showed that the YOLOv3 model maintained reliable detection performance—mean confidence of 77%, with all 30 test objects successfully identified across varied colors—at turbidity levels below 100 NTU. Once turbidity exceeded 120 NTU, light scattering became too severe for image-based detection to function reliably.

6.3 GPS Integration and Spatial Debris Mapping

The YOLOv12l deployment study introduces an important operational capability: GPS-based geolocation tagging of detection events through a UART-connected GPS module integrated into the Raspberry Pi payload. Each detection record is stored in a SQLite database, capturing the object class label, confidence score, bounding box coordinates, GPS latitude and longitude, and a timestamp. Once the vehicle surfaces, this data is synchronized to a cloud server over Wi-Fi or a 4G mobile connection. A web-based visualization dashboard then maps debris density against geographic coordinates, generating spatial intelligence that can be used to guide cleanup operations and track long-term changes in pollution distribution.

VII. PERSISTENT CHALLENGES AND RESEARCH GAPS

7.1 Turbidity and Variable Illumination

Water turbidity remains the most impactful and consistently identified environmental barrier to reliable underwater plastic detection. Multiple independent studies using different architectures converge on a threshold of approximately 100 NTU, above which detection performance begins to deteriorate substantially. At this level, suspended particles scatter incoming light strongly enough to reduce image contrast below the threshold at which convolutional feature extraction can operate effectively. Beyond 120 NTU, visual detection becomes practically impossible regardless of the model employed. Current architectures are not designed to self-adapt to turbidity changes encountered during deployment; models trained on clear-water imagery will not automatically generalize to murky conditions. Developing preprocessing pipelines that can dynamically correct for turbidity-induced color shifts and contrast degradation represents an underexplored but promising area of future work.

7.2 Microplastic Detection

All studies reviewed in this survey focus exclusively on macroplastic—items measuring centimeters or larger that are resolvable by standard underwater cameras at practical operating distances. Microplastics, defined as particles smaller than 5 mm, constitute the majority of plastic particles in ocean environments by count, and they present a fundamentally different detection challenge. Effective automated microplastic quantification in situ would require specialized sensing modalities—such as high-resolution microscopy-grade optics, hyperspectral imaging, or fluorescence spectroscopy—that go beyond the capabilities of conventional cameras. No study in the reviewed literature demonstrates reliable automated microplastic detection under in-situ underwater conditions.

7.3 Dataset Diversity and Standardization

The absence of a single, universally adopted benchmark dataset continues to impede rigorous cross-study comparison. Studies reviewed here use a variety of datasets with differing class taxonomies, annotation conventions, image resolutions, and geographic origins. While TrashCAN and DeepTrash provide a degree of standardization, the field still lacks a comprehensive benchmark covering the full operational spectrum—including coastal, riverine, deep-sea, high-turbidity, and low-light conditions. Most available labeled datasets remain small (typically fewer than 10,000 images), limiting the degree to which trained models generalize when deployed in environments that differ from training conditions.

VIII. DIRECTIONS FOR FUTURE RESEARCH

Based on findings across the reviewed literature, six priority research directions are identified for advancing the field. First, turbidity-adaptive architectures should be developed that incorporate real-time water quality sensor readings as auxiliary inputs, enabling dynamic adjustment of detection

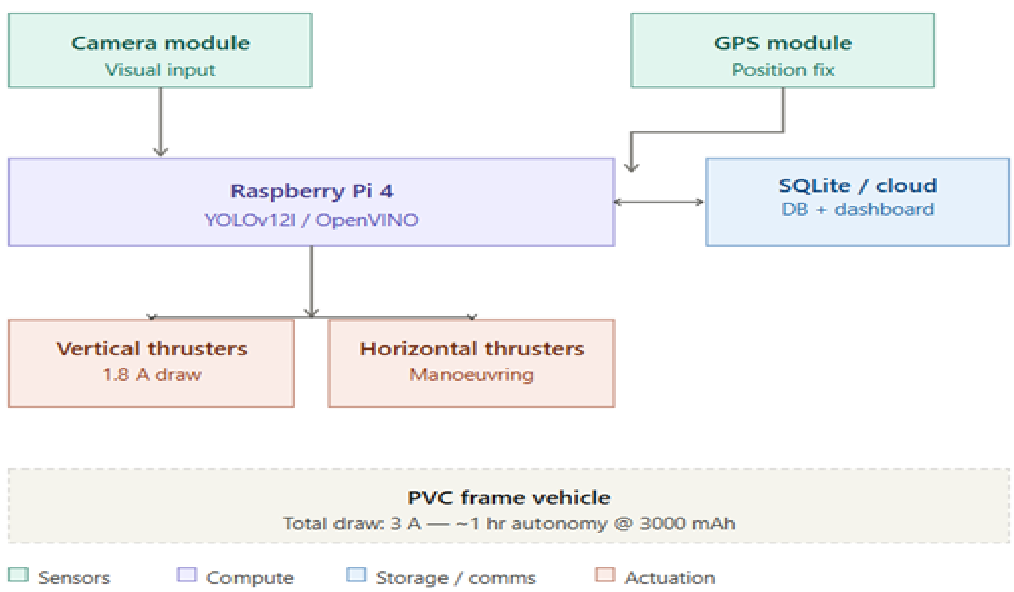

thresholds and preprocessing operations in response to measured turbidity, color temperature, and particle concentration. Synthetic data generation for high-turbidity conditions could support transfer learning pipelines and accelerate adaptation to low-visibility environments.

Second, multi-modal sensor fusion combining RGB cameras with forward-looking sonar for structure detection in zero-visibility conditions, hyperspectral sensors for material-level classification, and acoustic Doppler systems for tracking debris motion would substantially widen the operational envelope of automated detection systems. Feature-level fusion architectures, as opposed to simple decision-level combination, represent a particularly promising integration approach.

Third, a community-led benchmark dataset initiative modeled on COCO or ImageNet should be established for underwater debris—targeting at least fifty thousand annotated images spanning diverse environments, turbidity levels, and debris categories, with standardized splits enabling reproducible crosssystem evaluation. Fourth, real-time adaptive learning algorithms on edge devices could enable continuous model updating during long-duration deployments as domain characteristics shift. Fifth, extending from static frame analysis to spatiotemporal video analysis using recurrent or transformer-based temporal attention mechanisms would enable debris tracking across frames. Sixth, lightweight backbone architectures trained from scratch on domain-specific underwater corpora may yield accuracy improvements by eliminating the inherent domain mismatch that accompanies transfer learning from terrestrial image datasets.

IX. CONCLUSION

This survey has systematically reviewed the current state of deep learning-based automated underwater plastic detection, drawing on eleven studies published between 2023 and 2025. The results indicate a field that has achieved substantial technical maturity under controlled and shallow-water conditions, with leading architectures surpassing 90% mAP on standardized benchmarks and demonstrating viable real-time operation on low-cost edge hardware at ten frames per second. The superiority of newer YOLO-family architectures—particularly YOLOv9 and YOLOv12l—over earlier YOLO generations, Faster R-CNN, and SSD is well-established. YOLOv9's innovations in gradient flow management and feature aggregation deliver the most stable training dynamics and highest benchmark accuracy among the evaluated architectures. Notably, smaller model variants frequently matched or exceeded the performance of larger ones, underscoring that task-specific optimization carries more weight than raw model capacity in this application domain.

The most substantial barrier to real-world deployment remains water turbidity. The convergent finding—across multiple independent studies—that detection reliability degrades sharply above approximately 100 NTU identifies a critical constraint that architectural advances alone cannot overcome. Addressing it will require turbidity-adaptive preprocessing, multi-modal sensor integration, and training data that spans the full turbidity range of intended deployment environments. The ecological urgency motivating this research is clear: scaling automated underwater plastic detection is a prerequisite for meaningful ocean monitoring and remediation. The technical foundations documented in this survey confirm that practical systems are achievable with currently available technology; realizing this potential requires coordinated progress in standardized benchmarking, turbidity-adaptive methods, and multi-modal sensing architectures.

REFERENCES

[1] Chetana J., Jayashree R., Kamini M., and Siri D.H. (2024). Harnessing deep learning for underwater plastic trash identification. International Research Journal of Engineering and Technology (IRJET), 11(4), 2401–2406.

[2] Akshitha M., Yashwanth G., Srinath Reddy B., Boyapally R.R., Afzal M., and Ramesh M. (2025). Efficient deep learning-based marine plastic waste detection with YOLOv12. International Journal of Innovative Research in Technology (IJIRT), 11(12), 6465–6472.

[3] Rakshitha P.V. and Meghana B.S. (2024). Plastic identification underwater. International Advanced Research Journal in Science, Engineering and Technology (IARJSET), 11(8), 97–100.

[4] Khriss A., Elmiad A.K., Badaoui M., Barkaoui A.E., and Zarhloule Y. (2024). Exploring deep learning for underwater plastic debris detection and monitoring. Journal of Ecological Engineering, 25(7), 58–69.

[5] Delina M., Taryudi, Amin M.F., et al. (2024). The development of remotely operated underwater vehicle for plastic waste detection with Raspberry Pi. Journal of Physics: Conference Series, 2866(1), 012048.

[6] Balaga B.S. (2023). Detection of plastic in water using image processing. International Journal of Research Publication and Reviews, 4(12), 4799–4807.

[7] Yadav S., Singh H., Nishad N., and Shruthiba A. (2024). Underwater waste detection using YOLOv5. International Journal of Engineering Research in Computer Science and Engineering (IJERCSE), 11(7), 78–83.

[8] Aminurrashid A.R. and Sayuti M.N.S.M. (2024). A CNN plastic detection model for embedded platform of ROV. ITM Web of Conferences, 63, 01003.

[9] Sairamesh L., Selvakumar K., and Sindhu T. (2025). Deep learning models for trash detection in underwater through image processing. Global NEST Journal, 27(4), 06077.

[10] Madhavi A.T., Shriram K., Yugesh B., and Sanjay S. (2024). Plastic detection in underwater using image processing. International Journal of All Research Education and Scientific Methods (IJARESM), 12(4), 360–369.

[11] Bogisam S.S. (2024). Underwater plastic detection using YOLO v8 and v10 (MSc Research Project). National College of Ireland, School of Computing.

[12] Fulton M., Hong J., Islam M.J., and Sattar J. (2019). Robotic detection of marine litter using deep visual detection models. Proceedings of the IEEE ICRA, 5752–5758.

[13] Hong J., Fulton M.S., and Sattar J. (2020). TrashCan 1.0: An instance-segmentation labeled dataset of trash observations. arXiv preprint arXiv:2007.08097.

[14] Wang C.Y., Yeh I.H., and Liao H.Y.M. (2024). YOLOv9: Learning what you want to learn using programmable gradient information. arXiv preprint arXiv:2402.13616.

[15] Tata G., Royer S.J., Poirion O., and Lowe J. (2021). DeepPlastic: A novel approach to detecting epipelagic bound plastic using deep visual models. arXiv preprint arXiv:2105.01882.