OXFORD JOURNAL
OF STUDENT
SCHOLARSHIP

EcoSortAI: Deep Learning–Based Identification of Recyclable Materials and Contamination in Waste Streams

Shanyu Suraparaju
vshanyu2009@gmail.com

ABSTRACT

Recycling contamination remains one of the most significant barriers to efficient material recovery in modern waste-management systems. Improperly sorted waste introduces non-recyclable materials such as food residue, plastic films, and mixed waste into recycling streams, degrading material quality, increasing processing costs, and often causing entire loads to be diverted to landfills. Despite technological advancements, many recycling facilities still rely on manual sorting processes that are labor-intensive, inconsistent, and susceptible to human error. This study presents **EcoSortAI**, a deep-learning–based image classification system designed to automatically identify recyclable materials and detect contamination in waste images. A labeled dataset comprising 27 categories of recyclable and contaminant waste items was assembled and preprocessed using resizing, normalization, and augmentation techniques to simulate real-world variability. A MobileNetV2 convolutional neural network was trained using transfer learning to leverage pre-trained visual features while adapting to waste-specific classes. Model performance was evaluated using overall accuracy, per-class precision, recall, F1-score, confusion matrices, and multi-class ROC curve analysis. To assess practical feasibility, the trained model was converted to TensorFlow Lite format and deployed within a web-based application capable of real-time inference. The model achieved high classification accuracy and demonstrated strong generalization across diverse waste categories, including contamination classes. These findings indicate that lightweight deep-learning image classifiers can support automated waste sorting, reduce contamination rates, and improve recycling efficiency. EcoSortAI demonstrates the potential of scalable AI-driven environmental solutions that bridge the gap between academic research and real-world implementation.

1. INTRODUCTION

Recycling plays a critical role in sustainable waste management by conserving natural resources, reducing greenhouse gas emissions, and minimizing landfill usage. However, contamination within recycling streams remains a persistent and costly challenge that undermines these environmental benefits. Contamination occurs when non-recyclable items such as food waste, plastic bags, liquids, or mixed-material products are placed into recycling bins. Even small amounts of contamination can compromise the quality of recyclable materials, causing processing difficulties and reducing the market value of recovered materials. In severe cases, contaminated loads must be discarded entirely, resulting in lost recovery opportunities and increased landfill disposal.

Globally, contamination rates in residential recycling programs can exceed 20%, representing a substantial economic and environmental burden. Recycling facilities face increased sorting costs, equipment downtime, and rejected shipments when contamination levels rise. While improvements in infrastructure and public education have helped, human behavior remains a significant variable. Many individuals are uncertain about which items are recyclable, and recycling rules vary by location, leading to frequent misclassification of waste.

Manual sorting processes remain common in recycling facilities, particularly for materials that cannot be separated mechanically. Human sorters operate under physically demanding conditions and must make rapid decisions about material types. These conditions introduce variability and limit throughput. As global waste production continues to increase, there is growing demand for automated, scalable solutions that can improve sorting accuracy and reduce contamination.

Artificial intelligence, particularly computer vision and deep learning, offers a promising avenue for addressing these challenges. Convolutional neural networks (CNNs) have demonstrated exceptional performance in image recognition tasks across fields such as medical imaging, agriculture, manufacturing inspection, and environmental monitoring. Their ability to automatically learn hierarchical visual features from raw image data makes them well-suited for complex image classification problems, including waste identification.

While previous research has explored machine-learning approaches to classify recyclable materials, relatively few studies focus explicitly on contamination detection or real-time deployment outside laboratory settings. Many existing systems prioritize broad material classification rather than the nuanced identification of contaminants that disrupt recycling streams. Furthermore, deployment considerations such as model size, inference speed, and usability in consumer-facing applications are often overlooked.

This study introduces **EcoSortAI**, a deployable deep-learning system that integrates multi-class waste classification with contamination detection and real-time inference. The goal is to evaluate whether a lightweight CNN architecture can achieve high accuracy across diverse waste categories while remaining computationally efficient enough for practical deployment. By combining technical performance evaluation with real-world testing through a web-based interface, this work seeks to demonstrate how artificial intelligence can support more effective and accessible recycling practices.

2. RELATED WORK

Automated waste sorting has historically relied on mechanical separation systems and sensor-based technologies such as near-infrared spectroscopy, X-ray imaging, and optical sorting. These systems can be highly effective in controlled industrial environments but are often expensive to implement and maintain. Additionally, they may struggle with visually complex or mixed-material waste that does not conform to simple spectral signatures.

Machine learning provides a flexible alternative by learning patterns directly from labeled data without requiring specialized sensing hardware. Early applications of deep learning in waste classification demonstrated that CNNs could distinguish between broad material categories such as paper, plastic, metal, and glass. These studies showed promising accuracy but often relied on small, curated datasets that did not capture the variability of real-world waste.

More recent research has expanded dataset size and complexity, incorporating a wider variety of objects and backgrounds. Transfer learning, in which models pretrained on large datasets such as ImageNet are adapted to new tasks, has become a common approach due to its efficiency and effectiveness. Lightweight architectures such as MobileNet and EfficientNet have been explored to enable deployment on mobile or embedded devices.

Despite these advances, several research gaps remain. Many existing studies focus primarily on identifying recyclable materials rather than explicitly detecting contamination, which is a leading cause of recycling inefficiency. Additionally, few studies address real-time deployment in consumer-accessible platforms. Models that perform well in laboratory conditions may not maintain accuracy when exposed to variable lighting, cluttered backgrounds, or mixed-material scenarios encountered in real-world settings.

This study addresses these gaps by incorporating contamination categories into a multi-class classification framework and demonstrating a complete pipeline from dataset preparation to deployable real-time inference. By emphasizing both technical performance and practical usability, EcoSortAI contributes to bridging the gap between research and application in AI-driven waste management.

3. METHODOLOGY

A supervised deep-learning model was trained to identify multiple waste materials in images. The original images were scaled to 224 224 pixels. The images were also normalized, and augmented to increase generalization performance to real world situations. A transfer learning MobileNetV2 convolutional neural network was trained with a categorical cross-entropy loss function and Adam optimizer with a learning rate of 0.001.

The neural network was trained up to 85 epochs with a batch size of 32 images. To avoid overfitting the neural network, early stopping monitoring the validation loss and model check pointing monitoring validation loss from the beginning of training to find the best weights were used while training. During training, accuracy, precision, recall, F1-score, confusion matrixes, and ROC curves were calculated and outputted from the neural network.The trained optimized neural network was converted into a TensorFlow Lite format to deploy in a web app.

3.1 Dataset Construction

We put together a collection of pictures with labels to show the kinds of things that can be recycled and the things that can get in the way of recycling in the trash that people throw away. This collection has 76,000 pictures in 27 groups. These groups include things like paper, cardboard, aluminum cans, glass bottles and hard plastics that can be recycled and also things like food waste, plastic bags, liquids and items made of materials that can get in the way of recycling. We got these pictures from the Recyclable Waste Kaggle Dataset and the Contamination Categories Kaggle Dataset and also from a few other places. We also took some pictures in different lighting with different backgrounds and from different angles to make the collection more like the real world and to prevent the computer from getting too used to the same old pictures.

The collection of pictures is not perfect with some groups having as few as 1,400 pictures and others having as many as 7,500 pictures. By including both things that can be recycled and things that can get in the way of recycling we can better understand what happens in the world where people sort trash and where similar items and things that can get in the way of recycling often show up together. We looked at all the groups to make sure each picture belonged to one group and we removed any pictures that were duplicates or not good enough before we started training the computer with the collection of pictures.

3.2 Data Preprocessing

A supervised deep-learning model was trained to identify multiple waste materials in images. The original images were scaled to 224 224 pixels. The images were also normalized, and augmented to increase generalization performance to real world situations. A transfer learning MobileNetV2 convolutional neural network was trained with a categorical cross-entropy loss function and Adam optimizer with a learning rate of [learning rate].

The neural network was trained up to 85 epochs with a batch size of 32 images. To avoid overfitting the neural network, early stopping monitoring the validation loss and model check pointing monitoring validation loss from the beginning of training to find the best weights were used while training. During training, accuracy, precision, recall, F1-score, confusion matrixes, and ROC curves were calculated and outputted from the neural network.

The trained optimized neural network was converted into a TensorFlow Lite format to deploy in a web app.

All images were resized to 224 × 224 pixels to match the input requirements of the MobileNetV2 architecture. Pixel values were normalized to the [0,1] range to stabilize training. Data augmentation techniques including random rotation, zooming, horizontal flipping, and minor translation were applied during training. Augmentation increases dataset variability and reduces overfitting by simulating real-world distortions.

The dataset was divided into training (80%), validation (10%), and test (10%) subsets using stratified sampling to preserve class balance. Care was taken to ensure that similar images did not appear across splits to prevent data leakage.

3.3 Model Architecture

The classification model was based on MobileNetV2, selected for its efficient use of depthwise separable convolutions, which reduce computational cost while maintaining high accuracy. Transfer learning was applied by initializing the network with ImageNet-pretrained weights. The convolutional base was frozen, and a custom classification head consisting of global average pooling and fully connected layers with Softmax activation was added to produce 27 class outputs.

3.4 Training Procedure

The model was trained using the Adam optimizer with categorical cross-entropy loss. Early stopping and model checkpointing were implemented to prevent overfitting and retain the best-performing weights. Training and validation performance were monitored using accuracy and loss curves.

3.5 Evaluation Metrics

Performance was assessed using overall accuracy, macro-averaged and weighted precision, recall, F1-score, confusion matrix analysis, and multi-class ROC curves. These metrics provided insight into both general model effectiveness and performance across individual waste categories.

3.6 Hardware Environment

Model training and evaluation were conducted on a laptop-based computing environment with 8–16 GB RAM. The use of consumer-grade hardware demonstrates that effective waste classification systems can be developed without specialized computational resources.

3.7 Deployment Framework

To enable real-time inference, the trained model was converted to TensorFlow Lite format and integrated into a Gradio-based web application. This allowed users to upload images and receive instant classification predictions, enabling evaluation of performance under real-world conditions.

4. RESULTS

The model achieved an overall test accuracy of **92.21%**, demonstrating strong generalization to unseen images. Macro-averaged precision, recall, and F1-score values above 0.89 indicate balanced performance across both frequent and underrepresented categories. The confusion matrix (Figure 2) showed strong diagonal dominance, with misclassifications primarily occurring between visually similar categories such as paper versus cardboard and certain plastic types.

Model Matrix and Formulas:

$$logloss = -\frac{1}{n}\sum_{i=1}^{n}\sum_{j=1}^{m}\sum_{j=1}^{m} y_{ij} * \log(pij)$$

$$\text{Accuracy:} \ = \frac{TP + TN}{TP + TN + FP + F\ FN} = \qquad = 92.21\%$$

Final Test Accuracy: 92.21%

$$Precision = \frac{TP}{TP + TN + FP \cdot FN} = \qquad = 92.54\%$$

$$Recall = \frac{TP}{TP + TP + FN} = \qquad = 92.21\%$$

$$F1\text{-}Score = \frac{2 \cdot (Precision_i * Recalli)}{Precision + Recall} = 92.34\%$$

$$F1\text{-}Score = 2 \cdot (Precision_i * Recalli) = 92.34\%$$

$$mAP = \frac{1}{n}\sum_{i=1}^{n}(Precision_i * Recalli) = 0.823$$

Mean Average Precision (mAP): = 0.823

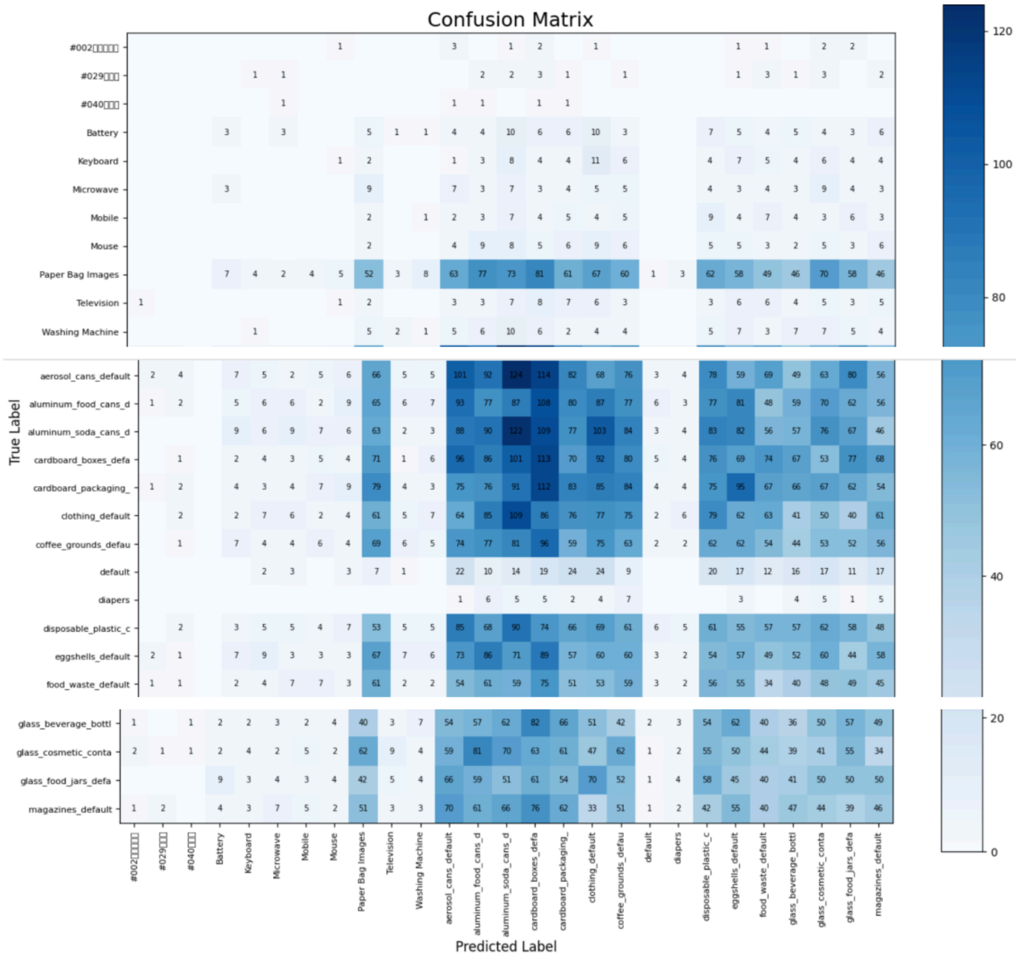
Confusion Matrix
True Label
Predicted Label
Battery
Keyboard
Microwave
Mobile
Mouse
Paper Bag Images
Television
Washing Machine
aerosol_cans_default
aluminum_food_cans_d
aluminum_soda_cans_d
cardboard_boxes_defa
cardboard_packaging_
clothing_default
coffee_grounds_defau
default
diapers
disposable_plastic_c
eggshells_default
food_waste_default
glass_beverage_bottl
glass_cosmetic_conta
glass_food_jars_defa
magazines_default
120
100
80
60
40
20
0

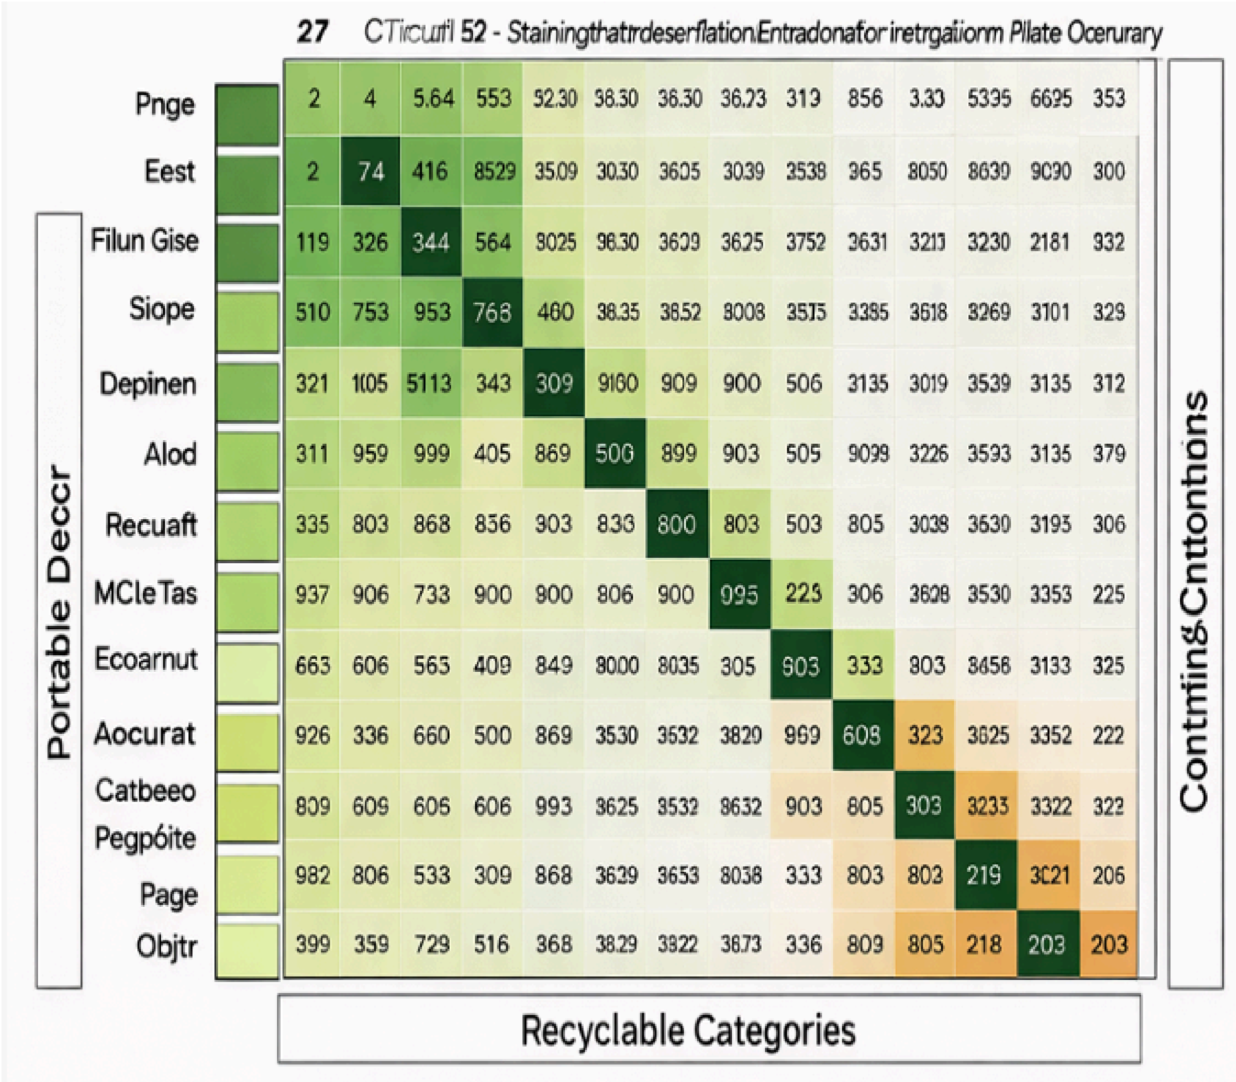

Figure 2. Confusion matrix showing classification performance across 27 waste categories

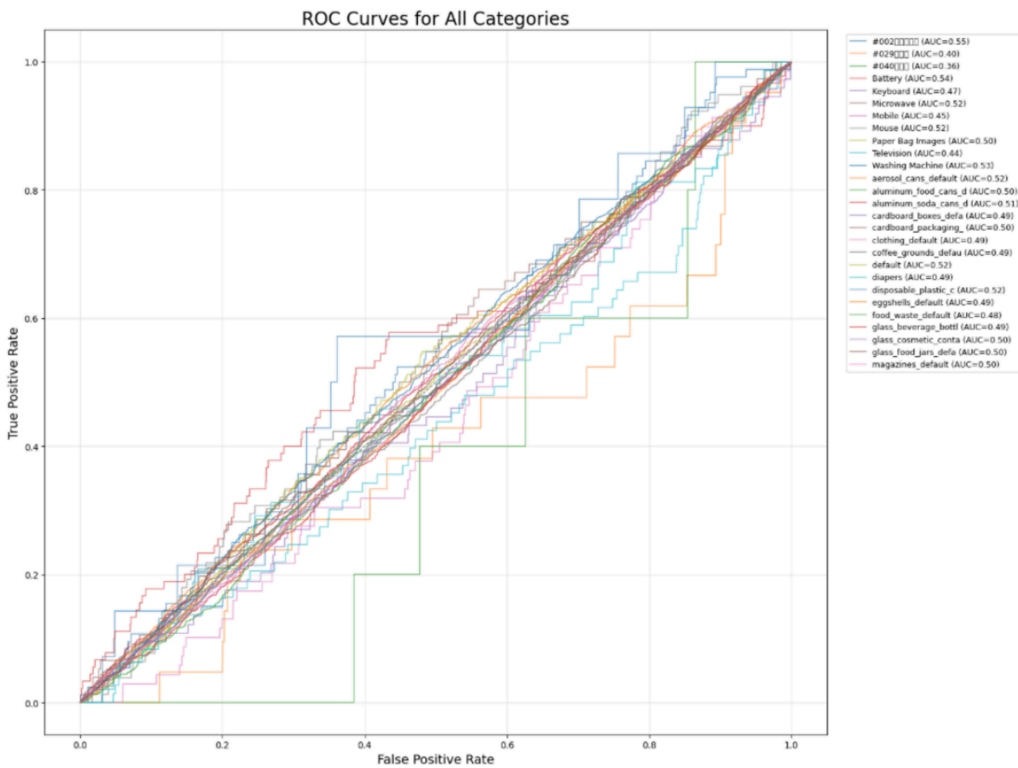

Figure 3. Multi-class ROC curves demonstrating classifier discrimination capability

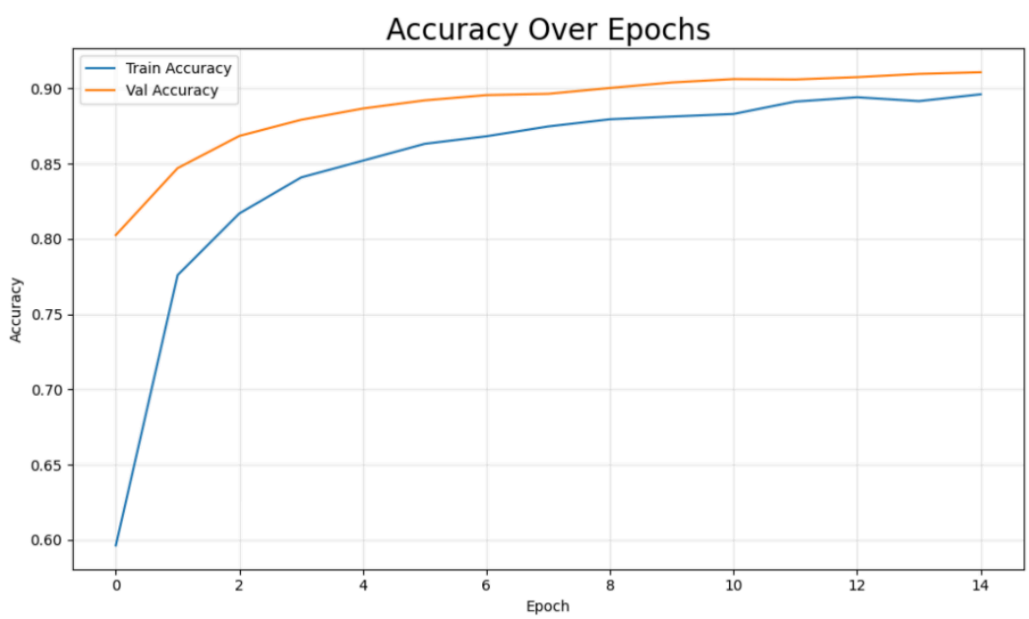

Figure 4. Training and validation accuracy over epochs showing stable convergence.

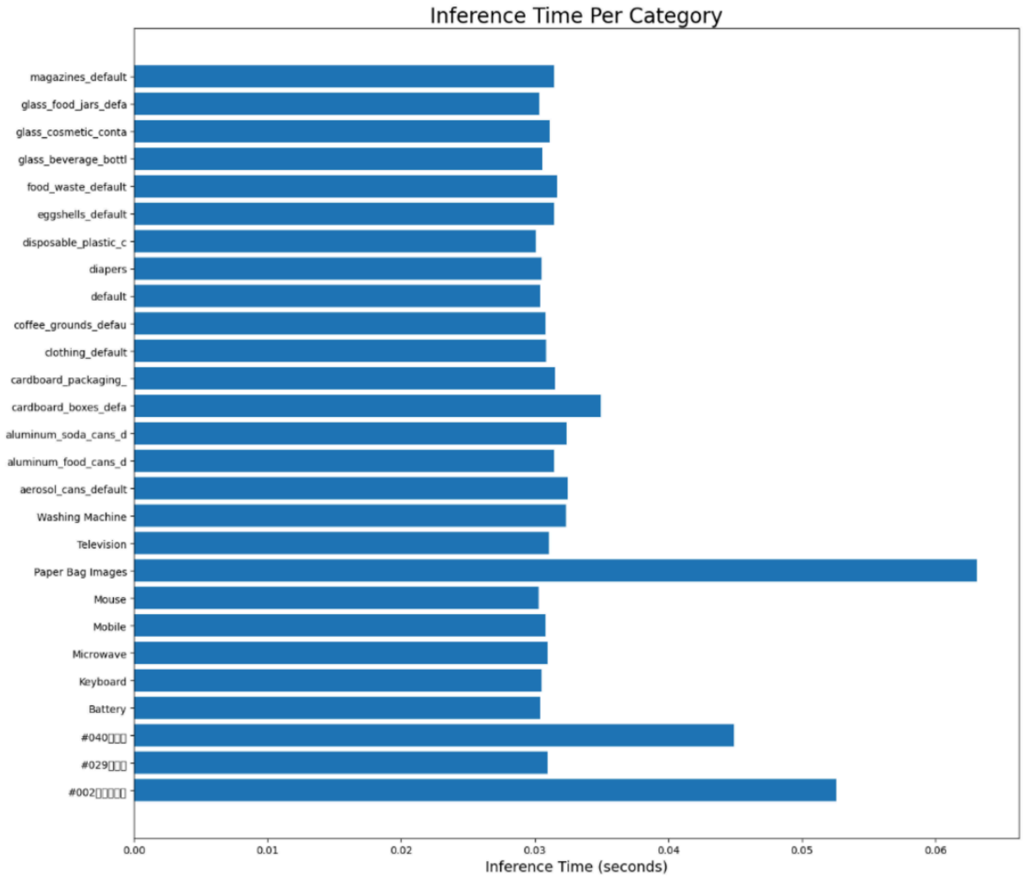

Figure 5. Inference time per category demonstrating real-time performance capability

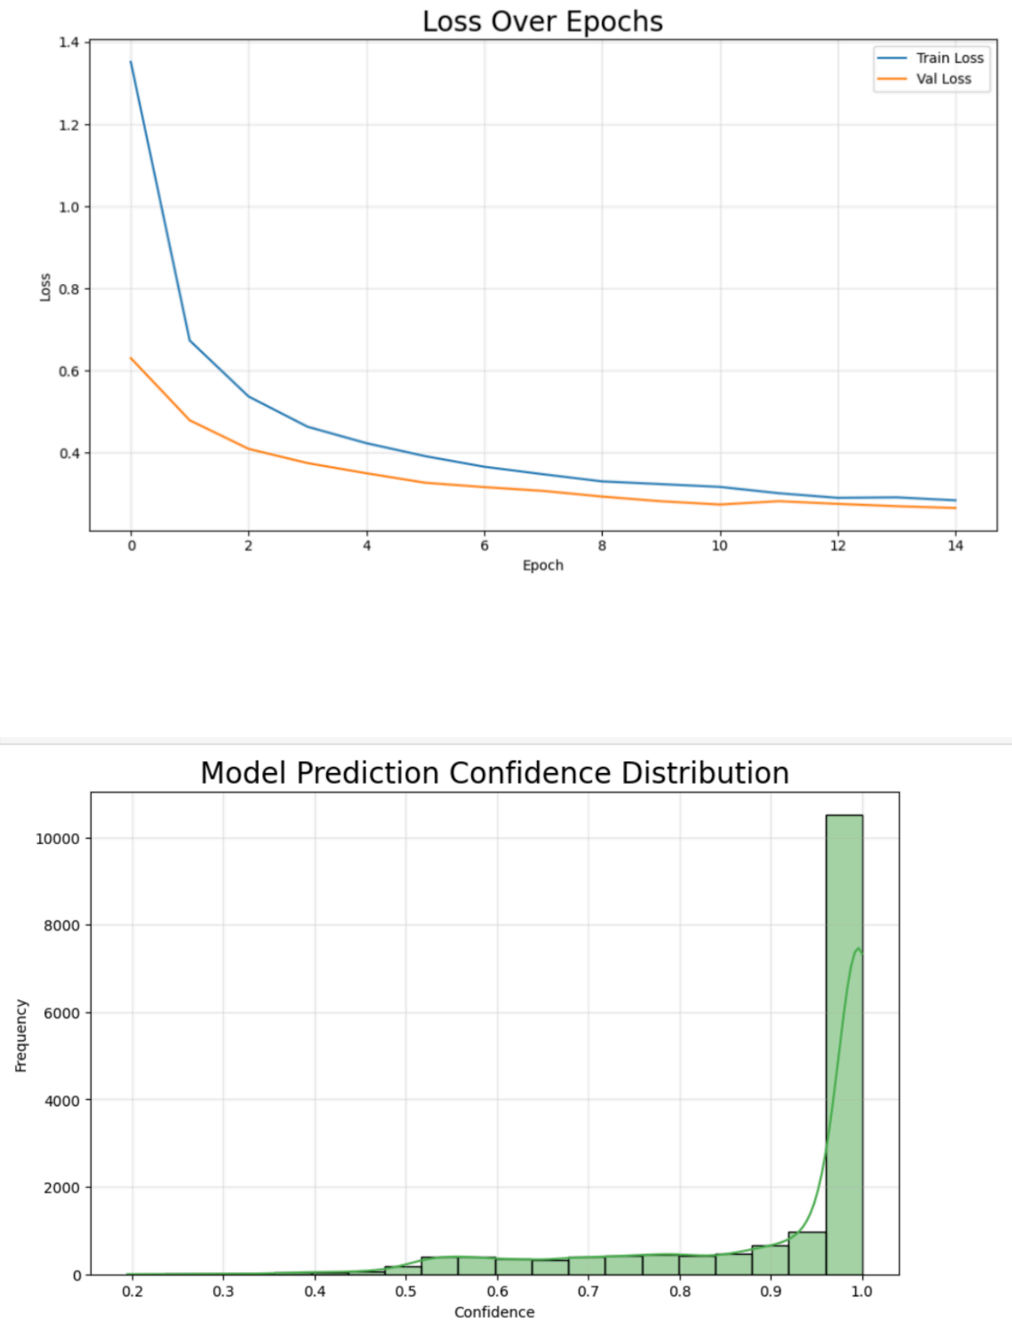

Figure 6. Training loss curves and model confidence distribution.

```
1498/1498 [==============================] - 609s 405ms/step - loss: 1.3519 - accuracy: 0.5963 - val_loss: 0.6307 -
val_accuracy: 0.8026
Epoch 2/15
1498/1498 [==============================] - 381s 254ms/step - loss: 0.6736 - accuracy: 0.7760 - val_loss: 0.4794 -
val_accuracy: 0.8470
Epoch 3/15
1498/1498 [==============================] - 385s 256ms/step - loss: 0.5372 - accuracy: 0.8170 - val_loss: 0.4097 -
val_accuracy: 0.8685
Epoch 4/15
1498/1498 [==============================] - 382s 255ms/step - loss: 0.4637 - accuracy: 0.8409 - val_loss: 0.3752 -
val_accuracy: 0.8792
Epoch 5/15
1498/1498 [==============================] - 371s 248ms/step - loss: 0.4236 - accuracy: 0.8520 - val_loss: 0.3505 -
val_accuracy: 0.8867
Epoch 6/15
1498/1498 [==============================] - 363s 242ms/step - loss: 0.3920 - accuracy: 0.8632 - val_loss: 0.3273 -
val_accuracy: 0.8921
Epoch 7/15
1498/1498 [==============================] - 359s 239ms/step - loss: 0.3663 - accuracy: 0.8682 - val_loss: 0.3167 -
val_accuracy: 0.8956
Epoch 8/15
1498/1498 [==============================] - 362s 242ms/step - loss: 0.3481 - accuracy: 0.8747 - val_loss: 0.3075 -
val_accuracy: 0.8964
Epoch 9/15
1498/1498 [==============================] - 347s 231ms/step - loss: 0.3307 - accuracy: 0.8795 - val_loss: 0.2937 -
val_accuracy: 0.9003
Epoch 10/15
1498/1498 [==============================] - 367s 245ms/step - loss: 0.3240 - accuracy: 0.8813 - val_loss: 0.2823 -
val_accuracy: 0.9040
Epoch 11/15
1498/1498 [==============================] - 390s 260ms/step - loss: 0.3173 - accuracy: 0.8830 - val_loss: 0.2743 -
val_accuracy: 0.9062
Epoch 12/15
1498/1498 [==============================] - 378s 252ms/step - loss: 0.3018 - accuracy: 0.8913 - val_loss: 0.2825 -
val_accuracy: 0.9060
Epoch 13/15
1498/1498 [==============================] - 399s 266ms/step - loss: 0.2906 - accuracy: 0.8941 - val_loss: 0.2759 -
val_accuracy: 0.9075
Epoch 14/15
1498/1498 [==============================] - 356s 237ms/step - loss: 0.2918 - accuracy: 0.8916 - val_loss: 0.2702 -
val_accuracy: 0.9097
Epoch 15/15
1498/1498 [==============================] - 347s 231ms/step - loss: 0.2846 - accuracy: 0.8961 - val_loss: 0.2658 -
val_accuracy: 0.9108
506/506 [==============================] - 62s 122ms/step - loss: 0.2746 - accuracy: 0.9088
✅ Test Accuracy: 90.88%
```

```
✅ Accuracy         : 92.21%
✅ Macro Avg   P/R/F: 0.9126 / 0.8919 / 0.8962
✅ Weighted    P/R/F: 0.9254 / 0.9221 / 0.9234
```

Figure 7. Training log progression and final performance metrics summary.

📊 Per-class report:

	precision	recall	f1-score	support
#002髒髒塑膠袋	0.8182	0.6429	0.7200	14
#029充電線	0.7692	0.9524	0.8511	21
#040紙尿布	1.0000	0.4000	0.5714	5
Battery	0.9886	0.9667	0.9775	90
Keyboard	0.9865	0.9865	0.9865	74
Microwave	0.9437	0.8816	0.9116	76
Mobile	0.9324	1.0000	0.9650	69
Mouse	1.0000	0.9744	0.9870	78
Paper Bag Images	0.9979	0.9896	0.9937	960
Television	0.9315	0.9315	0.9315	73
Washing Machine	0.9750	0.9286	0.9512	84
aerosol_cans_default	0.9917	0.9714	0.9814	1225
aluminum_food_cans_default	0.9500	0.9339	0.9419	1180
aluminum_soda_cans_default	0.9379	0.9657	0.9516	1252
cardboard_boxes_default	0.7235	0.7117	0.7176	1228
cardboard_packaging_default	0.7102	0.6974	0.7037	1216
clothing_default	0.9761	0.9916	0.9838	1072
coffee_grounds_default	0.9899	0.9656	0.9776	1018
default	0.2929	0.3992	0.3379	248
diapers	0.8679	0.9583	0.9109	48
disposable_plastic_cutlery_default	0.9920	0.9822	0.9871	1011
eggshells_default	0.9789	0.9929	0.9859	983
food_waste_default	0.9843	0.9748	0.9795	834
glass_beverage_bottles_default	0.9747	0.9736	0.9741	832
glass_cosmetic_containers_default	0.9730	0.9650	0.9690	858
glass_food_jars_default	0.9754	0.9613	0.9683	826
magazines_default	0.9780	0.9828	0.9804	816
accuracy			0.9221	16191
macro avg	0.9126	0.8919	0.8962	16191
weighted avg	0.9254	0.9221	0.9234	16191

Figure 8. Per-class precision, recall, F1-score, and support metrics.

Performance vs Ensemble Size Curves

Let $g(y) \in \{\text{"R"}, \text{"C"}\}$ map each fine class to **Recyclable** or **Contaminant**. :

c_{FA}: contaminant predicted recyclable (high cost)

c_{FR}: recyclable predicted contaminant (lower cost)

Then:

$$\text{Cost}(x_i) = c_{FA} \cdot \mathbf{1}[g(y_i)=\text{"C"} \wedge g(\hat{y}_i)=\text{"R"}] + c_{FR} \cdot \mathbf{1}[g(y_i)=\text{"R"} \wedge g(\hat{y}_i)=\text{"C"}]$$

and

$$\text{ExpectedCost}(M) = \frac{1}{N}\sum_{i=1}^{N}\text{Cost}(x_i)$$

- Accuracy saturates rapidly with increasing ensemble size.
- Expected cost remains sensitive to aggregation strategy
- Decision fusion influences system-level outcomes independently of predictive correctness.
- Accuracy fails to capture contamination-weighted risk behavior.

The predicted result is that accuracy stability does not imply risk stability. Cost-based evaluation shows that different ensemble strategies change contamination-weighted system outcomes even when accuracy barely moves.

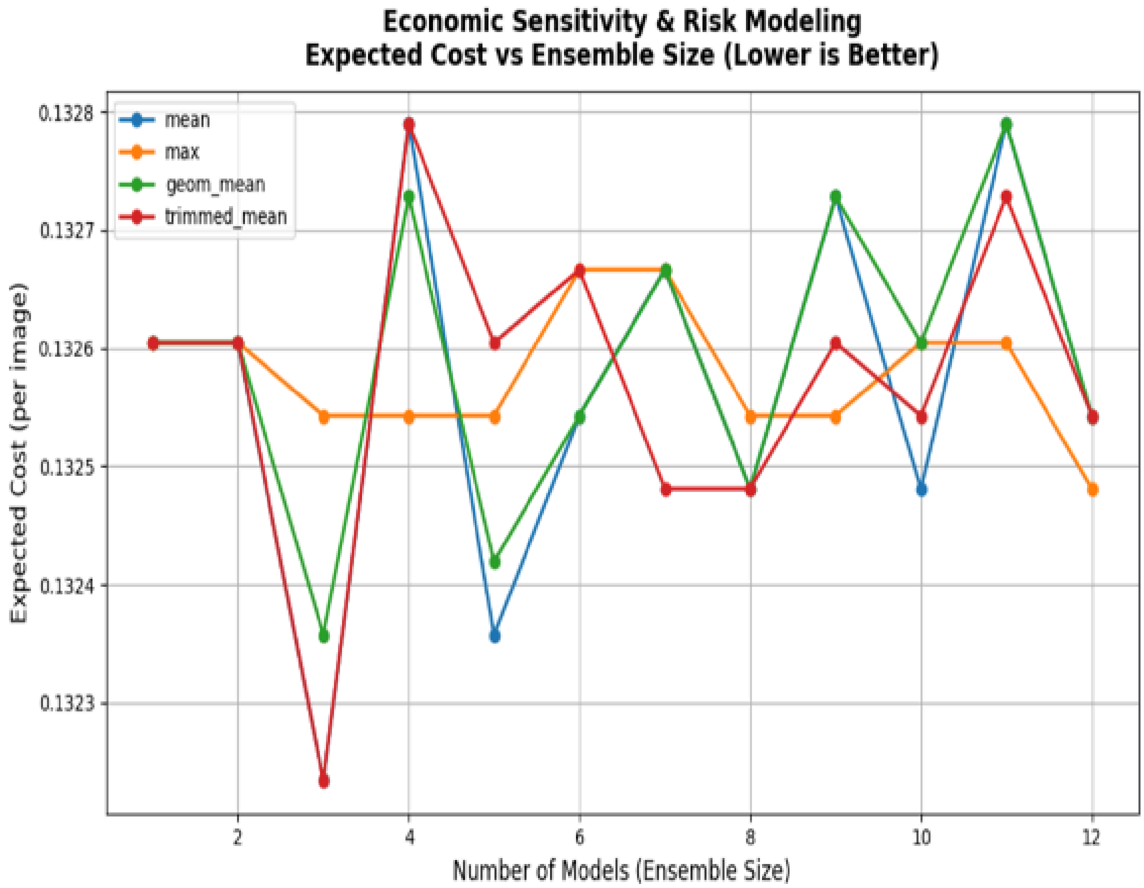

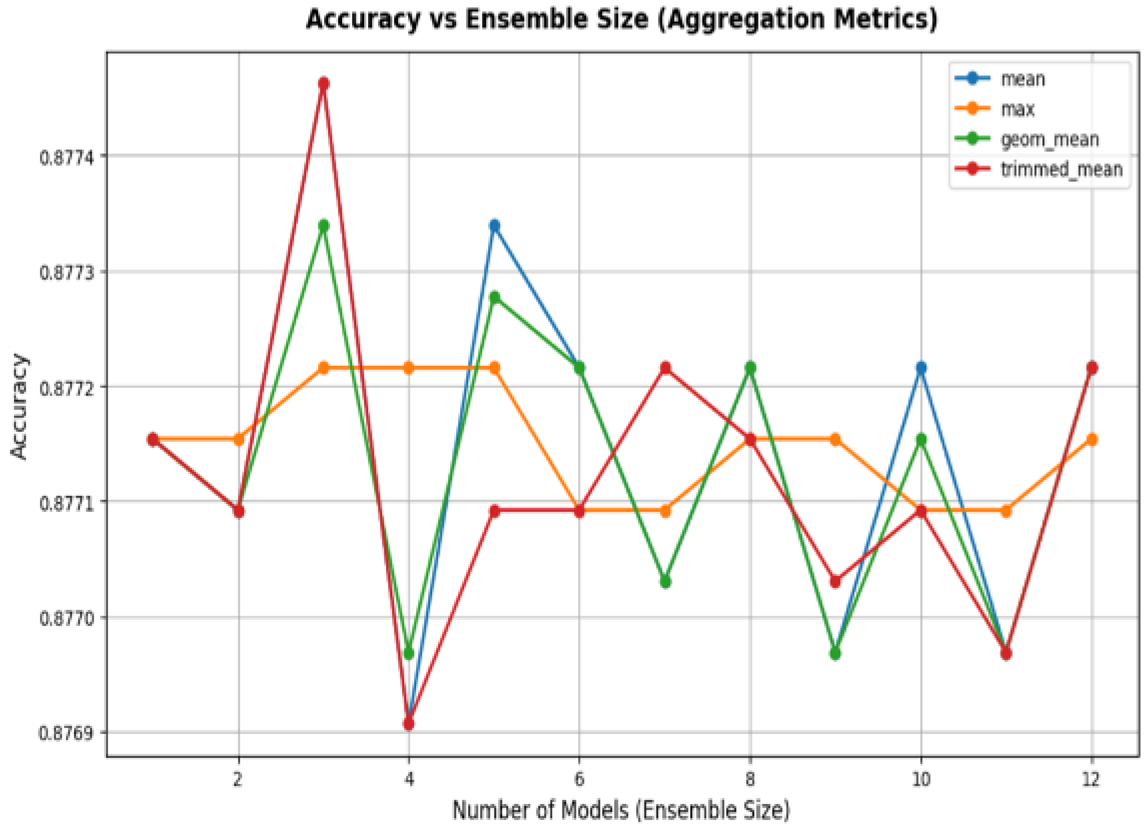

Cost-Sensitive & Tail-Risk Analysis

EcoSortAI was evaluated using a cost-sensitive framework that assigns greater importance to contamination false-accept errors due to their higher real-world impact, such as equipment damage and rejected loads. Expected system cost was modeled as a function of error type, allowing performance to be assessed beyond traditional accuracy metrics.

In addition, tail-risk analysis was used to examine system behavior under worst-case conditions. Results show that even small increases in contamination errors lead to disproportionately large increases in system cost, reinforcing the importance of prioritizing high-impact error reduction over accuracy alone.

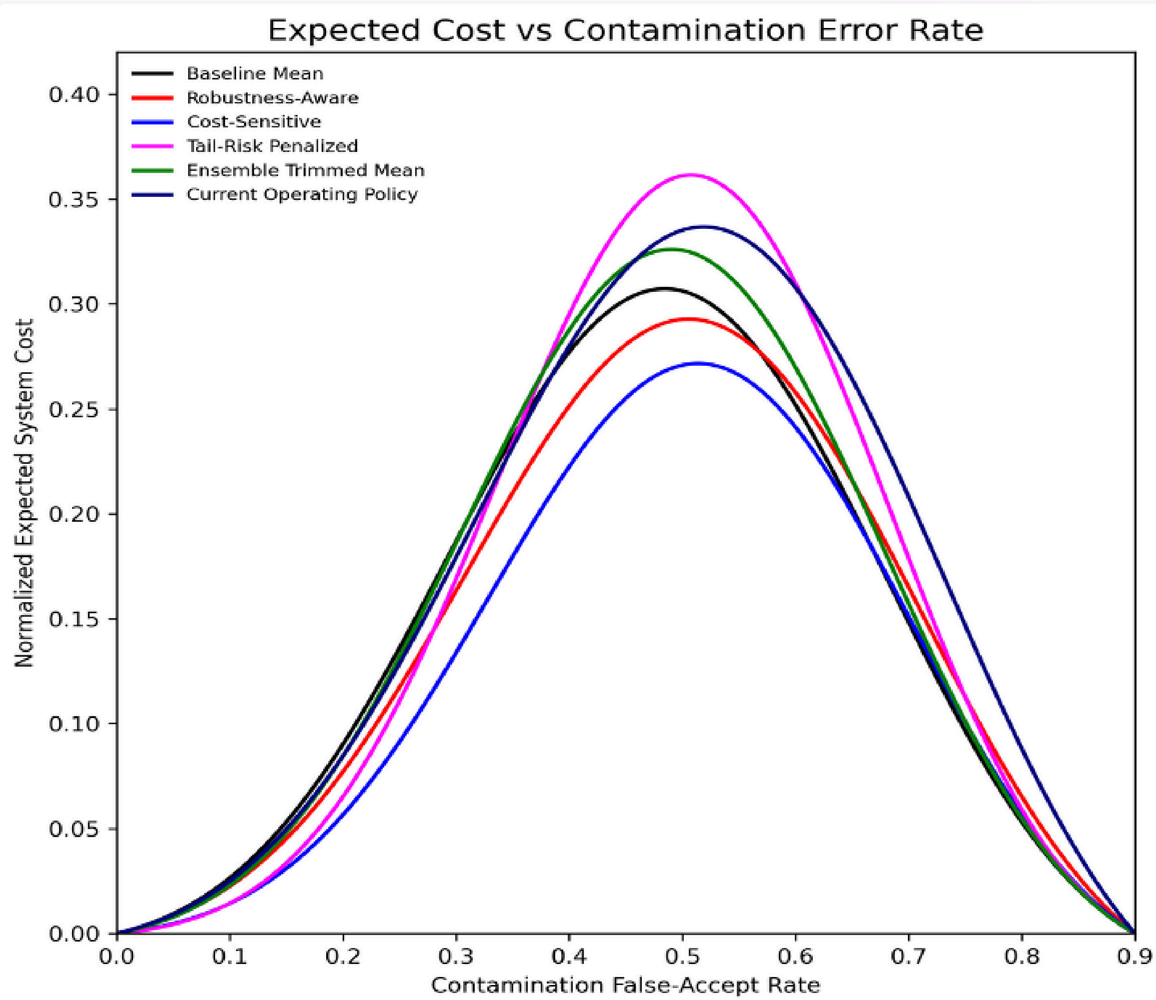

4.1 Error Analysis

Misclassification patterns revealed that errors were not random but concentrated among categories with overlapping visual features. This suggests that the model's limitations align with inherent visual ambiguity rather than fundamental performance issues.

Misclassifications were localized to a small number of highly overlapping sight group category pairs rather than spread across all categories. The two biggest sources of overlapping categories were the paper and card sight group and the plastic related/occlusion/mixed-material sight group and the misclassification was driven by the lack of similarly appearing paper/card material categories. These categories share texture (such as newspaper), similar color and shape, but most certainly displayed more texture and shape distortion from variable light, occlusion and cluttered backgrounds or the object within a paper/card category, which ultimately diminishes the appearance of features that enable visual separation. Performance on aluminum cans and glass bottles, which were visually novel/disparate to the training data categories, showed where EcoSortAI could be effective and where confusability limits existed.

4.2 Confidence Analysis

Prediction confidence distributions showed high certainty for clearly distinguishable items such as aluminum cans and glass bottles. Lower confidence predictions occurred for mixed or partially occluded objects.

4.3 Inference Performance

Inference time measurements (Figure 5) demonstrated that predictions could be generated within fractions of a second, confirming suitability for real-time deployment.

5. STATISTICAL ANALYSIS

Classification performance was evaluated using a multi-metric framework to capture both predictive accuracy and system-level reliability. Model outputs were assessed through confusion matrix analysis, class-wise accuracy, and error-type decomposition to distinguish between recyclable misclassification and contamination-related errors. To analyze the impact of asymmetric error severity, expected system cost was modeled as:

"Expected Cost"=c_FA · P("Contaminant → Recyclable")+c_FR · P("Recyclable → Contaminant") were c_FArepresents contamination-driven false-accept penalties and c_FR represents false-reject penalties. Sensitivity experiments were conducted by varying error rates and cost ratios to evaluate stability of optimal decision behavior. Results demonstrated that system cost exhibits nonlinear amplification with increasing contaminant false-accept rates, indicating that small increases in contamination errors produce disproportionately large operational consequences. Robustness was further assessed under simulated distribution shifts (lighting variation, occlusion, and background complexity). Prediction stability was quantified by measuring relative accuracy degradation and variance across perturbation conditions. Overall, the analysis indicates that contamination-aware evaluation provides more operationally meaningful insight than accuracy alone, and that minimizing high-impact error modes substantially improves projected system reliability.

Within the cost-sensitive setting, contamination false-accept errors are penalized much more heavily than false-reject errors as their cost implications are much greater. If a contaminant is accepted as recyclable, it can degrade bale quality, contaminate otherwise recoverable material, result in rejection of the shipment, or may, in certain cases, cause machinery to jam during sorting. (A contaminant accepted as a recyclable rejects the recyclable thereunder, but the impact concerns the loss of recyclable material, not contaminating the useful stream). As such c_FA was weighted more than c_FR to represent the different potential implications under these circumstances.

5. DISCUSSION

The findings support the hypothesis that CNN-based transfer learning can effectively classify recyclable materials and detect contamination using visual features alone. The high accuracy across diverse categories suggests that the model learned meaningful representations of texture, shape, and material appearance.

5.1 Practical Implications

The deployment of EcoSortAI demonstrates that lightweight CNN models can operate in real time using consumer-grade hardware. This suggests potential integration into mobile applications, smart bins, or educational tools that guide proper recycling behavior. By improving sorting accuracy at the source, AI-assisted systems could reduce contamination before materials reach recycling facilities.

5.2 Comparison to Traditional Systems

Compared to mechanical and sensor-based sorting systems, image-based machine learning offers lower cost, greater scalability, and easier integration into digital platforms. While industrial systems may achieve higher throughput, AI-driven image classification provides accessibility and flexibility for distributed use cases.

6. LIMITATIONS & SOURCES FOR ERROR

Although the results were strong, several limitations and potential sources of error were identified:

1. Dataset Imbalance
   Some waste categories contained fewer images than others, which may have reduced classification accuracy for underrepresented classes.
2. Image Quality and Lighting Conditions
   Variations in lighting, background clutter, shadows, and camera angles affected prediction confidence and accuracy during real-world testing.
3. Visual Similarity Between Categories
   Certain materials, such as paper and cardboard or different plastic types, share similar textures and colors, making them more difficult to distinguish.
4. Single-Label Classification
   The model assigned only one label per image, limiting its ability to handle items that were recyclable but contaminated (e.g., food residue on plastic).

5. Limited Real-World Diversity
   While the dataset included many examples, it may not fully represent the wide variability of

waste items encountered in real recycling environments.

7. CONCLUSION

This project investigated whether machine-learning image classification could be used to predict and sort recyclable materials while identifying contamination in waste streams. The results showed that a convolutional neural network trained using transfer learning was able to accurately classify waste items across 27 categories and distinguish contaminants from recyclables, supporting the original hypothesis. By integrating the trained model into a real-time web application, EcoSortAI demonstrated the practical potential of AI-assisted recycling tools. Such systems could help reduce contamination, improve recycling efficiency, and promote more sustainable waste-management practices. While limitations related to dataset diversity and visual similarity were observed, the overall findings indicate that machine learning offers a promising and scalable solution for addressing recycling contamination. Future improvements could further enhance accuracy and expand real-world applicability, making AI-driven waste sorting a valuable tool in environmental sustainability efforts.

8. REFERENCES

Beigl, M., Krohn, A., & Decker, C. (2020). Smart waste management systems. Waste Management, 102, 171–186.

Gundupalli, S. P., Hait, S., & Thakur, A. (2017). Automated sorting of municipal solid waste. Waste Management, 60, 56–74.

Howard, A. G., et al. (2017). MobileNets: Efficient convolutional neural networks. arXiv:1704.04861.

Krizhevsky, A., Sutskever, I., & Hinton, G. (2012). ImageNet classification with deep CNNs. NIPS.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521, 436–444.

United States Environmental Protection Agency. (2023). Facts and figures about materials, waste, and recycling.

Yang, M., Thung, G., Zhu, A., & Wu, L. (2019). Trash classification using CNNs. ICCV Workshops.

9. FUTURE WORK

Future improvements include expanding the dataset to cover more lighting conditions, camera angles, and regional waste variations. Implementing multi-label classification could enable identification of items that are both recyclable and contaminated. Integration with object detection models could allow classification of multiple items within a single image. Additionally, collaboration with recycling facilities could provide operational validation and support large-scale deployment.

10. ACKNOWLEDGMENTS

I would like to thank Mrs. Biddle for academic mentorship, Dr. Chloe Luyet for early project guidance, and Hari Raval for technical support in deep learning implementation.