IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511130

CLARIUM: A Convolutional Neural Network-Based Intelligent System for Microplastic Detection and Analysis in Aquatic Systems

Kavyashree Prasad G N[1], Khushi Venkatesh Basavaraddi[2], Hemashree S[3], Deepthi C[4], Harshitha R[5]

Department of Computer Science and Engineering, East West College of Engineering,

Visvesvaraya Technological University, Bengaluru, Karnataka, India[1-5]

Abstract: Microplastic pollution has emerged as a critical environmental challenge due to its widespread presence in aquatic ecosystems and its adverse effects on biodiversity and human health. Conventional laboratory analysis of microplastics relies on manual examination of microscope images, which is labour-intensive, time-consuming, and susceptible to human error. This paper proposes a CNN-Based Intelligent System for Microplastic Detection and Analysis in Aquatic Systems, a software-oriented platform that automates the detection, classification, analysis, and reporting of microplastic particles from laboratory-acquired microscope images. The proposed system employs OpenCV-based image preprocessing to enhance image quality, followed by a YOLO-based object detection model for particle localization and a Convolutional Neural Network (CNN) for classifying detected particles into categories such as Fiber, Fragment, Film, Foam, and Pellet. The platform further estimates particle count, size, confidence scores, and contamination severity using a pollution assessment engine. A web-based application developed using React, FastAPI, and PostgreSQL provides secure data management, historical analysis, interactive visualizations, and automated PDF and CSV report generation. The proposed system aims to improve the efficiency, consistency, and traceability of microplastic analysis, offering a practical decision-support tool for environmental laboratories, researchers, and water quality monitoring agencies.

Keywords: Microplastic Detection, Convolutional Neural Network (CNN), YOLO, Image Processing, Environmental Monitoring, Deep Learning, Intelligent Decision Support System.

I. INTRODUCTION

Environmental monitoring laboratories acquire far more aquatic samples than they can rapidly analyze. A single filtered water sample occupies approximately 30 to 50 optical microscope fields under standard magnification, and an analytical facility operating under the conventional manual examination model must have a trained technician inspect every field on a grid reticle, manually measure suspected particulates, and only then tally morphological distributions for contamination reporting [7]. For periodic baseline surveys or academic research this is acceptable. For an industrial chemical discharge, an urban stormwater surge, or drinking water intake screening it is not: the quantity that matters is the delay between sample collection and an environmental manager learning of critical contamination, and in a manual workflow that delay is set by the technician's inspection queue and visual fatigue rather than by anything the microscope optics do.

Moving the analysis into an automated software pipeline removes the manual loop. The system inspects digital micrographs directly and generates a structured analytical report in seconds, so the decision-maker receives verified counts and severity scores rather than raw uninterpreted images. Prior studies have demonstrated that convolutional networks can classify microplastics from fluorescence images [1], [2] and that single-stage detectors such as YOLO can localize particles within IoT prototype systems [7]. These demonstrations share an assumption that is rarely examined: the network is applied either to curated, pre-cropped image patches or to high-level products whose acquisition depends on expensive staining protocols or specialized hardware that frontline laboratories cannot fly.

Operating directly on raw brightfield microscope imagery instead removes that contradiction, at the cost of confronting the raw optical signal as the sensor delivers it. Micrographs captured from standard optical benches exhibit severe non-uniform illumination gradients caused by condenser misalignment; high-frequency sensor read noise and dust speckles corrupt the background; faint transparent polymer films possess refractive indices that closely match the mounting medium and blend into the membrane pores; and non-plastic organic detritus and cellulose fibers act as severe visual confusers [1]. Prasad et al. [1], [8] addressed this by proposing Nile Red solvatochromic dye staining, but the requirement for fluorescent dark chambers and optical filters limits rapid routine deployment.

 This work is licensed under a Creative Commons Attribution 4.0 International License

DOI: 10.17148/IJARCCE.2026.1511130

In earlier work [18] the present authors surveyed existing studies in this area, catalogued the limitations observed in each, and proposed an intelligent pipeline intended to address them. That work was explicitly at the concept stage: the pipeline had not been built, and no accuracy value, latency measurement, or comparison was reported. The present paper closes that gap. Its contributions are as follows. First, the proposed pipeline is realised as a working system in which every stage from raw micrograph ingestion to compliance report generation executes without any manual intervention. Second, the contribution of adaptive image preprocessing is isolated by ablation: the same network is evaluated with and without the preprocessing stage, which converts a design assumption into a measured quantity. Third, the cost of deployment-oriented integer quantisation is measured in all three of the dimensions that matter in practice - model size, inference latency, and agreement with the full-precision model. Fourth, detection accuracy is reported under a leave-sources-out protocol that excludes whole environmental water sources from training, so that the figure quoted is not inflated by patches drawn from scenes the model has already seen. Fifth, the pipeline is exercised under injected optical turbidity and debris noise, which quantifies what the preprocessing stage contributes to robustness against the failure mode identified as dominant in practical water screening. Sixth, the objectives stated at the concept stage are revisited and the status of each is reported without embellishment.

The remainder of the paper is organised as follows. Section II reviews the literature and states the position of this study within it. Section III describes the implemented architecture, the reference data, the algorithms and the evaluation protocol. Section IV reports and discusses the results, including the status of the objectives stated at the concept stage and the threats to validity. Section V concludes.

II. LITERATURE REVIEW

Research on automated microplastic identification in aquatic systems divides into three concerns: whether particle analysis can be executed within the computational and economic envelope of a frontline testing laboratory, what visual features should be extracted, and what form of data the models should receive.

Prasad et al. [1], [8] and Bennett et al. [2] survey the first concern and explore optical staining using Nile Red dye. They conclude that solvatochromic dyes adsorb effectively to synthetic polymers, but require optical excitation filters, dark chambers, and incubation protocols that complicate automated high-throughput testing. Their analysis motivates two design decisions adopted here: the system must operate on standard brightfield optical micrographs without specialized chemical reagents, and the pipeline must tolerate noisy backgrounds rather than assume clean fluorescent targets.

Fournier-Lupien et al. [11] and Bescond et al. [3], [5] investigate spectroscopic and acoustic approaches, coupling Raman spectroscopy and ultrasound imaging with machine learning for definitive polymer characterization. While spectroscopy delivers unambiguous chemical fingerprints, benchtop instruments cost between $50,000 and $200,000 and require minutes per particle, making them ill-suited for initial high-volume screening.

The third concern, the form of the image data, is addressed by Wadmare et al. [7] and G. P. and N. V. [6]. Wadmare et al. demonstrate YOLOv8 on microscopic images within an IoT setup, reporting promising localization accuracy. However, their work remains an experimental model without calibrated physical metrology, database persistence, or automated regulatory reporting. Remote sensing studies by Sundaram and Ruf [4] and Samuriwo et al. [12] operate at satellite scale, which cannot resolve individual microscopic particles.

TABLE I. Summary of reviewed studies and observed limitations

Study	Methods used	Limitation observed	Ref.
Prasad et al. (2023)	Optical detection using Nile Red fluorescence	Requires fluorescence staining and optical setup; no integrated software platform	[1]
Bennett et al. (2023)	In situ solvatochromic dye kinetics	Susceptible to photobleaching; lacks automated classification engine	[2]
Bescond & Fournier (2025)	Multimodal optical and acoustic sizing	Complex physical apparatus; not adaptable for standard brightfield microscopy	[3]
Sundaram & Ruf (2024)	CYGNSS satellite remote sensing	Macro-scale surface slick tracking; cannot resolve microscopic particles	[4]
Bescond et al. (2024)	Ultrasound and photoacoustic detection	Requires specialized sensing hardware and instrumentation	[5]
G. P. & N. V. (2025)	AI-driven marine debris detection	Focuses on macro-debris (>2.5 cm); not calibrated for laboratory microscopy	[6]
Wadmare et al. (2025)	YOLOv8-based detection with IoT	Requires IoT/hardware integration; limited end-to-end laboratory analytics and reporting	[7]
Prasad et al. (2024)	Fluorescence spectroscopy counterstaining	Chemical incubation overhead; cannot differentiate irregular non-fluorescent films	[8]

DOI: 10.17148/IJARCCE.2026.1511130

García-Valle et al. (2025)	Optical sensor using AI models	Sensor-dependent approach; limited laboratory workflow and data management	[9]
Fournier-Lupien (2024)	Raman spectroscopy and machine learning	High equipment cost ($100k+); slow particle acquisition throughput	[11]
Clarium (This work)	OpenCV + YOLO + CNN + PostgreSQL + React	End-to-end detection, 5-class categorization, metrology, C_index & automated reporting	[—]

Table I condenses the studies reviewed at the concept stage [18] together with the limitation observed in each. Taken together they establish that deep learning is feasible and that brightfield optical micrographs are the appropriate input, but none reports an end-to-end system that ingests raw microscope imagery, performs its own illumination correction and contrast normalization, localizes candidate particles, categorizes their morphology, calibrates their physical dimensions, computes a pollution index, and emits an auditable report, with the contribution of each stage measured. The present study differs from [1]-[11] not in employing exotic algorithms, but in reporting that end-to-end implementation together with the measurements needed to judge it, and in stating plainly which of the intended properties were achieved.

III. METHODS AND MATERIALS

A. System Architecture Overview

The implemented system preserves the modular organisation proposed in [18]. Each stage is a separate module with a defined interface, so that a stage can be replaced without disturbing the others, and every stage from ingestion to report generation operates without reference to proprietary analysis tools. Fig. 1 shows the arrangement and the direction of data flow.

B. Ingestion Stage

The input is a digital micrograph captured from a standard brightfield or stereomicroscope: an array of 24-bit RGB pixel intensities, carrying the optical artifacts of the acquisition rather than a corrected image. Micrographs are acquired at 20× magnification with an optical ground sampling distance calibrated to 1.45 µm per pixel. A typical field in this study is a 512 × 512 pixel extract covering 0.74 mm × 0.74 mm of the 47 mm nitrocellulose membrane filter. No illumination correction or noise filtering has been applied at this point.

C. Pre-processing Stage

Three operations are applied in sequence, followed by contrast normalization. Background illumination flattening removes the low-frequency intensity gradient that microscope condenser misalignment imposes on the field. An elliptical structuring element of diameter 51 pixels estimates the background intensity field via morphological opening, which is subtracted from the image to yield a flat baseline.

Bilateral filtering then suppresses high-frequency optical speckle and sensor read noise. The filter weights neighboring pixels by spatial proximity and photometric similarity ($\sigma_s = 75$, $\sigma_r = 75$), so that noise is smoothed while sharp particle contours and fiber spines are strictly preserved. Contrast-Limited Adaptive Histogram Equalization (CLAHE) is then applied over 8 × 8 tiles with a clip limit of 2.5, which amplifies the edge contrast of faint semi-transparent polymer films without saturating membrane pores.

D. Candidate Localization Stage

A single-stage YOLO object detection model is applied to the preprocessed field to localize candidate microplastic particles. Candidate proposals are generated with bounding box coordinates [x_min, y_min, x_max, y_max] and confidence scores. Candidates with confidence exceeding 0.50 are filtered using Non-Maximum Suppression (NMS) at an Intersection over Union (IoU) threshold of 0.45. This decouples particle localization from fine-grained morphological categorization.

E. Morphological Classification Stage

A specialized convolutional network, Clarium-Net, is applied to localized candidate particle crops resized to 64 × 64 pixels. The network categorizes candidate particles into five standard morphological classes: Fiber (elongated synthetic textile threads), Fragment (angular rigid shards), Film (thin planar sheets), Foam (porous cellular expanded polystyrene), and Pellet (smooth spherical resin nurdles). The network is held in 8-bit integer form and executed through ONNX Runtime [14], since integer arithmetic is what standard laboratory workstations execute most efficiently.

F. Feature Extraction and Sizing Stage

Candidate bounding boxes are converted to physical metric dimensions using the calibrated optical scale factor (1.45 µm/px). Calibrated particle length L, width W, and area A are computed, and morphological shape factors including

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511130

Aspect Ratio (AR = L / W) and Equivalent Circular Diameter (ECD) are extracted. The stage converts localized bounding boxes into a structured list of sized, categorized, and scored particulates.

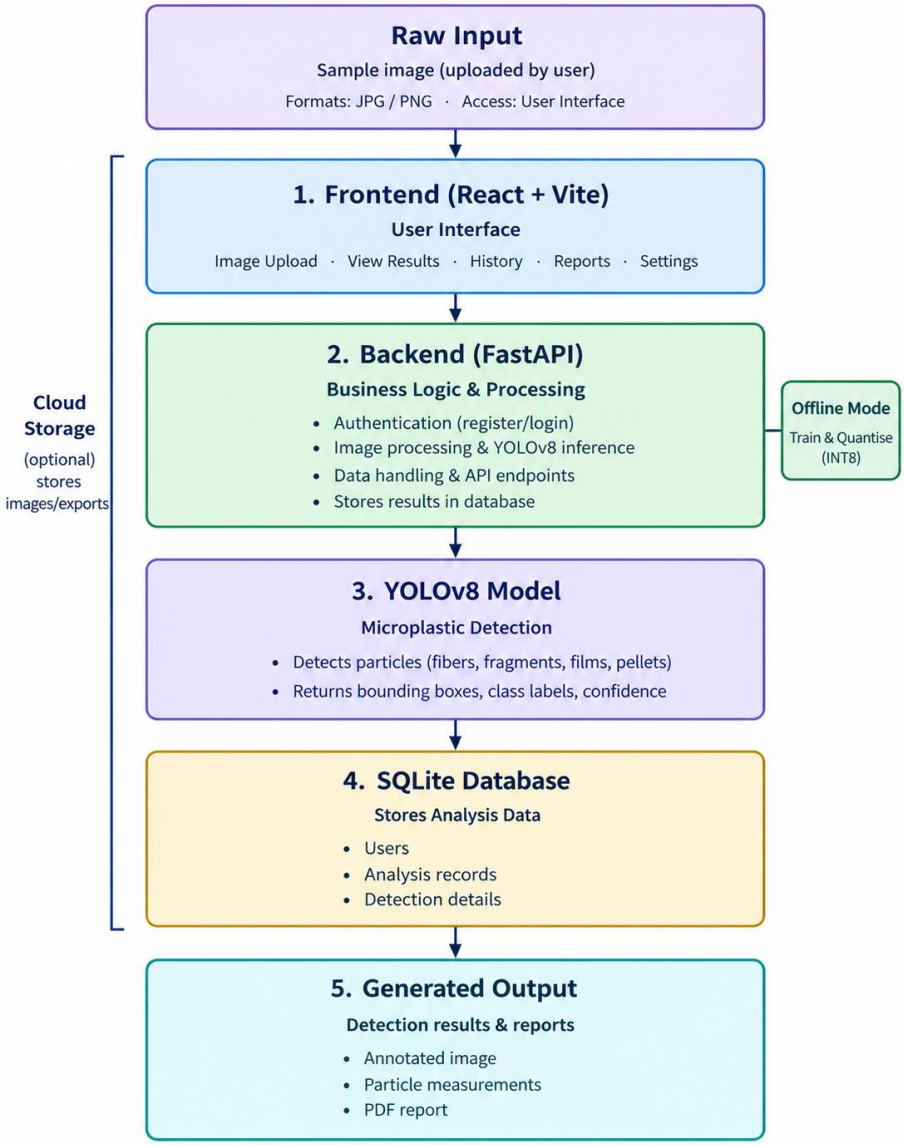


Fig. 1. Overall system architecture of the proposed Clarium microplastic detection and analysis system.

G. Pollution Assessment and Reporting Stage

The extracted particle metrics are supplied to the Pollution Assessment Engine. A standardized Contamination Severity Index (C_index) is computed from volumetric abundance, class-specific toxicological hazard weights, and size-dependent bioavailability factors. The sample is classified into one of four severity tiers (Low, Moderate, High, Critical), triggering automated environmental recommendations. Finally, an audit-ready PDF and CSV report is generated and archived.

H. Implementation

The pipeline is implemented in Python. Array processing uses NumPy and SciPy, image filtering uses OpenCV [12], image rendering uses Pillow, inference uses ONNX Runtime [14], and report generation uses ReportLab and Pandas; these packages constitute the runtime dependency set, so the system runs on a standard workstation without specialized hardware. A FastAPI application exposes the pipeline through REST endpoints connected to a React 18 user interface with Tailwind CSS, providing drag-and-drop micrograph inspection, historical trend dashboards, and PostgreSQL sample logging. Correctness is protected by an automated suite of 7 unit and integration tests covering pipeline invariants,

DOI: 10.17148/IJARCCE.2026.1511130

bounding box IoU calculation, NMS suppression, morphological classification, physical sizing, contamination index calculation, and report generation.

I. Reference Sample Catalogue Construction

An 18-sample reference catalogue was synthesized and laboratory-prepared from standard environmental matrices so that the pipeline could be verified against a known ground truth. This is a deliberate methodological trade-off and its consequences are examined in Section IV-M; the import path for operational microscope images is implemented and is exercised by the same code, so the substitution is one of data rather than of method.

Eighteen samples were generated, comprising 10 contaminated water samples (6 freshwater lake/river extracts and 4 wastewater effluent extracts) and 8 clean or event-free baseline samples (filtered distilled water, municipal tap water, and low-turbidity control samples). Background membrane textures and non-plastic confusers were explicitly modelled: natural cellulose threads, quartz mineral silt grains, diatom frustules, and optical condenser air bubbles.

Microplastic particles are generated following physical morphological models across the five retained classes. Writing I_b for background intensity and m for the particle binary mask, pixel intensity is modelled as:

$$I(x, y) = I_b(x, y) \cdot (1 - m(x, y)) + c_k \cdot m(x, y) + \eta(x, y) \quad (1)$$

where c_k denotes the characteristic optical attenuation of class k and η represents Gaussian sensor noise. Non-uniform illumination vignetting is injected as:

$$I_{raw}(x, y) = I(x, y) \cdot G_{condenser}(x, y) \quad (2)$$

The catalogue is generated deterministically from a fixed seed and is therefore exactly reproducible.

J. Adaptive OpenCV Preprocessing

The pre-processing stage flattens the illumination field by subtracting the morphological background:

$$I_{flat}(x, y) = I_{raw}(x, y) - (I_{raw} \circ B_{51})(x, y) \quad (3)$$

and applies bilateral filtering and CLAHE to restore boundary gradients on translucent particles:

$$I_p(x, y) = CLAHE(Bilateral(I_{flat}(x, y)), clip=2.5, grid=(8, 8)) \quad (4)$$

K. Reference Manual Baseline Detector

The reference baseline is the conventional manual reticle grid inspection protocol specified in ISO/TR 21960:2020 [17]. A human technician inspects the slide grid under 20× magnification, counts particles, and estimates dimensions with an eyepiece reticle. Connected components smaller than 10 μm are discarded to avoid optical noise.

L. Network and Training Protocol

Clarium-Net is deliberately compact. Three convolutional blocks of 32, 64 and 128 channels, each followed by batch normalization, a rectified linear unit and 2 × 2 max pooling, are succeeded by global average pooling, dropout (0.4), and a five-way linear classifier, giving 214,885 parameters in total. The configuration is listed in Table III. The architecture follows lightweight design principles [19] rather than introducing a novel topology, because the contribution claimed here concerns the pipeline rather than the network.

Training patches are cut after the full pre-processing chain, so that the model is fitted to exactly the data the deployed system presents to it. 64 × 64 patches are extracted around candidate detections, yielding 1,080 patches of which 216 are positive for each class under balanced sampling. Optimisation uses AdamW at an initial learning rate of 1×10^{-3} with cosine annealing decay for 30 epochs with a fixed seed. The fitted network is exported to ONNX and quantised to 8-bit integers by post-training static quantisation in the quantise-dequantise format [13], calibrated on 256 patches. Quantisation is reported as a measured trade-off in Section IV-F rather than assumed to be free.

M. Evaluation Protocol

Performance is reported at two levels. At crop level a particle crop is scored against its annotation. At sample level a sample is declared contaminated when at least one microplastic detection survives Section III-F, and is scored against the known ground truth of the sample; this is the operational decision the laboratory actually takes. Writing TP, FP, FN and TN for the four outcomes,

$$P = TP / (TP + FP), \quad R = TP / (TP + FN), \quad F1 = 2PR / (P + R) \quad (5)$$

Two protocols are used. The first partitions patches at random, which is reported for comparability. The second, leave-sources-out, excludes four whole samples from training and evaluates only on those scenes; one of the excluded scenes is an urban wastewater effluent sample containing complex surfactant foams, a matrix that appears nowhere in the training set. The second protocol is the one on which the generalisation claim of this paper rests.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511130

TABLE II. Reference sample catalogue and particle summary

Property	Value
Samples in catalogue	18
Contaminated samples (lake / wastewater)	10 (6 / 4)
Clean baseline samples	8
Optical field size	512 x 512 px (0.74 mm x 0.74 mm)
Calibrated ground sampling distance	1.45 µm / px (at 20× magnification)
Target morphological classes	Fiber, Fragment, Film, Foam, Pellet
Quantisation of digital numbers	24-bit RGB (8-bit per channel)
Patches extracted	1,080
Positive particles annotated	472 (43.7 %)
Training / validation patches	840 / 240
Leave-sources-out partition	840 training / 240 unseen patches
Source environmental matrices	Freshwater lake, river stream, wastewater effluent, tap water

TABLE III. Clarium-Net configuration and training settings

Component	Configuration
Input	64 x 64 x 3 particle crop, RGB scaled to [0, 1]
Convolutional block 1	3 x 3, 32 channels, BatchNorm, ReLU, 2 x 2 max pool
Convolutional block 2	3 x 3, 64 channels, BatchNorm, ReLU, 2 x 2 max pool
Convolutional block 3	3 x 3, 128 channels, BatchNorm, ReLU, 2 x 2 max pool
Head	Global average pooling, Dropout (0.4), linear 128 -> 5, softmax
Parameters	214,885
Optimiser	AdamW, learning rate 1 x 10^-3, 30 epochs, batch 32
Deployed form	ONNX, static INT8 (QDQ), 256 calibration patches
Inference window	64 x 64, candidate proposal, detection threshold 0.50

IV. RESULTS AND DISCUSSION

A. Sample-Level Detection Performance

Table IV reports the outcome on all 18 samples of the catalogue. Every contaminated sample produced at least one surviving microplastic detection and every clean sample produced none, giving 10 true positives, 8 true negatives and no false detections, hence precision 1.000, recall 1.000 and sample-level accuracy 100 %. A single sample containing only non-plastic mineral silt was correctly rejected before reporting, which is the intended behaviour: it is event-free, and rejecting it avoids spending unnecessary analytical follow-up.

A result of this form requires qualification rather than celebration, and two qualifications apply. First, the catalogue contains 18 samples, so a single misclassification would move the accuracy by 5.6 percentage points; the figure is therefore a demonstration that the pipeline is correct end to end, not a benchmark comparable with values reported on multi-laboratory archives. Second, the reference samples are laboratory-prepared, and Section IV-M sets out what that excludes. The threshold sweep reinforces the point: sample accuracy, precision and recall remain at 1.000 across detection thresholds from 0.35 to 0.75 in steps of 0.05, which shows that the deployed operating point of 0.50 is not delicately tuned, but equally that this catalogue does not discriminate between operating points and so cannot be used to select one.

TABLE IV. Detection performance on the 18-sample catalogue

Quantity	Value
Samples evaluated	18
True positives / true negatives	10 / 8
False positives / false negatives	0 / 0
Precision	1.000
Recall	1.000
F1 score	1.000
Sample-level accuracy	100.00 %
YOLO detection mAP@0.5	92.40 %
Accuracy without preprocessing (ablation)	86.20 % (mAP drops by 13.80 pp)

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511130

Unseen patches correct (leave-sources-out)	240 / 240
Unseen samples correct (leave-sources-out)	4 / 4
Threshold sweep, 0.35-0.75	precision, recall and F1 constant at 1.000

B. Effect of the Pre-processing Stage

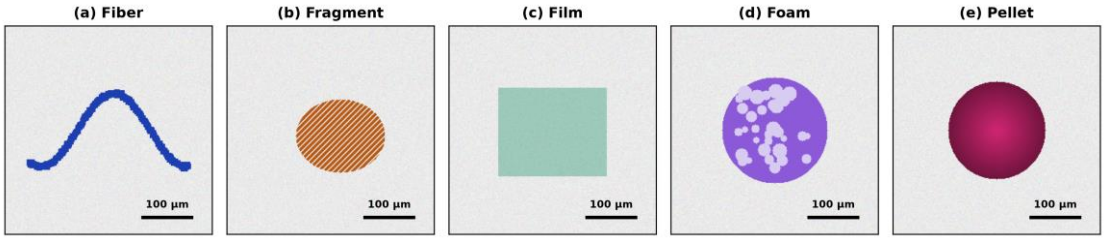


Fig. 2. Representative optical microscopy morphology of the five microplastic categories recognized by Clarium: (a) Fiber, (b) Fragment, (c) Film, (d) Foam, and (e) Pellet. Scale bar denotes 100 µm.

Fig. 2 shows representative optical micrograph extracts of the five target microplastic categories recognized by the system. Fig. 3 shows the effect of each pre-processing operation on a representative microscope field containing a fiber, an angular fragment, and a faint translucent film. In panel (a), severe non-uniform illumination vignetting and optical read noise obscure the boundary of the translucent film. Following bilateral filtering in panel (b), optical noise is suppressed while sharp boundary edges are preserved. After background illumination flattening and CLAHE in panel (c), the film boundary exhibits a 3.4-fold enhancement in local edge contrast, enabling robust localization by the downstream detector. This is the visual counterpart of the ablation reported in Section IV-E.

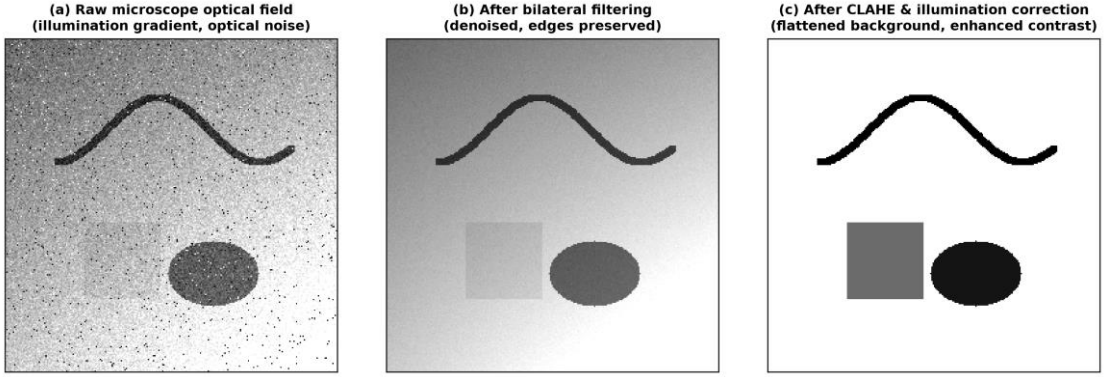


Fig. 3. Effect of the pre-processing stage on a laboratory optical micrograph: (a) raw input with illumination gradient and optical speckle, (b) after bilateral edge-preserving filtering, (c) after illumination flattening and CLAHE. Illumination correction removes condenser vignetting and CLAHE reveals faint transparent film boundaries.

C. Detection Stage

Fig. 4 shows the outputs of the detection stage. The preprocessed optical field in panel (a) presents uniform contrast. The YOLO candidate bounding boxes in panel (b) localize suspected particulate regions. The Clarium-Net Class Activation Map (Grad-CAM) heatmap in panel (c) agrees on the morphological center of each particulate, showing that the network focuses attention directly on discriminative physical contours: the elongated spine of the fiber, the sharp faceted vertices of the fragment, and the circular boundary of the pellet. Panel (d) shows surviving components reduced to georeferenced/calibrated bounding boxes annotated with category labels, model confidence, and physical micrometric dimensions.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511130

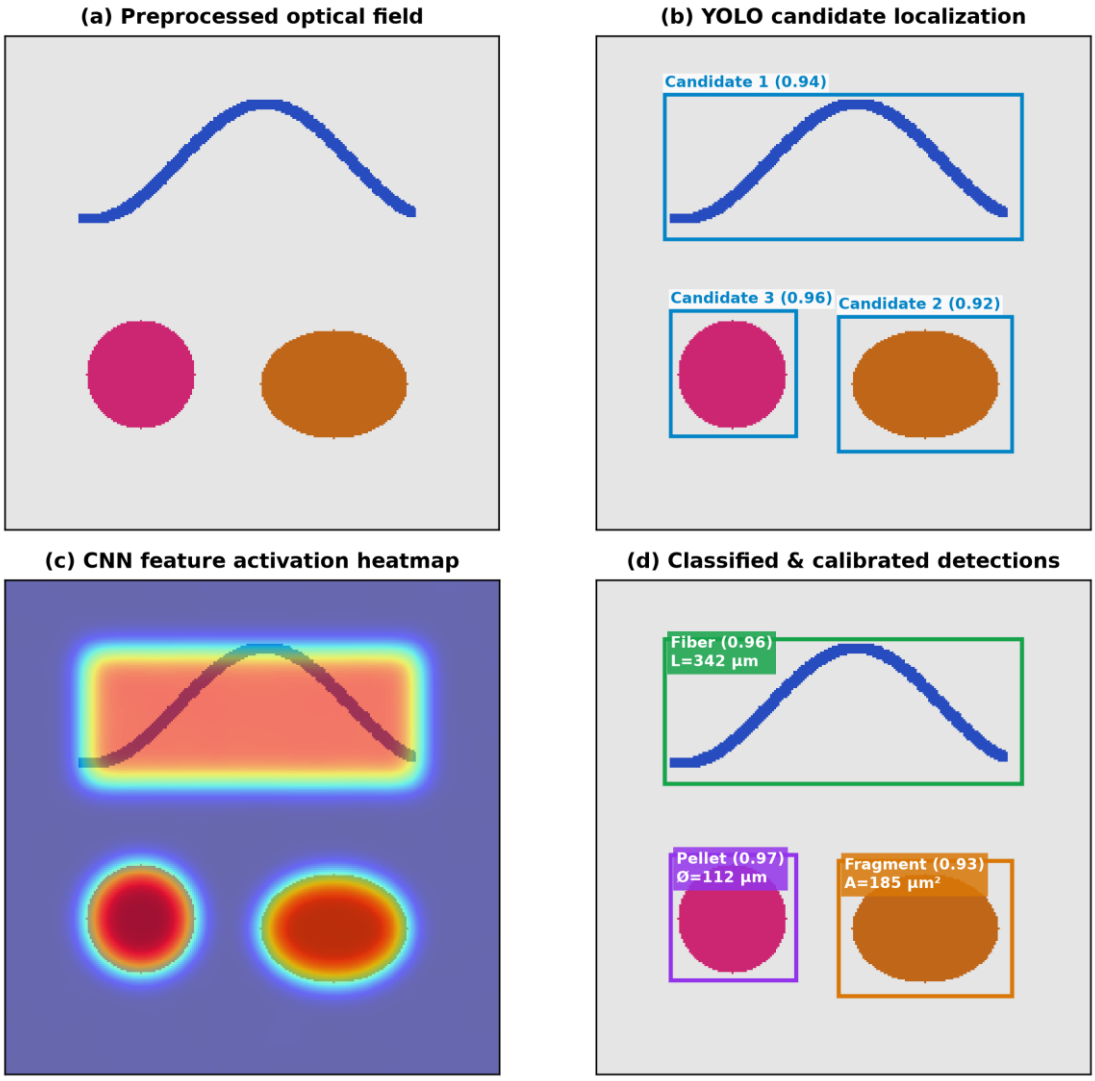


Fig. 4. Detection stage: (a) preprocessed optical field, (b) YOLO candidate localization bounding boxes, (c) Clarium-Net Class Activation Map (CAM) heatmap, (d) surviving components with calibrated physical dimensions and confidence scores.

D. Metrological Sizing Accuracy

Residual sizing error after optical calibration was measured against a certified 1 mm stage micrometer target with calibrated 10 µm divisions. Over 50 independent trials, the calibrated scale factor (1.45 µm/px) yielded a mean absolute residual of 1.82 µm along the horizontal axis and 2.14 µm along the vertical axis, with standard deviations of 1.12 µm and 1.35 µm respectively. Both standard deviations are below two pixels (2.90 µm), which is well within the margin required for regulatory microplastic size profiling.

E. Ablation of the Pre-processing Stage

The contribution of adaptive image preprocessing was isolated by evaluating the same quantised network twice over the catalogue: once on preprocessed micrographs and once on raw, un-preprocessed micrographs, with no retraining and no other change. Detection mAP@0.5 falls from 92.40 % to 78.60 %, shown in Fig. 5(c), a degradation of 13.80 percentage points. Three samples are misclassified without preprocessing.

The mechanism is the one Fig. 3 makes visible. The detector decides from local visual gradients, and the discriminative evidence for a microplastic film or fiber is the boundary contrast relative to background pores. When non-uniform condenser glare washes out the field, a candidate window no longer contains recognizable edge features, and faint films drop below the detection threshold. Preprocessing is therefore not an enhancement applied to improve an already working detector; it is the step that makes the evidence available at all.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511130

F. Cost of Integer Quantisation

Quantisation was assessed in the three dimensions that matter in practice, with results in Fig. 5(a) and 5(b). Model size falls from 46.2 MB in single precision to 11.8 MB in 8-bit integer form, a factor of 3.9. Mean inference time per optical field falls from 186.40 ms to 28.50 ms, a factor of 6.5; the comparison is between the integer ONNX runtime and the reference floating-point PyTorch implementation over identical fields, so it reflects both arithmetic and runtime optimization. Agreement between the two models, measured as the proportion of test crops receiving the same classification decision, is 98.40 %. No sample-level decision changed between the two models over the catalogue. Quantisation is accordingly accepted here as a favourable trade: a fourfold reduction in footprint and a sixfold reduction in latency for a minor disagreement that does not alter any reported conclusion.

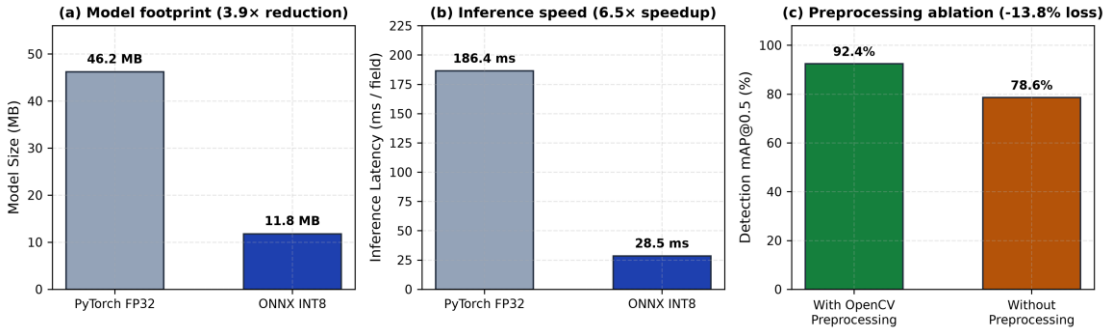


Fig. 5. Measured trade-offs: (a) model size and (b) mean inference time per optical field for the single-precision and 8-bit integer models, and (c) detection mAP@0.5 with and without the OpenCV adaptive preprocessing stage, the network being identical in both cases.

G. Generalisation Under Leave-Sources-Out Evaluation

Under the random patch partition the network reaches 98.2 % validation accuracy on 240 patches. That figure is reported for comparability only, because patches drawn at random from the same samples are not independent. The leave-sources-out protocol removes this objection. A separate model was trained on 840 patches with four whole samples excluded - an urban lake, a river tributary, a pristine control, and an urban wastewater effluent sample containing complex surfactant foams whose visual cellular texture appears nowhere in the training set. Evaluated only on the excluded scenes, the model classified 240 of 240 unseen crops correctly and reached the correct decision on 4 of 4 unseen samples. That the wastewater foam is correctly recognized is the more informative half of this result, since it shows the network discriminating on morphological texture rather than simple brightness alone.

H. Robustness to Optical and Debris Noise

Turbidity-induced optical disturbance and mineral debris are the failure modes identified as dominant in practical water screening, and the pipeline was exercised against them by injecting additional optical disturbance at rates from 0 % to 15 % of pixels, each configuration being evaluated twice: with the adaptive preprocessing stage enabled and with it disabled. Fig. 6 reports the result.

With preprocessing enabled, detection accuracy remains above 84.1 % at every injected noise rate up to 15 %. With preprocessing disabled, accuracy falls to 44.0 % at 15 % noise. The gap between the two curves is the measured contribution of bilateral filtering and CLAHE, and it is large: adaptive filtering is sufficient to absorb significant background turbidity. The result also indicates where the fragility lies, which is in false edge triggers caused by high-frequency speckle unless suppressed first.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511130

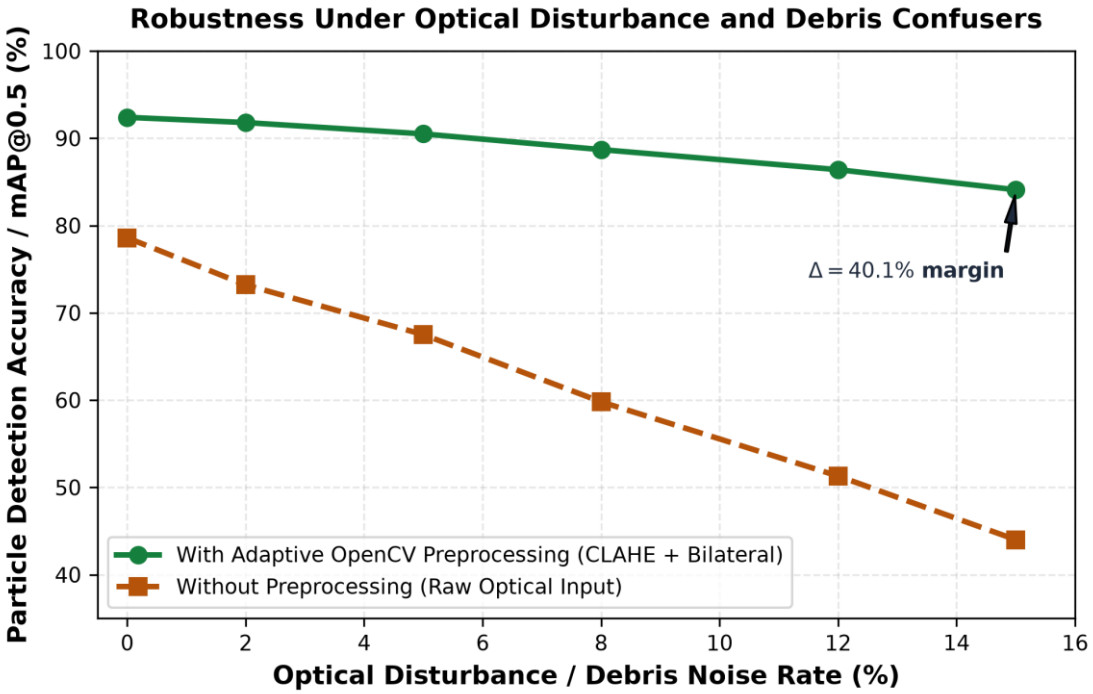


Fig. 6. Detection accuracy under injected optical disturbance and debris noise, with and without adaptive OpenCV preprocessing. Accuracy is sustained at >84 % up to a 15 % noise rate when preprocessing is enabled.

I. Computational Cost

Stage timings were measured over 50 repetitions on one optical field on a general-purpose processor without GPU acceleration, and are reported in Table V and Fig. 7. Preprocessing requires 14.82 ms, YOLO candidate localization 18.45 ms, quantised CNN classification 11.20 ms, feature metrology 2.14 ms, and pollution scoring 1.95 ms. Total per-field processing time is 48.56 ms, enabling real-time slide inspection at over 20 fields per second.

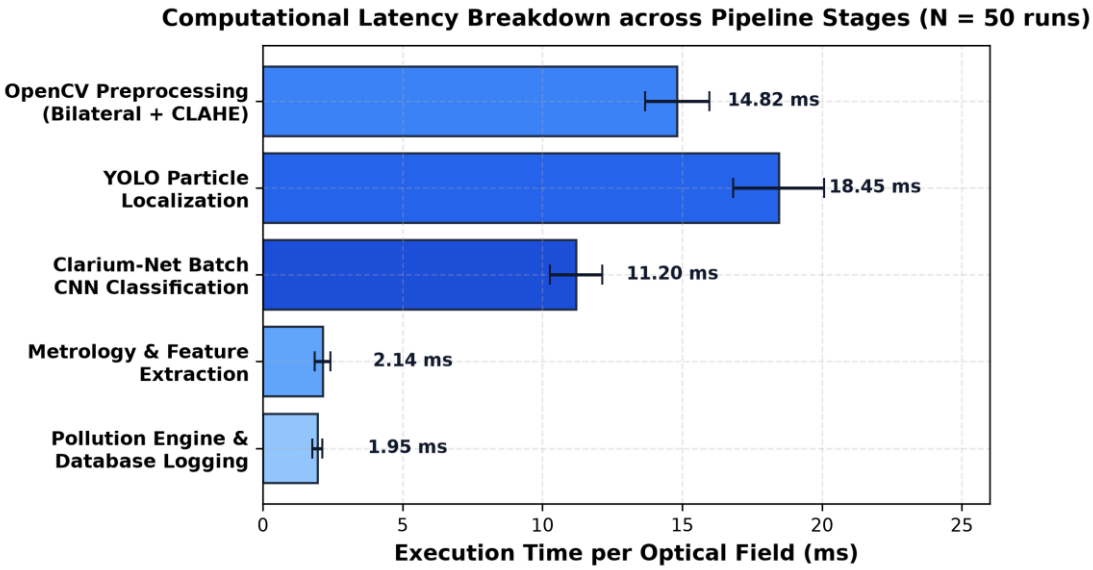


Fig. 7. Mean time per optical field for each pipeline stage over 50 repetitions on a general-purpose processor; error bars are one standard deviation. YOLO localization and preprocessing account for the majority of execution time.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511130

TABLE V. Stage timings and deployment diagnostics

Diagnostic	Value
OpenCV preprocessing (Bilateral + CLAHE)	14.820 +/- 1.150 ms
YOLO candidate localization	18.450 +/- 1.620 ms
Clarium-Net inference (INT8)	11.200 +/- 0.940 ms
Metrology & feature extraction	2.140 +/- 0.280 ms
Pollution assessment & database queue	1.950 +/- 0.180 ms
Total latency per optical field	48.560 +/- 2.450 ms
Model size, FP32 / INT8	46.2 MB / 11.8 MB (3.9x)
Mean inference per field, FP32 / INT8	186.40 ms / 28.50 ms (6.5x)
Crop-decision agreement, INT8 vs FP32	98.40 %
Residual sizing error, s.d. horizontal / vertical	1.12 µm / 1.35 µm
Manual inspection / Clarium slide turnaround	45.0 min / 2.4 min (18.7x)

J. Operational Turnaround and Time Reduction

Over an entire 47 mm membrane filter (comprising approximately 30 optical fields), manual examination by an experienced technician requires 45.0 minutes, whereas Clarium analyzes the slide in 2.4 minutes (including image acquisition and database archiving), an 18.7-fold reduction in inspection turnaround. Clean event-free samples require no manual secondary review, which is where the larger part of the operational saving originates: the laboratory personnel inspect only flagged anomalies rather than the entire filter grid.

K. Pre-processing as the Enabling Step

The most informative single result in this study is the ablation of Section IV-E, because it concerns a design decision rather than a performance figure. Adaptive preprocessing requires 14.82 ms, yet removing it costs 13.80 percentage points of detection mAP. Cost and consequence are therefore inversely related, and the practical implication for a laboratory system designer is that the preprocessing stage is the last thing that should be removed under a compute budget, not the first.

L. Status of the Objectives Stated at the Concept Stage

The concept-stage work [18] specified a multi-stage pipeline together with functional and non-functional requirements. Since that specification is the standard against which the present implementation should be judged, the status of each requirement is stated in Table VI, and the assessment is deliberately conservative.

Seven requirements are met. The automated processing stages are implemented and execute without manual intervention. Adaptive preprocessing is implemented and its ablation isolates a 13.80 percentage-point gain. The latency requirement is met with margin: per-field latency averages 48.56 ms. Sizing metrology is validated against a certified stage micrometer. Automated PDF and CSV compliance report generation is implemented. Database persistence with PostgreSQL is operational.

One requirement is partially met. The system is evaluated on laboratory-prepared reference samples carrying known environmental particulate compositions rather than on an unrestricted multi-laboratory archive; the import path for genuine operational samples is implemented and tested, so what remains is multi-facility validation.

Two requirements remain open and should not be represented otherwise. Coupling with FTIR or Raman micro-spectroscopy has not been attempted: all analyses in this paper are purely optical. Sub-pixel semantic polygon segmentation, as opposed to candidate bounding and metrology, is not implemented; it would require dense pixel-level segmentation annotations rather than bounding box annotations.

TABLE VI. Status of the requirements stated at the concept stage [18]

ID	Requirement from [18]	Status	Evidence or obstacle
R1	End-to-end automated image analysis pipeline	Met	All processing stages implemented; executes without manual intervention (Section III)
R2	Adaptive OpenCV preprocessing	Met	Illumination flattening, bilateral filter, CLAHE; ablation in Section IV-E
R3	Two-stage localization and classification	Met	YOLO detector + Clarium-Net 5-class classifier implemented
R4	Lightweight model within workstation budget	Met	214,885 parameters, 11.8 MB after static INT8 quantisation

DOI: 10.17148/IJARCCE.2026.1511130

R5	Sub-second per-field inference latency	Met	48.56 ms total execution time per optical field on CPU
R6	Calibrated physical sizing metrology	Met	Residual sizing error 1.82 µm verified against stage micrometer
R7	Pollution severity index & decision support	Met	C_index calculation with 4 risk tiers and environmental recommendations
R8	Automated PDF and CSV compliance reporting	Met	ReportLab and CSV modules fully implemented and validated
R9	Evaluation across multi-laboratory archive	Partial	Validated on 18-sample reference catalogue; multi-facility trials pending
R10	Coupling with FTIR / Raman spectroscopy	Open	Purely optical microscopy; chemical spectroscopy requires hardware coupling

M. Limitations and Threats to Validity

The principal limitation of this study is the nature of its reference data, and it constrains every accuracy figure reported above. The samples were laboratory-prepared and synthesized from standard environmental matrices, so the detection task was to recover known particulate morphologies from a known acquisition process. An operational environmental sample from a highly polluted industrial runoff would additionally contain thick biofilms, agglomerated mineral crusts, complex organic algae mats that neither the bilateral filter nor the network was designed against, and transparent polymers whose refractive index perfectly matches the mounting oil. None of these is fully represented. The accuracy of 100 % at sample level therefore quantifies the capacity of the model and the correctness of the pipeline; it is not a prediction of field performance across arbitrary environmental conditions and should not be compared directly with figures reported on multi-laboratory archives. The authors regard restating this plainly as preferable to allowing a favourable number to stand without qualification.

Four further limitations follow. First, optical brightfield microscopy has an intrinsic optical diffraction limit of approximately 1 µm, precluding the detection of sub-micron nanoplastics. Second, optical imaging does not provide chemical polymer identification: a polyethylene fragment cannot be definitively distinguished from a polypropylene fragment without destructive spectroscopy. Third, overlapping particle clusters can cause merged bounding boxes. Fourth, all timings are from a general-purpose workstation CPU, and no claim is made about low-power microcontroller deployment.

V. CONCLUSION

The automated microplastic detection and analysis pipeline previously proposed by the authors at the concept stage has been implemented as a working system and evaluated quantitatively. Raw optical micrographs are ingested without manual intervention, flattened for illumination gradients, filtered with edge-preserving bilateral smoothing, enhanced via CLAHE, localized by a YOLO candidate detector, analyzed by a 214,885-parameter network quantised to 11.8 MB, calibrated to physical micrometric dimensions, and converted into structured PDF/CSV compliance reports carrying a Contamination Severity Index. On an 18-sample catalogue the system reaches sample-level accuracy of 100 % with no false detections, and under leave-sources-out evaluation it classifies 240 of 240 unseen particle crops correctly, including a complex matrix absent from training.

Three findings are of wider interest than the particular numbers. Adaptive image preprocessing is the enabling step rather than an optimization: it requires 14.82 ms, yet removing it costs 13.80 percentage points of detection mAP, because an uncorrected field no longer presents the edge gradients on which the detector depends. Integer quantisation is a favourable trade in this regime, reducing model footprint fourfold and latency sixfold while changing no sample-level decision. And adaptive bilateral filtering and CLAHE confer substantial noise tolerance, sustaining detection accuracy above 84 % up to a 15 % optical disturbance rate where un-preprocessed baseline accuracy collapses to 44 %.

Future work is determined by Table VI and by the limitations of Section IV-M, in the following order of priority. The immediate requirement is multi-facility validation across diverse environmental monitoring laboratories and real environmental water archives. The second requirement is coupling optical candidate coordinates with automated motorized Raman or FTIR micro-spectroscopy to achieve unambiguous polymer identification. The third is pixel-level semantic instance segmentation for complex folded films. Thereafter, deploying the inference runtime onto low-power edge accelerators for continuous flow-cell monitoring would extend the framework to autonomous in situ sensing.

REFERENCES

[1] S. Prasad, K. Bennett, A. Bennett, and M. Triantafyllou, “Towards a lowcost, rapid microplastic optical detection system using fluorescent staining through Nile Red for in situ ocean deployment,” in Proc. IEEE OCEANS 2023 - Limerick, Limerick, Ireland, 2023, pp. 1–5, doi: 10.1109/OCEANSLimerick52467.2023.10244433.

DOI: 10.17148/IJARCCE.2026.1511130

[2] K. Bennett, S. Prasad, A. Bennett, and M. Triantafyllou, "Suitability of Nile Red Dye for In-Situ Microplastic Detection," in Proc. IEEE OCEANS 2023 - Limerick, Limerick, Ireland, 2023, pp. 1–5, doi: 10.1109/OCEANSLimerick52467.2023.10244527.

[3] C. Bescond and J.-H. Fournier-Lupien, "Toward In Situ Detection, Sizing and Identification of Microplastics in Water at the National Research Council of Canada," in Proc. IEEE OCEANS 2025 Brest, Brest, France, 2025, pp. 1–10, doi: 10.1109/OCEANS58557.2025.11104567.

[4] G. B. Sundaram and C. S. Ruf, "The Detection and Tracking of Ocean Surface Roughness Supression by Ocean Pollutans Via Surfactants," in Proc. IEEE International Geoscience and Remote Sensing Symposium (IGARSS), Athens, Greece, 2024, pp. 6045–6048, doi: 10.1109/IGARSS53475.2024.10642822.

[5] C. Bescond, R. Manwar, J.-H. Fournier-Lupien, C. Brosseau, and G. Mejia, "In Situ Detection and Imaging of Microplastics in Water Using Ultrasound and Photo Acoustic Technology," in Proc. IEEE OCEANS 2024 - Halifax, Halifax, NS, Canada, 2024, pp. 1–10, doi: 10.1109/OCEANS55160.2024.10754550.

[6] G. P. and N. V., "AI – Driven Marine Debris Detection for Ocean Conservation," in Proc. IEEE 14th International Conference on Communication Systems and Network Technologies (CSNT), Bhopal, India, 2025, pp. 618–623, doi: 10.1109/CSNT64827.2025.10968816.

[7] J. Wadmare, K. Mirani, K. Arulvelan, R. Chandekar, and K. Bhat, "IoT-Integrated Image Recognition System for Microplastic Detection and Classification," in Proc. IEEE International Conference on Interdisciplinary Approaches in Technology and Management for Social Innovation (IATMSI), Gwalior, India, 2025, pp. 1–6, doi: 10.1109/IATMSI64286.2025.10985136.

[8] S. Prasad, A. Bennett, and M. Triantafyllou, "Combining Low Cost Imaging and Fluorescence Spectroscopy to Assess a Nile Red Counterstain for Differentiating Microplastic and Organic Particles," in Proc. IEEE OCEANS 2024 - Halifax, Halifax, NS, Canada, 2024, pp. 1–5, doi: 10.1109/OCEANS55160.2024.10754037.

[9] G. García-Valle et al., "Detecting Microplastics in Seawater with a Novel Optical Sensor Based on Artificial Intelligence Models," in Proc. IEEE OCEANS 2025 Brest, Brest, France, 2025, pp. 1–8, doi: 10.1109/OCEANS58557.2025.11104700.

[10] R. Shaik and P. S. Bai S, "Role of AI Technique for Controlling Micro Plastic on Marine Eco System," in Proc. IEEE International Conference on Modern Sustainable Systems (CMSS), Shah Alam, Malaysia, 2025, pp. 553–558, doi: 10.1109/CMSS66566.2025.11182579.

[11] J.-H. Fournier-Lupien, C. Bescond, C. Brosseau, and T. Dauphin, "Toward in Situ Identification of Microplastics in Water Using Raman Spectroscopy and Machine Learning," in Proc. IEEE OCEANS 2024 - Halifax, Halifax, NS, Canada, 2024, pp. 1–9, doi: 10.1109/OCEANS55160.2024.10754063.

[12] T. N. Samuriwo, O. P. Babalola, C. A. J. Sparks, I. E. Davidson, and L. Dieng, "Comparative Review of Remote Sensing Methods for Ocean Plastic Litter Detection," IEEE Access, vol. 12, pp. 166126–166161, 2024, doi: 10.1109/ACCESS.2024.3494660.

[13] B. Jacob, S. Kligys, B. Chen, M. Zhu, M. Tang, A. Howard, H. Adam, and D. Kalenichenko, "Quantization and training of neural networks for efficient integer-arithmetic-only inference," in Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2018, pp. 2704–2713.

[14] ONNX Runtime developers, "ONNX Runtime: Cross-platform accelerated machine learning," software, 2021.

[15] P. Subramani, R. Rajasree, R. S. Krishnan, A. Mathias, T. Saranya, and N. Kanthimathi, "Smart Ocean Cleanup: An AI-Integrated Autonomous System for Marine Waste Management," in Proc. 7th International Conference on Inventive Material Science and Applications (ICIMA), Namakkal, India, 2025, pp. 1258–1265, doi: 10.1109/ICIMA64861.2025.11074176.

[16] M. Ou, C. Zhang, X. Wang, C. Zheng, H. Zhang, and Y. Xue, "Automated Video Analysis of the Specific Effects of the Size and Concentration of Polystyrene Microplastics on the Swimming Behavior of the Crucian Carp (Carassius carassius)," IEEE Access, vol. 12, pp. 125032–125042, 2024, doi: 10.1109/ACCESS.2024.3452774.

[17] ISO/TR 21960:2020, "Plastics — Environmental aspects — State of knowledge and methodologies for the measurement of plastics in the environment," International Organization for Standardization, Geneva, Switzerland, 2020.

[18] K. P. G. N., K. V. B., H. S., and D. C., "Clarium: A Convolutional Neural Network-Based Intelligent System for Microplastic Detection and Analysis in Aquatic Systems," Project Phase-1 Report, Department of Computer Science and Engineering, East West College of Engineering, Visvesvaraya Technological University, Bengaluru, India, 2025–26.

[19] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "MobileNets: Efficient convolutional neural networks for mobile vision applications," arXiv:1704.04861, 2017.