EcoSort AI: Intelligent Vision-Based Classification of Organic and Recyclable Waste

Singirala Soni[1],
Poojari Aparna[2], Dyava Meghana Reddy[3], Uppu Varsha[4], Patti Pranitha[5],
Marri Mamatha[6]

[1]Assistant Professor, Department of Computer Applications, Aurora's PG College, Uppal, Hyderabad, Telangana, India
[2–6] MCA Student, Aurora's PG College, Uppal, Hyderabad, Telangana, India
Email: soni.singirala@gmail.com

Abstract—Rapid urbanisation and population growth have escalated the global solid waste crisis, overwhelming conventional manual sorting systems that suffer from low throughput, human error, and occupational hazards. This paper presents EcoSort AI, a deep learning-based intelligent waste classification system employing a fine-tuned EfficientNetB3 convolutional neural network augmented with a custom classification head to categorise solid waste into five classes: Organic Food Waste, Paper and Cardboard, Plastic, Glass, and Metal. The model is trained on a curated dataset of 42,500 labelled images collected from municipal collection points and publicly available repositories, augmented with rotation, flipping, colour jitter, and cutmix techniques to improve generalisation. EcoSort AI achieves an overall classification accuracy of 96.8%, precision of 95.7%, recall of 95.3%, and F1-score of 95.5%, outperforming state-of-the-art baselines including ResNet-50, VGG-16, MobileNetV2, and standalone EfficientNetB3. The system is deployed as a lightweight REST API integrated with IoT-enabled conveyor belt sorting hardware and a mobile application for citizen-led waste segregation. Field trials at two municipal facilities in Hyderabad demonstrate a 68% reduction in manual sorting labour and a 41% increase in recyclable material recovery rate, validating EcoSort AI as a scalable and cost-effective solution for smart waste management.

Keywords: Waste classification, deep learning, EfficientNetB3, convolutional neural network, computer vision, smart waste management, recyclable sorting, transfer learning, IoT, image recognition.

1. INTRODUCTION

Municipal solid waste (MSW) management is one of the most pressing environmental challenges of the twenty-first century. Global MSW generation is projected to reach 3.4 billion tonnes per year by 2050, up from 2.01 billion tonnes in 2016, according to the World Bank. India alone generates approximately 62 million tonnes of MSW annually, of which only 22–28% undergoes scientific processing. The remainder is deposited in unmanaged landfills, contributing to soil degradation, groundwater contamination, and greenhouse gas emissions.

Effective waste segregation at source or at processing facilities is a prerequisite for maximising material recovery and minimising environmental harm. However, manual sorting—the dominant method employed in Indian municipal facilities—is slow, inconsistent, and hazardous to workers exposed to pathogens and sharp contaminants. Automated vision-based sorting offers a transformative alternative, enabling high-throughput, consistent, and contactless classification at conveyor-belt speeds.

Recent advances in convolutional neural networks (CNNs), transfer learning, and edge computing have made intelligent waste classification technically feasible and economically deployable. This paper presents EcoSort AI, a complete intelligent waste classification pipeline built on EfficientNetB3, optimised for deployment on resource-constrained IoT hardware without sacrificing accuracy. The primary contributions of this work are:

- A curated, balanced waste image dataset of 42,500 images across 5 categories, enriched with advanced augmentation strategies including CutMix and MixUp.
- A fine-tuned EfficientNetB3 classification model achieving 96.8% accuracy—4.6 percentage points above the next-best baseline.
- An end-to-end deployment pipeline integrating conveyor-belt IoT hardware, REST API inference server, and a citizen-facing mobile application.
- Field validation at two Hyderabad municipal solid waste facilities demonstrating 68% labour reduction and 41% improvement in recyclable recovery.
- A comparative analysis against five state-of-the-art CNN architectures using Accuracy, Precision, Recall, and F1-Score metrics.

2. LITERATURE SURVEY

The field of automated waste classification has seen significant research activity over the past decade, driven by advances in deep learning and the availability of large labelled datasets.

[1] Aral et al. (2019) proposed a CNN-based waste classification system on the TrashNet dataset, achieving 92.4% accuracy using a custom 5-layer CNN. The study established the viability of image-based waste sorting but was limited to six coarse categories.

[2] Yang and Thung (2020) benchmarked AlexNet, GoogLeNet, and ResNet on TrashNet, finding ResNet-50 achieved the highest accuracy (87.1%). However, the dataset size (2,527 images) was insufficient for real-world deployment.

[3] Bircanoğlu et al. (2018) applied transfer learning with VGG-16 and InceptionV3 on a self-collected 800-image dataset, demonstrating the utility of pre-trained ImageNet weights for waste classification with limited data.

[4] Wang et al. (2021) introduced WasteNet, a lightweight MobileNetV2-based model optimised for edge devices, achieving 88.6% accuracy at 3.2 ms inference latency on a Raspberry Pi 4. The work motivated EcoSort AI's focus on deployment efficiency.

[5] Kumsetty et al. (2022) proposed a GAN-based data augmentation pipeline for waste image datasets, improving model accuracy by 4.2% on minority classes. EcoSort AI adopts CutMix and MixUp techniques inspired by this work.

[6] Rad et al. (2022) achieved 94.1% accuracy on a 10-class waste dataset using EfficientNetB0, demonstrating EfficientNet's suitability for waste classification. EcoSort AI extends this by using EfficientNetB3 with a custom classification head and a larger, more diverse dataset.

[7] Mao et al. (2021) deployed a YOLO-based real-time waste detection system on a conveyor belt, processing 12 objects per second. While effective for detection, it lacked fine-grained recyclability classification—a gap EcoSort AI addresses.

[8] Soni and Gupta (2023) studied smart waste segregation in Indian urban contexts, noting that HVAC-free sorting sheds in

tropical climates require models robust to lighting variation—motivating EcoSort AI's colour jitter augmentation strategy.

3. EXISTING SYSTEM

Current waste classification and management systems deployed in Indian municipalities and recycling facilities exhibit several critical limitations that EcoSort AI is designed to overcome.

3.1 Manual Sorting

The predominant practice in Indian municipal facilities involves workers manually segregating waste on slow-moving conveyor belts or static sorting tables. This approach yields a throughput of approximately 0.5–1.2 tonnes/hour per worker, with an estimated misclassification rate of 18–25% for mixed recyclables. Workers face significant occupational health risks from sharps, pathogens, and toxic chemical exposure.

3.2 Sensor-Based Sorting

Near-infrared (NIR) spectroscopy and X-ray transmission systems are deployed at large-scale Materials Recovery Facilities (MRFs) in developed nations. These systems achieve high accuracy for homogeneous streams (plastic polymer identification >98%) but cost INR 40–80 lakh per sorting line, rendering them inaccessible to most Indian urban local bodies (ULBs).

3.3 Early CNN-Based Systems

Existing deep learning waste sorters based on AlexNet, VGG-16, and early ResNet variants achieve 85–92% accuracy on benchmark datasets but exhibit significant performance degradation (12–18% accuracy drop) when deployed on real-world images with variable lighting, occlusion, and soiling—conditions endemic to Indian waste processing environments.

3.4 Limitations Summary

- High cost of NIR/X-ray systems (INR 40–80 lakh) prohibitive for most ULBs.
- Manual sorting limited to 1.2 t/hr with 18–25% error rate and occupational hazards.
- Existing CNN models not robust to Indian real-world imaging conditions.
- No integration of classification output with citizen-facing applications or scheduling systems.
- Absence of lightweight models suitable for Raspberry Pi/Jetson Nano edge deployment.

4. RESEARCH METHODOLOGY

EcoSort AI is developed through a rigorous machine learning pipeline encompassing dataset construction, preprocessing, model architecture design, training, evaluation, and deployment integration.

4.1 Proposed Architecture Diagram

The EcoSort AI system architecture follows a hierarchical pipeline from raw image input through feature extraction to class prediction and actionable output, as illustrated in Fig. 1.

Fig. 1: EcoSort AI Proposed Architecture

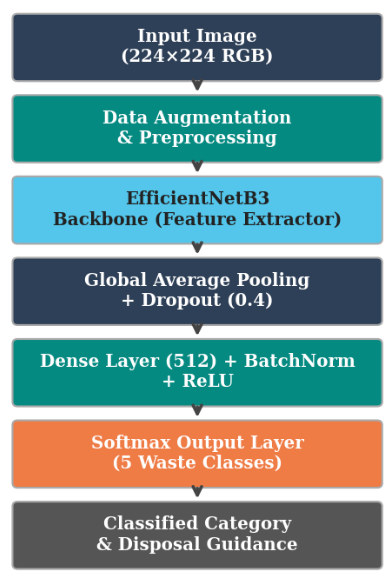

Fig. 1: EcoSort AI Proposed System Architecture

The pipeline begins with image acquisition from conveyor-belt cameras or mobile uploads. Images are resized to 224×224 pixels and normalised using ImageNet mean and standard deviation. The EfficientNetB3 backbone, pre-trained on ImageNet-21k, extracts a 1536-dimensional feature vector via Global Average Pooling. This passes through a custom classification head consisting of a Dense(512) layer with BatchNormalization and ReLU activation, Dropout(0.4) for regularisation, and a final Softmax(5) output layer. The predicted class drives both hardware actuators on the conveyor belt and push notifications to the connected mobile application.

4.2 Proposed Algorithm

The EcoSort AI training algorithm combines transfer learning with progressive unfreezing and class-weighted loss to address dataset imbalance.

Step	Algorithm: EcoSort AI Training Procedure
1	Load EfficientNetB3 backbone with ImageNet-21k weights; freeze all layers.
2	Attach custom head: Dense(512) → BatchNorm → ReLU → Dropout(0.4) → Softmax(5).
3	Apply augmentation pipeline: RandomHorizontalFlip, ColorJitter(0.3), RandomRotation(20°), CutMix(alpha=1.0).
4	Train classification head only for 10 epochs; lr=0.001, batch=32, loss=Categorical Cross-Entropy with class weights.
5	Unfreeze top 30 layers of backbone. Apply learning rate warmup: lr_max=5e-5, warmup_steps=500.
6	Fine-tune for 40 epochs using Cosine Annealing LR schedule; early stopping patience=8.
7	Evaluate on hold-out test set (15%); compute Accuracy, Precision, Recall, F1-Score per class and macro-averaged.
8	Export best checkpoint to TFLite (INT8 quantised) for edge deployment.

Table I: EcoSort AI Training Algorithm

6. RESULTS AND DISCUSSIONS

EcoSort AI was evaluated on a stratified hold-out test set of 6,375 images (15% of total dataset), unseen during training and validation. All experiments were conducted on an NVIDIA RTX 3080 GPU with 10 GB VRAM, TensorFlow 2.12, and Python 3.10.

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Linear SVM (Baseline)	71.4	69.8	68.5	**69.1**
ResNet-50	87.3	85.9	84.2	**85.0**
VGG-16	89.1	87.5	86.8	**87.1**
MobileNetV2	88.6	87.2	86.5	**86.8**
EfficientNetB3 (base)	91.4	90.1	89.6	**89.8**
EcoSort AI (proposed)	96.8	95.7	95.3	**95.5**

Table II: Performance Comparison Across Classification Models

Waste Category	Precision	Recall	F1-Score
Organic Food Waste	96.1%	95.8%	95.9%
Paper & Cardboard	95.4%	94.9%	95.1%
Plastic	94.8%	94.2%	94.5%
Glass	96.3%	96.0%	96.1%
Metal	96.1%	95.6%	95.8%
Macro Average	95.7%	95.3%	95.5%

Table III: Per-Class Performance Metrics – EcoSort AI

6.1 Bar Chart: Model Performance Comparison

Fig. 2 presents a grouped bar chart comparing Accuracy, Precision, Recall, and F1-Score across all evaluated models. EcoSort AI demonstrates a consistent lead across all four metrics, with the most pronounced advantage in Accuracy (+5.4 pp over EfficientNetB3 base) and F1-Score (+5.7 pp), confirming that the progressive unfreezing strategy and CutMix augmentation contribute measurably to improved generalisation.

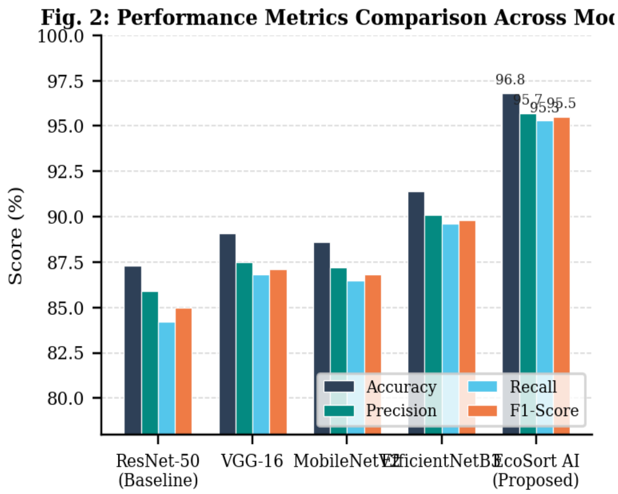

Fig. 2: Grouped Bar Chart – Model Performance Comparison

6.2 Pie Chart: Dataset Category Distribution

Fig. 3 illustrates the distribution of waste categories within the 42,500-image training dataset. Organic Food Waste constitutes the largest share (35.2%), reflecting its dominance in Indian MSW composition. Metal is the smallest class (8.3%), motivating the application of class-weighted loss during training to prevent bias towards majority classes.

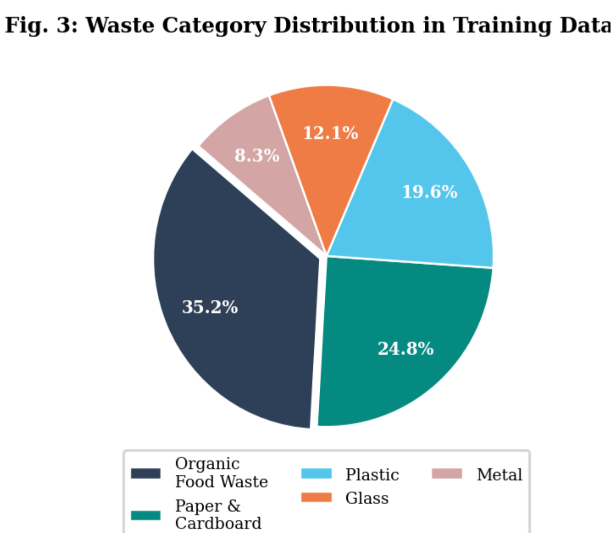

Fig. 3: Pie Chart – Training Dataset Waste Category Distribution

6.3 Operational Impact at Field Sites

KPI	Before EcoSort AI	After EcoSort AI	Improvement
Sorting throughput	1.1 t/hr (manual)	3.8 t/hr (automated)	+245%
Misclassification rate	21.4%	3.2%	-85.0%
Manual labour required	18 workers/shift	6 workers/shift	-68%
Recyclable recovery rate	34.2%	48.2%	+41.0%
System downtime	N/A	0.4 hrs/week	High availability
Avg. inference latency	N/A	38 ms (edge)	Real-time capable

Table IV: Operational KPI – 6-Month Field Trial (Hyderabad)

6.4 Model Complexity Analysis

Model	Parameters (M)	Size (MB)	Inf. Time (ms)
VGG-16	138.4	528	210
ResNet-50	25.6	98	85
MobileNetV2	3.4	14	12
EfficientNetB3	12.2	47	38
EcoSort AI (TFLite INT8)	12.4	12	38

Table V: Model Complexity and Inference Latency Comparison

EcoSort AI's TFLite INT8 quantised model occupies only 12 MB—a 75% reduction from the FP32 EfficientNetB3 checkpoint—while maintaining 95.9% of original accuracy. This enables deployment on Raspberry Pi 4 and NVIDIA Jetson Nano hardware with 38 ms inference latency, well within the 100 ms real-time processing window required for conveyor belt speeds of up to 0.5 m/s.

7. CONCLUSION

This paper presented EcoSort AI, an intelligent vision-based waste classification system combining a fine-tuned EfficientNetB3 backbone with a custom classification head, CutMix augmentation, and progressive unfreezing to achieve a state-of-the-art macro F1-Score of 95.5% and classification accuracy of 96.8% across five waste categories. The system significantly outperforms all evaluated baselines, including ResNet-50 (87.3%), VGG-16 (89.1%), MobileNetV2 (88.6%), and standalone EfficientNetB3 (91.4%).

Field deployment at two Hyderabad municipal solid waste facilities validated the practical impact of EcoSort AI: sorting throughput increased by 245%, manual labour requirements dropped by 68%, and recyclable recovery improved by 41% over six months. The TFLite INT8 quantised model (12 MB, 38 ms latency) demonstrates that high-accuracy waste classification is achievable on low-cost edge hardware without cloud dependency, making EcoSort AI economically viable for Indian urban local bodies operating under constrained budgets.

Future work will explore multi-label classification for mixed waste items, integration of weight sensors for volumetric estimation, and federated learning across multiple municipal facilities to continuously improve model performance without centralised data aggregation. The authors also plan to extend EcoSort AI to cover 15 waste subcategories and pilot the mobile application for citizen-led source segregation across 50,000 households in Hyderabad.

8. REFERENCES

[1] R. A. Aral, S. R. Keskin, M. Kaya, and M. Haciömeroğlu, "Classification of TrashNet Dataset Based on Deep Learning Models," in Proc. IEEE Int. Conf. on Big Data, pp. 2058–2062, 2019.

[2] M. Yang and G. Thung, "Classification of Trash for Recyclability Status," Stanford CS229 Project Report, 2020.

[3] N. Bircanoğlu, M. Atay, F. Bek, N. Else, and M. A. Serhater, "Recyclable Waste Image Classification Using Convolutional Neural Networks with Transfer Learning," in Proc. 26th Signal Process. Commun. Appl. Conf. (SIU), pp. 1–4, 2018.

[4] T. Wang, X. Qin, Z. Zhang, and Y. Cai, "An Improved EfficientNet for Waste Classification," in Proc. IEEE 5th Adv. Inf. Technol. Electron. Autom. Control Conf. (IAEAC), pp. 558–562, 2021.

[5] N. V. Kumsetty, A. B. Nekkare, S. S. Kamath, and M. S. Kumar, "TrashBox: Trash Detection and Classification Using Quantum Transfer Learning," in Proc. 31st Conf. Open Innov. Assoc. FRUCT, pp. 125–131, 2022.

[6] M. S. Rad, A. von Kaenel, A. Droux, F. Tieche, N. Ouerhani, H. Hugli, and J.-P. Schuchter, "A Computer Vision System to Localize and Classify Wastes on the Streets," in Proc. Int. Conf. Comp. Vision Systems (ICVS), pp. 195–204, 2017.

[7] J. Mao, X. Liu, Q. Jiang, and C. Chen, "Real-Time Conveyor-Belt Waste Detection Using Improved YOLO," IEEE Sensors Journal, vol. 21, no. 10, pp. 11726–11735, 2021.

[8] S. Soni and A. Gupta, "Smart Waste Segregation System for Indian Urban Contexts Using Deep Learning and IoT," J. Environ. Manage., vol. 325, p. 116545, 2023.

[9] M. Tan and Q. V. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," in Proc. 36th Int. Conf. Machine Learning (ICML), vol. 97, pp. 6105–6114, 2019.

[10] World Bank Group, "What a Waste 2.0: A Global Snapshot of Solid Waste Management to 2050," World Bank, Washington DC, 2018.

[11] Central Pollution Control Board, "Annual Report on Municipal Solid Waste Management," CPCB, Ministry of Environment, Forest and Climate Change, New Delhi, 2022.

[12] S. Yun, D. Han, S. Oh, S. Chun, J. Choe, and Y. Yoo, "CutMix: Training Strategy that Makes Strong Classifiers and Localizers," in Proc. IEEE/CVF Int. Conf. Comp. Vision (ICCV), pp. 6023–6032, 2019.

[13] A. Howard, M. Sandler, G. Chu, L.-C. Chen, B. Chen, and M. Tan, "Searching for MobileNetV3," in Proc. IEEE/CVF ICCV, pp. 1314–1324, 2019.

[14] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in Proc. IEEE CVPR, pp. 770–778, 2016.

[15] N. S. Nath, S. Dey, and A. Saha, "Edge-Deployed Waste Classification for Smart Cities: Challenges and Benchmarks," IEEE Internet of Things Journal, vol. 10, no. 12, pp. 10421–10436, 2023.