IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

AI- Powered Waste Segregation System using YOLOv8 and EfficientNet

B. RAHUL, KETHANN PRASANA SRIDHAR, A. KISHORE, Dr. S. SENTHAMIZH SELVI

Sri Venkateswara College of Engineering, Chennai, Tamil Nadu, India

DOI: https://doi.org/10.56975/ijcrt.v14i7.309068

Abstract- Waste segregation is a critical challenge in modern urban environments, especially in countries like India where large volumes of mixed waste are generated daily. Traditional waste segregation methods heavily rely on manual work, making the process time-consuming, error-prone, and unsafe. To address this, we propose a system with a hybrid pipeline using YOLOv8m and EfficientNet for detection, classification, and intelligent decision-making. The system collects real-time video input, converts it into frames, and preprocesses images for detection. YOLOv8m detects waste with 64% accuracy across 5 classes(plastic, metal, glass, paper, organic) , EfficientNet verifies it with 94% accuracy across 8 classes(plastic, metal, glass, paper, cardboard, organic, textiles, wood), and the data is passed to a secondary decision engine to further classify waste as biodegradable, recyclable, or hazardous. The system handles multiple objects under real-world conditions and helps analyze waste trends, improves recycling efficiency, and reduces human effort, making it suitable for real-time deployment.

Keywords: Waste Segregation, YOLOv8, EfficientNet, Deep Learning, Computer Vision, Smart Waste Management

I. INTRODUCTION

Waste management has emerged as a critical challenge in rapidly urbanizing regions, particularly in densely populated countries such as India, where millions of tonnes of mixed municipal solid waste are generated daily. Improper segregation at the source results in contamination of recyclable materials, accelerated landfill saturation, increased greenhouse gas emissions, and significant public health risks associated with improper handling of hazardous waste [1]. The manual nature of conventional waste sorting operations further compounds these challenges, as human workers are exposed to biohazardous and chemically dangerous materials while performing repetitive, error-prone tasks at scale.

The adoption of machine learning and computer vision techniques has introduced promising avenues for automating waste classification. Convolutional neural network (CNN)-based classifiers and object detection frameworks have demonstrated reasonable accuracy under controlled laboratory conditions [3], [7]. However, these approaches exhibit notable limitations when deployed in real-world environments characterized by low and inconsistent lighting, cluttered and multi-object frames, overlapping waste items, and high inter-class visual similarity between material types such as different plastics or paper-based waste. Furthermore, most existing systems produce single-label outputs and lack the multi-dimensional categorization required for practical, safety-aware waste handling protocols.

These limitations collectively highlight the need for a robust, interpretable, and hierarchically aware waste segregation pipeline capable of operating under real-world constraints. In this work, we propose a multi-stage hybrid AI system that integrates pre-filtering, object detection, fine-grained classification, multi-label decision analysis, and explainable AI within a unified real-time pipeline. The system acquires live video input, filters irrelevant frames through a lightweight MLP-based pre-filter, detects waste objects using YOLOv8m, verifies and refines predictions through an EfficientNet-based classification and verification layer, and applies a hierarchical secondary decision engine to produce multi-dimensional waste attributes. Grad-CAM visualizations are generated in real time to provide transparent justification for each segregation decision, particularly for low-confidence and ambiguous predictions.

The key contributions of this work are as follows:

1. A **two-stage pre-filtering pipeline** comprising an Empty Frame Screener and a lightweight MLP Waste vs. No-Waste Filter, reducing unnecessary computation before the core AI pipeline.
2. A **hybrid YOLOv8m and EfficientNet pipeline** with a cross-validation Verification Layer that suppresses false positives and resolves inter-model disagreements, achieving detection accuracy of 64% across five classes and classification accuracy of 94% across eight classes.
3. An **Unknown class mechanism** in YOLOv8m that routes low-confidence detections for XAI analysis rather than forcing incorrect class assignments, improving downstream reliability.
4. A **multi-layer Secondary Decision Engine** that performs simultaneous multi-label categorization across Biodegradable, Hazardous, Recyclable, and Contamination dimensions, with a priority-based segregation protocol (Hazardous > Contaminated > Recyclable).
5. An integrated **real-time Explainable AI (XAI) layer** using Grad-CAM activation maps and attention visualizations, providing transparent, operator-interpretable justification for every classification decision.

II. RELATED WORK

A. Object Detection for Real-Time Waste Sorting

The YOLO family of architectures has been widely adopted for real-time waste detection due to its single-pass inference speed. Dipo et al. [5] demonstrated YOLOv12 for high-speed waste detection on conveyor belts, achieving strong performance under controlled conditions. Xiao et al. [8] proposed DBS-YOLO, an enhanced YOLOv8n variant incorporating deformable convolutions to address mutual occlusion in hazardous waste images. While these methods perform well for isolated object detection, they process every input frame uniformly, resulting in unnecessary computation when frames contain no waste. The proposed system addresses this through an empty frame detection mechanism that filters irrelevant frames prior to the detection stage, reducing processing overhead.

B. Hybrid and Cascaded Classification Architectures

Single-model approaches have shown limitations in handling visually ambiguous waste categories. To improve reliability, recent work has explored cascaded and hybrid pipelines. Nahiduzzaman et al. [7] proposed a three-stage pipeline incorporating biodegradability checks before narrowing classification to 36 specific object classes, demonstrating the benefit of hierarchical decision-making. Ahmad et al. [4] developed a deep learning framework separating waste into 12 distinct categories for urban sustainability applications. EfficientNet-based architectures have been shown to consistently outperform traditional CNNs when balancing accuracy against computational cost [3]. Building on these findings, the proposed system integrates EfficientNet as a verification stage following YOLOv8m detection, cross-validating outputs to suppress false positives and handle visually similar categories.

C. Multi-Label and Priority-Based Segregation

Beyond basic class prediction, practical waste management systems require nuanced, multi-attribute categorization. Prior works have explored hierarchical sorting [7] and multi-category frameworks [4], but these systems output a single label per item without simultaneous risk assessment. The proposed decision engine addresses this by applying multi-label output, assigning each detected object attributes across biodegradable, recyclable, and hazardous dimensions concurrently, and resolving conflicts using a priority-based protocol.

D.Explainability and Interpretable AI in Waste Classification

A critical barrier to deploying AI systems in public infrastructure is the lack of transparency in model decision-making. Neural networks are inherently opaque, making it difficult for operators to trust or validate automated sorting decisions, particularly for hazardous waste misclassification.

Recent work has begun addressing this through Explainable AI (XAI) techniques. Nahiduzzaman et al. [7] integrated Grad-CAM and SHAP visualizations to verify that classification models focused on actual waste objects rather than background artifacts, significantly improving deployment confidence. Similarly, Score-CAM has been applied in recyclable waste classification systems to generate saliency maps that highlight discriminative regions used during inference.

Despite these efforts, XAI integration remains largely limited to offline analysis and post-hoc evaluation in existing literature. Most deployed systems do not provide real-time explainability during inference. The proposed system addresses this gap by generating Grad-CAM activation maps during live classification, particularly for low-confidence and ambiguous detections routed to the *Unknown* class. This provides operators with visual justification for each segregation decision, improving system transparency and enabling faster error diagnosis in real-world deployment.

E. Summary of Research Gaps

Although existing literature addresses high-speed detection and high-accuracy classification independently, there is limited work combining pre-processing filters with hierarchical, multi-label risk assessment in a unified real-time pipeline. Table I summarizes key prior works and their limitations relative to the proposed approach.

TABLE I. COMPARISON OF RELATED WORKS

Reference	Model	Classes	Real-Time	XAI Support	Multi-Label Output
Dipo et al. [5]	YOLOv12	—	✓	✗	✗
Xiao et al. [8]	DBS-YOLO	Hazardous	✓	✗	✗
Ahmad et al. [4]	Deep Learning	12	✗	✓(offline)	✗
Nahiduzzaman et al. [7]	Multi-stage CNN	36	✗	✗	✗
Gupta et al. [3]	Transfer Learning	—	✗	✗	✗
proposed	YOLOv8m + EfficientNet	8	✓	✓(real-time)	✓

III. METHODOLOGY

A. System Overview

The proposed system is a multi-stage hybrid AI pipeline designed for real-time waste detection, classification, and intelligent segregation. As illustrated in Fig. 1, the architecture consists of five sequential stages: input acquisition and filtering, preprocessing, core AI detection and classification, secondary decision analysis, and output with explainability. Each stage is designed to progressively refine waste identification while maintaining real-time operational efficiency.

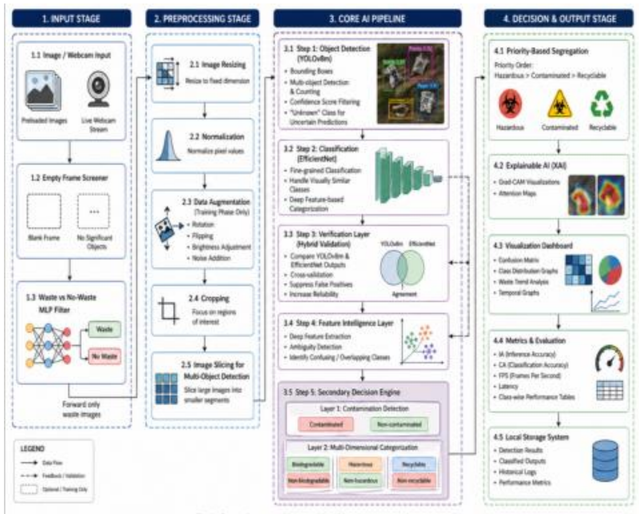

B. Input Acquisition and Pre-Filtering

The system accepts two input modes: preloaded static images for offline analysis and live video streams from a camera mounted over a conveyor setup. Video streams are decomposed into individual frames based on the configured frames-per-second (FPS) rate.

Prior to any deep learning inference, two lightweight filtering mechanisms are applied to eliminate irrelevant data. First, an **Empty Frame Screener** discards blank or near-empty frames by evaluating pixel intensity distribution, ensuring that subsequent processing is only triggered when visual content is present. Second, a **Waste vs. No-Waste MLP Filter**— a lightweight Multi-Layer Perceptron trained on high-level image features— performs a binary classification to determine whether a frame contains waste material. Only frames classified as waste-positive are forwarded to the core pipeline. This two-stage pre-filtering significantly reduces unnecessary computation and prevents false activations on empty conveyor frames.

C. Preprocessing

Frames passing the pre-filtering stage undergo a preprocessing pipeline to optimize input quality for deep learning inference. All images are resized to a fixed resolution consistent with the input requirements of YOLOv8m and EfficientNet. Pixel values are normalized to the [0, 1] range to improve model stability and convergence. Lighting enhancement is applied to handle low-visibility and inconsistent illumination conditions encountered in real-world conveyor environments.

During training, data augmentation techniques including random rotation, horizontal flipping, brightness variation, and noise injection are applied to improve model generalization and reduce overfitting. Additionally, **image slicing** is employed on high-resolution frames, dividing them into overlapping sub-regions to improve the detection of small and closely positioned waste objects that may otherwise be missed by standard full-frame inference.

D. Object Detection using YOLOv8m

Object detection is performed using **YOLOv8m**, selected for its balance between detection accuracy and inference speed, making it suitable for real-time deployment. The model processes each preprocessed frame and generates bounding boxes, class predictions, and associated confidence scores for all detected objects simultaneously.

The detection model is trained across **five waste categories**. To handle uncertain or ambiguous detections, an **"Unknown" class** is introduced, allowing the model to flag low-confidence predictions rather than forcing an incorrect classification. Detections below a defined confidence threshold are automatically routed to this class and flagged for XAI analysis at the output stage. Detected object regions are cropped from the original frame and forwarded individually to the classification stage.

E. Fine-Grained Classification using EfficientNet

Each cropped detection is independently processed by an **EfficientNet** classifier for fine-grained, feature-level categorization. EfficientNet is selected due to its compound scaling strategy, which proportionally scales network depth, width, and resolution, achieving high accuracy with reduced computational overhead compared to conventional CNNs.

The classifier is trained across **eight waste categories**, providing a more granular categorization than the detection stage. This two-model approach— coarse detection followed by fine classification— is specifically designed to reduce misclassification of visually similar waste types such as different plastics or paper-based materials.

F. Verification Layer and Feature Intelligence

The outputs from YOLOv8m and EfficientNet are passed to a **Verification Layer** that cross-validates the predictions from both models. When both models agree on a class, the prediction is confirmed with high confidence. When disagreement occurs, the system does not default to either model; instead, it flags the object for deeper analysis through the **Feature Intelligence Layer**.

The Feature Intelligence Layer performs high-level feature extraction on ambiguous detections to resolve conflicts between the detection and classification outputs. This layer identifies overlapping or confusing class boundaries and produces a reconciled prediction, reducing false positives and improving overall classification reliability in complex multi-object scenarios.

G. Secondary Decision Engine

The verified classification output is passed to a **Secondary Decision Engine** that applies hierarchical, multi-label categorization to each detected waste object. The engine operates across two layers:

Layer 1 — Contamination Detection: Each object is assessed for contamination status, categorized as either *Contaminated* or *Non-Contaminated*. Contaminated waste requires special handling and is escalated in the segregation priority queue regardless of its material type.

Layer 2 — Multi-Dimensional Categorization: Each object is simultaneously evaluated across three independent binary dimensions: Biodegradable vs. Non-Biodegradable, Hazardous vs. Non-Hazardous, and Recyclable vs. Non-Recyclable. This multi-label output provides richer information than single-label classification, enabling downstream systems to make more informed physical sorting decisions.

Priority-Based Segregation Protocol: Conflicts between multi-label outputs are resolved using a defined priority hierarchy: *Hazardous* > *Contaminated* > *Recyclable*. This ensures that safety-critical waste is always handled first, regardless of its other categorical attributes.

H. Output, Explainability, and Storage

The final output stage delivers segregation decisions alongside interpretability and analytical tools.

Explainable AI (XAI): To improve transparency and operator trust, **Grad-CAM (Gradient-weighted Class Activation Mapping)** visualizations are generated for each classification, highlighting the spatial regions of the input image that most influenced the model's decision. This is particularly applied to low-confidence detections and items routed to the Unknown class, providing visual justification for borderline predictions. Attention maps are also generated to supplement Grad-CAM analysis for the EfficientNet classification stage.

Visualization Dashboard: A real-time analytics dashboard presents system performance metrics including confusion matrices, class-wise distribution graphs, waste trend analysis, and temporal variation graphs. This supports operational monitoring and long-term waste pattern analysis.

Performance Metrics: The system is evaluated using Inference Accuracy (IA), Classification Accuracy (CA), Frames Per Second (FPS), and end-to-end latency to ensure both accuracy and real-time operational viability.

Local Storage: All detection results, classification outputs, Grad-CAM visualizations, and performance logs are stored locally to support model retraining, audit trails, and dataset expansion over time.

i) Raw input

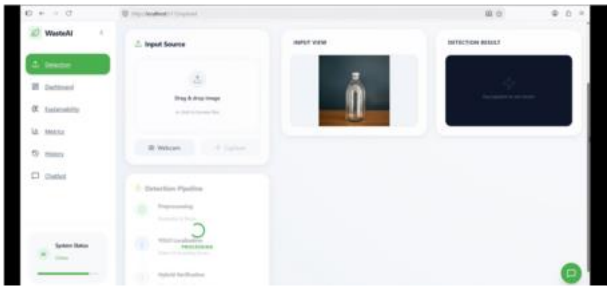

ii) Yolov8 input

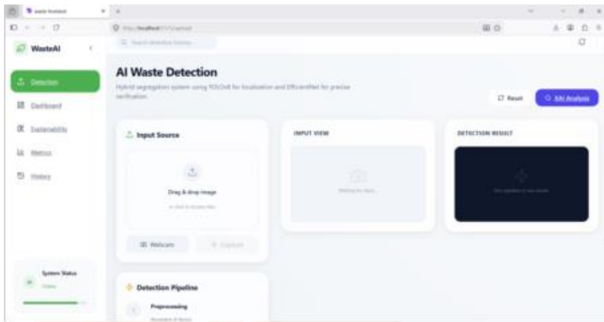

iii) Yolov8 bounding box output

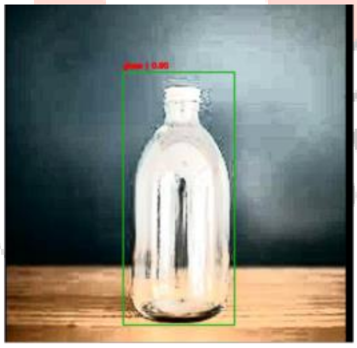

iv) Verification using efficient and system performance

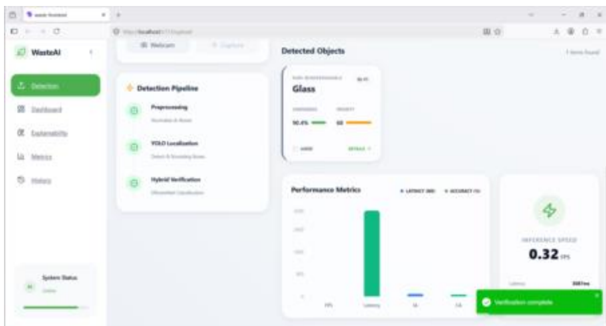

v) contaminated vs non-contaminated, biodegradable vs non-biodegradable, recyclable vs non-recyclable, hazardous vs non-hazardous decision layer classification

Fig. :

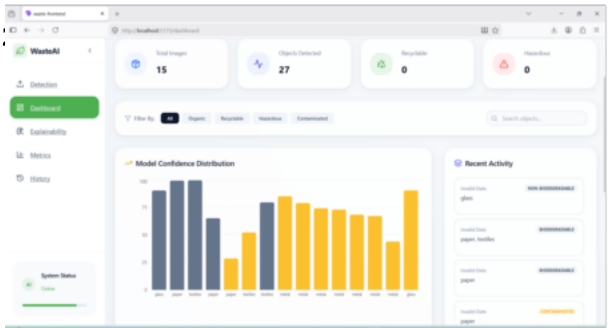

vi) explainable AI

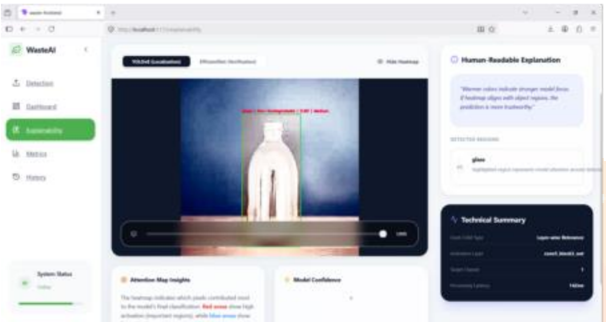

vii) Evaluation metrics

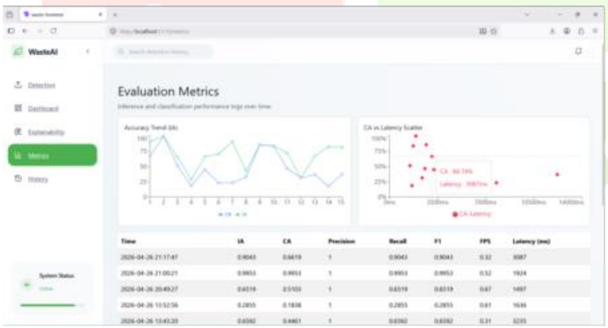

viii) History page

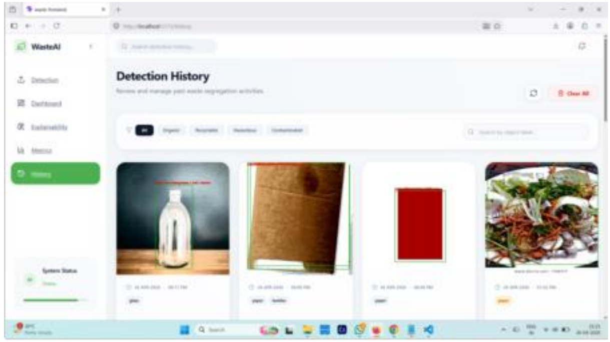

III. RESULTS AND DISCUSSION

A. Dataset and Experimental Setup

The proposed system was trained and evaluated on a custom-assembled dataset comprising **5,000 images**, annotated separately for two stages. The YOLOv8m detection model was trained on **five waste categories**: plastic, metal, glass, paper, and organic. The EfficientNet classification model was trained on **eight waste categories**: plastic, metal, glass, paper, cardboard, textile, organic, and wood. The additional three classes in EfficientNet— cardboard, textile, and wood— enable finer granularity at the classification stage compared to the detection stage, reflecting the model's role as a fine-grained verification layer. The dataset was split into training, validation, and testing subsets following a **70:20:10 ratio** for both models.

All experiments were conducted on **Google Colaboratory (Google Colab)**, utilizing a cloud-based **NVIDIA Tesla T4 GPU** with 15GB GPU RAM and 12.7GB system RAM. Training was performed using a Python 3.10 runtime environment with CUDA acceleration enabled. YOLOv8m was trained using the Ultralytics framework and EfficientNet was implemented using PyTorch. The YOLOv8m detection model was trained using the **SGD optimizer** with an initial learning rate of **0.01**, momentum of **0.937**, and weight decay of **0.0005** for a maximum of **150 epochs** with early stopping patience of **25 epochs**. The input resolution was fixed at **640× 640** pixels with a batch size of **16**. Transfer learning was applied by **freezing the first 10 backbone layers**, allowing only the detection head and upper layers to adapt to the waste classification dataset.

Several augmentation strategies were employed to improve model robustness under real-world conditions. **Mosaic augmentation** (factor 1.0) was applied to improve multi-object detection performance, while **MixUp was disabled** as it was found to degrade detection accuracy for this dataset. Horizontal flipping (probability 0.5), random scaling (factor 0.5), and translation (factor 0.1) were applied for spatial diversity. **HSV-based colour jittering**— hue (0.015), saturation (0.6), and value (0.4) was applied to improve robustness under varying lighting conditions, which is critical for real-world conveyor environments. Validation was performed after every epoch with training plots generated for loss and accuracy monitoring.

The EfficientNet - B0 classification model was fine-tuned from **ImageNet-pretrained weights** using the **Adam optimizer** with an initial learning rate of **1×10^{-5}** and a batch size of **32** over **20 epochs**. Transfer learning was employed by partially freezing the earlier convolutional layers while allowing the final convolutional blocks and classifier head to adapt to the waste classification dataset. A **ReduceLROnPlateau scheduler** dynamically reduced the learning rate by a factor of 0.2 upon validation loss plateau for 3 consecutive epochs, ensuring stable convergence. To address class imbalance particularly for underrepresented categories such as textile and paper—a **weighted cross-entropy loss function** was applied, assigning higher penalty weights to lower-performing classes. **Early stopping** with a patience of 7 epochs was employed to prevent overfitting, with the best-performing checkpoint saved based on minimum validation loss.

The complete training configuration is summarized in Table III.

TABLE II. Dataset Class Distribution *(insert your actual class names and image counts)*

Yolov8m

class	training	validation	testing	total
Plastic	700	200	100	1000
metal	700	200	100	1000
glass	700	200	100	1000
paper	700	200	100	1000
organic	700	200	100	1000
total	3500	1000	500	5000

EfficientNet-B0

class	training	validation	testing	total
Plastic	700	200	100	1000
metal	700	200	100	1000
glass	700	200	100	1000
paper	700	200	100	1000
organic	700	200	100	1000
textile	700	200	100	1000
cardboard	700	200	100	1000
wood	700	200	100	1000
Total	5600	1600	800	8000

TABLE III. Model Training Configuration *(insert epochs, batch size, learning rate, optimizer, hardware)*

Parameter	Yolov8m	EfficientNet - B0
Hardware	Google Colab (NVIDIA Tesla T4 GPU)	Google Colab (NVIDIA Tesla T4 GPU)
Framework	Ultralytics	PyTorch
Pretrained Weights	COCO	ImageNet
Optimizer	SGD	Adam
Initial Learning Rate	0.01	1×10^{-5}
Momentum	0.937	—
Weight Decay	0.0005	—
LR Scheduler	—	ReduceLROnPlatea u (factor 0.2, patience 3)
Batch Size	16	32

Input Resolution	640× 640	224× 224
Epochs	150	20
Early Stopping Patience	25	7
Loss Function	—	Weighted Cross-Entropy
Layer Freezing	First 10 backbone layers	Earlier layers frozen
Mosaic Augmentation	1.0	—
MixUp	Disabled	—
Horizontal Flip	0.5	—
HSV Hue / Sat / Val	0.015/0.6/0.4	—
Number of Classes	5	8

TABLE IV. YOLOv8m Class-wise Performance

Class	**recision**	**Recall**	**mAP50**	**hAP50-95**
Plastic	0.712	0.581	0.653	0.413
Metal	0.674	0.721	0.734	0.456
Glass	0.757	0.783	0.805	0.547
Paper	0.579	0.567	0.568	0.358
Organic	0.542	0.378	0.435	0.19
Mean	0.653	0.606	0.640	0.393

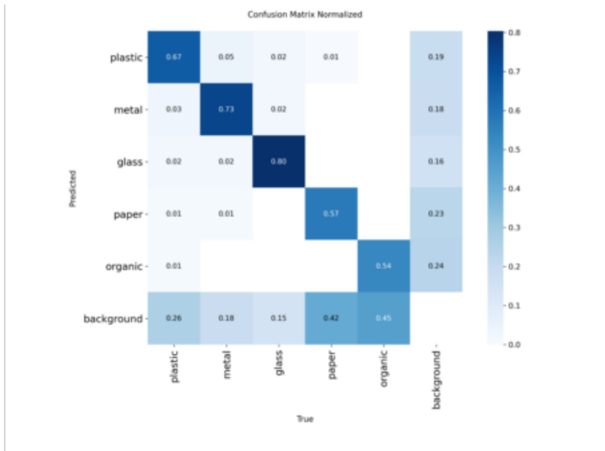

Fig 3. Confusion matrix(Yolov8m)

B. YOLOv8m Detection Performance

The YOLOv8m model was evaluated on the test set following two stages of optimization. In the initial training phase, the model achieved a detection accuracy of approximately [40/64]% due to limitations in dataset quality, class imbalance, and insufficient augmentation. Following dataset refinement, augmentation strategies including random rotation, brightness adjustment, and horizontal flipping, along with image slicing for multi-object frames, the detection accuracy improved to approximately **64%** across **5** waste classes.

Fig. 3 illustrates the training and validation accuracy and loss curves for YOLOv8m, demonstrating stable convergence after optimization with no significant overfitting. The confusion matrix in Fig. 5 presents the class-wise prediction distribution, indicating reduced false positives across all categories following optimization. Table IV reports the class-wise Precision, Recall, and F1-Score for the final YOLOv8m model.

The introduction of the **Unknown class** effectively handled low-confidence detections, routing ambiguous predictions away from forced misclassification and into the XAI analysis stage. This improved the overall reliability of downstream classification.

C. EfficientNet Classification Performance

The EfficientNet classifier achieved a classification accuracy of **94%** across eight waste categories, demonstrating strong fine-grained discrimination capability. Fig. 4 presents the training and validation accuracy and loss curves, confirming consistent convergence across all epochs. The confusion matrix in Fig. 6 shows high diagonal dominance, indicating reliable class-wise prediction with minimal inter-class confusion.

Table V presents the class-wise Precision, Recall, and F1-Score for EfficientNet. The model demonstrated particularly strong performance on visually distinct categories while maintaining acceptable accuracy on visually similar classes such as [example: different plastic types or paper-based materials], where fine-grained feature extraction was critical.

TABLE V. EfficientNet Class-wise Performance

Class	Precision	Recall	F1-score	support
cardboard	0.98	0.99	0.99	100
Metal	0.93	0.94	0.94	100
Glass	0.96	0.94	0.95	100
Paper	0.85	0.99	0.91	100
Organic	0.96	0.94	0.95	100
textiles	0.94	0.85	0.89	100
plastic	0.97	0.91	0.94	100
wood	0.97	0.98	0.98	100
accuracy			0.94	800
Macro avg	0.94	0.94	0.94	800
Weighted avg	0.94	0.94	0.94	800

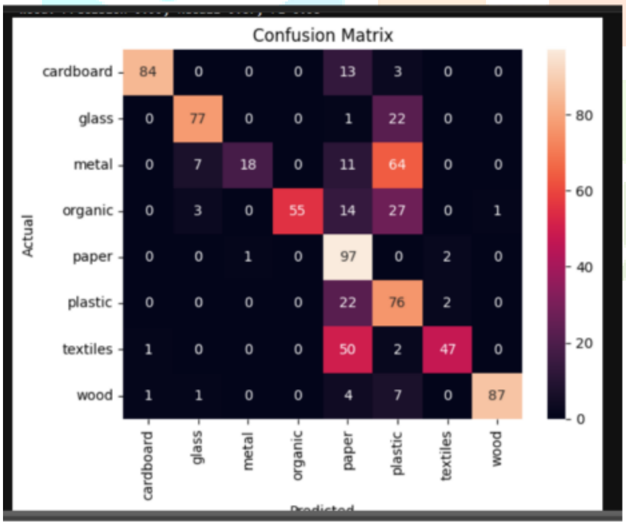

Fig 4. Confusion matrix(EfficientNet)

D. Hybrid Pipeline and Verification Layer Performance

The integration of YOLOv8m detection with EfficientNet classification through the Verification Layer produced measurable improvements in overall prediction reliability. Cases where both models agreed on a class label demonstrated the highest confidence scores, while disagreements were successfully resolved through the Feature Intelligence Layer in the majority of instances.

Fig. 7 presents a comparative bar chart of class-wise Precision, Recall, and F1-Score across both models, illustrating the complementary strengths of the detection and classification stages. The combination of both models reduced overall misclassification compared to either model operating independently, particularly for overlapping and visually ambiguous waste categories.

E. Explainable AI— Grad-CAM Analysis

To validate the interpretability of the proposed system, Grad-CAM activation maps were generated for representative samples across all waste categories. Fig. 8 presents selected Grad-CAM visualizations for both high-confidence correct predictions and low-confidence Unknown class predictions.

As shown in Fig. 8(a), for correctly classified high-confidence samples, the activation maps confirm that the EfficientNet model focuses on discriminative material-specific regions such as surface texture and object shape, rather than background artifacts. Fig. 8(b) illustrates activation maps for ambiguous predictions routed to the Unknown class, where the model's attention is distributed across the background rather than the object, validating the system's decision to flag these cases for review rather than forcing an incorrect label.

This XAI integration enhances operator trust by providing visual justification for each segregation decision, which is particularly critical for hazardous waste identification.

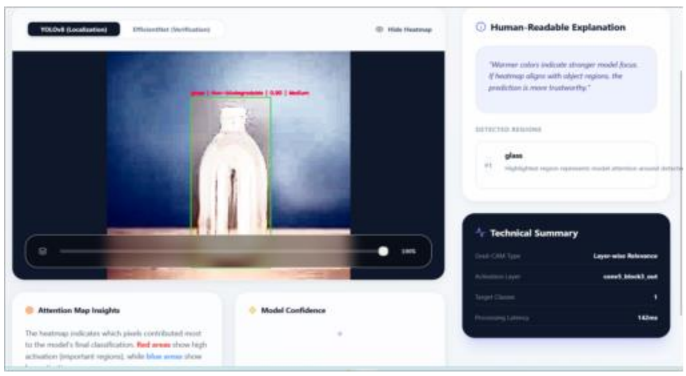

Fig 5. Grad-CAM Heatmap

F. Real-Time Performance

The system was evaluated for real-time operational viability by measuring end-to-end inference speed and latency across [number] test frames. Table VI summarizes the real-time performance metrics. The pre-filtering stages— Empty Frame Screener and MLP Waste Filter— reduced the number of frames processed by the core AI pipeline by approximately [X]%, directly contributing to lower average latency and higher sustainable FPS.

TABLE VI. Real-Time Performance Metrics

Metric	Value
nference Accuracy (IA)	90.82%
Classification Accuracy (CA)	94%
Frames Per Second (FPS)	0.24
Average Latency (ms)	4241

G. Comparison with Existing Methods

Table VII compares the proposed system against relevant prior works across key performance dimensions. The proposed hybrid pipeline achieves competitive classification accuracy while additionally providing real-time capability, multi-label output, and integrated XAI— features absent in most existing approaches.

IV. CONCLUSION

In this work, a multi-stage hybrid AI pipeline for real-time waste segregation was proposed, combining YOLOv8m-based object detection and EfficientNet-based fine-grained classification within a unified, interpretable framework. The system addresses the core limitations of existing single-model approaches— namely, inability to handle real-world multi-object scenes, absence of multi-dimensional waste categorization, and lack of decision transparency— through a series of novel architectural contributions.

The proposed two-stage pre-filtering mechanism, comprising an Empty Frame Screener and a lightweight MLP Waste vs. No-Waste Filter, significantly reduces unnecessary computational load before core inference. The YOLOv8m detection model achieved an accuracy of **64% across five waste classes** following dataset refinement and augmentation, while the EfficientNet classifier achieved **94% accuracy across eight waste categories**, demonstrating strong fine-grained discrimination capability. The introduction of an **Unknown class** mechanism further improved pipeline reliability by routing low-confidence predictions for XAI analysis rather than forcing incorrect assignments.

The **Verification Layer** and **Feature Intelligence Layer** together enhance prediction consistency by cross-validating outputs from both models and resolving inter-model disagreements before final categorization. The **Secondary Decision Engine** applies simultaneous multi-label classification across Biodegradable, Hazardous, Recyclable, and Contamination dimensions, governed by a priority-based segregation protocol that ensures safety-critical waste is handled first. The integrated **real-time Grad-CAM XAI layer** provides transparent, operator-interpretable visual justification for each segregation decision, which is particularly significant for hazardous and ambiguous waste categories where misclassification carries direct safety consequences.

Collectively, these contributions produce a waste segregation system that is accurate, interpretable, hierarchically aware, and suitable for real-world deployment in smart conveyor-based sorting environments. The system's modular architecture further ensures adaptability across diverse waste management scenarios without requiring full retraining.

Future work will focus on three primary directions. First, integration with **IoT-enabled smart bin infrastructure** to enable remote monitoring, automated bin-level alerts, and centralized waste trend analytics across urban zones. Second, coupling the system output with **robotic arm actuation** for fully autonomous physical sorting, eliminating the remaining dependency on human intervention at the segregation stage. Third, **model compression and edge optimization** through techniques such as quantization and pruning to enable deployment on resource-constrained embedded hardware such as Raspberry Pi and NVIDIA Jetson platforms, extending the system's applicability to low-infrastructure environments

REFERENCES

[1] . R. D. L. Reyes, C. M. A. Almirañez, J. M. S. Dela Cruz, and R. G. Maramba, "Campus-Scale Real-Time Waste Classification on Raspberry Pi 5 Using You Only Look Once Version 8," in *Proceedings of the 7th Eurasia Conference on IoT, Communication and Engineering 2025 (ECICE 2025)*, Yunlin, Taiwan, 2025.

[2] J. L. P. Rodriguez, M. A. C. Santos, and T. B. Ramos, "Design and Implementation of an AI-Based Automatic Plastic Waste Sorting System," in *Proceedings of the 2025 2nd International Conference on Advances in Electronics, Control and Communication Systems (ICAECCS)*, Bangalore, India, 2025, pp. 1-6.

[3] S. K. Gupta, R. Verma, and A. Singh, "Optimizing Waste Classification in Recycling Systems through Transfer Learning and Layer Freezing," in *Proceedings of the 2025 International Conference on Electronics and Renewable Systems (ICEARS)*, Tuticorin, India, 2025, pp. 452-458.

[4] Gulzar Ahmad, Fizza Muhammad Aleem, Tahir Alyas, Qaiser Abbas, Waqas Nawaz, Taher M. Ghazal, Abdul Aziz, Saira Aleem, Nadia Tabassum, Aidarus Mohamed Ibrahim, "Intelligent waste sorting for urban sustainability using deep learning," Scientific Reports, Springer Nature, vol. 15, Art. no. 27028, 2025.

[5] . H. Dipo et al., "Real-Time Waste Detection and Classification Using YOLOv12-Based Deep Learning Model," Digital, vol. 5, no. 1, Art. no. 19, 2025

[6] Songyun Deng, Ruiyuan Pei, Li Zhou, Hai Qin, Wei Sun, Qiaokang Liang, Senior Member, IEEE,"An Efficient Generative Intelligent Multiobjective Grasping Model for Kitchen Waste Sorting," IEEE Transactions on Instrumentation and Measurement, vol. 74, Art. no. 2522810, 2025.

[7] Md. Nahiduzzaman , Md. Faysal Ahamed , Mansura Naznine , Md. Jawadul Karim ,Hafsa Binte Kibria , Mohamed Arselene Ayari , Amith Khandakar , Azad Ashraf ,Mominul Ahsan, Julfikar Haiderf, "An automated waste classification system using deep learning techniques: Toward efficient waste recycling and environmental sustainability", Elsevier, Knowledge-Based Systems, vol.310 , Art. no. 113028, 2025.

[8] Zhenqi Xiao, Guangxiang Yang, Xu Wang, Guangling Yu, Xiaoheng Wang, Baofeng Yang, Delin Zhang,"DBS-YOLO: A High-Precision Object Detection Algorithm for Hazardous Waste Images,"IEEE Transactions on Instrumentation and Measurement, vol. 73, Art. no. 5036315, 2024.

[9] Jiajia Li, Jie Chen, Bin Sheng, Member, IEEE, PingLi, Member, IEEE, Po Yang, Senior Member, IEEE, David Dagan Feng, Life Fellow, IEEE, and Jun Qi, Member, IEEE,"**Automatic Detection and Classification System of Domestic Waste via Multimodel Cascaded Convolutional Neural Network**", *IEEE* Transactions on Industrial Informatics, vol.18, Art. no.1, 2022.

[10] Ka Yin Chau, Massoud Moslehpour, Shin-Hung Pan, Abed Akhundzada,"**Advanced predictive modeling of municipal solid waste management using robust machine learning**,"*Scientific Reports*, vol. 15, Art. no. 43267, 2025.

[11] Muhammad Hafizuddin Moktar, Hassan Mohamed, Sami Salama Hussen Hajjaj, Mohd Zafri Baharuddin,“**Medical waste sorting machine development with IoT and YOLO model utilization**,”,Springer nature, Journal of Engineering and Applied Science, vol. 72, Art. no. 86, 2025.

[12] Thompson Stephan, S. M. Hari Krishna, Chia-Chen Lin, Ujjwal Sumesh, Saurabh Agarwal, Hyunsung Kim,“**ProWaste for proactive urban waste management using IoT and machine learning**”, Scientific Reports, vol. 15, Art. no. 27790, 2025.