Image Classification-Based Intelligent Recycling Bin

Samadhi K. Chandrasena

Senior Project in Computer Science | Earlham College

Introduction

Improper waste sorting is a widespread problem that contaminates recyclable materials, increases landfill waste, and raises processing costs. In many settings, people dispose of items incorrectly due to confusion, lack of awareness, or insufficient labelling on bins. Manual sorting can address this issue but is labour-intensive, inconsistent, and impractical in high-traffic environments such as schools, offices, and public spaces. This project aims to automate the waste sorting process by developing a smart bin that can identify different types of waste materials by taking an image, classifying them, and physically directing them into the appropriate collection compartment without requiring human intervention.

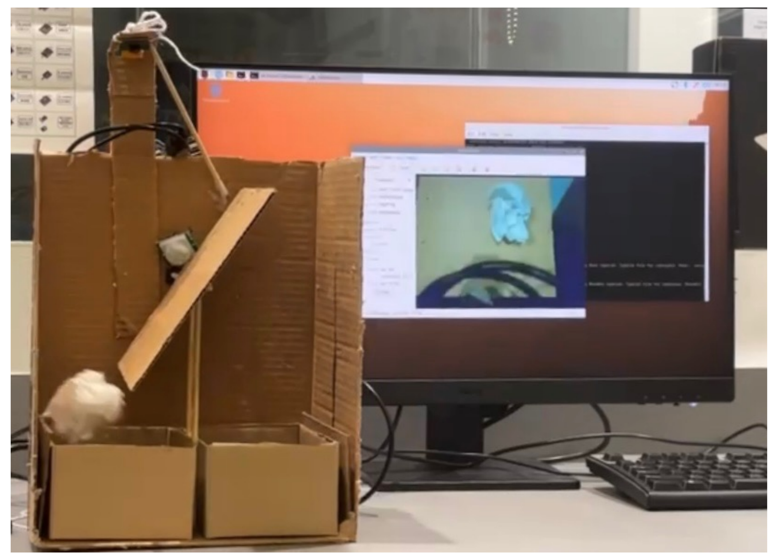

Figure 1: Final Product

Graphical Abstract

The process can be visually depicted in the following four steps:

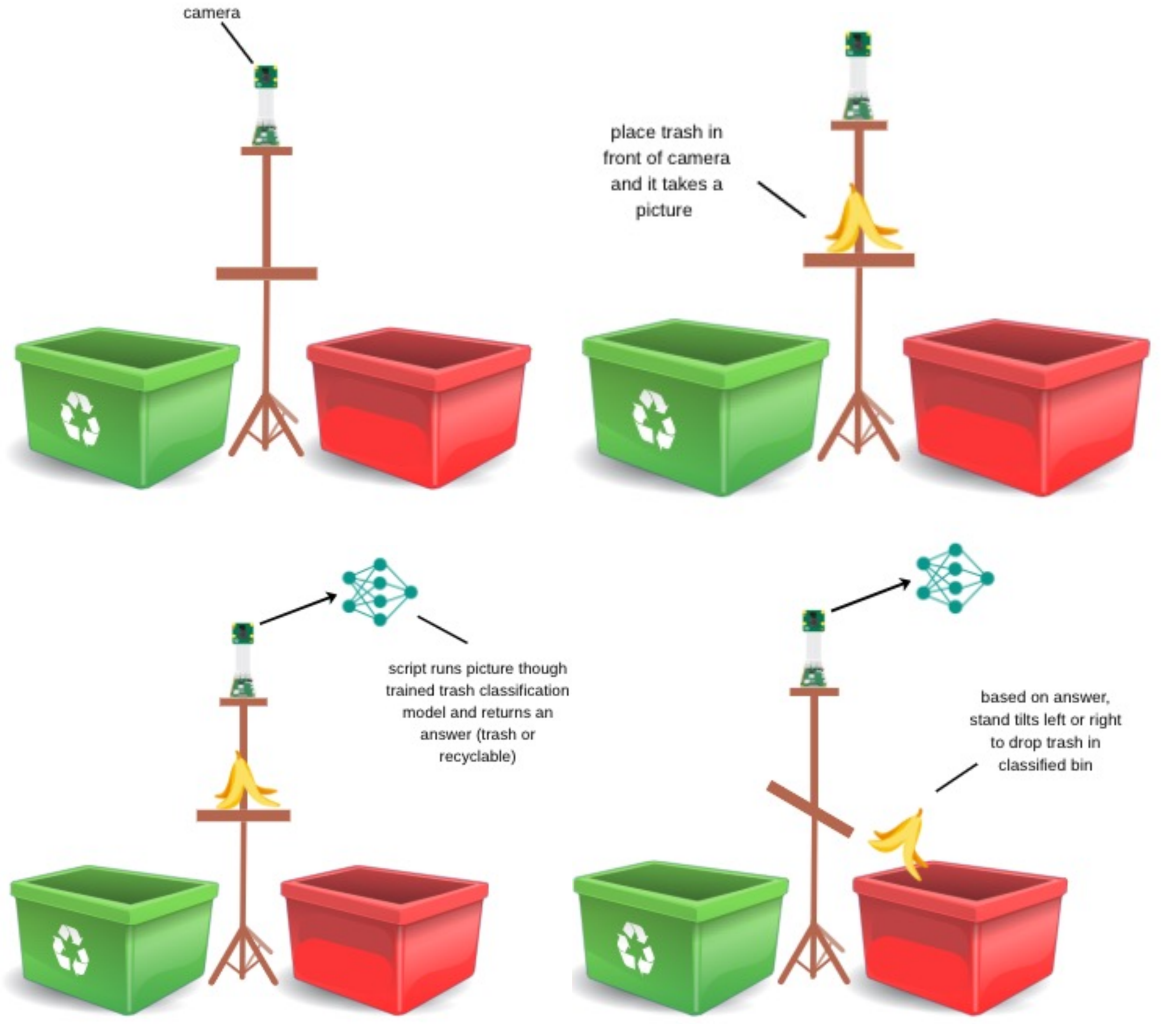

Figure 2: Graphical Abstract

Background

Convolutional Neural Networks (CNNs)

- AI systems designed to recognise visual patterns, inspired by the human visual system.
- Widely used in waste classification, with advanced models achieving over 90% accuracy.
- High-performance CNNs often require powerful hardware, making them unsuitable for low-cost devices (like the Raspberry Pi I used).

Choosing MobileNetV2

- A lightweight CNN optimized for mobile and embedded systems.
- Uses efficient computational techniques to maintain good accuracy with low processing requirements.
- Ideal for deployment on a Raspberry Pi, enabling affordable, practical waste-sorting solutions.

Transfer Learning Approach

- Avoids training from scratch by starting with a model already trained on millions of general images.
- Only the final layers are retrained to classify specific waste materials.
- Requires far less data and computation, making it suitable for resource-limited projects.

Dataset Usage (TrashNet)

- TrashNet includes 2,527 images across six categories: cardboard, glass, metal, paper, plastic, and trash.
- Limited size and environmental diversity often require augmentation or additional datasets.
- This project trains on all six TrashNet classes, then maps predictions to two physical bins - recyclable and non-recyclable - during hardware actuation to combine detailed classification with simple sorting.

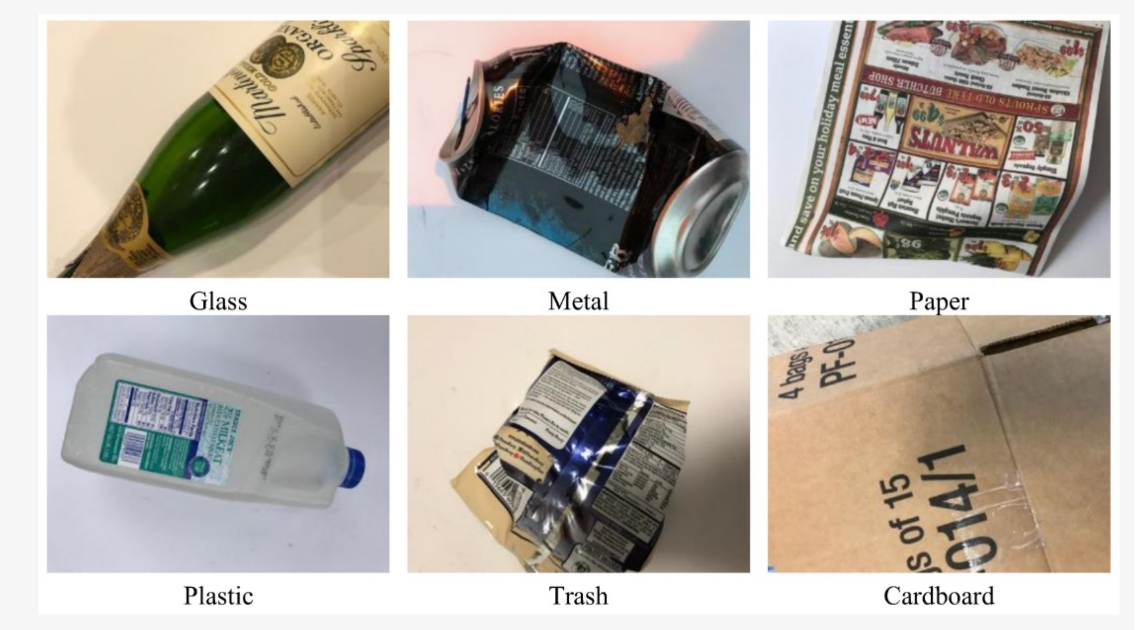

Figure 3: Sample Images from the TrashNet Dataset

Methodology

The system's data architecture diagram (see Figure 4) follows a straightforward pipeline. The Python script running has the IR sensor actively waiting to detect a piece of trash. Once detected, the Pi camera takes a picture and runs it through the trained model. Once it determines its decision (Trash or Recycling), the servo which is attached to a piece of cardboard tilts the trash left or right to slide it into the correct bin. It then returns to its original position, awaiting the next piece of trash.

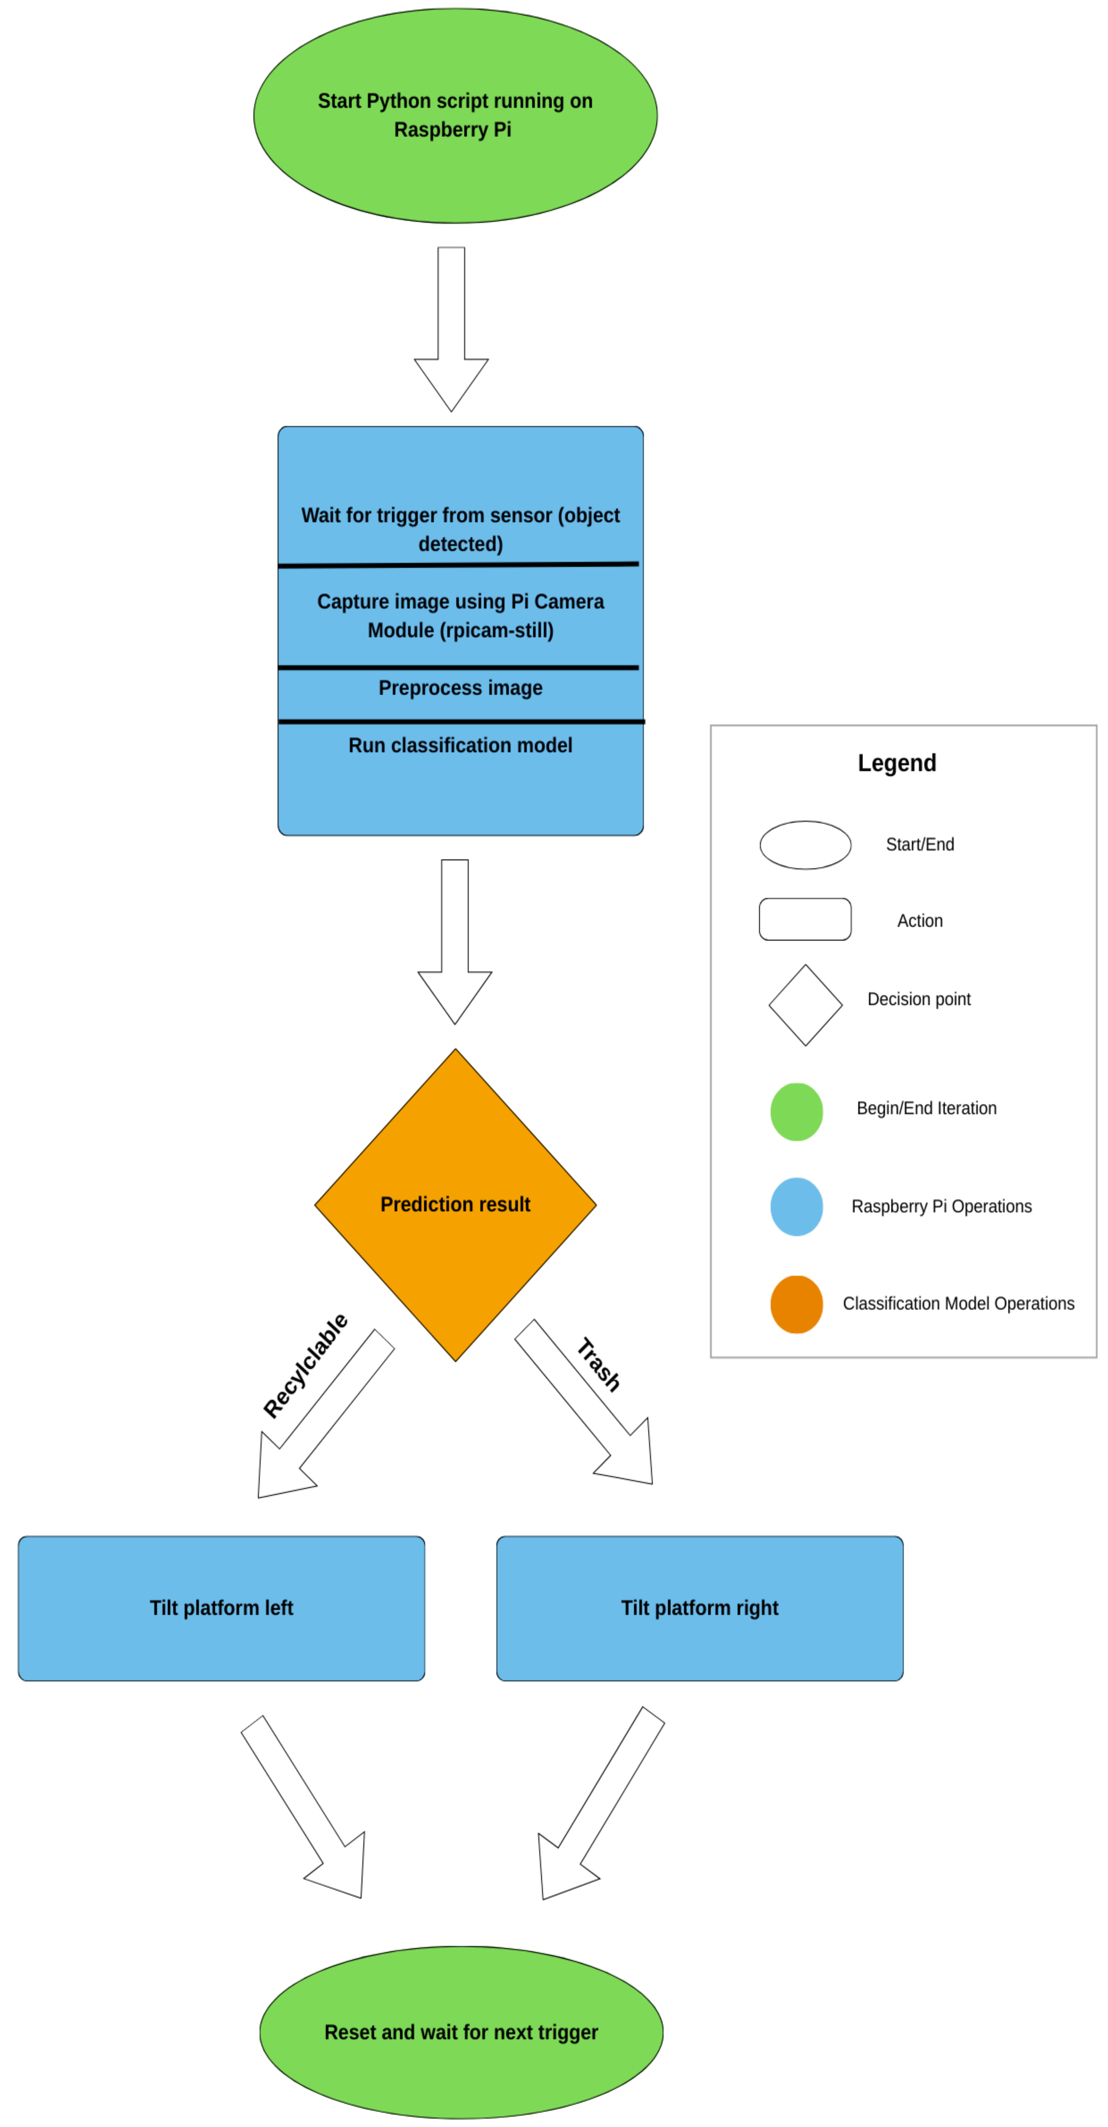

Figure 4: Data Architecture Diagram

Results

As shown in the accuracy graph (Figure 5), both training and validation accuracy improved over the first few epochs, with validation accuracy stabilising at 72.56%, triggering early stopping at epoch five. The loss graph (Figure 5) shows a consistent drop in training and validation loss before the validation curve plateaus, indicating only mild overfitting. In a practical sense, due to limitations that will be discussed below, there was more difficulty in sorting the trash into the "Trash" bin, with more items being classified as "Recycling", even when they were not.

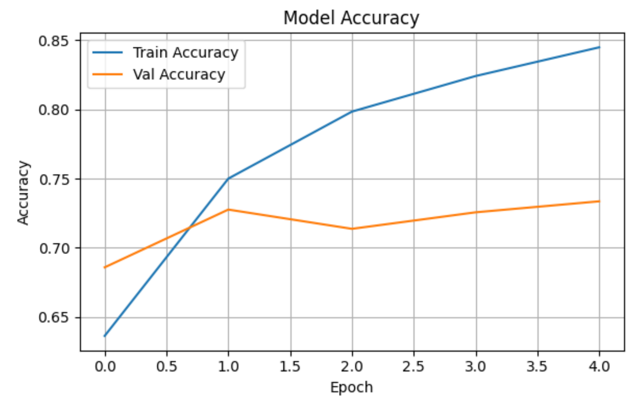

Figure 5: Model Accuracy

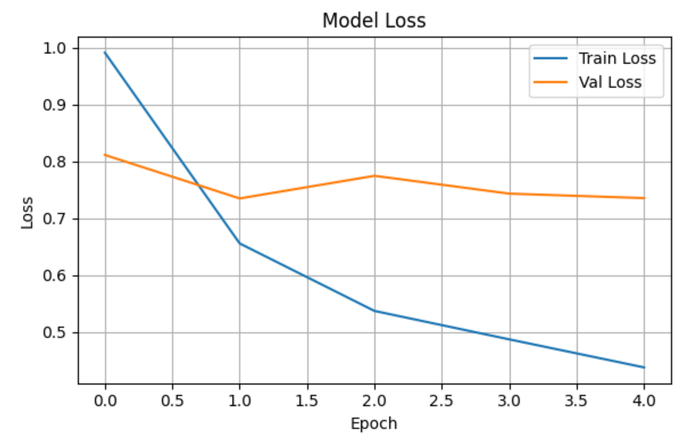

Figure 6: Model Loss

Issues & Future Work

The project faced several challenges, including limited model accuracy due to the small, non-diverse TrashNet dataset and a strong bias toward recyclable predictions, which weakened binary sorting performance. Physical construction added further difficulty, with issues in camera alignment, servo calibration, structural rigidity, and reliable sensor timing. Integrating software and hardware required careful tuning to coordinate image capture, classification, and actuation. Future work should focus on building a large, custom Pi-Camera dataset, improving model architecture and class balance, redesigning the physical system with stronger materials and more reliable mechanisms, and developing a more user-friendly, fully autonomous bin workflow. With a better, more diverse dataset, I am confident the accuracy would be much higher, and the sorting would be more reliable.

Acknowledgements

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