IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

Comprehensive Analysis Of Automated Garbage Classification Using MobileNetV2 And Transfer Learning Strategies

1. Rushikesh Ramesh Kharage, 2.Tejas Pintu Kharat, 3.Shivam Suresh Kumbhar, 4. Manthan Prasanna Kore

Department of Computer Science and Engineering

DKTE Society's Textile And Engineering Institute, Ichalkaranji, Maharashtra

Abstract

The rapid rise in municipal solid waste is placing enormous pressure on urban infrastructure, recycling systems, and the natural environment. Proper waste separation at the source remains one of the most effective strategies for reducing landfill burden and improving recyclability, but manual sorting is slow, inconsistent, and hazardous for workers. This paper describes an automated garbage classification system built on a deep learning backbone, designed to sort waste in real time. The system uses MobileNetV2 with transfer learning to classify waste into twelve distinct categories. Training was supported by aggressive data augmentation and a two-phase learning process that first trains a custom classification head, then fine-tunes the deeper feature extraction layers. On a dataset of 15,515 images, the system reached a validation accuracy of 96.97%, which held up well across precision, recall, F1-score, and confusion matrix evaluations. The results confirm that lightweight deep learning architectures can approach the performance of far heavier models while remaining deployable on edge hardware, smart waste bins, and IoT-connected systems.

Keywords: Garbage Classification, Deep Learning, MobileNetV2, Transfer Learning, Image Classification, Convolutional Neural Networks (CNN)

1. Introduction

Waste generation has grown into a defining challenge of the modern world. As cities expand and consumption patterns intensify, the volume of garbage produced every day has outpaced the capacity of traditional disposal and recycling methods. The consequences include overflowing landfills, contaminated water sources, increased greenhouse gas emissions, and a significant drain on municipal budgets. Recycling can address many of these issues, but its effectiveness depends heavily on waste being correctly sorted at the source or at sorting facilities.

In practice, waste sorting is still performed manually in most regions. Workers at sorting stations must physically identify and separate materials such as plastic, metal, glass, cardboard, and organic waste. This is not only slow and expensive, but it also puts workers in contact with sharp objects, chemical residues, and biological material on a daily basis. Attempts to automate this using traditional image processing and machine learning, relying on handcrafted features and classifiers like SVM or KNN, showed early promise but fell short in real-world conditions where garbage is crumpled, dirty, overlapping, or only partially visible.

Deep learning has changed the equation considerably. Convolutional neural networks can learn to extract relevant visual features directly from raw image data, without the need for manual feature engineering. High-performing architectures like ResNet, DenseNet, and NASNet have demonstrated strong classification accuracy on waste datasets. However, these models are computationally expensive and impractical for real-time use on low-cost hardware embedded in smart waste bins.

This work addresses that gap by proposing a garbage classification system based on MobileNetV2, a lightweight CNN architecture designed with computational efficiency in mind. By combining pretrained ImageNet weights with transfer learning and targeted fine-tuning, the system achieves high accuracy while remaining suitable for edge device deployment.

2. Literature Review

The field of automated waste classification has evolved considerably over the past decade, moving from rule-based image processing methods toward sophisticated deep learning pipelines. Understanding this progression helps clarify where the present work fits and what gaps it aims to address.

2.1 Traditional Approaches

Early research leaned heavily on handcrafted features combined with classical machine learning classifiers. SVM and KNN were paired with SIFT and HOG feature extractors to describe the shape, color, and texture of waste items. Under controlled laboratory conditions these systems worked reasonably well, but they broke down when applied to real-world garbage that is frequently deformed, soiled, overlapping, or only partially visible.

2.2 Deep Learning Advances

The introduction of deep convolutional networks shifted the field substantially. VGG16, ResNet, and DenseNet demonstrated the power of hierarchical feature learning from raw pixel data. ResNet's skip connections and DenseNet's dense connectivity both yielded strong accuracy improvements for waste classification. However, large parameter counts made these models unsuitable for edge deployment.

2.3 Lightweight Models and Transfer Learning

Research attention eventually shifted toward architectures that balance accuracy and efficiency. MobileNetV2 emerged as a strong candidate through depthwise separable convolutions and inverted residual blocks. Transfer learning, pretraining on large datasets and fine-tuning on task-specific data, has become the standard technique for overcoming small dataset limitations in applied domains such as waste classification.

Table 1: Summary of Related Work

Author(s)	Method	Dataset / Classes	Result
Thung & Yang	SVM + HOG	TrashNet / 6	63% Accuracy
Simonyan et al.	VGG16	ImageNet / Large	89.1% Accuracy
He et al.	ResNet-50	TrashNet / 6	87.4% Accuracy
Kruthika et al.	NASNet	Custom / 12	99.7% Accuracy
Proposed	**MobileNetV2+TL**	**Kaggle GC / 12**	**96.97% Val. Acc.**

3. Objective

The primary goal of this research is to design, implement, and evaluate a deep learning-based garbage classification system that operates reliably under real-world conditions and remains lightweight enough for edge deployment. The following specific objectives guide the work:

- Collect and prepare the Kaggle Garbage Classification Dataset (15,515 images, 12 classes) with resizing, normalization, and real-time data augmentation.
- Develop a transfer learning model using MobileNetV2 as the backbone with a custom 12-class classification head.
- Train using a two-phase approach: head training with frozen backbone layers, followed by fine-tuning of deeper layers with appropriate callbacks.
- Evaluate using accuracy curves, confusion matrix, and per-class precision, recall, and F1-score.
- Compare performance against state-of-the-art architectures and assess feasibility for smart waste management deployment on edge hardware.

4. Theoretical Framework

MobileNetV2's efficiency over standard CNNs stems from two mathematical innovations: depthwise separable convolutions and inverted residual blocks with linear bottlenecks.

4.1 Standard Convolution Cost

In a conventional convolutional layer, an input feature map of size H x W x M is processed by N filters of size K x K. The computational cost in multiply-accumulate operations is: Cost = H x W x M x N x K^2. For a 3x3 kernel applied to a 64-channel feature map generating 128 output channels at 28x28 spatial resolution, this amounts to tens of millions of operations per layer.

4.2 Depthwise Separable Convolution

MobileNetV2 factorizes this into a depthwise convolution applying one K x K filter per input channel, followed by a pointwise 1x1 convolution combining outputs to produce N feature maps. Combined cost = H x W x M x (K^2 + N). The reduction ratio simplifies to 1/N + 1/K^2, yielding roughly 8 to 9 times fewer operations for K=3. This fundamental gain makes MobileNetV2 attractive for edge deployment without significant accuracy trade-offs.

5. Implementation

The garbage classification system was developed in Python using TensorFlow 2.x and the Keras API on Google Colab with GPU acceleration. The implementation covered data loading, preprocessing, augmentation, model construction, two-phase training, evaluation, and export for deployment.

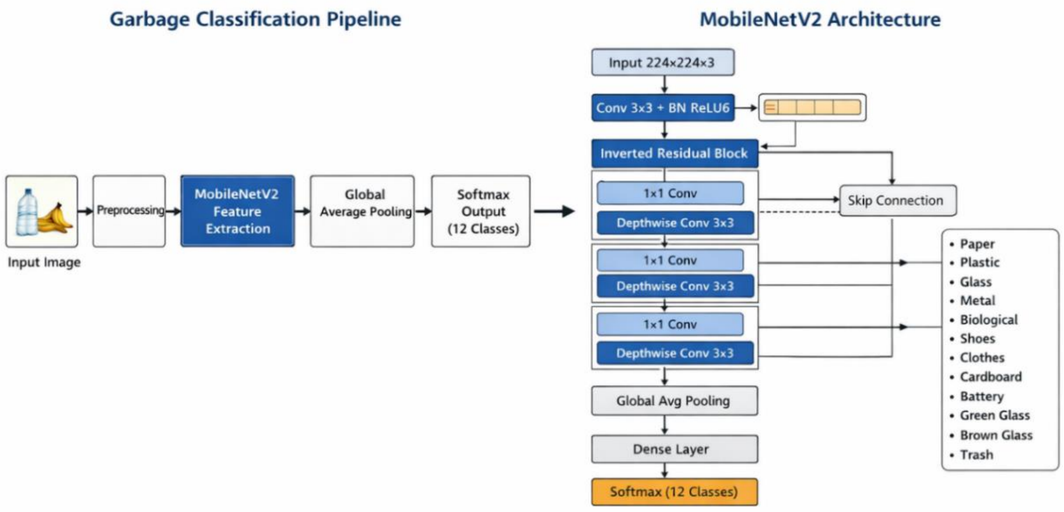

Fig. 1. Overall workflow of the proposed garbage classification system.

5.1 Dataset and Preprocessing

The publicly available Kaggle Garbage Classification Dataset contains 15,515 RGB images distributed across twelve waste categories: cardboard, brown-glass, metal, paper, plastic, biological, battery, clothes, green-glass, shoes, white-glass, and trash. Images were loaded using TensorFlow's image_dataset_from_directory at an 80:20 train-validation split, yielding approximately 12,400 training images and 3,100 validation images. All images were resized to 224 x 224 pixels and normalized using MobileNetV2's preprocess_input to match the distribution expected by the ImageNet-pretrained weights.

Fig. 3. Sample waste images from each of the twelve dataset categories.

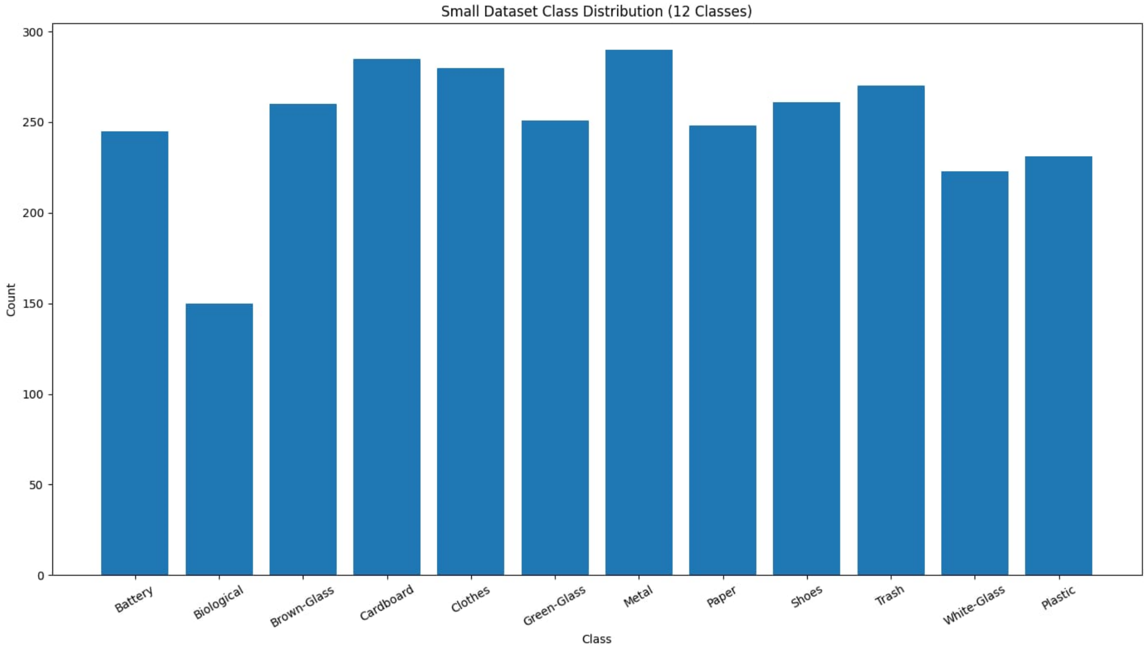

Fig. 4. Class-wise distribution of samples in the small dataset (12 classes).

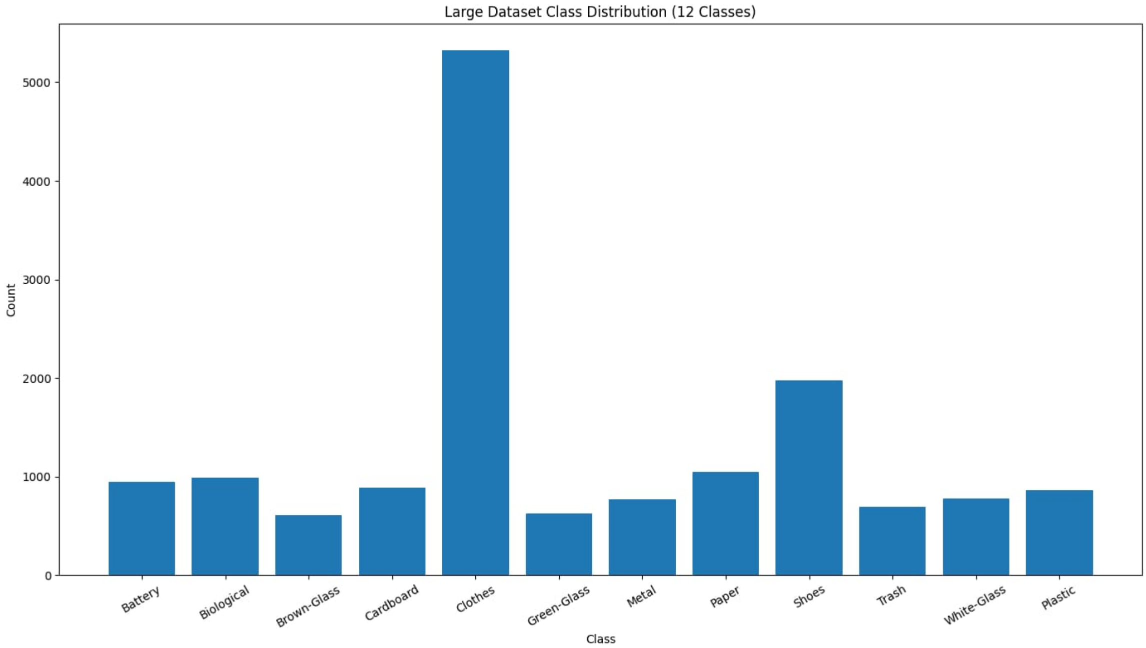

Fig. 5. Class-wise distribution of samples in the large dataset (12 classes).

5.2 Data Augmentation

Real-time augmentation was applied to the training pipeline including random horizontal flipping, rotation up to 15 degrees, zoom up to 10%, and contrast adjustment. These operations simulate real garbage disposal environment variability such as different orientations, lighting conditions, and partial occlusions, without requiring additional data collection.

5.3 Model Architecture

MobileNetV2, pretrained on ImageNet, served as the backbone feature extractor. Its final layers were replaced with a custom classification head consisting of: Global Average Pooling, Batch Normalization, Dense layer with 256 neurons and ReLU activation, Dropout at a rate of 0.5, and a Softmax output layer with 12 neurons.

5.4 Training Procedure

Training proceeded in two phases. Phase one froze all backbone layers and trained only the custom head for 10 epochs using Adam optimizer with sparse categorical cross-entropy loss. Phase two unfroze approximately 70% of the deeper backbone layers and fine-tuned the full network for 20 additional epochs at a reduced learning rate. Callbacks including Early Stopping, ReduceLROnPlateau, and Model Checkpointing ensured optimal weight preservation.

6. Results

Performance was assessed using training and validation accuracy curves, loss convergence plots, confusion matrix analysis, and per-class precision, recall, and F1-score metrics, providing a comprehensive view of system behavior across all twelve waste categories.

6.1 Training Dynamics

During phase one, validation accuracy exceeded training accuracy from the very first epoch, a clear indicator that effective transfer learning was occurring. The pretrained backbone immediately provided generalizable visual patterns useful for waste domain classification. Fine-tuning produced further gains, with smooth accuracy and loss curves throughout both phases.

6.2 Final Performance

At the conclusion of fine-tuning, the model achieved approximately 99% training accuracy and 96.97% validation accuracy with a validation loss of around 0.15. The small gap between training and validation metrics indicates effective generalization rather than memorization of training samples.

6.3 Confusion Matrix Analysis

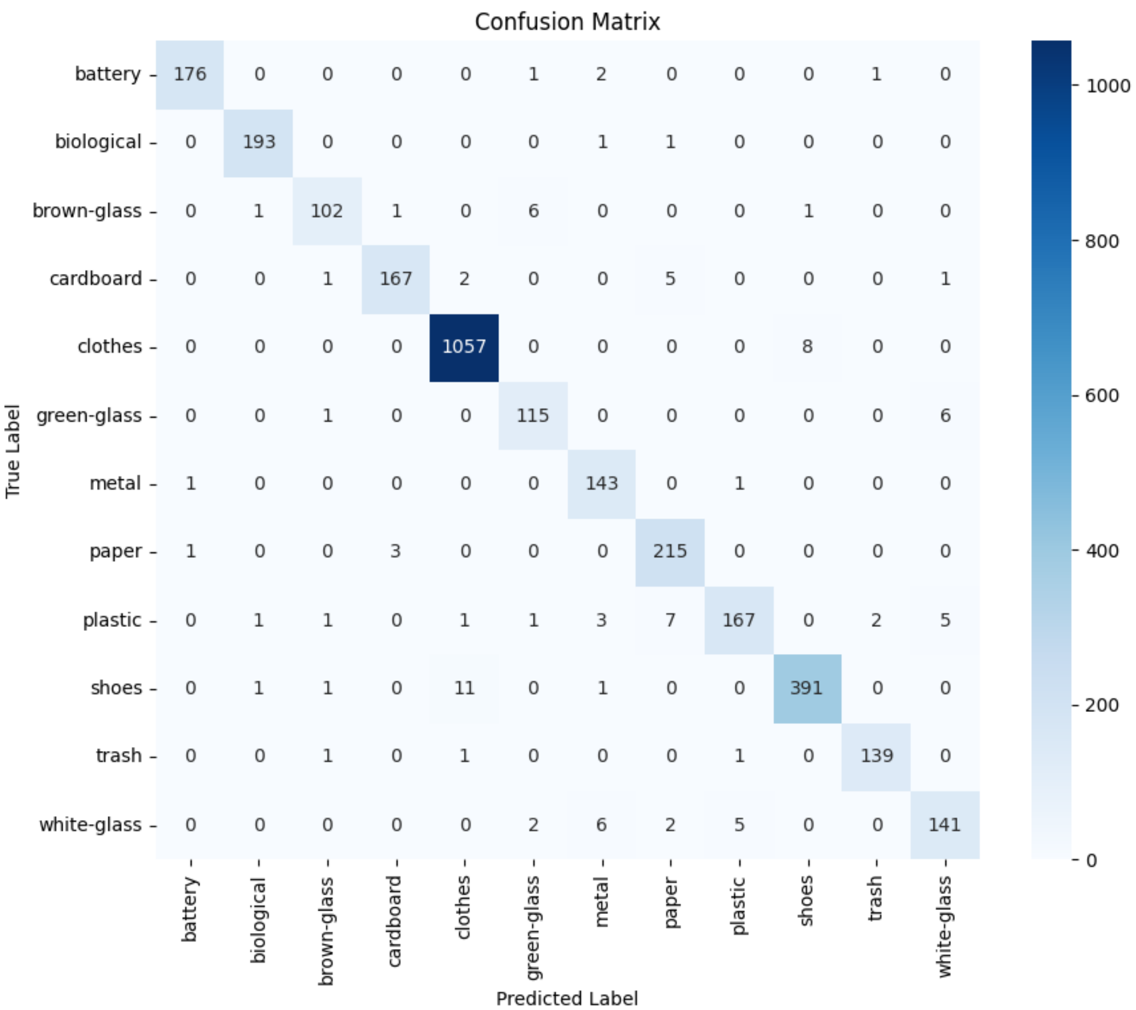

Fig. 6. Confusion matrix for 12-class garbage classification.

The confusion matrix showed a strongly dominant diagonal, confirming correct predictions in the vast majority of cases. Off-diagonal values were small and concentrated in visually similar class pairs such as the three glass subtypes and between plastic and metal items. No single category showed systematic misclassification, indicating that the model has learned meaningful representations for all twelve waste types.

6.4 Classification Report Analysis

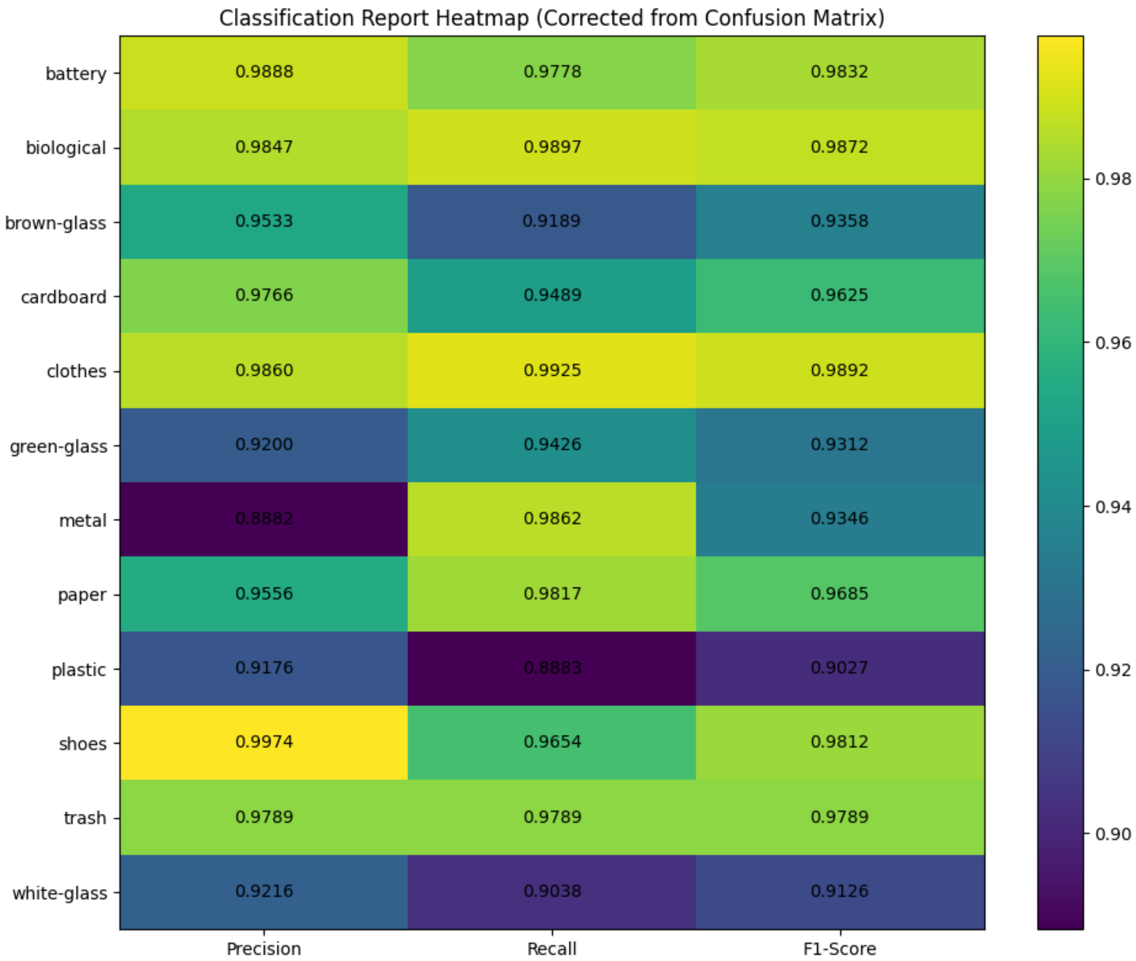

Fig. 7. Classification report heatmap showing precision, recall, and F1-score for 12 waste classes.

Most waste categories achieved F1-scores above 0.95, with clothes, shoes, battery, paper, and cardboard performing particularly well. Some reduction in scores was observed for the glass subtypes and plastic vs metal, which share surface texture and reflective properties. The heatmap confirms robust and balanced multi-class classification performance suitable for real-world deployment.

7. Comparative Discussion

To situate the proposed system within the broader research landscape, its performance is compared against well-known architectures on similar waste classification tasks.

Table 2: Comparison with State-of-the-Art Models

Model	Accuracy	Dataset Size	Deployment
NASNet	99.7%	15,555 images	Cloud Server Only
InceptionV3	99.6%	15,555 images	Cloud / Desktop
MobileNet (Ref.)	95.3%	15,555 images	Edge / Mobile
MobileNetV2 (Prop.)	**96.97%**	**15,515 images**	**Edge / IoT / Smart Bin**

NASNet and InceptionV3 achieve slightly higher raw accuracy but require cloud-grade computational infrastructure and are impractical for real-time edge deployment. MobileNetV2 reaches 96.97% validation accuracy on a comparable dataset while remaining deployable on single-board computers. The two-phase fine-tuning strategy applied here yields meaningful accuracy gains over the reference MobileNet model without additional hardware cost.

Limitations: The dataset does not fully capture the diversity of real-world garbage, particularly soiled, heavily deformed, or partially occluded items in outdoor environments. The glass subcategories remain the most challenging, and further work is needed to improve discrimination among materials with similar reflective surface properties.

9. Conclusion

This research presented the design, implementation, and evaluation of a deep learning-based automated garbage classification system using MobileNetV2 with transfer learning. Trained and evaluated on 15,515 waste images spanning twelve categories, the model achieved a validation accuracy of 96.97% through a combination of data augmentation, two-phase training, and targeted regularization, demonstrating that a lightweight architecture can deliver performance approaching far more resource-intensive alternatives.

The per-class analysis and classification report heatmap confirmed strong and balanced classification capability across all twelve waste categories. The confusion matrix reinforced these findings, showing a dominant diagonal with only small and expected off-diagonal values, primarily among visually similar materials.

From a practical standpoint, the proposed system holds real promise for deployment in smart waste management contexts including smart bins, IoT monitoring systems, and edge computing sorting facilities. Future work should focus on expanding the dataset, exploring attention mechanisms or transformer-based hybrid architectures, and investigating pruning and quantization for lower-power hardware deployment.

10. References

[1] S. Kruthika M., R. Rajadevi, D. Sathya, et al., Garbage Classification: A Deep Learning Perspective, IRJAEH, vol. 02, no. 12, pp. 2774-2780, Dec. 2024.

[2] Kaggle, Garbage Classification Dataset, Available: https://www.kaggle.com

[3] M. Sandler, A. Howard, M. Zhu, et al., MobileNetV2: Inverted Residuals and Linear Bottlenecks, Proc. IEEE CVPR, 2018.

[4] G. Thung and M. Yang, Classification of Trash for Recyclability Status, Stanford University Project Report, 2016.

[5] S. J. Pan and Q. Yang, A Survey on Transfer Learning, IEEE Trans. Knowledge and Data Engineering, vol. 22, no. 10, pp. 1345-1359, 2010.

[6] K. Simonyan and A. Zisserman, Very Deep Convolutional Networks for Large-Scale Image Recognition, ICLR, 2015.

[7] B. Zoph et al., Learning Transferable Architectures for Scalable Image Recognition (NASNet), IEEE CVPR, 2018.

[8] G. Huang et al., Densely Connected Convolutional Networks, IEEE CVPR, 2017.