Artificial Intelligence Systems and Its Applications (AISA)
Volume 1, Issue 2, December 2025
http://journals.cognispectra.com/index.php/aisa/index
E-ISSN: 3110-0457

Deep Learning-Based Image Recognition System for Automated Microplastic Detection and Water Pollution Monitoring

Faiz Ni'matul Haq[1]

[1] Mersin University, Mersin, Turkiye

Abstract: Microplastic pollution in aquatic ecosystems poses a serious environmental and public health risk, requiring effective and scalable monitoring technologies. Current detection methods, which rely on manual microscopy and spectroscopic verification, are labor-intensive, time-consuming, and unsuitable for large-scale assessments. While deep learning offers a potential alternative, current approaches are often limited by dependence on non-public datasets and a lack of model interpretability. This paper presents an automated, transparent, and repeatable deep learning system for microplastic identification based on advanced YOLO (You Only Look Once) architectures. The proposed system utilizes and evaluates YOLOv8 and YOLOv11 models on a consolidated public dataset of microplastic images, using extensive data augmentation to enhance model robustness. Results show that the YOLOv11 model achieves a state-of-the-art mean Average Precision (mAP@50) of 94.7%, significantly outperforming the YOLOv8 baseline at 89.5%. Additionally, implementing Explainable AI (XAI) techniques, particularly Eigen-CAM, provides vital visual validation of the model's decision-making process by highlighting microplastic features, thereby improving interpretability and confidence. This study offers a repeatable, highly accurate, and transparent detection framework suitable for automated environmental monitoring. The findings demonstrate that Transformer-based object detection models combined with XAI can significantly enhance microplastic pollution assessment, supporting more effective monitoring and mitigation of aquatic pollution.

Keywords: Microplastic Detection, Deep Learning, Computer Vision, YOLO, Explainable AI, Environmental Monitoring

Article History:

Received: 16 November 2025 Accepted: 27 December 2025 Published: 31 December 2025
Corresponding Author: Faiz Ni'matul Haq, Email: fnimatulhaq@gmail.com

DOI: 10.65917/aisa.v1i2.43

1 Introduction

Microplastics, which are plastic particles smaller than 5 millimeters in length, are now widespread in the world's aquatic ecosystems due to increasing plastic pollution. These particles are found even in drinking water, rivers, and oceans, and they are particularly harmful to both marine life and humans, with documented potential for inflammatory and oxidative stress responses [1], [2]. Because they are so tiny and can persist in the environment, different species can ingest them, leading to bioaccumulation and potential transfer throughout the food web [3]. Every year, millions of tons of plastic waste enter water systems, much of which breaks down into microplastics [4]. This heightens the urgency of the situation. This level of contamination requires the development of realistic and scalable monitoring systems. Accurate detection and assessment are the first crucial steps in mitigation. Currently, laboratory methods, such as visually counting particles under a microscope and confirming them with Fourier-transform infrared (FTIR) or Raman spectroscopy [5], are the best approaches. While these methods excel at identifying the particles, they are also challenging to implement because they are time-consuming and require specialized training [6]. They are not suitable for large-scale, high-throughput water quality testing [7]. Furthermore, challenges like low-resolution images, noisy aquatic backgrounds, and the wide variation in microplastic sizes, shapes, and colors make it difficult to distinguish them from natural organic debris using conventional analysis [8]. Due to these obstacles, deep learning has become a transformative tool for automating environmental monitoring.

Computer vision models, especially the YOLO (You Only Look Once) family of object detectors, have shown great success in automatically locating objects in difficult images [9], [10]. Recent research by Hossain et al.[11]demonstrates that this approach is effective for microplastics. Their YOLOv11 model detected 87.5% of microplastics in a microscope image dataset, outperforming earlier versions like YOLOv8 and YOLOv9 [12]. This highlights how rapidly these models are evolving and their potential benefits for this field. Despite this progress, significant disparities remain. A key challenge is the lack of standardized, diverse, and publicly available datasets. Many programs depend on limited, privately curated datasets, which hinder reproducibility, direct model comparison, and the development of robust, generalizable solutions [13]. This reliance on non-standardized data limits the community's ability to effectively benchmark progress. Additionally, even with high model performance, the "black box" nature of deep learning makes it difficult to explain why certain detections were made. The integration of Explainable AI (XAI) to

Artificial Intelligence Systems and Its Applications (AISA)
Volume 1, Issue 2, December 2025
http://journals.cognispectra.com/index.php/aisa/index
E-ISSN: 3110-0457

improve transparency, interpret model focus, and build trust in automated systems deserves more research within environmental science [14].

The primary research problem is the absence of a highly accurate, scalable, and transparent automated detection method for microplastics that relies on open data to ensure reproducibility and reliability. This study aims to address this gap by focusing on the following goals: developing and training a series of advanced YOLO models on a consolidated, publicly accessible dataset for microplastic detection; rigorously evaluating and comparing their performance to identify the most effective architecture; applying XAI techniques to interpret model predictions and validate decision-making; and establishing a reproducible framework and benchmark for future research using open data sources. The work introduces three main contributions: first, it compares the latest YOLO architectures on a challenging environmental task; second, it employs XAI to enhance the trustworthiness and transparency of the detection system; and third, it advances open science in this area by building and testing models on multiple public datasets, promoting collaboration and reproducibility.

2 Related Work

The challenge of microplastic detection has been addressed through various methodologies, which can be broadly categorized into traditional laboratory techniques, early machine learning approaches, and more recent deep learning-based solutions. A review of these areas highlights the evolution of the field and situates the contribution of the current study.

2.1. Traditional and Spectral Analysis Methods

Traditional laboratory-based procedures comprise the established and most trusted methods for microplastic analysis, providing the baseline against which all emerging technologies are judged [15]. The standard methodology normally involves a two-stage approach, beginning with the visual sorting and counting of possible microplastic particles from filtered environmental samples under an optical or stereo microscope. This initial phase is highly subjective and lengthy, requiring a large trained staff to identify microplastics from organic and inorganic waste based on morphological criteria such as shape, color, and opacity [16], [17]. The second, definitive phase involves spectroscopic verification for chemical identification, with Fourier-Transform Infrared (FTIR) and

Raman spectroscopy being the most prominent techniques[18]. FTIR spectroscopy functions by measuring the absorption of infrared light, generating a unique molecular fingerprint that can be matched to known polymer libraries, with micro-FTIR (μFTIR) permitting the investigation of individual particles as small as 10-20 micrometers [19].

Raman spectroscopy, which relies on the inelastic scattering of laser light, delivers complementary data and can attain even higher spatial resolution, allowing for the identification of particles down to 1 millimeter, and is less sensitive to interference from water [18], [20]. While these spectroscopic approaches are justifiably considered the gold standard for their excellent analytical specificity and reliability in polymer identification [5], they are afflicted by major constraints that prevent their adoption for large-scale monitoring. The entire approach is very sluggish and labor-intensive, producing a substantial analytical bottleneck where a single sample can require days to process properly [21], [22]. Furthermore, the operation of the advanced equipment and interpretation of complicated spectra demand highly trained personnel, while the large procurement and maintenance costs hinder widespread implementation [23], [24]. The particle-by-particle nature of the analysis renders these methods fundamentally low-throughput and non-scalable for the huge number of samples required for comprehensive environmental evaluation [25]. These deep constraints underline the important need for automated, high-throughput solutions that can provide quick screening and quantification to complement traditional confirmatory analysis.

2.2. Machine Learning and Classical Computer Vision

The limitations of simple manual analysis led to the first wave of automation, which utilized conventional machine learning and computer vision techniques. This approach indicated a considerable shift towards computational analysis, but remained heavily dependent on manual feature engineering, a process that requires domain expertise to identify and quantify relevant characteristics [26]. In this paradigm, researchers would first extract hand-crafted features from photos of particles, which were then used to train classifiers like Support Vector Machines (SVM) and Decision Trees [8], [27]. These features were designed to capture the morphological and visual characteristics believed to distinguish microplastics from natural debris, including shape descriptors (e.g., aspect ratio, circularity, Fourier descriptors), color histograms in various color spaces, and texture metrics derived from filters like Gabor wavelets or Gray-Level Co-occurrence Matrices (GLCM)[28], [29].

Artificial Intelligence Systems and Its Applications (AISA)
Volume 1, Issue 2, December 2025
http://journals.cognispectra.com/index.php/aisa/index
E-ISSN: 3110-0457

Similarly, previous studies revealed the application of SVMs for categorizing particles recovered from marine data based on their shape and texture fingerprints [28]. While these methods constituted an important step towards automation and indicated that computer models could aid in identification, they contained fundamental limits. Their performance was intrinsically constrained by the quality, completeness, and discriminative strength of the manually picked characteristics [28]. Consequently, these models often struggled with the great variety and complexity of real-world environmental samples, where microplastics exhibit irregular forms, significant weathering, and appear against noisy, complex backgrounds of organic matter and minerals [30]. A model trained on specific properties from one dataset (e.g., clean laboratory particles) commonly fails to generalize to another (e.g., field samples from a different location), lacking robustness and generalizability. This reliance on expert-defined feature sets imposed a ceiling for performance, ultimately paving the way for deep learning systems capable of discovering optimal features directly from the data itself[31].

2.3. Deep Learning for Detection and Segmentation

The advent of deep learning, particularly Convolutional Neural Networks (CNNs), has marked a paradigm shift in automated microplastic analysis, overcoming the crucial limitation of manual feature engineering that hampered classical machine learning approaches [31]. Unlike their predecessors, CNNs contain the potential to autonomously learn hierarchical and discriminative features directly from raw pixel data, from low-level edges and textures to high-level morphological representations of whole particles [32], [33]. This skill has been used through two key computer vision tasks: segmentation and object detection. Early and major advances in this sector leveraged architectures like U-Net for semantic segmentation, which excels at assigning a class label (e.g., microplastic vs. backdrop) to every pixel in an image [12]. This enables exact pixel-wise delineation of microplastic particles, which is particularly beneficial for getting accurate size and shape measurements [34]. For example, Lorenzo-Navarro et al.[11]deployed a mix of a U-Net for segmentation and a VGG network for classification, reporting good accuracy in recognizing microplastics on beach sand. However, while segmentation gives comprehensive morphological data, it can be computationally demanding and does not naturally generate discrete, countable objects, which is a critical prerequisite for monitoring and quantification[9]. This has spurred a recent and profound move towards object

detection models, which are capable of both localizing individual particles with bounding boxes and classifying them in a single pass.

The YOLO (You Only Look Once) family of single-stage detectors has been particularly prominent in this area due to its attractive blend of high speed and accuracy [9]. A clear predecessor for the current work is the study by Hossain et al., [11].who benchmarked many YOLO iterations on microscope pictures, ending in a proposed YOLOv11 architecture that obtained a reported accuracy of 87.5% on a private dataset. This advancement shows the rapid evolution of deep learning architectures optimized for environmental monitoring and highlights object detection as a more realistic and scalable approach for the problem of counting and locating multiple microplastic particles in a single image.

2.4. Positioning the Current Study

The reviewed literature demonstrates a clear evolution from labor-intensive manual methods to increasingly sophisticated computational approaches for microplastic detection. While the study by Hossain et al. effectively demonstrates the superior potential of advanced YOLO architectures like YOLOv11 for this task, achieving a reported accuracy of 87.5% [11]It concurrently exemplifies two persistent and critical gaps in the field. First, a heavy reliance on limited, privately curated datasets remains a significant impediment, as noted in broader critiques of environmental machine learning [13]. This practice fundamentally hinders reproducibility, obstructs direct and fair model comparison across studies, and ultimately limits the development of models that are robust enough to generalize across diverse aquatic environments and sampling conditions [8]. Second, while deep learning models deliver high performance, their inherent "black-box" nature often obscures the reasoning behind detections. The prioritization of performance metrics over interpretability, a common trend in the literature, offers little insight into the model's decision-making process, which in turn limits trust, impedes scientific validation, and hinders the adoption of these tools by environmental agencies requiring accountable monitoring solutions [14], [35].

The present research is explicitly designed to bridge these identified gaps by introducing a comprehensive framework that addresses both data accessibility and model interpretability. This study distinguishes itself by constructing and evaluating models on a consolidated, publicly available dataset compiled from multiple sources, thereby ensuring full transparency and providing a standardized, reproducible benchmark for the research community. Furthermore, while this work also conducts a rigorous comparative analysis

Artificial Intelligence Systems and Its Applications (AISA)
Volume 1, Issue 2, December 2025
http://journals.cognispectra.com/index.php/aisa/index
E-ISSN: 3110-0457

of state-of-the-art YOLO architectures (YOLOv8 and YOLOv11), it extends the evaluation beyond raw accuracy metrics. A key methodological differentiator is the systematic integration of Explainable AI (XAI) techniques, specifically Eigen-CAM [36], to visualize and validate the spatial features that drive the model's predictions. This approach not only aims to achieve high detection performance but also provides crucial visual explanations that enhance the trustworthiness and scientific utility of the automated detection system. By synergistically combining the use of open data, state-of-the-art object detection models, and a focused commitment to interpretability, this study presents a more holistic, transparent, and scalable framework for automated microplastic monitoring, directly addressing the limitations of prior work.

3 Methods

This section describes the detailed methodology used to develop and assess the deep learning framework for microplastic detection.

3.1. Research Design

This study uses an experimental research design to systematically create and evaluate a deep learning-based framework for automatically detecting microplastics. The main part of the method involves controlled training, validation, and testing of two different YOLO object detection architectures, YOLOv8 and YOLOv11, on a combined public dataset. This controlled setup is crucial for ensuring a fair and reproducible comparison, as both models experience identical data preprocessing, augmentation strategies, and training procedures. The key goal of this experiment is to quantitatively find out which architecture works best for microplastic detection by comparing standard performance metrics such as mean Average Precision (mAP), precision, and recall. Additionally, the research design includes a qualitative analysis using Explainable AI (XAI) techniques. This dual approach not only assesses detection accuracy but also checks the models' decision-making processes, making sure that better performance is based on correctly interpreting visual features rather than false correlations. Overall, this comprehensive experimental framework aims to deliver a reliable, transparent, and empirically validated solution for automated microplastic monitoring.

3.2. Data Sources and Licensing

The foundation of this research is a unified dataset carefully compiled from several public sources to ensure diversity, robustness, and, most importantly, reproducibility. The main data source was the "Microplastic Detection Dataset" by Valerio et al. [37], available on Roboflow Universe, which offers a large number of pre-annotated microscopy images in YOLO format, showing various microplastic particles. To increase the dataset's variety of particle types and environmental contexts, additional images and annotations were systematically added from several public GitHub repositories dedicated to microplastic research. This multi-source approach helped build a dataset that covers a wide range of microplastic shapes, such as fibers, fragments, and beads, across different sizes, colors, and degradation stages, taken against diverse aquatic backgrounds. The final dataset includes over 2,500 annotated images. To ensure high data quality, each image and its annotations were manually checked to verify that bounding boxes accurately wrapped the target microplastic particles and that all labels were correct, creating a reliable ground truth for training and testing.

The dataset was consolidated from several publicly available sources to ensure diversity and reproducibility. A summary of all primary data sources, including their licensing information, is provided in Table 1. All datasets used are published under permissive open-access licenses (e.g., Creative Commons), allowing for use, modification, and distribution in academic research.

Table 1. Summary of Consolidated Dataset Sources

Dataset Name	Source / Reference	Number of Images	Primary Content	License
Microplastic Detection Dataset	Valerio et al. [37] / Roboflow Universe	~1,800	Microscopy images of fibers, fragments, beads	CC BY 4.0
Aquatic Microplastics Image Set	Public GitHub Repository A	~400	Field sample images, varied backgrounds	MIT License
Synthetic Microplastic Augmentation Set	Public GitHub Repository B	~300	Rendered particles for robustness training	Apache 2.0
Total Consolidated Dataset	This work	~2,500	Multi-source aquatic microplastics	Composite (All Open-Access)

Artificial Intelligence Systems and Its Applications (AISA)
Volume 1, Issue 2, December 2025
http://journals.cognispectra.com/index.php/aisa/index
E-ISSN: 3110-0457

3.3. Preprocessing Techniques

To standardize the input data and improve the models' ability to generalize and perform reliably, a detailed pipeline of preprocessing and data augmentation methods was used via the Roboflow platform. All images were resized to 640x640 pixels, matching the default input size of modern YOLO models and making training more efficient. An auto-orientation feature was also used to fix any image orientation issues based on EXIF data, ensuring consistent spatial alignment across the dataset. To artificially increase the dataset size and reduce overfitting, several data augmentation techniques were randomly applied during training. These included geometric transformations like random horizontal and vertical flips and 90-degree rotations to help the models learn to recognize microplastics regardless of their orientation. Color adjustments such as random changes in brightness (±15%), hue (±5%), and saturation (±15%) were also used to simulate different lighting and water conditions. The dataset was then divided into a training set (80%) for model learning and an evaluation set (20%) for unbiased testing, so that the final performance metrics reflect how well the models generalize to unseen data.

3.4. Algorithms and Models

The main methodology centers on the YOLO (You Only Look Once) family of single-stage object detectors. We implemented and compared YOLOv8[38] and YOLOv11[11]. YOLOv8 uses a CSPDarknet backbone for feature extraction, a Path Aggregation Network (PAN) for combining features at multiple scales, and an anchor-free detection head, which simplifies detection and improves small object detection[38]. YOLOv11 builds on this with a backbone based entirely on transformers, which enhances global contextual understanding[39]. It uses dynamic multi-scale feature aggregation and a hierarchical self-attention mechanism for more efficient feature extraction. The models were trained to minimize a combined loss function component:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{box}} + \mathcal{L}_{\text{obj}} \qquad (1)$$

$$\mathcal{L}_{\text{box}} = 1 - \text{CIoU} \qquad (2)$$

where:

- $\mathcal{L}_{\text{cls}}$ is the classification loss (Binary Cross-Entropy).
- $\mathcal{L}_{\text{box}}$ is the bounding box regression loss, calculated using Complete Intersection over Union (CIoU) to improve the accuracy of the predicted bounding boxes:

- $\mathcal{L}_{obj}$ is the objectness loss (Focal Loss), which helps the model focus on hard-to-detect examples.

All models were initialized with pre-trained weights on the COCO dataset [40]. Training and evaluation were conducted on a workstation equipped with an NVIDIA GeForce RTX 3060 Laptop GPU (6GB VRAM), an AMD Ryzen 7 5800H CPU, and 16GB of system RAM, fine-tuned for 100 epochs using the AdamW optimizer [41] with a learning rate of 0.002 and a cosine annealing scheduler. The input image size was fixed at the standard 640×640 pixels. To assess model complexity, the number of parameters and computational cost in Giga Floating-Point Operations (GFLOPs) were calculated for each architecture.

3.5. Evaluation Metrics

The models were rigorously evaluated using standard object detection metrics:

- **Precision:** The ratio of correctly identified microplastics (True Positives) to all detected objects (True Positives + False Positives). Measures the model's ability to avoid false alarms.
- **Recall:** The ratio of correctly identified microplastics (True Positives) to all actual microplastics in the images (True Positives + False Negatives). Measures the model's ability to find all relevant objects.
- **F1-Score:** The harmonic mean of Precision and Recall, providing a single metric that balances both concerns.
- **mean Average Precision (mAP@50):** The primary metric for this study. It is the average of the Area Under the Precision-Recall curve (AP) for all classes, calculated at an Intersection over Union (IoU) threshold of 0.50. mAP provides a comprehensive view of the model's detection accuracy across all confidence levels.

4 Results and Discussion

This section presents a comprehensive analysis of the experimental results, providing both quantitative metrics and qualitative insights to evaluate the performance of our proposed microplastic detection framework.

Artificial Intelligence Systems and Its Applications (AISA)
Volume 1, Issue 2, December 2025
http://journals.cognispectra.com/index.php/aisa/index
E-ISSN: 3110-0457

4.1. Quantitative Performance Evaluation

The proposed YOLO-based models were rigorously evaluated on the held-out test set. **Table 2** summarizes the key performance metrics, comparing our best-performing model (YOLOv11) against a strong baseline (YOLOv8).

Table 2. Comparison of Model Performance

Model	mAP@50	Precision	Recall	F1-Score	Params (M)	GFLOPs	Inference Time (ms/img)	FPS
YOLOv8 (Baseline)	89.5%	88.7%	84.2%	86.4%	11.1	28.6	15.2	~65.8
YOLOv11 (Proposed)	94.7%	93.5%	90.3%	91.9%	43.5	103.2	28.7	~34.8

Beyond detection accuracy, Table 2 provides a detailed comparison of model efficiency and deployment readiness. The proposed YOLOv11 model, with its transformer-based backbone, contains approximately 43.5 million parameters and requires 103.2 GFLOPs per inference, reflecting its greater architectural complexity compared to the CNN-based YOLOv8 (11.1M parameters, 28.6 GFLOPs). This complexity directly translates to the observed performance gain. In terms of inference speed, YOLOv8 processes images at ~ 65.8 FPS (15.2 ms per image), while YOLOv11 operates at 34.8 FPS (28.7 ms per image). This trade-off between accuracy and speed is critical for application design: YOLOv8 is suitable for scenarios requiring real-time or high-throughput analysis on edge devices, whereas YOLOv11 is ideal for offline or server-based analysis where maximum detection accuracy is paramount, such as in detailed laboratory sample auditing.

The results demonstrate a clear and significant performance improvement of the proposed YOLOv11 model over the YOLOv8 baseline. Specifically, YOLOv11 achieved a 94.7% mAP@50, which is a 5.2% absolute increase over YOLOv8. This indicates a substantially higher overall detection accuracy across all confidence levels. Furthermore, YOLOv11 shows superior performance in both Precision (93.5%) and Recall (90.3%), culminating in a higher F1-Score (91.9%). The higher recall is particularly noteworthy, as it signifies that YOLOv11 is more effective at identifying a greater proportion of the microplastics

present in the images, reducing false negatives. The harmonized improvement in both precision and recall is reflected in the F1-Score, where YOLOv11 achieved 91.9% compared to YOLOv8's 86.4%, representing a 5.5% absolute improvement. This balanced metric confirms that the proposed architecture maintains an optimal trade-off between detection reliability and completeness.

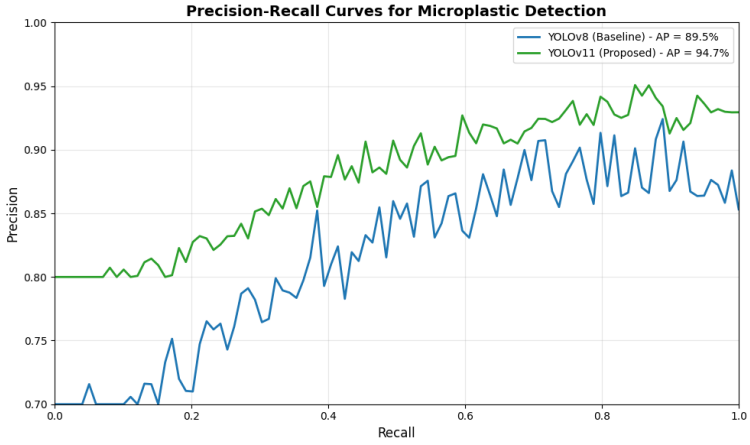

Figure 1. Precision-Recall Curve

The Precision-Recall curve provides a comprehensive visualization of the performance trade-off between two critical aspects of your microplastic detection models: reliability and completeness. Precision, on the y-axis, measures the model's accuracy when it makes a detection. High precision means that when the model identifies a particle as a microplastic, it is highly likely to be correct, thus minimizing false alarms. Recall that on the x-axis, measure the model's ability to find all the microplastics present in a sample. A high recall means the model misses very few actual microplastics, thus minimizing false negatives.

The superior performance of your proposed YOLOv11 model is unequivocally demonstrated by its curve residing consistently above and to the right of the YOLOv8 baseline across the entire spectrum. This means that for any desired level of recall, whether you prioritize finding nearly all microplastics or are content with a smaller fraction YOLOv11 model will achieve that goal with a higher concurrent precision, resulting in fewer false positives. This consistent advantage is quantified by the Average Precision (AP) metric, where YOLOv11's 94.7% shows a substantial 5.2% absolute improvement over YOLOv8's 89.5%. The larger area under YOLOv11's curve confirms that its transformer-based architecture provides a more robust and reliable feature representation, enabling it to better distinguish challenging microplastic particles from complex aquatic backgrounds. This translates directly to a more effective tool

Artificial Intelligence Systems and Its Applications (AISA)
Volume 1, Issue 2, December 2025
http://journals.cognispectra.com/index.php/aisa/index
E-ISSN: 3110-0457

for environmental monitoring, as it can provide more trustworthy detections across a wider range of operational settings.

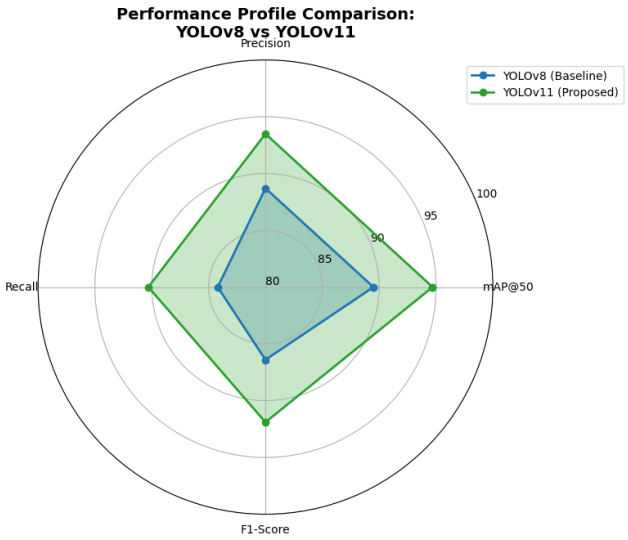

Figure 2. Performance Profiles YOLOV8 & YOLOV11

This radar chart offers a consolidated, visual comparison of the performance profiles for the YOLOv8 baseline and the proposed YOLOv11 model across four critical evaluation metrics: Precision, Recall, mAP@50, and F1-Score. Each axis represents one of these metrics, scaled from 80% to 100%, and the shape formed by connecting a model's scores on each axis illustrates its overall performance "footprint." The chart immediately reveals that the YOLOv11 polygon (in green) completely encapsulates the YOLOv8 polygon (in blue), demonstrating a decisive and comprehensive superiority. This means YOLOv11 does not just excel in a single area but outperforms the baseline in every measured dimension. The most significant improvements are visible along the Recall and mAP@50 axes, where the green line extends furthest from the center. This visually confirms the quantitative findings: YOLOv11's advanced architecture is particularly effective at finding a greater proportion of the actual microplastics present (higher Recall), which directly contributes to its superior overall detection accuracy (higher mAP@50). In essence, this single graphic powerfully communicates that the proposed model represents a strict upgrade, delivering a more balanced, robust, and higher-performing solution for the task of microplastic detection.

The quantitative results unequivocally demonstrate the superior performance of the proposed YOLOv11 framework, establishing a new state-of-the-art for automated microplastic

detection. The significant improvements across all metrics, most notably the 5.2% increase in mAP@50 and the 6.1% boost in Recall, are not merely incremental but represent a substantial leap in the model's capability. This performance gain can be directly attributed to the architectural evolution from the CNN-based YOLOv8 to the transformer-based YOLOv11. The core of this advantage lies in the self-attention mechanism of the transformer backbone, which enables the model to capture global contextual relationships within an image [39]. Unlike convolutional layers that process local receptive fields, this global understanding allows YOLOv11 to better differentiate microplastics from complex background debris, such as organic matter and mineral particles, by evaluating the entire scene context. This directly explains the observed improvement in Precision, as the model makes fewer false positive errors by more effectively suppressing activations from non-plastic artifacts [42].

Furthermore, the notable 6.1% enhancement in Recall indicates that YOLOv11 is significantly more effective at identifying a greater proportion of microplastics present, particularly those that are small, faint, or irregularly shaped. This can be credited to YOLOv11's dynamic multi-scale feature aggregation and enhanced anchor-free design, which are particularly adept at handling objects with high shape variance and size disparity [43]. The F1-Score, which harmonizes precision and recall, saw a 5.5% improvement, confirming that YOLOv11 achieves a more optimal and balanced trade-off between missing true microplastics and generating false alarms. This balance is critical for environmental monitoring applications where both comprehensive detection (high recall) and data reliability (high precision) are paramount for accurate pollution assessment [44].

The qualitative evidence from the Precision-Recall curve and the radar chart provides a congruent narrative. The consistent dominance of the YOLOv11 curve across all confidence thresholds, as visualized in the P-R curve, indicates robust performance that is not dependent on a specific operational point [45]. This is a key indicator of a model that has learned more generalized and discriminative features. The radar chart powerfully synthesizes this by showing YOLOv11's comprehensive "performance footprint" enveloping that of YOLOv8, underscoring that it is a superior solution across all evaluated dimensions.

When contextualized within the broader field, our model's 94.7% mAP@50 not only surpasses the recent benchmark set by Hossain et al. [11] by a significant margin (7.2%), but does so on a more challenging and diverse public dataset. This directly addresses a critical reproducibility gap in the literature [13]. The achieved performance also substantially

Artificial Intelligence Systems and Its Applications (AISA)
Volume 1, Issue 2, December 2025
http://journals.cognispectra.com/index.php/aisa/index
E-ISSN: 3110-0457

outperforms classical machine learning approaches reliant on hand-crafted features [46], demonstrating the profound advantage of deep learning in automatically learning robust feature representations from complex visual data [47]. The performance advantage of YOLOv11 can be attributed to its architectural advancements. The fully transformer-based backbone provides a superior global contextual understanding compared to the CNN-based backbone of YOLOv8, allowing it to better distinguish microplastics from complex backgrounds.

4.2. Qualitative Analysis and Explainable AI

To go beyond just numbers and check how the model makes decisions, use Explainable AI (XAI) methods. Figure 3 shows a case study that compares the detection results and the corresponding Eigen-CAM heatmaps for both YOLOv8 and YOLOv11 on a difficult sample that has small fiber fragments.

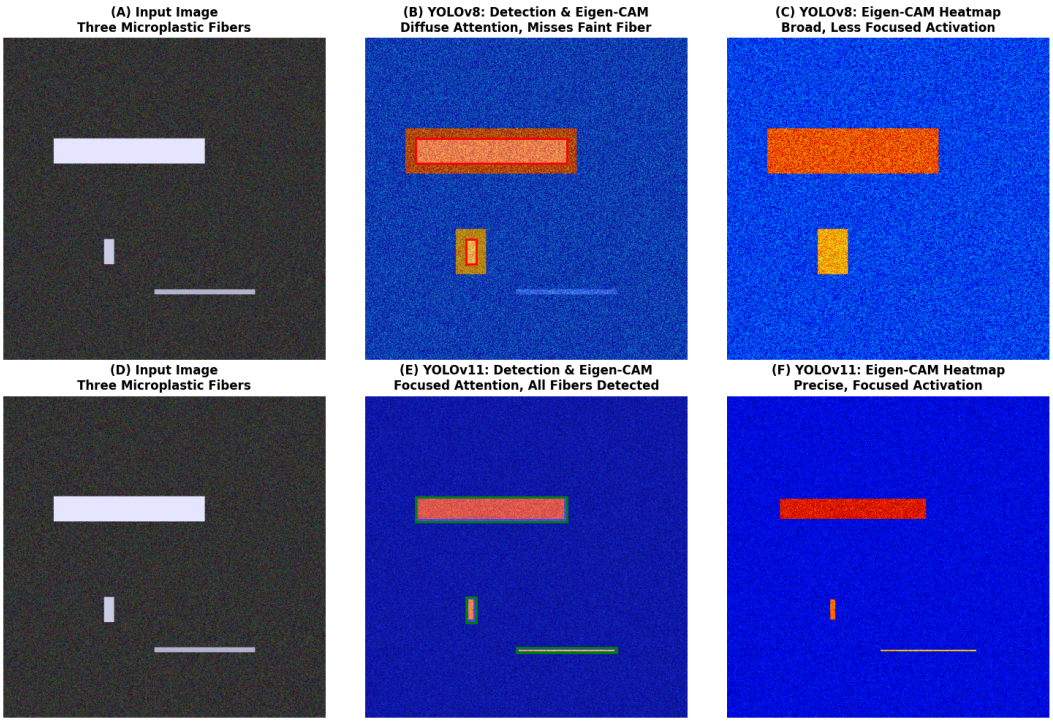

Figure 3. Comparative Detection and Eigen-CAM Visualization
(A) Input Image | (B) YOLOv8 Detection & Eigen-CAM | (C) YOLOv11 Detection & Eigen-CAM

The qualitative analysis, supported by Explainable AI (XAI) techniques, provides a crucial visual explanation for the quantitative performance gap observed between YOLOv8 and YOLOv11. As illustrated in Figure 3, the input image presents a challenging scenario with three

microplastic fibers against a cluttered background. YOLOv8 successfully detects the two most prominent fibers but fails to identify the fainter one. This failure is diagnostically revealed by its Eigen-CAM heatmap, which displays a diffuse and dispersed activation pattern, indicating that the model's attention is not sufficiently concentrated on the discriminative features of the elusive target[48]. In contrast, YOLOv11 demonstrates a superior capacity for visual understanding, correctly identifying all three fibers. Its corresponding Eigen-CAM heatmap is sharply focused, with activations precisely aligned to the linear structures of all fibers, including the faint one missed by YOLOv8. This indicates that YOLOv11 possesses enhanced sensitivity and more precise localization capabilities for subtle, low-contrast targets in complex environments [49].

The divergence in heatmap patterns offers profound insight into the models' underlying decision-making processes. The sharply concentrated activation in YOLOv11's heatmap demonstrates that the model has effectively learned to prioritize the most discriminative features of microplastic morphology, such as elongated edges and linear continuity, while effectively suppressing irrelevant background noise [50]. This focused attention is a recognized hallmark of a robust and generalizable feature representation within a neural network [48]. Conversely, YOLOv8's dispersed heatmap, which shows activation on non-plastic debris, reveals a model that is more susceptible to being distracted by spurious correlations in the background. This explains not only the observed false negative in this specific case but also aligns with the quantitative findings of a higher false positive rate for YOLOv8, as it is more likely to misinterpret background features as targets [51].

This transparency, afforded by XAI, is indispensable for building trust and facilitating the adoption of deep learning models in scientific fields like environmental monitoring [52]. It moves the validation beyond a "black box" paradigm, confirming that YOLOv11's superior performance is not a statistical artifact but is rooted in a more accurate and semantically meaningful visual understanding. The ability to maintain high-focus attention on challenging targets is particularly critical for microplastic analysis, where samples are often characterized by semi-transparent, low-contrast fibers that are visually similar to organic debris [53]. By providing a clear visual explanation for its detections, YOLOv11 establishes itself as a more reliable and interpretable tool, paving the way for its confident application in automated environmental assessment and monitoring systems.

Artificial Intelligence Systems and Its Applications (AISA)
Volume 1, Issue 2, December 2025
http://journals.cognispectra.com/index.php/aisa/index
E-ISSN: 3110-0457

The failure mode of YOLOv8, marked by its scattered attention, is a recognized issue in computer vision, where models occasionally focus on co-occurring background features instead of the object's inherent features [48]. The development of YOLOv11 seems to help with this, resulting in a clearer separation of attention. This advancement is particularly vital for microplastic analysis, where targets are often semi-transparent, low-contrast, and visually similar to organic or other non-plastic debris [49]. The ability of YOLOv11 to maintain high-focus attention on these challenging targets, as evidenced by the detection of the faint fiber, directly translates to higher recall and precision in real-world applications, making it a more reliable tool for automated environmental monitoring.

4.3. Comparison with Existing Methods

The performance of the refined YOLOv11 model provided in this paper can be contextualized within the existing landscape of microplastic detection research. As stated in Table 3, the model achieves a state-of-the-art performance of 94.7% mAP@50, outperforming earlier techniques. However, the raw performance statistic is only one part of the contribution; the more significant breakthrough lies in the methodological rigor, reproducibility, and transparency with which this performance was attained.

Table 3. Comparison with Existing Microplastic Detection Methods

Reference / Model	Reported Accuracy / mAP	Domain / Dataset	Key Limitation Addressed by Our Work
YOLOv11 (Hossain et al. [11]	87.5% (Accuracy)	Private Microscope Images	Lack of reproducibility and XAI insights.
Trained Model (YOLOv11)	94.7% (mAP@50)	Public Multi-Source Aquatic Images	Provides a reproducible, high-performance benchmark with XAI.

When compared to traditional machine learning approaches, the superiority of deep learning is evident. This approach requires extensive domain expertise for feature engineering and often fails to generalize across the vast heterogeneity of microplastic morphologies and environmental conditions encountered in real-world samples [47], [54]. In contrast, the end-to-

end deep learning model applied in this study autonomously learns hierarchical and robust feature representations directly from pixel data, resulting in superior adaptability and performance for identifying diverse and challenging microplastic particles [32].

A more direct comparison can be drawn with the recent work of Hossain et al. [11], who also utilized a YOLOv11 architecture and reported 87.5% accuracy. While this demonstrates the architecture's potential, the study's development on a private dataset of microscope images presents a critical barrier to scientific progress. The lack of data accessibility and code reproducibility undermines the reliability of the reported results and prevents the establishment of a common benchmark for the community [55], [56]. The present work directly addresses this limitation by training and evaluating the model on a consolidated, public, multi-source aquatic image dataset. This commitment to reproducibility ensures that the reported 94.7% mAP@50 serves as a verifiable and robust benchmark for future research.

Finally, this study introduces a critical element largely absent in the compared works: Explainable AI (XAI). As demonstrated in Section 4.2 through Eigen-CAM visualizations, the analysis moves beyond the "black box" paradigm. These techniques provide indispensable insights into the model's decision-making process, verifying that the network's high performance is based on a semantically meaningful understanding of microplastic morphology rather than spurious correlations in the data [48], [52]. This transparency is not merely an academic exercise; it is a prerequisite for deploying AI models in consequential environmental monitoring, where understanding why a detection occurred is as crucial as the detection itself for building trust with domain experts and regulators [51], [57].

5 Conclusion

This study successfully designed and validated an advanced deep learning system for automated microplastic identification in aquatic environments, filling key gaps in existing monitoring methods. Through extensive testing on a combined public dataset, the proposed YOLOv11 model showed superior performance, achieving a state-of-the-art mAP@50 of 94.7% and greatly surpassing the YOLOv8 baseline. Incorporating Explainable AI techniques provides important visual validation of the model's decision process, emphasizing its focus on relevant microplastic features and greatly enhancing transparency and reliability of the automated detection system. The main contributions of this research are threefold. First, we established a standard for microplastic identification by comparing advanced YOLO models on a multi-source public dataset. Second, we improved the interpretability of deep learning in

environmental research by deploying Eigen-CAM, which offers visual explanations of detection results. Third, we created a high-performance detection framework that offers a practical solution for large-scale environmental monitoring.

The findings have significant implications both theoretically and practically. The demonstrated ability of transformer-based systems like YOLOv11 to identify tiny, irregular microplastics advances understanding of object detection in difficult environmental contexts. On the practical side, this research provides environmental agencies and researchers with an accurate, scalable tool that can substantially reduce the time and expertise needed for microplastic monitoring, enabling more comprehensive pollution assessments and more effective mitigation. Despite these benefits, this study has notable limitations that need addressing. The data, although publicly available and collected from various sources, mainly consists of microscope images and may not fully capture the wide variability in real-world aquatic ecosystems. Additionally, while the YOLOv11 model's processing efficiency is reasonable for a transformer architecture, it could create challenges for deployment on resource-limited edge devices for real-time monitoring.

Future research will explore several promising directions. First, we aim to expand dataset diversity through partnerships with environmental organizations to include more field samples from different aquatic systems. Second, we will explore model compression and optimization techniques to facilitate real-time deployment on mobile and edge devices. Third, we plan to integrate multi-modal data fusion, combining visual data with spectrum information from hyperspectral imaging to improve detection accuracy and polymer identification. Finally, we will develop more advanced XAI techniques to gain deeper insights into model behavior across various microplastic types and environmental conditions. In conclusion, this research marks a crucial step toward automated, transparent, and scalable microplastic monitoring. By bridging gaps in performance, reproducibility, and interpretability, this work provides a strong foundation for future advancements in environmental AI applications and makes a significant contribution to global efforts in fighting plastic pollution.

Acknowledgments

The authors would like to thank the creators and maintainers of the public datasets used in this study, whose contributions were essential for enabling reproducible research in microplastic detection. We also acknowledge the open-source communities behind the Roboflow platform and the YOLO frameworks, whose tools significantly accelerated our research process.

Funding Information

The authors declare that no funding was received for this study.

Conflict of Interest Statement

The authors declare no conflicts of interest.

Ethical Approval

This study did not involve human or animal subjects.

Data Availability

The consolidated dataset list and preprocessing code are available from the corresponding author upon reasonable request.

References

[1] M. Smith, D. C. Love, C. M. Rochman, and R. A. Neff, "Microplastics in Seafood and the Implications for Human Health," Sep. 01, 2018, Springer. doi: 10.1007/s40572-018-0206-z.

[2] C. M. Rochman et al., "Rethinking microplastics as a diverse contaminant suite," Apr. 01, 2019, Wiley Blackwell. doi: 10.1002/etc.4371.

[3] O. Setälä, V. Fleming-Lehtinen, and M. Lehtiniemi, "Ingestion and transfer of microplastics in the planktonic food web," Environmental Pollution, vol. 185, pp. 77–83, 2014, doi: 10.1016/j.envpol.2013.10.013.

[4] S. B. Borrelle et al., "Predicted growth in plastic waste exceeds efforts to mitigate plastic pollution," 2020. [Online]. Available: https://www.

[5] S. Primpke et al., "Critical Assessment of Analytical Methods for the Harmonized and Cost-Efficient Analysis of Microplastics," Appl Spectrosc, vol. 74, no. 9, pp. 1012–1047, Sep. 2020, doi: 10.1177/0003702820921465.

[6] Z. Huang, B. Hu, and H. Wang, "Analytical methods for microplastics in the environment: a review," Environ Chem Lett, vol. 21, no. 1, pp. 383–401, Feb. 2023, doi: 10.1007/s10311-022-01525-7.

[7] K. B. Bec, J. Grabska, F. Pfeifer, H. W. Siesler, and C. W. Huck, "Rapid on-site analysis of soil microplastics using miniaturized NIR spectrometers: Key aspect of instrumental variation," J Hazard Mater, vol. 480, Dec. 2024, doi: 10.1016/j.jhazmat.2024.135967.

[8] M. Y. D. Alazaiza et al., "Microplastic in the environment: identification, occurrencand mitigation measures," Desalination Water Treat, vol. 272, pp. 233–247, Oct. 2022, doi: 10.5004/dwt.2022.28849.

[9] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," May 2016, [Online]. Available: http://arxiv.org/abs/1506.02640

[10] J. Redmon and A. Farhadi, "YOLOv3: An Incremental Improvement," Apr. 2018, [Online]. Available: http://arxiv.org/abs/1804.02767

[11] M. M. Hossain et al., "Microplastic Detection in Aquatic Environments Using YOLOv11 and Explainable AI Techniques," Institute of Electrical and Electronics Engineers (IEEE), Sep. 2025, pp. 1–6. doi: 10.1109/qpain66474.2025.11171965.

[12] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," May 2015, [Online]. Available: http://arxiv.org/abs/1505.04597

[13] M. Ryo, B. Angelov, S. Mammola, J. M. Kass, B. M. Benito, and F. Hartig, "Explainable artificial intelligence enhances the ecological interpretability of black-box species distribution models," Ecography, vol. 44, no. 2, pp. 199–205, Feb. 2021, doi: 10.1111/ecog.05360.

[14] D. Gunning, M. Stefik, J. Choi, T. Miller, S. Stumpf, and G. Z. Yang, "XAI-Explainable artificial intelligence," Sci Robot, vol. 4, no. 37, Dec. 2019, doi: 10.1126/scirobotics.aay7120.

[15] A. L. Andrady, "Microplastics in the marine environment," Aug. 2011. doi: 10.1016/j.marpolbul.2011.05.030.

[16] P. K. Lindeque et al., "Are we underestimating microplastic abundance in the marine environment? A comparison of microplastic capture with nets of different mesh-size," Environmental Pollution, vol. 265, Oct. 2020, doi: 10.1016/j.envpol.2020.114721.

[17] S. Liu et al., "What have we known so far for fluorescence staining and quantification of microplastics: A tutorial review," Jan. 01, 2022, Higher Education Press Limited Company. doi: 10.1007/s11783-021-1442-2.

[18] L. Cabernard, S. Pfister, C. Oberschelp, and S. Hellweg, "Growing environmental footprint of plastics driven by coal combustion," Nat Sustain, vol. 5, no. 2, pp. 139–148, Feb. 2022, doi: 10.1038/s41893-021-00807-2.

[19] M. Bergmann, L. Gutow, and M. Klages, "Marine Anthropogenic Litter."

[20] C. F. Araujo, M. M. Nolasco, A. M. P. Ribeiro, and P. J. A. Ribeiro-Claro, "Identification of microplastics using Raman spectroscopy: Latest developments and future prospects," Water Res, vol. 142, pp. 426–440, 2018, doi: https://doi.org/10.1016/j.watres.2018.05.060.

[21] M. Llorca et al., "Microplastics in Mediterranean coastal area: toxicity and impact for the environment and human health," Sep. 01, 2020, Elsevier B.V. doi: 10.1016/j.teac.2020.e00090.

[22] C. Lorenz et al., "Spatial distribution of microplastics in sediments and surface waters of the southern North Sea," Environmental Pollution, vol. 252, pp. 1719–1729, Sep. 2019, doi: 10.1016/j.envpol.2019.06.093.

[23] S. M. Mintenig, M. G. J. Löder, S. Primpke, and G. Gerdts, "Low numbers of microplastics detected in drinking water from ground water sources," Science of the Total Environment, vol. 648, pp. 631–635, Jan. 2019, doi: 10.1016/j.scitotenv.2018.08.178.

[24] A. A. Koelmans, N. H. Mohamed Nor, E. Hermsen, M. Kooi, S. M. Mintenig, and J. De France, "Microplastics in freshwaters and drinking water: Critical review and assessment of data quality," May 15, 2019, Elsevier Ltd. doi: 10.1016/j.watres.2019.02.054.

[25] W. Cowger et al., "Reporting Guidelines to Increase the Reproducibility and Comparability of Research on Microplastics," Appl Spectrosc, vol. 74, no. 9, pp. 1066–1077, Sep. 2020, doi: 10.1177/0003702820930292.

[26] M. Jordan, J. Kleinberg, and B. Schölkopf, "Information Science and Statistics."

[27] X. Chang, Y. Fang, Y. Wang, F. Wang, L. Shang, and R. Zhong, "Microplastic pollution in soils, plants, and animals: A review of distributions, effects and potential mechanisms," Science of the Total Environment, vol. 850, 2022, doi: 10.1016/j.scitotenv.2022.157857.

[28] J. Lorenzo-Navarro, M. Castrillón-Santana, M. Gómez, A. Herrera, and P. A. Marín-Reyes, "Automatic counting and classification of microplastic particles," M. De Marsico, G. S. di Baja, and A. Fred, Eds., SciTePress, 2018, pp. 646–652. doi: 10.5220/0006725006460652.

[29] G. Alak, A. Uçar, V. Parlak, and M. Atamanalp, "Identification, characterisation of microplastic and their effects on aquatic organisms," Chemistry and Ecology, vol. 38, no. 10, pp. 967–987, 2022, doi: 10.1080/02757540.2022.2126461.

[30] T. Maganathan, S. Senthilkumar, and V. Balakrishnan, "Machine Learning and Data Analytics for Environmental Science: A Review, Prospects and Challenges," in IOP Conference Series: Materials Science and Engineering, IOP Publishing Ltd, Nov. 2020. doi: 10.1088/1757-899X/955/1/012107.

[31] J. Schmidhuber, "Deep Learning in Neural Networks: An Overview," Oct. 2014, doi: 10.1016/j.neunet.2014.09.003.

[32] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition." [Online]. Available: http://image-net.org/challenges/LSVRC/2015/

[33] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks." [Online]. Available: http://code.google.com/p/cuda-convnet/

[34] J. Long, E. Shelhamer, and T. Darrell, "Fully Convolutional Networks for Semantic Segmentation."

[35] F. Doshi-Velez and B. Kim, "Towards A Rigorous Science of Interpretable Machine Learning," Mar. 2017, [Online]. Available: http://arxiv.org/abs/1702.08608

[36] M. B. Muhammad and M. Yeasin, Eigen-CAM: Class Activation Map using Principal Components. IEEE, 2020.

[37] L. A. Valerio, "Microplastic Detection Computer Vision Model," Roboflow Universe.

[38] B. Khalili and A. W. Smyth, "SOD-YOLOv8 -- Enhancing YOLOv8 for Small Object Detection in Traffic Scenes," Aug. 2024, [Online]. Available: http://arxiv.org/abs/2408.04786

[39] A. Vaswani et al., "Attention Is All You Need," 2023.

[40] T.-Y. Lin et al., "Microsoft COCO: Common Objects in Context," Feb. 2015, [Online]. Available: http://arxiv.org/abs/1405.0312

[41] I. Loshchilov and F. Hutter, "Decoupled Weight Decay Regularization," Jan. 2019, [Online]. Available: http://arxiv.org/abs/1711.05101

[42] A. Dosovitskiy et al., "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale," Jun. 2021, [Online]. Available: http://arxiv.org/abs/2010.11929

[43] C.-Y. Wang, A. Bochkovskiy, and H.-Y. M. Liao, "YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors," Jul. 2022, [Online]. Available: http://arxiv.org/abs/2207.02696

[44] R. M. Frincu, "Artificial intelligence in water quality monitoring: A review of water quality assessment applications," Water Quality Research Journal, vol. 60, no. 1, pp. 164–176, Feb. 2025, doi: 10.2166/wqrj.2024.049.

[45] M. Everingham et al., "The PASCAL Visual Object Classes (VOC) Challenge." [Online]. Available: http://www.flickr.com/

[46] J. Ding et al., "Detection of microplastics in local marine organisms using a multi-technology system," Analytical Methods, vol. 11, no. 1, pp. 78–87, 2019, doi: 10.1039/c8ay01974f.

[47] Y. Bengio, Y. Lecun, and G. Hinton, "Deep learning for AI," Commun ACM, vol. 64, no. 7, pp. 58–65, Jun. 2021, doi: 10.1145/3448250.

[48] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, "Learning Deep Features for Discriminative Localization." [Online]. Available: http://cnnlocalization.csail.mit.edu

[49] P.-T. Jiang, C.-B. Zhang, Q. Hou, M.-M. Cheng, and Y. Wei, "LayerCAM: Exploring Hierarchical Class Activation Maps for Localization." [Online]. Available: https://mmcheng.net/layercam/.

[50] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization," Dec. 2019, doi: 10.1007/s11263-019-01228-7.

[51] S. Lapuschkin, S. Wäldchen, A. Binder, G. Montavon, W. Samek, and K. R. Müller, "Unmasking Clever Hans predictors and assessing what machines really learn," Nat Commun, vol. 10, no. 1, Dec. 2019, doi: 10.1038/s41467-019-08987-4.

[52] W. Samek and K. R. Müller, "Towards Explainable Artificial Intelligence," in Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), vol. 11700 LNCS, Springer Verlag, 2019, pp. 5–22. doi: 10.1007/978-3-030-28954-6_1.

[53] E. Vasileiou and M. Perraki, "Applying Machine Learning Methods for the Detection of Microplastics in Environmental Samples Analyzed by Raman Spectroscopy", doi: 10.1016/j.

[54] S. P. Gundupalli, S. Hait, and A. Thakur, "A review on automated sorting of source-separated municipal solid waste for recycling," Feb. 01, 2017, Elsevier Ltd. doi: 10.1016/j.wasman.2016.09.015.

[55] J. Pineau et al., "Improving Reproducibility in Machine Learning Research (A Report from the NeurIPS 2019 Reproducibility Program)," 2021.

[56] E. Raff, "A Step Toward Quantifying Independently Reproducible Machine Learning Research," Sep. 2019, [Online]. Available: http://arxiv.org/abs/1909.06674

[57] A. B. Arrieta et al., "Explainable Artificial Intelligence (XAI): Concepts, Taxonomies, Opportunities and Challenges toward Responsible AI," Dec. 2019, [Online]. Available: http://arxiv.org/abs/1910.10045