Using Image Classification to Create an Intelligent Recycling Bin

Samadhi Chandrasena[1]

Earlham College, Richmond, Indiana

skchand22@earlham.edu

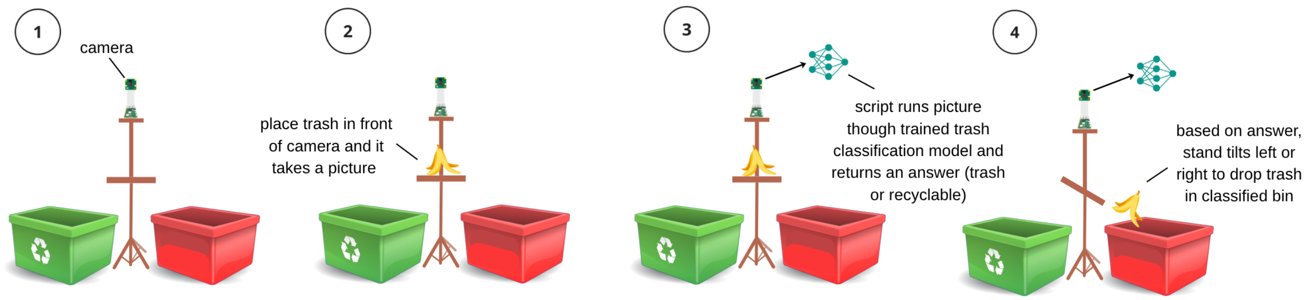

Graphical Abstract

Abstract :

Improper waste sorting leads to contaminated recyclables, increased landfill waste, and higher waste management costs. To address these issues, this project proposes the design and development of a smart trash sorting system that combines image classification with physical actuation to automatically categorise and direct waste into the appropriate bins. A Raspberry Pi Camera captures real-time images of incoming waste, which are then processed by a pre-trained MobileNetV2 convolutional neural network to identify the material type across six categories (cardboard, glass, metal, paper, plastic, and trash). Based on the classification result, a Raspberry Pi 5 controls a servo-actuated tilting platform to sort items into two bins: recyclable and non-recyclable. The system achieved 73.36 percent validation accuracy on the TrashNet dataset and demonstrated real-time inference capabilities with approximately 0.5-1 second response time.

Index Terms: Smart trash sorting system, Machine learning, Waste management automation, Edge computing, Transfer learning

1 Introduction

In many public and household settings, people often dispose of items incorrectly due to confusion, lack of awareness, or insufficient infrastructure, which reduces the effectiveness of recycling efforts. Manual sorting can help, but it is time-consuming, inconsistent, and often impractical in high-traffic environments. This project aims to address these challenges by developing a small-scale, automated waste sorting system that uses image classification and embedded hardware control. The system detects and identifies six common waste material types (cardboard, glass, metal, paper, plastic, and trash) using a Raspberry Pi Camera and a MobileNetV2 convolutional neural network trained via transfer learning. Based on the classification results, the system activates a servo-controlled tilting platform managed by a Raspberry Pi 5 to direct waste items into the appropriate bin—either recyclable or non-recyclable. The proposed system could be adapted for use in smart homes, schools, or public waste stations, particularly in areas where environmental education and recycling habits are still developing.

1.1 Contribution

This paper has the follow contributions:

- A trained and tested MobileNetV2 image classifier optimised for small datasets using transfer learning, achieving 73.36 percent validation accuracy on six waste material categories.
- A fully integrated hardware-software prototype deployed on a Raspberry Pi 5 that classifies and physically sorts waste items into two separate bins in real-time.

2 Related Work

2.1 Convolutional Neural Networks

Convolutional Neural Networks (CNNs) have proven highly effective in automated waste classification systems. Several studies demonstrate that CNN-based models like YOLOv8, VGG16, and EfficientNet consistently achieve high accuracy in classifying different waste materials. Akbar, 2023 tested multiple CNN models, reporting that YOLOv8 achieved the highest precision at 96.5 percent on the TrashNet dataset. YOLOv8's strong performance in object detection and its efficiency in real-time applications make it especially suitable for waste sorting systems where speed and resource constraints are critical. Similarly, VGG16 and EfficientNet have shown strong classification capabilities in related contexts. These results highlight the importance of choosing a deep learning model that balances accuracy and processing efficiency for real-world deployments.

2.2 Lightweight Models for Low-Resource Deployment

Many real-world waste classification systems are designed to run on embedded devices like Raspberry Pi, which have limited memory and processing power. To make this possible, researchers use lightweight CNN models that require less computation but still deliver robust results. For example, Deeluea, 2022 used EfficientNet-Lite0 - a compact model designed for mobile or edge devices - and achieved 93 percent accuracy on a dataset of over 2,800 images. Similarly, Mendy et al., 2020 demonstrated that even small models can be effective for classifying waste when deployed in smart bins. MobileNetV2, developed by Google, is another widely-adopted lightweight architecture optimised for mobile and embedded deployment through techniques like depthwise separable convolutions that reduce computational complexity while maintaining classification performance. These lightweight models are ideal for projects requiring on-device inference without relying on cloud computing or external servers.

2.3 Transfer Learning

Transfer learning has emerged as a powerful technique for training effective models with limited domain-specific data. Rather than training a CNN from scratch, transfer learning leverages pre-trained models that have already learned generalisable features from large datasets like ImageNet. Studies in waste classification frequently employ this approach by using pre-trained architectures such as MobileNetV2, ResNet, or EfficientNet as feature extractors, then fine-tuning only the final classification layers on waste-specific datasets. This significantly reduces training time and data requirements while often improving accuracy. Rad et al., 2017 demonstrated that transfer learning combined with data augmentation yielded substantial performance gains when working with small training sets. For resource-constrained applications, freezing the pre-trained layers and training only custom classification heads provides an optimal balance between model performance and training efficiency.

2.4 Data Augmentation

A recurring challenge in waste classification is the limited availability of large, high-quality datasets. Data augmentation has emerged as a key technique to address this problem by artificially expanding training datasets through transformations that preserve label integrity. Studies such as Akbar, 2023 and Deeluea, 2022 applied augmentation strategies like random rotation (typically ±20°), horizontal flipping, and width/height shifts to prevent overfitting and improve model generalization. Rad et al., 2017 similarly reported a 12 percent increase in accuracy using augmentation techniques when training on a small dataset. These enhancements allow models to handle variations in lighting, positioning, and object orientation - all of which are common in real-world waste sorting scenarios. Incorporating data augmentation is therefore a practical and essential step for improving robustness when training on limited datasets, particularly when deploying models to edge devices where inference must handle diverse real-world conditions.

2.5 Datasets

The TrashNet dataset is frequently used in studies involving AI-based waste classification. It contains 2,527 images of waste items across six categories: cardboard, glass, metal, paper, plastic, and trash. While useful for benchmarking and educational purposes, researchers often note its limitations, particularly its relatively small size and limited environmental diversity (controlled lighting, clean backgrounds, centered objects). This makes data augmentation or custom dataset creation a common necessity in research projects. Some researchers supplement TrashNet with additional images or curate their own datasets to increase model robustness and better represent real-world deployment conditions. However, TrashNet's limited diversity has led researchers to explore more complex datasets like TACO (Trash Annotations in Context), which includes annotated images of litter found in real-world environments with cluttered backgrounds and varying lighting conditions. Other datasets such as WasteNet and EcoWasteNet offer larger and more heterogeneous collections, allowing for the training of more generalised models. These datasets enable comparative analysis of different models' performance and help advance the development of more resilient waste-sorting systems capable of handling the visual complexity of uncontrolled environments. For this project, the complete TrashNet dataset with all six material categories is used for model training. The six-class predictions are subsequently mapped to two physical bins (recyclable and non-recyclable) at the hardware actuation stage, allowing the system to perform practical binary sorting while leveraging the model's fine-grained classification capabilities.

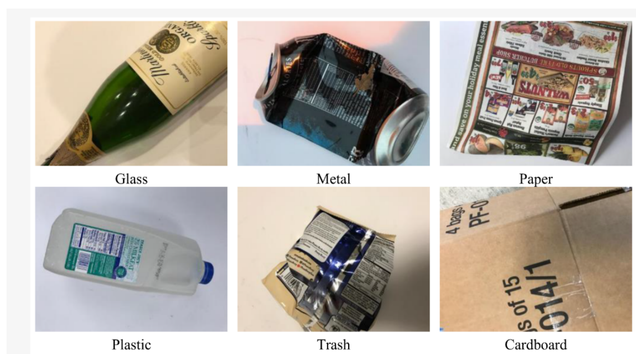

Figure 1. Sample images with labels from the TrashNet Dataset

2.6 Hardware Integration

In addition to software advances, several projects have explored hardware implementations for automated waste sorting. Many of these systems use actuators or servo motors connected to embedded computers such as Raspberry Pi to physically redirect waste into appropriate bins. While some systems rely on basic metrics like weight or sensor-based triggers, fewer integrate visual classification with mechanical actuation. This combination presents challenges, particularly due to the limited processing power of embedded devices, which can make real-time image classification difficult to perform. Some systems adopt a hybrid approach - using a PC or cloud service to handle classification while a microcontroller performs the physical sorting. However, advances in model optimisation techniques like TensorFlow Lite have made deployment increasingly feasible. Some notable implementations of hardware-based waste sorting systems demonstrate the practicality of edge deployment. For example, projects like those by Mendy et al., 2020 combined lightweight CNN models with servo motors to sort items into recyclable and non-recyclable categories. Other designs, such as those outlined in Deeluea, 2022, prioritised low-cost and portable systems by deploying image classification on Raspberry Pi units with simple mechanical actuation mechanisms such as rotating arms or tilting platforms. These designs emphasise accessibility and affordability, making them suitable for educational settings or small-scale deployment. However, limitations such as slow sorting speeds, inference latency on embedded devices, and sensitivity to lighting and camera positioning remain common challenges across many hardware-oriented solutions.

2.7 *Approach for this Project*

Building on these prior studies, this project implements MobileNetV2, a lightweight convolutional neural network architecture optimised for mobile and embedded deployment. The model is trained using transfer learning on the TrashNet dataset to classify images across six waste material categories: cardboard, glass, metal, paper, plastic, and trash. Unlike high-complexity models like YOLOv8, MobileNetV2 prioritiset ales deployability on resource-constrained devices like the Raspberry Pi 5, balancing classification accuracy with low computational requirements. The base MobileNetV2 layers, pre-trained on ImageNet, are frozen to serve as feature extractors, while custom classification layers are trained on the TrashNet dataset. Data augmentation techniques including random rotation (±20°), horizontal flipping, and width/height shifts are applied to address the limited size and environmental diversity of the dataset. The trained model is converted to TensorFlow Lite format and deployed on the Raspberry Pi 5, which handles image capture via the Pi Camera Module, preprocessing, and real-time inference. An IR obstacle sensor detects the presence of waste items, triggering the classification pipeline. Based on the six-class prediction output, the system maps the result to a binary category (recyclable or non-recyclable) and controls a servo motor driving a tilting platform to physically direct waste items into the appropriate collection bin. This approach demonstrates a fully integrated, end-to-end autonomous waste sorting system operating entirely on embedded hardware without reliance on cloud services or external computing resources.

3 Methodology

The proposed system is an end-to-end, autonomous waste classification and sorting unit powered entirely by a Raspberry Pi 5. It combines real-time image classification using a trained convolutional neural network (CNN) model with physical actuation via a servo motor to sort objects into predefined bins based on material type. This system is built to demonstrate the practical deployment of machine learning in resource-constrained environments, with the entire pipeline—from image capture to mechanical sorting—handled on-device.

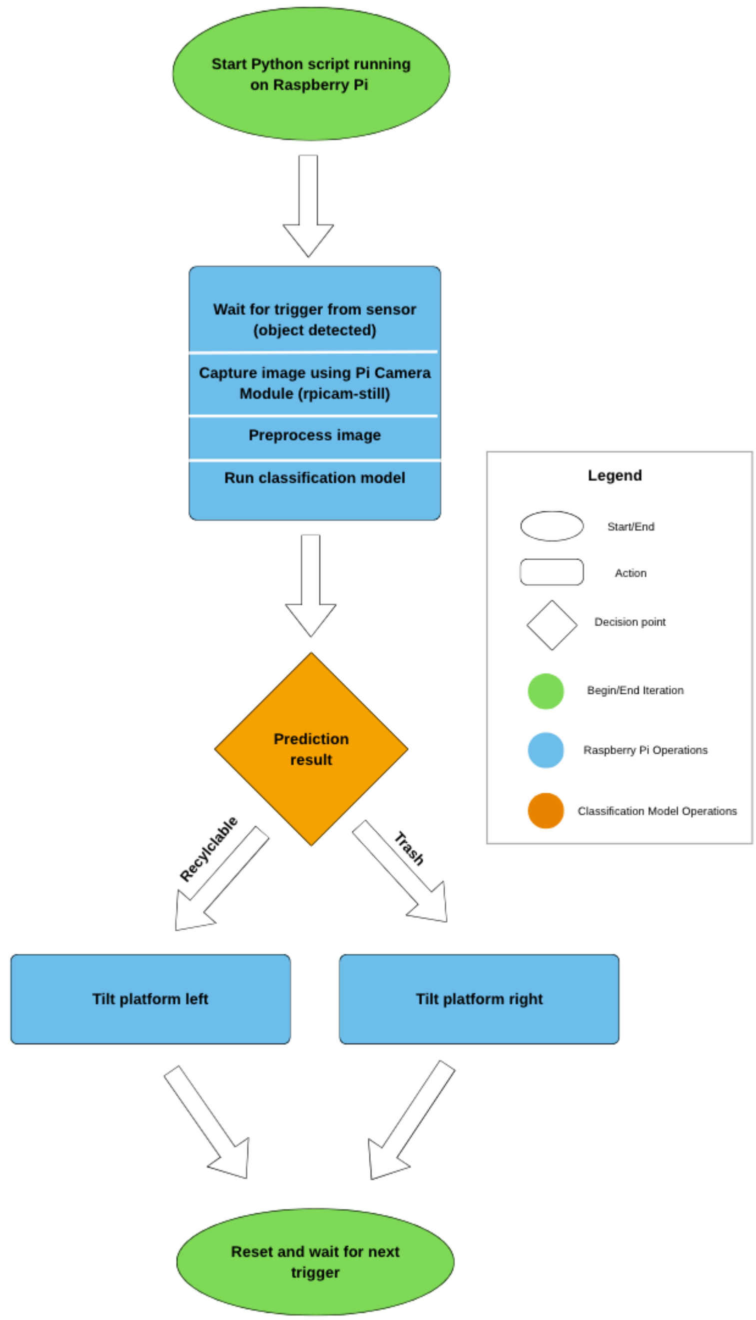

Figure 2. Data Architecture Diagram

3.1 System Initialisation and Monitoring

The process begins by initializing a Python script on the Raspberry Pi, which serves as the central control unit for both the software and hardware subsystems. On startup, the script configures the GPIO pins and sets up the necessary peripherals, including the IR obstacle sensor (GPIO27) and the Pi Camera Module. The IR sensor continuously monitors for object presence, and when an object is placed on the sorting platform, the sensor's output signal transitions from HIGH to LOW, triggering the classification and sorting routine.

3.2 Image Capture and Preprocessing

Upon detecting an object, the Pi captures an image using the rpicam-still command via the Pi Camera Module. The captured image is stored as a NumPy array representing RGB pixel values. The image then undergoes a preprocessing pipeline tailored to the model's input specifications, including resizing to 224×224 pixels (MobileNetV2's required input size) and normalizing pixel values to a range of 0-1 for consistency with the training set. This preprocessing ensures the model receives input in the same format and quality as it was trained on, thereby improving classification reliability..

3.3 Model Loading and Inference

Once preprocessed, the image is passed to a trained CNN model deployed in TensorFlow Lite format (.tflite) for optimized performance on the Raspberry Pi. The MobileNetV2-based model performs inference and outputs a class prediction across six categories: cardboard, glass, metal, paper, plastic, and trash. The output class index is then mapped to a binary classification (Recyclable or Trash) where cardboard, glass, metal, paper, and plastic are categorized as recyclable materials, while trash remains as non-recyclable. This binary classification is passed to the sorting mechanism for physical actuation.

3.4 Hardware Control for Sorting

The physical sorting mechanism comprises a single servo-controlled tilting platform. A positional servo motor (GPIO17) is mounted vertically, with the platform attached to the servo arm. Depending on the predicted bin type:

Recyclable → Platform tilts to maximum left position Trash → Platform tilts to maximum right position

The platform maintains the tilted position for 2.5 seconds, allowing the item to slide off into the appropriate collection bin positioned below. After sorting, the servo returns the platform to the center position, and the system resumes monitoring the IR sensor for the next item.

3.5 System Reset and Readiness

After the item is sorted, the system returns the servo to its default center position, clears the camera buffer, and resumes monitoring the IR sensor. This closed-loop cycle continues indefinitely, allowing for continuous autonomous operation without manual intervention.

3.6 Resource Optimisation

The Raspberry Pi 5 successfully handles the lightweight MobileNetV2 model deployed in TensorFlow Lite format, achieving inference times of approximately 0.5-1 second per image. The use of TensorFlow Lite optimization techniques, including model quantization, ensures efficient edge deployment without requiring external computational resources. The system operates entirely on-device, demonstrating the feasibility of real-time machine learning inference on resource-constrained embedded systems.

4 Preliminary Evaluation

To evaluate the system's performance, several experiments were conducted across three key metrics:

Model Accuracy: Classification performance was measured on the TrashNet validation set comprising 503 images (20 percent of the total dataset). The MobileNetV2 model achieved a validation accuracy of 72.56 percent, correctly classifying approximately 365 out of 503 unseen images across six material categories.

Inference Latency: The total time from image capture through preprocessing to classification output was measured on the Raspberry Pi 5 hardware. The system achieved an average inference latency of approximately 0.5-1 second per item, demonstrating real-time classification capability suitable for autonomous deployment.

System Integration: The complete end-to-end pipeline - from IR sensor object detection through image classification to servo-controlled platform actuation - was tested to verify functional integration of all hardware and software components. The system successfully detected waste items, classified them into material categories, mapped predictions to binary sorting decisions (recyclable or non-recyclable), and actuated the tilting platform to physically direct items into the appropriate collection bin.

The 72.56 percent validation accuracy represents reasonable performance given several constraining factors: the relatively small dataset size (2,527 images total), visual similarity between certain material categories (particularly glass and metal), limited environmental diversity in the TrashNet dataset (controlled lighting and backgrounds), and the use of a lightweight model architecture optimized for resource-constrained edge deployment rather than maximum accuracy. The inference latency of 0.5-1 second significantly exceeded the target threshold of under 2 seconds, demonstrating the effectiveness of TensorFlow Lite optimization and the MobileNetV2 architecture for real-time inference on embedded systems.

5 Key Results

The training process generated two key performance metrics visualized in Figure 3 and Figure 4: model accuracy and model loss across training epochs, respectively. The accuracy graph demonstrates steady improvement in both training and validation performance throughout the training duration. Training accuracy increased from an initial 64 percent at epoch zero to 85 percent by epoch four, indicating successful feature learning by the custom classification layers. Validation accuracy exhibited a similar upward trend, rising from 68 percent to a peak of 73 percent at epoch two, before stabilizing around 72-73 percent for subsequent epochs. This convergence pattern triggered the early stopping mechanism at epoch five, as the validation accuracy showed no further improvement within the patience window of three epochs. The final validation accuracy of 72.56 percent represents reasonable performance for a six-class classification problem with a relatively small dataset, particularly considering the inherent visual similarity between certain material categories such as glass and metal or different types of paper products.

The loss graph provides complementary insights into model optimization. Training loss decreased consistently from approximately 1.0 to 0.45 over the five epochs, demonstrating effective gradient descent and weight optimization. Validation loss similarly declined from 0.8 to approximately 0.74, where it plateaued after epoch two. The gap between training and validation metrics indicates mild overfitting - the model performs better on seen training data than unseen validation data - however, this gap remains acceptable and does not suggest severe memorization. The plateau in validation loss after epoch two, combined with the stagnant validation accuracy, justified the early termination of training. Continuing beyond epoch five would likely have resulted in increased overfitting without improving generalization performance, as evidenced by the diverging training and validation curves. The implementation of early stopping therefore successfully prevented unnecessary computational expense while preserving optimal model weights. These results validate the transfer learning approach, demonstrating that pre-trained MobileNetV2 features combined with targeted fine-tuning of classification layers can achieve functional performance even with limited domain-specific training data.

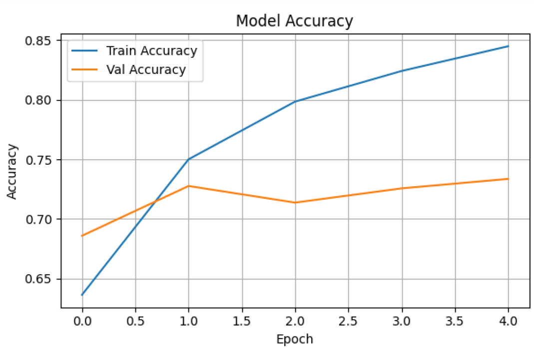

Figure 3. Model Accuracy

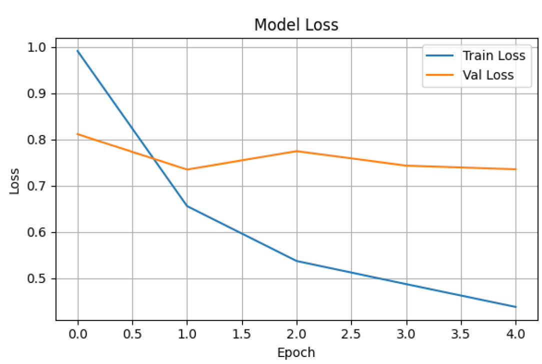

Figure 4. Model Loss

6 Discussion of Challenges

While the availability of existing training scripts for the TrashNet dataset facilitated the development of the model, and the Raspberry Pi components were well supported by documentation, the project encountered significant challenges in both model performance and physical implementation.

Dataset Limitations and Classification Accuracy: A primary challenge was achieving reliable classification accuracy suitable for autonomous sorting. The TrashNet dataset, while widely used for benchmarking, presents several limitations that directly impacted real-world performance. With only 2,527 images total across six categories, the dataset size is relatively small for training robust convolutional neural networks. More critically, the dataset lacks the environmental diversity present in real-world deployment scenarios - TrashNet images feature controlled lighting, clean backgrounds, and centered objects, whereas the sorting system operates under variable lighting conditions with diverse object orientations and backgrounds. This mismatch between training conditions and deployment environment contributed to the 72.56 percent validation accuracy, which proved challenging for reliable autonomous sorting.

Additionally, the decision to map six material categories to two physical bins (recyclable and non-recyclable) introduced an implicit assumption that all six-class predictions would translate cleanly to binary sorting decisions. In practice, the model exhibited strong bias toward recyclable categories (plastic, paper, cardboard, glass, metal), rarely predicting the "trash" class. This imbalance likely stems from the underrepresentation of the trash category in the TrashNet dataset and the visual diversity within that class. As a result, the physical sorting system predominantly directed items to the recyclable bin, limiting the demonstration of true binary classification capability. Creating a custom dataset using the actual Pi Camera in the deployment environment with balanced representation of both recyclable and non-recyclable items would likely have yielded significantly improved real-world performance.

Physical Integration Challenges: Beyond software limitations, the transition from model development to hardware integration introduced substantial mechanical and construction difficulties. Constructing a stable physical sorting system - particularly securing the camera at the correct angle and height, mounting the servo motor to enable consistent tilting motion, and building a platform structure rigid enough to support waste items while remaining responsive to servo actuation—proved significantly more challenging than anticipated. The camera positioning required multiple iterations to achieve the correct top-down viewing angle that matched the training data perspective while maintaining sufficient field of view. The servo-controlled tilting platform presented additional complications: ensuring the platform tilted at sufficient angles to reliably direct items into collection bins, preventing items from sliding prematurely, and achieving consistent return-to-center positioning all required careful calibration and structural reinforcement. Cardboard construction, while accessible and low-cost, lacked the rigidity needed for reliable repeated operation, necessitating additional support structures and adhesive reinforcement. These physical construction challenges represented the most time-consuming aspect of the project, requiring multiple design iterations and substantially more effort than the software development and model training phases.

System Integration and Timing: Integrating all components into a cohesive autonomous system also revealed timing and coordination challenges. Ensuring the IR sensor reliably detected objects without false triggers, synchronising image capture with object placement, allowing sufficient time for the servo to complete its motion before accepting the next item, and managing the inference latency within the overall sorting cycle all required careful tuning of delays and thresholds in the control script. These integration issues highlighted the complexity of deploying machine learning models in physical robotic systems, where software timing must coordinate with mechanical response times and sensor behavior.

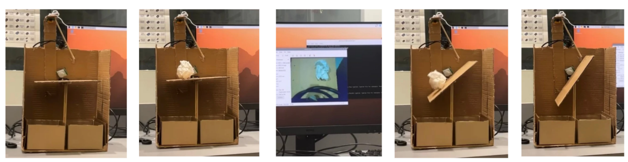

(a) (b) (c) (d) (e)

Figure 5. Physical system demonstration: (a) IR sensor waits for object, (b) User places waste item underneath camera, (c) Pi Camera captures image and model classifies, (d) Platform tilts based on prediction, (e) Will return to centre and wait for next item

7 Future Work

7.1 Dataset and Model Improvements

The most critical priority for future work is addressing the fundamental limitations in classification accuracy through comprehensive dataset improvements. Creating a custom dataset captured with the actual Pi Camera Module under the deployment environment's lighting conditions, camera angles, and background settings would dramatically improve model performance and generalisation. This custom dataset should include substantially more images—ideally 10,000+ samples rather than the current 2,527 - with balanced representation across all material categories, particularly addressing the severe underrepresentation of the trash class that caused the current system's bias toward recyclable predictions. Collecting diverse examples of each material type (different brands, colors, conditions, and orientations) and capturing images under various lighting conditions (natural daylight, indoor lighting, shadows) would enable the model to handle real-world variability more robustly. Beyond dataset expansion, model architecture improvements could yield significant accuracy gains. Fine-tuning unfrozen layers of MobileNetV2 or experimenting with larger architectures like EfficientNet may improve classification performance. Implementing class weighting or oversampling techniques to address the trash class imbalance would help ensure more reliable binary sorting. Establishing a target accuracy of 85-90 percent on a properly representative validation set should be achieved before attempting physical deployment, as the current 72.56 percent accuracy proved insufficient for reliable autonomous operation. Only after achieving robust classification performance in controlled testing should hardware integration proceed, as premature physical construction without adequate model performance wastes significant development time on a system that cannot function reliably.

7.2 User Experience and Deployment Design

A key limitation of the current system is the requirement for users to position waste items deliberately in front of the camera to ensure clear and usable images. This interaction is not representative of typical waste-disposal behavior, where items are usually dropped or thrown into a bin without precise placement. Future work will focus on developing a more natural and user-independent workflow. One direction is to redesign the system so that image capture occurs inside the bin rather than at the entrance. In this model, users could dispose of items normally, and an internal camera - potentially paired with controlled lighting to ensure consistent image quality - would capture images as the object enters the bin. This approach would allow the classification and sorting mechanism to operate autonomously without requiring user cooperation or awareness of the camera system. Additionally, integrating a more robust internal sorting mechanism, such as a motorised rotating platform with defined bin positions or a multi-stage chute system with actuated gates, could enable more reliable sorting than the current tilting platform approach. Using more durable construction materials such as acrylic, wood, or 3D-printed components rather than cardboard would improve mechanical stability and longevity. Implementing multiple servos for independent control of platform rotation and trapdoor actuation (as originally envisioned in the design) would provide more precise sorting control and reduce the risk of items failing to fall into the correct bin.

7.3 System Robustness and Real-World Deployment

For practical deployment, several additional enhancements would be necessary. Implementing feedback mechanisms such as weight sensors or secondary cameras to verify that items successfully reached the intended bin would enable the system to detect and correct sorting failures. Adding a manual override or "uncertain" category for items the model classifies with low confidence would prevent misclassification of ambiguous items. Developing a user interface - such as an LED indicator system or small display - to communicate system status, classification results, and error conditions would improve user trust and adoption. Finally, conducting extensive real-world testing in target environments (households, schools, public spaces) with diverse users and waste items would reveal usability issues and edge cases not apparent in controlled laboratory testing, enabling iterative refinement toward a truly deployable autonomous waste-sorting system.

Acknowledgements

I would like to thank Dr. Charlie Peck, Dr. Yunting Ying, Dr. Debanjali Banerjee, and Porter Libby for their assistance in this project.

References

Akbar, Mohammad Zubair (2023). *Intelligent Solid Waste Classification System Using Deep Learning*. URL: https://hammer.purdue.edu/articles/thesis/INTELLIGENT_SOLID_WASTE_CLASSIFICATION_SYSTEM_USING_DEEP_LEARNING/23786484?file=41773260.

Chen, H. (2021). "Optimization of an Intelligent Sorting and Recycling System for Solid Waste Based on Image Recognition Technology". In: *Hindawi*.

Cheng, Tianhao et al. (2023). "Smart waste management system based on image recognition and machine learning: A case study of the recycling process in Japan". In: *Sustainability* 16.23, p. 10155.

Deeluea, Saran (2022). "Automatic Waste Sorting Trash Bin by Image Classification". In: *Medium*. URL: https://medium.com/@saran.deeluea/automatic-waste-sorting-trash-bin-by-image-classification-79f22d96498e.

Feng, Zhicheng et al. (2022). "An Intelligent Waste-Sorting and Recycling Device Based on Improved EfficientNet". In: *International Journal of Environmental Research and Public Health* 19.23, p. 15987. URL: https://www.mdpi.com/1660-4601/19/23/15987.

Hani Abu-Qdais, Nawras Shatnawi and Esra'a AL-Alami (2023). "Intelligent System for Solid Waste Classification Using Combination of Image Processing and Machine Learning Models". In: *Journal of Environmental Science and Technology*.

Kang, Xiaoqing, Ying Li, and Liang Zhang (2021). "Domain adaptation techniques for waste classification tasks". In: *2021 International Conference on Machine Learning and Cybernetics (ICMLC)*. IEEE, pp. 98–103.

Mendy, Gabriel et al. (2020). "Smart waste bin: Smart waste management system based on embedded systems, image classification, and IoT". In: *International Journal of Advanced Computer Science and Applications* 11.7.

Olawade, David B. et al. (2023). "Smart Waste Management: A Paradigm Shift Enabled by Artificial Intelligence". In: *Sustainability* 16.23, p. 10155.

Rad, Mahsa Rezanejad et al. (2017). "A deep learning approach for waste classification". In: *2017 IEEE Global Humanitarian Technology Conference (GHTC)*. IEEE, pp. 1–5.

Ruíz, V. et al (2019). "Automatic Image-Based Waste Classification". In: *Procedia Computer Science* 153, pp. 422–431.

Zhu, Junya, Rui Zhang, and Jiajun Jiang (2019). "Classification of Garbage Waste Images Using Convolutional Neural Networks". In: *International Journal of Environmental Research and Public Health* 16.17, p. 3148.