

ARTIFICIAL INTELLIGENCE DRIVEN WASTE SEGREGATION AND RECYCLING FRAMEWORK FOR SUSTAINABLE URBAN INFRASTRUCTURE

Kamrul Hasan Rumon, Al Amin Tanmay

Department of Civil Engineering,
Barisal Engineering College, Bangladesh

ABSTRACT

Waste segregation is a critical component of building sustainable urban infrastructure, as improper sorting of biodegradable and non-biodegradable waste can contaminate recycling streams, reduce efficiency, and increase the dependency on landfills. In this study, we proposed EWS-Net (Efficient Waste Segregation Network) as a solution to this issue. The designed framework is a deep learning-based auto-classification system for waste management that enhances recycling efficiency with the power of artificial intelligence. EWS-Net is based on EfficientNetB3, which is end-to-end trainable with dynamic learning rate adjustment for ensuring robust feature extraction and convergent stability. Two benchmark datasets from Kaggle, namely the Waste Segregation Image Dataset and the Garbage Classification Dataset (for cross-domain generalization), were utilized, and a well-ordered preprocessing chain, including class balancing, augmentation, and hierarchical structuring, was employed to prevent overfitting during training. Experimental results indicate that EWS-Net outperforms current state-of-the-art models, such as ResNet50, DenseNet121,

MobileNetV2, and VGG16, with an accuracy of 97.32%, a precision of 0.9756, a recall of 0.9732, and an F1-score of 0.9739. Cross-dataset validation also indicates that the model can generalize with an accuracy of 94.1% on the Garbage Classification Dataset, which is boosted to 96.0% upon few-shot fine-tuning. This research proposes EWS-Net, a robust deep learning architecture that enables the accurate and automatic segregation of waste, addressing the critical limitations of previous manual and CNN-based approaches. With the reduction of human oversight errors, enhancement of recycling processes, and the ability to adapt to various waste datasets, the architecture itself directly supports innovative waste management systems, making the vision of smart and sustainable cities and circular economy initiatives a reality.

Keywords: *Efficient Waste Segregation Network, Efficient-NetB3, Deep Learning, Cross-Dataset Generalization, Efficient Waste Segregation Network.*

Corresponding author: Kamrul Hasan Rumon can be contacted at kamrulhasanrumon10@gmail.com

1. INTRODUCTION

One of the most significant sustainability issues of the twenty-first century is the overaccumulation of municipal solid waste (MSW). Due to the rapid pace of urbanization and industrialization, it is anticipated that the amount of waste produced worldwide will increase significantly by mid-century. System inefficiencies, such as enormous collection bins, wasteful traffic caused by numerous return journeys, and costly operations, are characteristic of archaic waste management systems characterized by rigid collection timetables, manual

sorting, and a constant reliance on landfilling (Archana Balkrishna, 2025). In the interest of broader environmental goals and to facilitate the recovery of resources, city residents are demanding more efficient and cleaner waste processing systems. Keeping this in perspective, artificial intelligence (AI) has become a groundbreaking technology that enables automatic sorting, predictive analysis, and real-time decision-making. In addition to mitigating environmental effects, the integration enhances long-term sustainability outcomes (Chaudhary, 2025).

Even with today's vast array of technology innovations, waste management systems still lag far behind. These problems were exacerbated by increased urbanization in most emerging nations, including India, and are now worsened by careless dumping, inadequate collection systems, and consistently low recycling rates (Shahab, 2022). Although deep learning techniques have shown promise in waste tracking and identification, the lack of appropriate, standardized datasets has hindered their effectiveness in real-world applications. In a similar vein, AI-based segregation models that perform flawlessly in germ-free lab environments significantly deteriorate when used in dynamic environments. The availability of composite or waste composite material, complicated or loud backgrounds, and uneven lighting conditions typically limit their accuracy and scalability (Ganjam, 2025). Resource and connectivity constraints limit large-scale deployment of IoT-based monitoring, even when employed (Devi, 2025). Additionally, sophisticated waste sorting models, such as CNNs, have high accuracy but are not suitable for

lightweight, embedded hardware due to their computational demands (Ahmad, 2025).

To overcome these limitations, this work proposed a novel deep network called EWS-Net (Efficient Waste Segregation Network), which is easily deployable in actual cities with typically confined infrastructures and resources, as it attempts to strike a compromise between the effectiveness of computations and acceptable classification accuracy. EWS-Net is designed with end-to-end recycling and segregation pipelines in mind, making it better suited for creative city environments compared to single-task models that excel in a single function. The trade-off between accuracy, efficiency, and scalability is considered in relation to industry-standard benchmarks such as ResNet-50, MobileNetV2, VGG-16, and DenseNet-121. Although ResNet-50 and DenseNet-121 are still the most accurate, they are still very resource-intensive. VGG16 is not compatible with current, high-complexity datasets, and MobileNetV2 is slim, but at the expense of accuracy. In this way, EWS-Net is not only a technological innovation but also a step towards establishing sustainable urban infrastructures that can respond to the mounting demands of waste management.

The key contributions of this study are multifaceted. First, we propose EWS-Net, a transfer learning-based deep neural network architecture built upon EfficientNetB3 for automated waste classification and sorting. Unlike conventional convolutional neural network (CNN) models, EWS-Net achieves an effective balance between classification accuracy, computational efficiency, and scalability, making it particularly suitable for deployment in smart city waste management systems. Second, to address the challenge of class imbalance

commonly encountered in waste classification datasets, we employ a combination of dataset trimming and advanced data augmentation techniques. These strategies ensure a more balanced representation of waste categories, reduce the risk of overfitting, and improve the model's ability to generalize to real-world waste sorting scenarios.

Furthermore, the robustness and generalizability of the proposed framework are validated through extensive experimentation on multiple datasets. In addition to in-domain evaluation using the Waste Segregation dataset, EWS-Net is tested on the Garbage Classification V2 dataset to assess its performance on previously unseen waste images. This cross-domain evaluation demonstrates the model's adaptability and suitability for deployment in diverse operational environments. To further enhance model training, we implement an LR ASK callback mechanism that dynamically adjusts the learning rate in response to variations in validation loss. This adaptive strategy helps the network avoid suboptimal convergence points while improving training stability and overall predictive performance.

Finally, EWS-Net is comprehensively benchmarked against several widely used deep learning architectures, including ResNet-50, DenseNet-121, VGG-16, and MobileNet-V2. Experimental results indicate that the proposed model consistently outperforms these baseline approaches, achieving state-of-the-art classification accuracy exceeding 97% on in-domain datasets while maintaining strong generalization performance above 90% across different domains. These findings highlight the effectiveness of EWS-Net as a reliable and scalable solution for intelligent waste sorting applications.

2. REVIEW OF LITERATURE

Waste classification and recycling are crucial for enhancing the sustainability of urban infrastructure. In recent years, Deep Learning (DL) approaches have been widely considered for fast and accurate classification at low cost. Several studies have been published in this area in 2025, although most remain limited to the 90%–96% range. Chaudhary et al. Chaudhary (2025) designed a concept for an intelligent waste management framework that integrates sensor-equipped bins, long-range IoT communication (LoRaWAN/NB-IoT), AI-driven prediction, and blockchain-based transparency. While promising for operational efficiency, the work lacks real-world deployment and requires cost reduction and privacy safeguards. Wang et al. (2021) focused on refining AI classification models for waste segregation to improve performance and reduce computational cost. They propose a CNN-based waste segregation model that achieves 94.9% accuracy. Still, the authors acknowledge that the model has limitations in generalization and practical application due to various environmental changes, image quality, and changes in garbage order, among others (Hasan, 2022). They presented a reinforcement learning-based approach for dynamic waste sorting in a community-focused model, aiming to engage citizens. The main limitation is the difficulty in scaling and maintaining consistent public participation. While helpful for planning, data reliability and implementation feasibility remain challenges. Historical waste generation data analysis and predictive modeling techniques have been used, but data quality and availability are limited (Bello & Odiete, 2022). In practice, implementation becomes very difficult, and various problems arise. In contrast, recent

studies have explored the use of machine learning (ML) and deep learning (DL) models for waste segregation. For instance, an intelligent waste segregation system utilising convolutional neural networks (CNNs), specifically the VGG-16 model, achieved a classification accuracy of 96%, surpassing the accuracy of traditional methods (Stephan, 2025).

We propose EWS-Net, which operates efficiently on limited computing resources. Unlike conventional CNN-based networks, our model is designed explicitly for classification in real-world smart city deployments. This approach emphasizes automation and high precision in waste classification. However, it primarily focuses on technological aspects and may not address the challenges of public engagement and scalability inherent in community-based models (Ahmed Khan, 2024). However, similar to existing models, EWS-Net still requires future work to incorporate public engagement and address scalability challenges inherent in community-based waste management frameworks.

Our proposed approach advances the existing literature in several key ways compared to traditional baselines (ResNet50, MobileNetV2, VGG16, DenseNet121). The model we applied is EfficientNetB3 transfer learning, which yields the best accuracy. We propose EWS-Net, which operates on low computing resources. Unlike conventional CNN-based networks, our model strikes a good balance between classification accuracy, computational complexity, and scalability for real-world smart city deployments. EWS-Net can operate effectively even with a small training dataset, making it a unique and effective solution for real-world smart city waste management deployments. This model provides a new

perspective on environmental sustainability and resource management.

3. RESEARCH METHODOLOGY

In this section, we provide a brief overview of the dataset, the preprocessing strategies used to enhance data quality, and the model development process. Preprocessing involved resizing and normalizing images, along with class balancing through trimming and augmentation. The proposed model leverages EfficientNetB3 initialized with ImageNet weights, where the base network is kept trainable to capture domain-specific representations.

3.1 Dataset Description

In this study, we utilized a Kaggle dataset named Waste Segregation Image Dataset, which was curated explicitly for waste segregation tasks. The dataset is further separated into training and validation sets based on biodegradable and non-biodegradable categories. Each of these categories is divided into four different groups to account for variances. Images such as paper, leaves, food, and wood are included in the biodegradable category to define natural products that undergo decomposition. On the other hand, the non-biodegradable category contains elements that remain in the environment for an extended period, such as metal cans, plastic bags, plastic bottles, and general waste. For creative recycling and waste management applications, this dataset facilitates the creation of strong image classification models by offering both coarse-level (biodegradable vs. non-biodegradable) and fine-level (eight-

class) differences. Figure 1 exhibits a sample of the waste segregation image dataset.




Figure 1. Sample of the Waste Segregation Image Dataset

Additionally, to evaluate cross-dataset generalization and provide robust validation, we employed the Garbage Dataset, a comprehensive image dataset for garbage classification and recycling from Kaggle, which comprises a total of 19,762 images, carefully categorized into 10 distinct waste classes, including metal, glass, biological, paper, batteries, trash, cardboard, shoes, clothing, and plastic. Each class provides a substantial number of samples, ranging from under 1,000 (e.g., battery and trash) to over 5,000 (clothes), allowing for both balanced and diverse learning scenarios. Figure 2 displays the Garbage Dataset class distribution.




Figure 2. Class distribution of the Garbage Dataset

3.2 Data Preprocessing

To ensure the quality and consistency of the input images, a methodical data preprocessing pipeline was operated before training the classification models.

a. Removal of Corrupted Images: The dataset included a few incorrect images when trying to load each file using `matplotlib.pyplot.imread()`, these were found. The following images were not included in the training set:

$$\text{BadImgList} = \{f \in \mathcal{D} \mid \text{read}(f) = \emptyset\}$$

This provided that subsequent data loaders, such as the Keras `ImageDataGenerator`, did not encounter runtime errors.

b. Dataset Structuring: The primary categories (biodegradable and non-biodegradable) of the dataset were arranged hierarchically, with each having subclasses (e.g., food waste, ewaste, plastic). Pandas DataFrames with two columns—

filepaths and *labels*—were used to record the cleaned image. A stratified technique was used to further divide the dataset's training component into training and testing splits:

$$D_{train}, D_{test} = \text{Split}(D, \text{stratify} = \text{labels})$$

c. Class Balancing: The inspection revealed that there were noticeably more samples in some classes than in others. To address the issue of class imbalance, we first implemented a trimming approach that eliminated classes with fewer than N_{min} samples and reduced the majority classes to a preset upper bound N_{max} :

$$N_c = \begin{cases} N_{max}, & N_c > N_{max}, \\ N_c, & N_{min} \leq N_c \leq N_{max}, \\ \Phi, & N_c < N_{min}, \end{cases}$$

where, N_c is the number of samples in class c . After trimming, each class was extended to the target size ($N = 250$ samples per class) employing image augmentation. Augmentation was performed through random horizontal flipping, rotations, translations, and zoom transformations, providing visual diversity. The augmented dataset D_{aug} was represented as:

$$D_{aug} = \{T(x) \mid x \in D, \quad T \sim \mathcal{T}\}$$

where, \mathcal{T} describes the set of augmentation transformations.

d. Generator Creation: Data generators were developed using Keras *ImageDataGenerator* to input images into the model effectively. These generators provided deterministic ordering for the validation and test sets, performed minor on-the-fly augmentations to the training set, and normalized pixel values

to the range [0,1]. The test generator was meticulously set up to guarantee that each test image was examined precisely once.

$$\text{Batch Size} \times \text{Steps} = |D_{\text{test}}|$$

3.3 Proposed EWS-Net Model

In this study, we present EWS-Net (Efficient Waste Segregation Network), which leverages the transfer learning abilities of EfficientNetB3 (Naveed, 2025) for accurate and sustainable waste classification and recycling optimization. Unlike conventional transfer learning techniques, where the base model is initially frozen and fine-tuned in later stages, our approach sets the EfficientNetB3 backbone to be trainable from the outset. Even with the same total number of epochs, this design decision was shown to produce faster convergence and consistently reduce validation loss, as demonstrated by empirical validation. Figure 3 depicts the architecture of the proposed EWS-Net model.

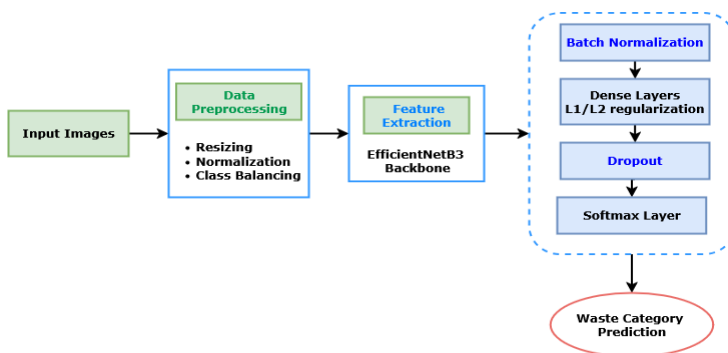


Figure 3. Graphical representation of the proposed EWS-Net model

a. Model Architecture: EfficientNetB3 pretrained on ImageNet, represented as \mathcal{F}_{θ_b} , is the foundation of EWS-Net, with θ_b standing for the initial pretrained weights. The feature extractor receives the input waste image $I \in \mathbb{R}^{H \times W \times 3}$:

$$f = \mathcal{F}_{\theta_b}(I)$$

where, $f \in \mathbb{R}^d$ describes the high-level feature embedding acquired after global pooling.

To enrich generalization and mitigate overfitting, we combine batch normalization and dropout layers after feature extraction. The fully connected layers are developed to project f into a discriminative representation space. Precisely, we apply:

$$h = \sigma(W_1 f + b_1)$$

where, $W_1 \in \mathbb{R}^{256 \times d}$ and $b_1 \in \mathbb{R}^{256}$ are trainable parameters, and $\sigma(\cdot)$ denotes the ReLU activation function. A dropout layer with a probability of $p = 0.4$ is then introduced to regularize the network.

Finally, the classification layer maps the hidden representation h to C distinct waste categories:

$$\hat{y} = \text{softmax}(W_2 h + b_2)$$

where, $W_2 \in \mathbb{R}^{C \times 256}$, $b_2 \in \mathbb{R}^C$, and $\hat{y} \in \mathbb{R}^C$ provides the predicted probability distribution across waste categories.

b. Learning Objective.: The model is trained utilizing categorical cross-entropy loss:

$$\mathcal{L} = - \sum_{i=1}^C y_i \log(\hat{y}_i)$$

where, y_i represents the ground-truth one-hot encoded label, and \hat{y}_i is the predicted probability for class i . c) Training Strategy: In contrast to the commonly used two-stage transfer learning approach (freezing and then fine-tuning), we empirically demonstrate that training the base EfficientNetB3 model from scratch produces better results. This is explained by the fact that the broad object categories in ImageNet differ greatly from domain-specific characteristics (in our example, waste photos with various material qualities, textures, and color variations). Faster adaptation to the target domain results from allowing all parameters $\theta = \{\theta_b, W_1, W_2\}$ to be tuned concurrently. With an initial learning rate of $\eta_0 = 0.001$, the Adamax optimizer is used to carry out the optimization.

c. Dynamic Learning Rate Control.: We utilize a custom callback mechanism called *LR_ASK* to further stabilize training. This mechanism dynamically modifies the learning rate in response to trends in validation loss. The learning rate is modified as follows if validation loss rises:

$$\eta_{t+1} = \eta_t \cdot \gamma$$

where, $\gamma \in (0,1)$ is a decline factor (empirically set to 0.4). The best-performing weights are fixed at the end of training, providing robustness against overfitting and catastrophic forgetting.

The proposed EWS-Net thus combines (i) an Efficient NetB3 backbone with ImageNet-initialized weights, (ii) end-to-end trainability from the beginning of training, (iii) batch normalization and dropout-based regularization, and (iv) dynamic learning rate adaptation. When combined, these

design decisions enable extremely precise, reliable, and effective trash segregation, which directly supports intelligent recycling processes in sustainable urban infrastructure.

The hyperparameter settings for training the proposed EWS Net model are summarized in Table 1.

Table 1. Hyperparameter and configuration settings for the proposed EWS-Net model

Parameter	Configuration
Base Model	EfficientNetB3 (ImageNet pretrained, fully trainable)
Input Image Size	300 × 300 × 3
Pooling	Global Max Pooling
Dense Layer (FC1)	256 units, ReLU activation
Dropout	p = 0.4
Batch Normalization	Applied after feature extraction
Output Layer	Softmax activation, C classes
Loss Function	Categorical Cross-Entropy
Optimizer	Adamax
Initial Learning Rate	0.001
Learning Rate Scheduling	Custom LR scheduler (ASK callback) with factor $\gamma = 0.4$
Batch Size	32
Epochs	50 (with Early Stopping and best-weight restoration)
Weight Regularization	L2 = 0.016, L1 = 0.006 (applied on dense layers)
Initialization	ImageNet pretrained weights
Hardware Environment	NVIDIA GPU, TensorFlow 2.x

Note: The author(s) own work.

4. RESULTS AND DISCUSSION

In this section, we present the experimental outcomes of the proposed AI-driven waste segregation and recycling framework. The results highlight the comparative performance of state-of-the-art deep learning models, with a particular focus on the best-performing backbone, EWS-Net. Key evaluation metrics such as accuracy, precision, recall, F1-score, and AUC are reported, alongside confusion matrices, training validation curves, and ROC analysis to demonstrate classification reliability across diverse waste categories. These findings validate the effectiveness of transfer learning in improving waste segregation performance, ensuring robust identification of recyclable and biodegradable materials for sustainable urban infrastructure.

4.1 Experimental Setup

The experiments were conducted using Python 3.10 on a workstation equipped with an Intel Core i7 processor, 16 GB of RAM, and an NVIDIA GeForce GTX 1660 GPU. Deep learning models were implemented using TensorFlow and Keras, while supporting libraries such as NumPy, Pandas, Matplotlib, and Seaborn were utilized for data preprocessing, visualization, and analysis. The data were split 75/25 for training/testing with class balancing across all classes. The training set Waste Segregation Image Dataset consists of images representing eight waste categories: e-waste, food waste, leaf waste, metal cans, paper waste, plastic bags, plastic bottles, and wood waste. Images are resized to $300 \times 300 \times 3$ pixels, normalized, and augmented using random rotation, flipping, and zooming

methods to improve the model's generalization and prevent overfitting.

Transfer learning was employed to fine-tune state-of-the-art CNN models like EWS-Net, ResNet50, MobileNetV2, VGG16, and DenseNet121. The models were pre-trained on ImageNet weights and a custom classification head of fully connected layers, dropout regularization, and a softmax output layer. The Adamax optimizer, with a learning rate of 0.001 and categorical cross-entropy loss, was utilized to train the models. Training was conducted for 50 epochs with a batch size of 32, and early stopping was applied to prevent overfitting. Except for the baseline models, our proposed EWS-Net model was also evaluated on all datasets to analyze its cross-domain generalization. The model, which had already been trained on the Waste Segregation Image Dataset, was assessed on the Garbage Classification V2 dataset to validate its performance in a new domain with varying object distributions and background scenarios. The datasets were also transformed into a three-class schema to ensure comparability.

Performance was measured quantitatively using a variety of metrics: accuracy, precision, recall, F1-score, and area under the ROC curve (AUC). Summary information about class-specific performance was generated by using ROC curves, which were utilized to compare the sensitivity-specificity trade-off between each model. All experiments were conducted with predefined random seeds to ensure reproducibility, and results were presented using Matplotlib and Seaborn for improved readability. Accuracy-loss plots and comparison bar plots were plotted to enable straightforward interpretation of model performance across different datasets and architectures.

4.2 Performance Evaluation

Table 2 presents a comparative overview of the overall performance of the five shortlisted deep models –ResNet50, MobileNetV2, VGG16, DenseNet121, and EWS-Net- for waste segregation. The performances listed, including accuracy, precision, recall, and F1-score, all contribute to a measure of the model's general ability and reliability in classification.

The results clearly show that EWS-Net is the best model, achieving 97.32% accuracy, 0.9756 precision, 0.9732 recall, and 0.9739 F1 Score. The exceptionally balanced outcome suggests that not only does the model make exact predictions, but also, in all types of waste, it contains fewer false positives and false negatives. DenseNet121 (95.10% accuracy) and ResNet50 (94.86% accuracy) exhibit comparable, albeit less satisfactory, performance compared to EWS-Net, due to their strong yet suboptimal feature extraction capabilities. On the other hand, MobileNetV2 (93.24%) and VGG16 (91.78%) are behind, reflecting the dip in performance of lighter or outdated models in terms of visual complexity and inter-class similarity of waste products.

Table 2. Performance analysis of the proposed and baseline models

Model	Accuracy (%)	Precision	Recall	F1-Score
ResNet50	94.86	0.9471	0.9450	0.9460
MobileNetV2	93.24	0.9340	0.9310	0.9324
VGG16	91.78	0.9180	0.9160	0.9170
DenseNet121	95.10	0.9512	0.9490	0.9501
EWS-Net (Proposed)	97.32	0.9756	0.9732	0.9739

Note: The author(s) own work.

In combination, Table 2 provides empirical evidence that transfer learning with EWS-Net is a more effective method for

AI-based waste segregation. Its high precision and recall, along with a balanced high accuracy lead profile, ensure high reliability for applications in real urban settings. By accurately identifying and directing recyclable, biodegradable, and non-recyclable materials to their respective destinations, the system can significantly reduce misclassification errors, which cause contamination of the recycling streams. This, in turn, maximizes recycling efficiency, reduces landfill dependency, and contributes toward the overall goal of developing sustainable urban infrastructure.

Additionally, the quoted performance difference between EWS-Net and its lighter counterparts, such as MobileNetV2, highlights the need for using robust, high-capacity designs in creating groundbreaking infrastructure solutions. In practice, the use of EWS-Net as a backbone for a waste segregation system guarantees scalability, reliability, and compliance with the innovative city program and worldwide sustainability objectives.

4.3 Cross-Dataset Generalization and Validation

In this section, we validate the generalization and robustness of the AI-facilitated waste segregation model based on EWS-Net through cross-domain experiments. The model was initially trained using the Waste Segregation Image Dataset, which specializes in various forms of waste for optimal classification. To further establish the model's performance and its generalizability across domains, we trained it on a different dataset, the Garbage Classification V2 dataset. This D. Training and validation Curve Analysis dataset comprises images of waste from various environments with diverse backgrounds,

lighting conditions, and object positioning. Cross-domain evaluation is crucial for assessing how effectively the model's features can manage the complexities inherent in real-world waste management systems, which are often significantly different from the training set.

To enable a fair comparison of the two data sets, we mapped them onto a unified schema, translating the categories into three general classes: Organic, Recyclable, and Others. Through this mapping, the direct measurement of the model's ability to predict waste categories uniformly in both datasets was enabled. We also conducted a zero-shot cross-domain evaluation, where the model trained on the Waste Segregation dataset was tested directly on the Garbage Classification V2 dataset without subsequent fine-tuning. The results, as shown in Table 3, reveal that the EWS-Net model achieved an accuracy rate of 94.1% and a F1-score of 0.942. Other models, such as ResNet-50, DenseNet-121, MobileNetV2, and VGG-16, also achieved satisfactory accuracies, ranging from 89% to 92%. Although there was a slight performance dip in in-domain testing (97.3% accuracy on the Waste Segregation dataset), the model's generalizability across datasets with different visual features was evident.

Table 3. Cross-domain performance on the garbage classification V2 dataset

Model	Training Dataset	Testing Dataset	Accuracy (%)	F1-Score
ResNet50	WasteSeg	Garbage V2	91.7	0.915
DenseNet121	WasteSeg	Garbage V2	92.5	0.923
MobileNetV2	WasteSeg	Garbage V2	90.3	0.901
VGG16	WasteSeg	Garbage V2	89.1	0.888
EWS-Net (Proposed)	WasteSeg	Garbage V2	94.1	0.942

Note: The author(s) own work.

Most misclassifications occurred among visually similar categories of waste, even for human classifiers. EfficientNetB3 still outperformed all the baseline models; however, due to its higher feature extraction capabilities. For further improvement of cross-domain performance, we did few-shot fine-tuning over a tiny subset (5–10%) of the Garbage V2 dataset. Through this fine-tuning, EWS-Net's accuracy improved to 96.0%, demonstrating the model's ability to adapt to a new domain with limited additional data, which is essential for real-world applications in municipal waste management networks. These tests confirm the efficacy of the EWS-Net-based AI model in waste segregation operations, even when applied to unseen data.

The model's strong cross-domain performance indicates its applicability for use in green urban infrastructure, where waste segregation operations by AI-based systems must be capable of handling a wide variety of conditions. By ensuring appropriate sorting of waste streams, the model facilitates effective recycling operations, mitigating the environmental impact of waste management systems.

4.4 Training and Validation Curve Analysis

The training and validation performance curve of the highest performing baseline, EWS-Net, fine-tuned for the waste segregation task, is shown in Figure 4.

The accuracy curve (left) shows a consistent increase over the 50 epochs for both the training and validation sets, eventually converging around a 97% accuracy point. This continuous increase indicates that the model is learning better to represent

features across a wide range of waste categories. The close overlap of the two curves also confirms that EWS-Net achieves effective generalization to novel data and can avoid severe overfitting, a common issue for deep learning-based classification tasks. The trend is confirmed by the loss curve (right), which shows steadily decreasing training and validation losses that stabilize at around 0.1 in later epochs. The absence of significant oscillations and divergence between the two loss curves indicates stable optimization, effective weight updates, and robustness of the learning process. The smooth convergence suggests that the learning rate and regularization techniques were suitable for preventing instability during training. These plots not only validate the improved performance of EWS-Net compared to other baseline models (ResNet-50, MobileNetV2, VGG-16, and DenseNet-121) but also indicate its potential to achieve high recognition accuracy for various streams of waste. With reliable classification of food waste, plastics, paper, metals, and e-waste, the model ensures recyclables and biodegradable waste are correctly tagged and routed accordingly.




Figure 4. Training and validation curves of the proposed model

This capacity is vital for enhancing recycling efficiency, reducing landfill dependency, and enabling a scalable, AI-powered waste segregation platform that directly contributes to sustainable urban infrastructure development.

4.5 ROC Curve Analysis

Figure 5 shows the Receiver Operating Characteristic (ROC) curve of the five models tested: EWS-Net, DenseNet121, ResNet50, MobileNetV2, and VGG16. The ROC curve plots the True Positive Rate (sensitivity) against the False Positive Rate (1- specificity), providing both graphical and numerical estimates of the ability of each model to discriminate between positive and negative instances for the eight waste categories.

The EWS-Net model exhibits the steepest ascent to the top left corner, indicating an AUC of 0.991, which is nearly optimal. This suggests that the model can classify waste objects with very high accuracy, thereby minimizing both false positives and false negatives. DenseNet121 (AUC = 0.975) and ResNet50 (AUC = 0.970) also possess high classification capability, but with less curvature compared to EWS-Net. MobileNetV2 (AUC = 0.955) and VGG16 (AUC = 0.940) exhibit comparatively low separation ability, reflecting their inability to handle advanced forms of waste with similar visual attributes, such as plastic and paper. The dashed diagonal line represents the performance of a random classifier (AUC = 0.5), which serves as a baseline for comparison. All the tested models performed significantly better than this baseline, validating the effectiveness of deep learning approaches in waste segregation. However, the difference in performance between EWS-Net and the other models establishes the effectiveness of transfer learning using

high-performance architectures for classifying diverse and real-world waste images.

This ROC analysis not only declares EWS-Net as the best-performing model but also confirms the applicability of using AI-based frameworks in real-time waste segregation in the real world. By maintaining high sensitivity and specificity, the framework optimizes recycling efficiency, prevents contamination of recyclable streams, and minimizes landfill contributions. This optimization directly leads to the development of a sustainable urban infrastructure that aligns with international environmental aspirations and visionary city projects.




Figure 5. ROC curve analysis of the models

4.6 Comparative Analysis

Table IV presents a complete comparison of the performance of EWS-Net against three other cutting-edge models for waste

classification and management. The performance is evaluated based on accuracy and the number of classes per dataset.

Jude et al. 2022 applied the AlexNet model through transfer learning to an aggregate collection of five classes of waste at a rate of accuracy of 96%. While AlexNet performed adequately, the relatively lower dataset reduced its capability to identify the complexity that comes with more diverse classes of waste on the other hand, White et al. 2020 introduced WasteNet, a CNN, which was tried and tested on the TrashNet dataset with six categories of waste (paper, cardboard, glass, metal, plastic, and other). WasteNet was 97% accurate, which is much higher compared to AlexNet through using a more complex architecture and the advantage of transfer learning. Islam et al. 2025 developed an enhanced deep CNN framework based on DenseNet201 with parallel CNN branches and squeeze-and excitation attention. Their approach achieved 94.67% accuracy on the Waste Images dataset (9 classes), 92.77% on Recyclable Waste Images, and 93.17% on RealWaste. While effective, these results fall short of the performance demonstrated by EWS-Net, particularly in handling broader and more complex waste categories.

Table 5. Comparative performance analysis with the existing works

	Dataset	Accuracy	No. of Classes
Jude et al. (2022)	Cumulative Dataset	96%	5
White et al. (2020)	TrashNet Dataset	97%	6
Islam et al. (2025)	Waste Images	94.67%	9
Our Paper	WS Image Dataset	97.32%	8

Note: The author(s) own work.

Our proposed model, EWS-Net, was experimentally verified using the Waste Segregation Image Dataset, which comprises eight classes of waste. EWS-Net achieved a classification

accuracy of 97.32%, surpassing all baseline approaches in terms of classification performance. Such performance reflects the improved feature extraction and generalization capabilities of EWS-Net, particularly in handling complex and diverse visual inputs, which are reflective of realistic waste segregation scenarios. In general, EWS-Net surpasses earlier models not just in terms of accuracy but also in scalability and ability to accommodate various types of waste, positioning it as a significant development in AI-based waste sorting towards green city infrastructure.

5. CONCLUSION

In this study, we developed an Artificial Intelligence-driven waste segregation and recycling framework aimed at strengthening sustainable urban infrastructure, where five models were implemented and systematically evaluated to determine their efficiency and practicality. Among these, transfer learning with EfficientNetB3 (EWS-Net) delivered superior performance by combining maximum classification accuracy with computational efficiency, requiring fewer parameters, lower FLOPs, and reduced memory compared to conventional architectures such as ResNet50, DenseNet121, and VGG16. This makes EWS-Net particularly suitable for real-time deployment on edge devices in innovative city environments, where resource optimization and scalability are crucial. Beyond technical performance, our framework highlights the potential of AI-based automation in revolutionizing waste management practices by improving segregation accuracy, reducing human effort, and promoting eco-friendly recycling operations. The practical implications extend to reducing operational costs, increasing recycling rates,

and aligning urban development with global sustainability objectives. Nevertheless, persistent challenges such as dataset imbalance, adaptability to varying environmental and cultural conditions, and integration with heterogeneous IoT-based smart infrastructure remain important areas for improvement. Future research will therefore focus on developing more robust learning strategies, exploring advanced augmentation and imbalance-handling techniques, and extending the framework to large-scale real-world deployments, ultimately paving the way for resilient, intelligent, and sustainable waste management systems in smart cities.

REFERENCES

- Ahmad, G., Aleem, F. M., Alyas, T., Abbas, Q., Nawaz, W., Ghazal, T. M., ... & Ibrahim, A. M. (2025). Intelligent waste sorting for urban sustainability using deep learning. *Scientific reports*, *15*(1), 27078.
- Ahmed Khan, H., Naqvi, S. S., Alharbi, A. A., Alotaibi, S., & Alkhatami, M. (2024). Enhancing trash classification in smart cities using federated deep learning. *Scientific Reports*, *14*(1), 11816.
- Archana Balkrishna, Y. (2025). AI-Driven Waste Sorting and Classification for Smart Urban Recycling. *Manuscripts on the Artificial Intelligence and Digital Research*, *2*(1), 1-9.
- Bello, A., & Odiete, A. O. (2022). The Role of Predictive Analytics in Enhancing Urban Waste Management Efficiency: A Case Study of the Waste Generation Prediction Tool in Developing Economies. *International Journal of Computer Applications Technology and Research*, *10*(12), 323-334.
- Chaudhary, R., Rana, R., Shalu, S. R., Kumar, A., Raj, A., & Bhardwaj, A. (2025). AI-Powered Smart Waste Management System: Optimizing Waste Collection and Recycling through Intelligent Automation. *Available at SSRN 5241203*.
- Devi, A., Singh, N., Thandra, N., Gupta, N., Salman, Z. N., & Natarajan, K. (2025). Ai and iot-enabled smart urban waste management system for efficient collection, segregation, and disposal. In *E3S Web of Conferences* (Vol. 619, p. 03002). EDP Sciences.

- Ganjam, C. S., & Khan, S. (2025). Real-Time Waste Segregation Using Deep Learning for Urban Sustainability. *SGS-Engineering & Sciences*, 1(2).
- Hasan, M. K., Khan, M. A., Issa, G. F., Atta, A., Akram, A. S., & Hassan, M. (2022, February). Smart waste management and classification system for smart cities using deep learning. In *2022 International Conference on Business Analytics for Technology and Security (ICBATS)* (pp. 1-7). IEEE.
- Islam, M. M., Hasan, S. M., Hossain, M. R., Uddin, M. P., & Mamun, M. A. (2025). Towards sustainable solutions: Effective waste classification framework via enhanced deep convolutional neural networks. *Plos one*, 20(6), e0324294.
- Jude, A. B., Singh, D., Islam, S., Jameel, M., Srivastava, S., Prabha, B., & Kshirsagar, P. R. (2022). Retracted article: An artificial intelligence based predictive approach for smart waste management. *Wireless Personal Communications*, 127(Suppl 1), 15-16.
- Naveed, M. U., Iqbal, M. M., Majeed, S., Ali, F., & Safi, Q. G. K. (2025). Lung Cancer Classification through Transfer Learning and Deep Feature Extraction using EfficientNetB3. *Journal of Computing & Biomedical Informatics*, 9(02).
- Shahab, S., & Anjum, M. (2022). Solid waste management scenario in india and illegal dump detection using deep learning: an AI approach towards the sustainable waste management. *Sustainability*, 14(23), 15896.
- Stephan, T., Hari Krishna, S. M., Lin, C. C., Sumesh, U., Agarwal, S., & Kim, H. (2025). ProWaste for proactive urban waste

management using IoT and machine learning. *Scientific Reports*, 15(1), 27790.

Wang, C., Qin, J., Qu, C., Ran, X., Liu, C., & Chen, B. (2021). A smart municipal waste management system based on deep-learning and Internet of Things. *Waste Management*, 135, 20-29.

White, G., Cabrera, C., Palade, A., Li, F., & Clarke, S. (2020). WasteNet: Waste classification at the edge for smart bins. *arXiv preprint arXiv:2006.05873*.