IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

The V: Autonomous Waste Collection Robot

Bassim Abdu Razak PK[1], Irfan Fathan M[2], Mohammed Shamil KP[3], Muhammed Adhil A[4], Neethu PM[5], Thasneem P[6]

1. B.Tech Student (Electronics and Communication Engineering), Eranad Knowledge City Technical Campus, Manjeri, India.
2. B.Tech Student (Electronics and Communication Engineering), Eranad Knowledge City Technical Campus, Manjeri, India.
3. B.Tech Student (Electronics and Communication Engineering), Eranad Knowledge City Technical Campus, Manjeri, India.
4. B.Tech Student (Electronics and Communication Engineering), Eranad Knowledge City Technical Campus, Manjeri, India.
5. H.O.D and Professor (Electronics and Communication Engineering), Eranad Knowledge City Technical Campus, Manjeri, India.
6. Assistant Professor (Electronics and Communication Engineering), Eranad Knowledge City Technical Campus, Manjeri, India.

Abstract: Waste collection in busy indoor spaces like hospitals, offices, and schools has become a real operational problem. Staff spend considerable time manually sorting and collecting mixed waste, and the results are often inconsistent. This paper describes The V, a robot we built to take over that job autonomously. It uses a webcam and OpenCV to identify waste in real time, a three-joint servo arm to physically pick items up, and a Raspberry Pi 4 to run everything. We wanted to see if this kind of system could work on cheap, everyday hardware without needing expensive processors or custom components. Testing showed a classification accuracy of 90.0%, an arm grasping success rate of 92.6%, and an end-to-end segregation accuracy of 83.3% at around 8 to 10 frames per second. We think these numbers show that practical, affordable automation for indoor waste handling is genuinely achievable.

***Index Terms* -** autonomous waste management, computer vision, OpenCV, outdoor robotics.

I. Introduction

Walk into any large office building, hospital ward, or university cafeteria, and the waste situation is usually the same: overflowing bins, mixed materials in the wrong containers, and cleaning staff constantly having to sort through what other people threw together. It is not a glamorous problem, but it is a persistent one. As indoor environments get busier and waste volumes climb, doing this entirely by hand becomes less and less sustainable.

The appeal of using a robot for this job is straightforward. A robot does not get tired, does not skip items, and can apply the same classification logic every single time without variation. The challenge is making one that is actually practical. A lot of research in this space ends up building systems that depend on expensive hardware or work only under narrow, pre-set conditions. We wanted to try a different approach: build something useful from low-cost parts and see how far it could actually go.

The V is the result of that. It combines an OpenCV vision pipeline on a Raspberry Pi 4 with a servo-driven robotic arm and a basic chassis for indoor movement. The robot detects waste objects through its camera, classifies them as plastic, metal, or paper, drives toward them, picks them up, and drops them into the correct slot. The whole thing runs on a single-board computer that costs well under a hundred dollars. This paper covers how we designed it, how we built it, and what we found when we tested it.

II. LITERATURE REVIEW

Several research groups have worked on automated waste handling over the past few years, and looking at what they built helped us figure out where the real gaps are.

One study from 2025 used YOLOv8 running on a Raspberry Pi, connected to a servo arm designed to grab plastic bottles [1]. Getting 0.90 precision at 10 FPS on that hardware is genuinely good, and the work showed that real-time detection on a Pi is feasible. The limitation was scope: only one waste type, no mobility, no way for the robot to go find what it detects. For a real deployment you need all of those things together.

A different approach from 2023 skipped cameras entirely and used physical sensors instead: IR, moisture, metal detection, and ultrasonic distance [2]. The ESP8266 module let it report data remotely, which was a nice touch. It worked reasonably well when waste was introduced in a predictable way, but the moment conditions shifted slightly, the sensor logic struggled. Physical sensors are brittle in that way; they need the object to behave in a very specific manner to produce a clean reading.

On the pure classification side, a CNN-based system from the same year hit accuracy numbers as high as 98% across several waste categories, with sub-0.1 second inference [3]. Those are excellent figures on paper, but the system never left the lab. There was no arm, no chassis, no real interaction with physical objects. Classification that cannot trigger any physical response is only half the problem solved.

Two hardware-focused papers gave us useful reference data. One examined a two-joint aluminium arm and validated it for loads up to 0.8 kg, though it had no sensing of its own [4]. Another structurally tested a wheel-track hybrid chassis design using CAE and FEA methods, which confirmed the mechanical concept without connecting it to any detection or manipulation system [5].

What all of this shows is that the pieces exist, but they rarely come together. Vision, manipulation, and mobility have mostly been developed in isolation. The V is our attempt to combine them into one working system.

III. SYSTEM ARCHITECTURE

The V was designed to handle waste detection, pickup, and sorting in one continuous cycle, without stopping for human input at any stage. The architecture is divided into sensing, processing, actuation, and communication, with the Raspberry Pi 4 sitting at the center of all four.

3.1 Block Diagram Description

Three sensors connect directly to the Pi. The camera feeds visual data into the classification pipeline. A GPS module tracks the robot's location within the space it is operating in. A GSM module keeps a communication channel open to a remote monitoring interface. Two motor driver circuits receive control signals from the Pi: one for the drive motors that move the chassis around, and one for the arm servos. Waste collected by the arm goes into one of three slots on the chassis depending on its category.

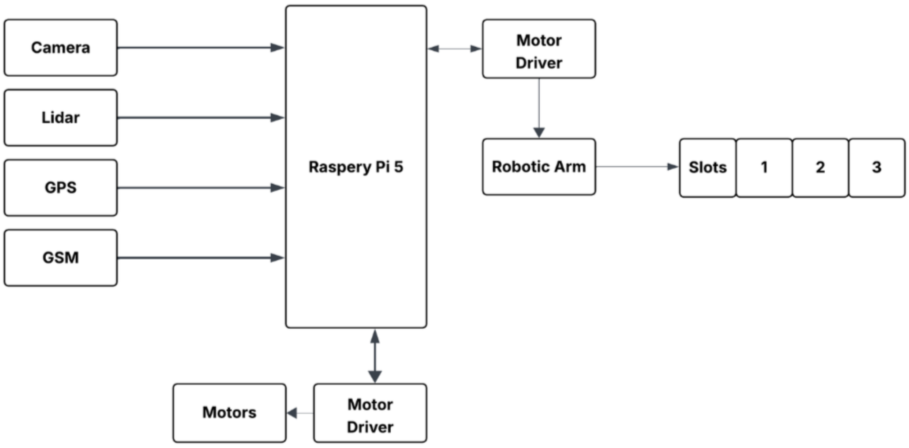

Fig 3.1 Block diagram

3.2 Hardware Architecture

The Raspberry Pi 4 manages everything centrally. It reads sensor data, runs the image processing pipeline, makes decisions, and sends out control signals, all at the same time. The hardware breaks down into three layers based on function.

The sensing layer handles data collection. The USB webcam streams frames for waste detection and classification. The GPS module records position information, which is used to track coverage during operation. The GSM module handles outbound status reports and inbound commands from whoever is monitoring the system remotely.

The mobility layer moves the robot around. A motor driver converts PWM signals from the Pi into the voltage levels needed to drive the locomotion motors. The robot can navigate to wherever a waste object has been detected, stop, complete the collection task, and then continue its patrol.

The manipulation layer handles physical contact with waste objects. A second motor driver controls the three servo joints in the arm. When a classification result comes in, the arm extends to the target position, the gripper closes around the object, and the arm carries it to the correct deposit slot before returning to its resting position.

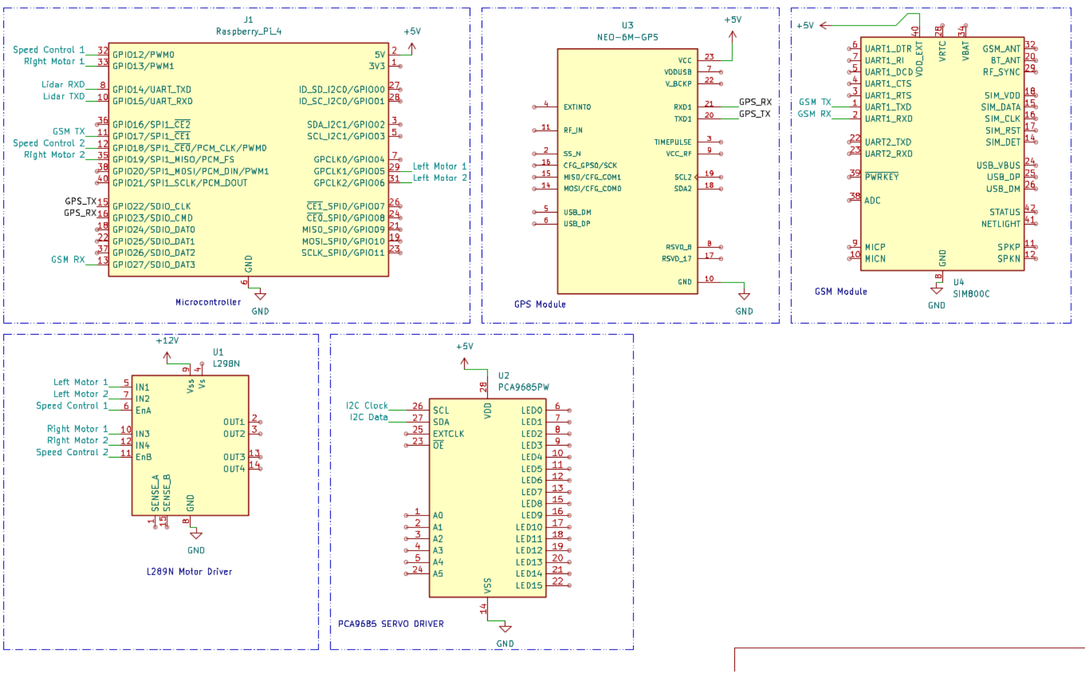

Fig. 3 .2 Schematic design of the overall system

3.4 Software Architecture

The software runs as a continuous loop from the moment the system starts up. Frames from the camera flow into the OpenCV classification module. When waste is detected, its category and screen position go to the path planning module, which calculates a route to the object. The robot drives to the location, the arm controller computes the joint angles needed and actuates accordingly, and the item gets deposited into its slot. The GSM module sends out a status update after each completed cycle.

IV. METHODOLOGY

The V works through two main subsystems that operate in sequence. First the camera and classification software identify a waste object and determine its type. Then the arm controller takes that result and physically retrieves and deposits the object. Everything else in the system exists to support those two operations.

4.1 Waste Detection and Classification

Each frame captured by the webcam starts in BGR format, which is how OpenCV reads camera data by default. Before any analysis happens, we convert it to HSV. The reason for this is practical: HSV separates colour information from brightness in a way that BGR does not, which makes the classifier much less sensitive to lighting changes. An object that looks different under fluorescent light versus natural light will still produce a recognisable HSV profile.

We then apply colour range masks that we tuned for each waste category. Contour detection runs on the masked regions to extract the outline shapes of candidate objects, and morphological operations clean up noise in the result caused by reflections or surface irregularities. From the cleaned contours, we extract both shape and colour features, and a decision tree assigns the object to a category.

Plastic objects tend to produce smooth, rounded contours. Depending on the specific item, they may appear transparent or show a solid saturated colour. Metal objects produce sharper, more angular contours and a grey, low-saturation surface appearance that is fairly consistent across different metal items. Paper and cardboard produce flat rectangular outlines with white or lightly toned colouring in the HSV space. These visual signatures are distinct enough for the rule-based approach to work well in practice, though they do have edge cases, particularly between plastic and paper as we found during testing.

Once a category is determined, the label and position coordinates go to the arm controller.

4.2 Robotic Arm Control

The arm has three joints (shoulder, elbow, and wrist) and a two-finger gripper at the end. To reach a detected object, we need to know what angle each joint should be at, which is what the inverse kinematics solver calculates. It takes the estimated three-dimensional position of the target object and works backwards through the arm geometry to produce the shoulder, elbow, and wrist angles that will bring the gripper to that point.

Those angles get converted into PWM duty cycle values and sent through the motor driver to each servo. The motor driver steps up the signal to the voltage the servos need. When the gripper reaches the target position, it closes, lifts the object, swings to the correct deposit slot, opens, and then the arm returns home. The whole pick-and-place cycle takes a few seconds per object.

4.3 Waste Segregation

Slot 1 on the chassis receives plastic, Slot 2 receives metal, and Slot 3 receives paper. The classification result from the vision pipeline determines which slot the arm targets; there are no mechanical switching parts involved. The target object's three-dimensional coordinates are estimated by combining the screen-space position from the camera with a depth reading from an ultrasonic sensor mounted at the front of the chassis. This gives the IK solver the spatial input it needs to accurately direct the arm.

V. Results and Discussion

We tested The V in a controlled indoor space using physical waste samples, running each subsystem through a structured evaluation before assessing the complete integrated pipeline.

5.1 Waste Classification

We ran the classifier on 90 samples, 30 from each category. Metal classification came out highest at 93.3%, plastic at 90.0%, and paper at 86.7%, for an overall accuracy of 90.0%. The metal result made sense to us given how visually distinctive metal surfaces are in HSV space. The slightly lower paper score was expected: during testing we noticed that a few of our plastic samples had cream-coloured surfaces that fell just inside the HSV range we had defined for paper. It was a boundary case that a colour-only classifier will always struggle with, and it accounted for most of the paper misclassifications.

5.2 Robotic Arm Performance

The arm achieved a mean end-effector positional error of plus or minus 5 mm, which was accurate enough for every object size we tested. The grasping success rate was 92.6%. Most of the failures happened when objects had settled into orientations that the wrist servo could not compensate for fully, so the gripper could not get a solid grip before attempting the lift. It was not a common occurrence but happened often enough to show up in the numbers. Adding another degree of freedom at the wrist joint would likely resolve it.

5.3 Overall System Performance

Running the full pipeline end to end, including detection, navigation, arm actuation, and deposit, gave us an overall segregation accuracy of 83.3% at 8 to 10 FPS. The drop from 90.0% classification accuracy to 83.3% segregation accuracy reflects the combined effect of occasional misclassifications and the arm failures described above. Neither was dominant; both contributed about equally.

The frame rate confirmed that real-time operation on the Raspberry Pi 4 is practical. We did not need any additional compute hardware, which was one of the things we specifically wanted to validate. The robot has no difficulty processing frames fast enough to support reliable decision-making at its current operating speed. Replacing the rule-based classifier with a CNN and adding a wrist rotation joint would be the two changes most likely to push overall segregation accuracy meaningfully higher in the next iteration.

Table 1: Waste Classification Accuracy by Category

SI NO	Waste Category	Samples Tested	Correctly Classified	Accuracy (%)
1	Plastic	30	27	90.0
2	Metal	30	28	93.3
3	Paper	30	26	86.7
4	Overall	90	81	90.0

VI. Conclusion and Future Work

The V is a working autonomous waste collection robot built on a Raspberry Pi 4. It detects, classifies, picks up, and deposits waste into the correct category slot without requiring any human input during operation. In our testing it achieved a classification accuracy of 90.0%, an arm grasping success rate of 92.6%, and an end-to-end segregation accuracy of 83.3% at 8 to 10 FPS. We built it from inexpensive, off-the-shelf components to demonstrate that effective waste automation does not require specialist or high-cost hardware.

There is a clear path for improving on these results. The highest-impact next step would be swapping the current rule-based classifier for a trained convolutional neural network, which would handle edge cases between categories much more reliably. Adding a LiDAR sensor and implementing a SLAM algorithm would allow the robot to map and navigate unknown environments on its own rather than operating in a pre-defined space. Migrating the software to ROS 2 would make the system more modular and much easier to extend as new subsystems are added. We also plan to extend the segregation system beyond three categories to cover organic waste, glass, and electronic materials, which would make The V applicable to a much broader set of indoor environments.

VII. Acknowledgment

We gratefully acknowledge the support of the management of Eranad Knowledge City Technical Campus, Manjeri, and sincerely thank Dr. Adharsh T K, Principal, for providing the resources that made this project possible. Our deepest gratitude goes to Asst. Prof. Neethu P M, Head of the Department of Electronics and Communication Engineering and our Project Guide and Coordinator, for her guidance, mentorship, and consistent encouragement from start to finish. We also thank the faculty of the Department of Electronics and Communication Engineering for their support throughout. Finally, we appreciate the patience and backing of our families and colleagues, who made it easier to see this work through to completion.

References

[1] A. Kumar and R. Sharma, "YOLOv8-based plastic waste detection using Raspberry Pi and servo-controlled robotic arm," Journal of Robotics and Automation, vol. 12, no. 3, pp. 145–153, 2025.

[2] S. Patel, M. Nair, and T. Raj, "IoT-enabled waste segregation system using IR, moisture, and metal sensors with ESP8266 communication," International Journal of Embedded Systems and Applications, vol. 13, no. 4, pp. 67–75, 2023.

[3] L. Zhang, H. Chen, and Y. Liu, "CNN-based robotic waste classification with high-accuracy inference for indoor environments," IEEE Access, vol. 11, pp. 23014–23025, 2023.

[4] R. Fernandez and J. Torres, "Design and kinematic analysis of a 2-DoF aluminium robotic arm for educational applications," Proceedings of the International Conference on Mechanical and Robotics Engineering (ICMRE), pp. 112–118, 2020.

[5] B. Singh and P. Mehta, "Wheel-track hybrid mobility mechanism for waste collection robots: Structural validation using CAE/FEA," International Journal of Mechanical Engineering and Robotics Research, vol. 9, no. 2, pp. 88–95, 2020.