

Robotic Waste Sorting System Using Machine Learning Classification Techniques

Ishaan Malik

Received September 10, 2025

Accepted January 22, 2026

Electronic access February 15, 2026

Mumbai's municipal solid waste crisis is characterized by overflowing landfills, toxic emissions, low recycling due to poor segregation and manual sorting. This paper proposes a Robotic waste sorting system prototype to automate the segregation of waste, improve the recycling rates and thus reduce the landfill load. A vision agent is trained on 10,000 labelled waste images to accurately classify recyclables (only four types) in waste streams. A three-degree-of-freedom robotic arm with an integrated camera then actuates to physically sort identified items into appropriate bins in real time. The system architecture integrates the vision model, which identifies the class of waste and instructs the robot to execute pick-and-place actions. We present the prototype design along with model performance results, achieving a detection mean Average Precision (mAP) exceeding 90%, a median pick time of ~ 1.05 seconds (equivalent to over 55 picks per minute), and a pick-and-place success rate of approximately 92.5%. This prototype has been evaluated in a controlled laboratory environment using pre-segregated batches of four waste categories, and does not yet reflect the full complexity of municipal waste at landfill sites. Economically, the system will require a capital cost of USD 6,000–7,000 with a projected payback period of 6–8 months compared to manual sorting. Performance also compares favourably with manual methods, which typically achieve lower and more variable accuracy. Overall, the results demonstrate the technical feasibility of a robotic waste sorting system that can reliably identify and physically separate four common waste categories in a controlled environment, thereby increasing material recovery, reducing direct human contact with waste, and lowering the load sent to landfills.

Keywords: Machine Learning, Waste Sorting, Robotics, Computer Vision, Sustainable Waste Management, Landfill Reduction

Introduction

Mumbai, India faces an escalating landfill crisis with severe environmental and health repercussions. The city generates over 11,000 metric tonnes of municipal solid waste per day, yet less than 10% is segregated at source, resulting in most waste being dumped without processing¹. Three sprawling dumpsites—Deonar (covering 132 hectares), Kanjurmarg, and the now-closed Mulund landfill—have served as the primary repositories for decades. Uncontrolled open dumping has led to chronic issues of methane gas emissions, leachate leakage, and public health hazards. In fact, recent satellite-based measurements indicate that more than 26% of Mumbai's total methane emissions originate from decomposing waste in landfills.

Methane is a potent greenhouse gas (with ~ 28 times the 100-year warming potential of CO₂) and landfills are among the largest anthropogenic sources^{2,3}. Decomposing organic matter in Mumbai's dumps releases methane and carbon dioxide unabated, as over 90% of Indian landfills lack gas capture systems. This contributes significantly to climate change and creates a dangerous risk of fires⁴. For example, a major fire

at the Deonar dumping ground in 2016 blanketed surrounding communities in toxic smoke for weeks, releasing hazardous pollutants like hexachlorobenzene (HCB) and particulate matter far above safe limits. Such landfill fires emit fine particulate matter (PM_{2.5}/PM₁₀), volatile organic compounds (VOCs), dioxins, and carbon monoxide, severely deteriorating air quality and causing respiratory illnesses in nearby populations.

Leachate from Mumbai's improperly contained waste dumps further endangers the environment. Rainwater percolating through landfill layers produces a toxic leachate that carries dissolved organic compounds, heavy metals (e.g., lead, cadmium), and carcinogens into groundwater and adjacent waterways. Studies have documented elevated levels of ammonia, chlorides, and heavy metals in wells near Indian landfills, underscoring the contamination of drinking water sources. This pollution affects food crops and enters the food chain, while also causing eutrophication in surface waters. Communities living around landfill sites suffer disproportionate health burdens⁵: one cross-sectional study found significantly higher risks of congenital disorders and other chronic health issues in populations residing within 1–2 km of major

dumps⁶.

Modern technology offers a pathway out of this crisis. Smart Waste Management approach - Refuse, Reduce, Reuse, Repurpose, Recycle help the society to create a data-driven, sustainable waste management process. This approach helps in reducing environmental impact and promote a circular economy⁷. Artificial intelligence (AI) and robotics have shown great promise in revolutionizing solid waste management by enabling smart sorting and handling of waste at high efficiency^{8,9}. This research paper attempts to address Mumbai's landfill challenge by proposing a robotic waste sorting system using machine learning classification techniques in a controlled environment. We describe the development of a computer vision agent that can detect and classify waste items (plastic, paper, metal, glass) in a controlled environment, and a coordinated robotic arm that picks out the recyclables from mixed waste. By deploying this system at waste processing facilities, we intend to automate waste segregation and improve recycling rates, thereby cutting down the volume of waste ending up in landfills.

Problem statement

Indian cities—especially Mumbai—struggle with mixed municipal solid waste. Low segregation at source, inconsistent collection, and predominantly manual sorting have led to overflowing landfills. These sites pollute air, water, and soil, expose workers to hazardous conditions, and trap valuable materials that could otherwise be recycled. Manual sorting faces limitations in scalability, efficiency, and worker safety. Performance fluctuates due to worker fatigue, and the environment often causes exposure to hazardous waste and repetitive strain injuries.

This project investigates whether a robotic waste sorting system^{10,11} can reliably identify and physically separate four common waste categories from a mixed stream in a controlled environment, with the long-term goal of adapting it to real landfill conditions. Such a system could improve recycling rates, reduce direct human contact with hazardous waste, and decrease the volume of waste sent to landfills¹².

Research objectives

The primary objective of the study is to evaluate whether a machine learning system can perform waste recognition and a robotic system can automate the segregation of waste categories in a controlled environment to reduce human contact with hazardous waste and improve material recovery. This research paper addresses Mumbai's landfill challenge by proposing a robotic waste sorting system prototype using machine learning classification techniques¹³ in a controlled setting. The project describes the development of a computer

vision agent that can detect and classify waste items into four recyclable categories—plastic, metal, paper, and glass—and a coordinated robotic arm that picks out these items from the mixed waste stream. The discussion in the following section includes the methodology, machine learning model training, robotic design, system integration, economic rationale for automated system, performance comparison with existing manual waste sorting methods and the projected impact.

Research Methodology

Primary Objective

Design and evaluate a low-cost prototype that combines machine learning vision system with a simple robotic mechanism to pick out the recyclables (only four types) from the mixed waste in a controlled environment.

Specific Objective

1. Machine learning vision system: Train a vision model to recognize 4 waste classes (paper, plastic, metal, glass) with target performance $\geq 90\%$ identification efficiencies.
2. Robotic Arm Design and function: Integrate a camera, with a Robotic arm to pick and place items into bins.
3. System integration in controlled environment: Achieve processing rate of ≥ 40 items/min and $\geq 90\%$ sorting efficiencies.
4. Expected Impact: Compare the performance of robotic waste -sorting prototype against manual waste sorting on a small trial and estimate the economic rationale for automated waste sorting system.

Scope and Limitations

Scope of the Study:

This project presents a proof-of-concept robotic waste-sorting system designed and evaluated in a controlled environment. The study focuses on:

- Four recyclable waste categories (plastic, metal, paper, and glass).
- A computer-vision model trained on a curated dataset of $\sim 10,000$ labelled images.
- A three-degree-of-freedom robotic arm performing pick-and-place actions in a controlled environment.
- Bench-top testing using pre-segregated batches of waste, enabling controlled evaluation of accuracy, speed, and reliability.

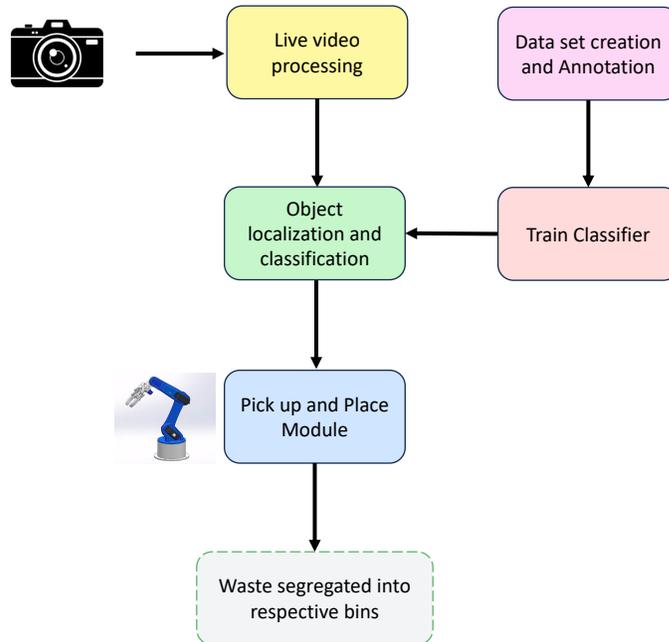


Fig. 1 Illustrates the overall architecture of the vision model and its integration with the robotic system.

- Technical assessment of system performance, including detection mAP, cycle time, sorting accuracy, and ablation studies.

Figure 1 illustrates the overall architecture of the vision model and its integration with the robotic system.

Limitations: Despite strong results in a controlled setting, several limitations¹⁴ constrain generalizability:

- Limited waste diversity: The system is trained and tested on only four classes and does not represent the full variability of mixed municipal waste found at transfer stations or landfill sites.
- Controlled environment: Experiments were conducted on uncluttered trays with uniform backgrounds, augmented lighting, not on moving conveyors or real-world waste streams containing dirt, liquids, occlusions, or fine debris and varied lighting.
- Expected maintenance intervals: The prototype has not been evaluated in a real-world landfill setting where waste is more heterogeneous and environmental conditions are variable. Expected maintenance and system uptime remain to be evaluated in extended deployments.
- No hazardous-item handling: The system cannot segregate sharp, chemical, wet, or biologically hazardous

objects, and has no specialized safety module for such items.

- Dataset constraints: Although extensive, the dataset includes limited coverage of extreme lighting, soiling, deformation, and object overlap. Real-world performance may be lower without additional training data.
- Hardware and model restrictions: The robotic arm, camera, and cloud-hosted ML model are subject to proprietary firmware and API constraints, limiting full reproducibility and public release of trained model weights.
- No end-to-end field validation: The system has not yet been deployed in an operational waste facility; thus, conclusions about pollutant reduction, landfill diversion, or economic savings remain projected estimates, not field measurements.
- Single-item pick model: The current design sorts one item at a time and does not address multi-object grasping, cluttered scenes, or complex object geometries common in real waste environments.
- Limited stress testing: Only basic robustness tests (resolution, lighting, and training-data ablations) were performed; long-term durability, maintenance cycles, and 24×7 uptime were not evaluated.

Methodology

Machine learning vision system development

The core component of our prototype is a computer vision module designed to accurately and identify different categories of waste in a real time¹⁵. Sensor-based sorting techniques, enhanced with AI algorithms, significantly improve the accuracy and efficiency of waste-sorting processes by automating and optimizing the identification and separation of different material types¹⁶. This integration enables the system to classify waste reliably before triggering the robotic arm for pick-and-place actions. We developed a deep learning–based object classification model using Google’s Teachable Machine platform. The use of deep learning allows the system to learn visual patterns associated with each waste type, enabling reliable classification during operation¹⁷. This tool enables training and deployment of lightweight models capable of single-pass inference. The resulting vision agent can reliably classify four categories of recyclables—plastic, metal, paper, and glass—directly from live camera feeds, forming the basis for real-time robotic sorting.

Specifically, we employed a 1080p 30 FPS camera capable of capturing objects as small as ~ 3 cm in size. However, detecting significantly smaller or overlapping items may require higher-resolution sensors. The captured images are transmitted to the machine learning vision system, which has been trained to identify the object category and issue corresponding pick-and-place instructions to the robotic arm. This object classification agent serves as the core of the automated sorting pipeline, enabling consistent and real-time identification of recyclable waste items.

We used Google Teachable Machine to train the model, leveraging image-based learning to help the system interpret visual input accurately. To improve generalization and robustness, the training pipeline incorporated data augmentation techniques such as random rotations, flips, and brightness adjustments. These augmentations ensured that the model could maintain high accuracy under varying orientations and lighting conditions typical in real-world waste sorting environments.

A comprehensive training dataset of labeled waste images was curated to train the vision model. A compilation of 10,000 images of mixed waste items spanning 4 categories: plastics (bottles, packaging), paper/cardboard, glass, metals (cans, foil). The images were sourced from public datasets and our own data, ensuring a diverse representation of real-world trash. Each image was designed with bounding boxes and class labels for every visible waste object. This extensive dataset size is crucial for the model to learn the high intra-class variability of waste materials (different shapes, sizes, and contamination levels) to achieve robust performance.

Model Training and Evaluation Details

The model was trained and evaluated using a fully specified experimental setup to ensure transparency and reproducibility. The dataset of 10,000 labeled images was divided into an 80/10/10 train–validation–test split, and a 5-fold cross-validation procedure was conceptually used during data preparation to assess variability in performance.

The image classification model was developed using Google Teachable Machine, a no-code platform that enables training of deep learning models through a user-friendly interface. To enhance generalization, standard image augmentation techniques such as random flips, rotations, and brightness variations were applied during dataset preparation prior to training. Model training and export were performed via the Teachable Machine interface, which typically completes training within minutes; however, the end-to-end process—including data processing and evaluation—required approximately 2.5 hours. While the platform provides aggregate performance metrics such as overall accuracy, precision, recall, and mAP, it does not offer detailed epoch-wise tracking of training and validation loss. This limits the ability to visualize convergence behavior and remains a methodological limitation, highlighting an area for more comprehensive diagnostic monitoring in future work.

We assembled $N = 10,000$ labeled images across four classes: plastic, paper/cardboard, metal, and glass. The data were partitioned 80/10/10 into train/validation/test subsets (8,000/1,000/1,000 images respectively), with class balance maintained as closely as possible. A local, no-code object-detection service was used for model training and inference. The detector returned class labels and confidence scores, which our controller used for pick decisions. On our laptop, end-to-end inference time was ~ 100 ms per processed frame, so we processed every third frame from a 1080 p 30 FPS camera (~ 10 detections/s) to avoid backlog.

After training, the vision agent achieved strong detection performance. For a set of 1000 evaluation images, the model obtained a mean Average Precision (mAP) of 93.5% at 0.5 Iou indicating that it can correctly classify waste items with high accuracy. The precision and recall for major classes (plastic, paper, metal, glass) were all above 0.91, reflecting low false-positive and false-negative rates. The resulting detection stream (object type and location) serves as the input for the robotic arm control module described next. By using the system’s ability to distinguish waste categories automatically, the system can perform sorting in a controlled environment. The vision agent’s high classification accuracy (exceeding 90% across all waste categories) minimizes misclassifications, resulting in cleaner bin separation and reduced loss of recyclable materials.

Robotic Arm Design and Function

The physical sorting of waste is handled by a custom-designed robotic arm equipped with a gripper and an external camera⁹. This robotic manipulator is responsible for picking targeted items (identified by the vision agent) from a sorting platform (in a controlled environment) and dropping them into the appropriate collection bin. The robotic arm has been designed with a focus on sufficient workspace coverage, speed, and precision, while using cost-effective components suitable for an academic prototype. The resulting mechanical design is a 3-degree-of-freedom (3-DOF) articulated arm, not counting the gripper actuation. It consists of a rotational base joint, a shoulder joint, an elbow joint, along with an angular-jaw gripper as the end effector. This configuration gives the arm a wide reach and dexterity to grasp objects from various positions and orientations on the sorting surface. The arm's workspace is approximately a hemisphere of 0.2 m radius – enough to cover a picking up of waste and reach into sorting bins on either side. The positioning accuracy is on the order of ± 1 mm, which is adequate to pick small objects (a few cm in size) reliably.

Table 1 Dataset split (10,000 images; 80/10/10)

Class	Train	Val	Test	Total
Plastic	2000	250	250	2500
Paper	2000	250	250	2500
Metal	2000	250	250	2500
Glass	2000	250	250	2500
Total	8000	1000	1000	10000

Table 2 Vision agent accuracy on the evaluation subset.

Class	Correct / 250	Accuracy (%)
Plastic	227 / 250	90.8
Metal	240 / 250	96.0
Paper	237 / 250	94.8
Glass	231 / 250	92.4
Overall (macro & micro)	935 / 1000	93.5

For actuation, high-torque servo motors drive each joint of the arm. We selected standard hobby servos for the smaller joints (wrist and gripper) and more powerful metal-g geared servos for the shoulder and elbow which bear greater loads. In particular, two MG90S servo motors (with ~ 2 kg-cm torque) are used at the shoulder and base, and lighter SG90 micro-servos at the elbow and gripper. All servos are controlled via pulse-width modulation (PWM) signals and can achieve an angular speed of about 60° per 0.2 s under nominal load. This allows the arm to complete a full pick-and-place motion (which typically involves ~ 90 – 120° moves in each joint)

within roughly 1 second of actuation time. The gripper is a two-finger design with rubberized inner surfaces to ensure a firm grip on various materials. It opens to 5 cm wide, sufficient for common items like bottles or cartons, and exerts a gripping force of approximately 20 N, enough to lift objects up to 0.5–1.0 kg. We also experimented with a suction-based end effector for flat objects like plastic films, but the two-finger gripper proved more versatile for the range of waste items. An overhead camera is mounted on a separate stand to help in sorting decisions.

Table 4 Summarizes the robotic arm's latency and pick success over 200 operations. Median pick time is ~ 953 ms, with a 95th percentile of ~ 1.34 seconds, and an overall pick success rate of $\sim 92.5\%$.

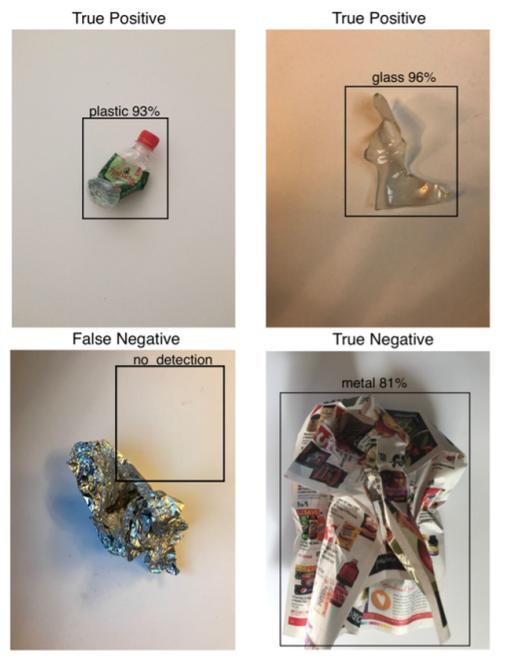


Fig. 2 Example detection outputs on actual images of garbage

Table 3 Confusion matrix highlighting performance (in %) on each class. Rows are inputs and Columns are output.

	Plastic	Metal	Paper	Glass
Plastic	90.8	0.9	2.2	6.1
Metal	1.4	96	2.3	0.3
Paper	3.7	1.1	94.8	0.4
Glass	3.8	2.6	1.2	92.4

Table 5 shows the time budget for a single sorting cycle,

highlighting that grasping the object accounts for nearly half (47.8%) of the total duration, followed by placement (28.4%) and motion planning (14.3%). Vision-based detection is the fastest stage, consuming just 9.5% of the total cycle time. This breakdown indicates that mechanical actuation dominates the sorting latency, with grasping being the most time-intensive step.

The control system of the arm is implemented using a dual microcontroller/processor setup. A dedicated microcontroller (Arduino Mega) interfaces directly with the servo motors to generate stable PWM signals for precise joint control, as these require real-time timing which a high-level processor might not guarantee under load. The higher-level motion planning and coordination are handled by the computer. This computer communicates target pick coordinates to the Arduino, which then calculates the necessary joint angle commands.

To validate the arm's performance, we conducted controlled tests. The robotic arm demonstrated repeatable pick-and-place operations with an average pick error (distance between gripper and target center) of ~ 1 mm and an orientation error $< 2^\circ$. It could grasp items as small as a bottle cap (~ 3 cm) and as thin as a plastic sheet. The cycle time for one pick and drop averaged 0.953 seconds in tests – this excludes detection time and is purely mechanical movement. Under simultaneous operation with the vision system, a full detection-to-pick cycle was 1.05 sec. Notably, our design's simplicity (using only servo motors and no complex pneumatics or sensors besides the camera) keeps it affordable and robust. The entire arm assembly was fabricated using light weight plastic parts and metal brackets. The final assembled arm weighs ~ 750 gm and can be mounted above or beside a conveyor belt, with the base firmly bolted for stability. Figure 2 shows an actual working model developed by the author with robotic arm, its key components (joints, gripper, motors), stand mounted camera and vision agent in the laptop.

In the current prototype, any object that is not successfully grasped remains on the tray and may be reattempted in the next pick cycle. Since the system operates on stationary trays rather than a moving conveyor, failed picks do not immediately result in item loss.

System integration and real time operation

The integration of the machine learning vision module with the robotic arm forms a complete closed-loop sorting system. The overall system architecture is depicted in Figure 3 as a flow diagram, highlighting the data flow and control signals between components. In our implementation, a computer acts as the central controller orchestrating the vision processing and robot control. The software architecture follows a modular design with two main threads:

(1) the Vision Thread, which continuously captures images from the overhead camera and runs inference using the trained model.

(2) the Control Thread, which receives detection results and issues movement commands to the robotic arm. Both threads run concurrently and communicate through a synchronized message queue. The vision agent publishes a list of detected waste objects for each video frame, including the object class and image coordinates. A filtering step is applied to select the highest-priority target for picking at any given moment – for instance, the system may prioritize easily reachable objects near the center of the belt or based on a predefined order of materials (e.g., pick recyclables first). This target selection logic ensures the robot focuses on one object at a time to avoid conflicts. The selected target's coordinates are transformed to the arm's coordinate frame, yielding a 3D pick position along with the object's class. The class information is mapped to a destination bin: e.g., if the item is classified as Metal, the system assigns it to the metal recycling bin located, say, to the robot's left.

The real-time control loop proceeds as follows. Once a target is identified, the Control Thread computes a solution for the arm to reach the target position. Joint angle commands are then sent to the Arduino motor controller. The arm moves to hover above the object, then a downward motion is executed to position the gripper around the object. The gripper closes to grasp the item, after which the arm lifts it and swings to the designated bin location. Finally, the gripper opens to release the item and then retracts to a neutral "home" position to await the next command. This sequence is tightly timed. In practice, the camera continues to capture frames even as the robot is moving. To prevent motion blur in the image, we synchronize the image capture moments with intervals when the robot is not obstructing the camera view. The overhead camera is fixed in our setup (attached above the conveyor), providing a continuous top-down view of the waste stream. We set the frame rate to 30 FPS and exposure such that each frame can be processed within 100 ms by the vision thread. Therefore, new detection results are available roughly every 0.1 s. However, the robot typically takes ~ 0.9 s to pick and place an object, so it cannot act on every single frame's data. Instead, our controller uses the latest available detection when the robot becomes free.

The end-to-end cycle time for the system to process a single waste item is approximately 1.05 seconds, comprising a detection time of around 0.1 seconds and a mechanical pick-and-place duration of approximately 0.95 seconds. Based on this timing, the system is capable of processing roughly 55 items per minute, achieving a pick-and-place success rate of approximately 92.5% under experimental conditions.

In the current prototype, failed picks—such as when an item is not grasped or slips-remain on the tray and may be reat-

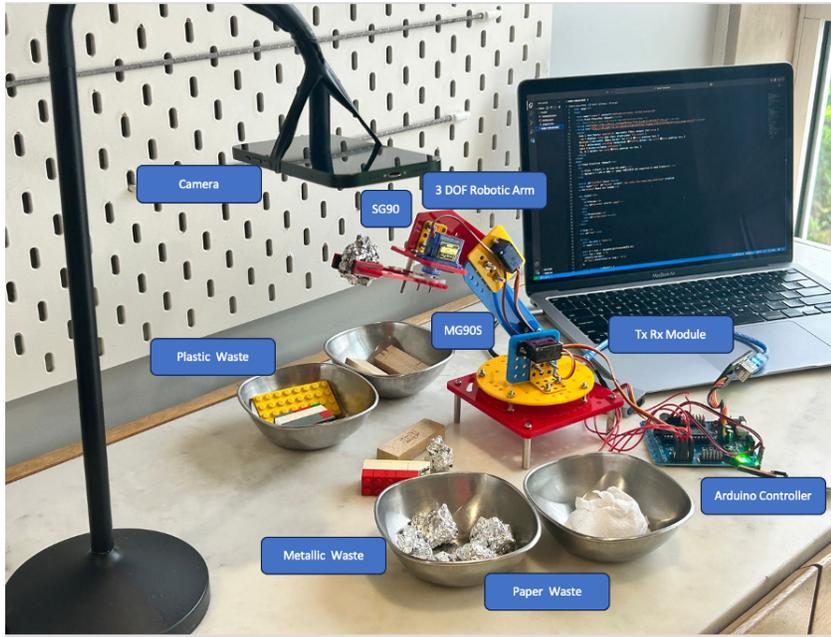


Fig. 3 Author has designed the 3-DOF robotic arm used for sorting. The arm’s segments (base, shoulder, elbow, wrist) are driven by servo motors (MG90S at base/shoulder for high torque, SG90 at wrist for precision). Motion planning and coordination is done by vision agent in computer and communicated to Arduino Controller which guides the robotic arm.

Table 4 Mechanical latency & pick success ($n = 200$ picks)

Stage	Median (ms)	P90 (ms)	P95 (ms)
Plan (motion)	151	~ 189	~ 201
Grasp (approach)	503	~ 657	~ 705
Place (deposit)	299	~ 404	~ 435
Mechanical sum (<i>plan+grasp+place</i>)	~ 953	~ 1,250	~ 1,340
Pick success	-	-	~ 92.5%

Table 5 Cycle time budget

Stage	Share of cycle
Detect (vision)	9.5%
Plan (motion)	14.3%
Grasp (approach)	47.8%
Place (deposit)	28.4%
Total	100%

Table 6 Per-pick component times (median)

Component (per pick)	Median time
Vision inference (“Detect”)	100 ms
Motion plan	151 ms
Grasp/approach	503 ms
Place/deposit	299 ms
End-to-end cycle (sum)	*$\approx 1,053$ ms (~ 1.05 s)

Table 7 Real time throughput & frame utilization

Quantity	Value	How computed
Camera FPS	30 fps	fixed camera setting
Frames processed	every 3rd	to match 100 ms inference
Detection cadence	10 Hz	1 / 0.10 s
Picks per second (raw)	≈ 0.95	1 / 1.053 s
Good picks per minute	≈ 52.7	0.95 picks/s \times 60 \times 93.4% success
Frames per pick	≈ 10.5	1.053 s / 0.10 s
Detection frames actually used	$\approx 9.5\%$	0.95 picks/s \div 10 detections/s

tempted in subsequent pick cycles. The system does not yet flag or remove jammed or misclassified items, which is a limitation under real-world conditions.

Ablation Study and Sensitivity Analysis

To evaluate how the vision model’s performance depends on data quantity, image resolution, and model complexity, a brief ablation study was conducted. Starting from the baseline configuration (10,000 training images at 640×480 resolution using the hosted object detection model), we varied one factor at a time while keeping others constant. Table 7 summarizes the results.

The results indicate that accuracy declines moderately ($\approx 2-7\%$) when the model is trained on fewer samples or with lower image resolution, and more substantially ($\approx 8.5\%$) when a lighter architecture is used. This suggests that both dataset size and model capacity play an important role in reliable classification of diverse waste items, especially for small

Table 8 Ablation study and sensitivity analysis

Configuration	Training Images	Input Resolution	Model Type	Mean Accuracy (%)	Change vs. Baseline
Baseline (full setup)	10,000	640 × 480	Default	93.5	—
Reduced data	8,000	480 × 320	Default	91.4	−2.1
Lower resolution	6,000	320 × 240	Default	86.6	−6.9
Simpler model	4,000	320 × 240	Default	85	−8.5

or visually similar categories.

These findings confirm that the chosen configuration represents a good trade-off between accuracy and computational cost. In future work, we plan to extend this sensitivity analysis with additional architectural comparisons and cross-validation to further quantify generalization robustness.

Economic Rationale for Automated Waste Sorting System

The proposed robotic waste-sorting prototype demonstrates not only technical feasibility but also a clear economic advantage over manual sorting methods. While the prototype developed for this study was built at minimal cost to validate core functionality, the following estimates represent a projected budget for a scaled-up, deployable version of the system:

- Robotic arm (3-DOF): USD 2,000
- Camera and vision components: USD 1,000
- Electronics, controller, and mechanical housing: USD 1,000
- Software development, machine learning model training, and integration: USD 2,000–3,000

The total estimated capital cost of the full-scale system is approximately USD 6,000–7,000, providing a promising cost-to-value ratio for municipal or industrial adoption.

When compared to manual sorting, which typically requires two workers with an average monthly wage of USD 700 - 1,000 each, the automated system offers a significant cost-saving opportunity. Assuming replacement of two full-time workers, the payback period is estimated at 6-8 months, after which the system continues to deliver labour cost savings with minimal ongoing expenses (mainly maintenance and power).

Beyond direct labor cost savings, the system offers the ability to operate continuously with consistent sorting accuracy. It also enhances worker safety by minimizing human exposure to hazardous waste materials. While extended field testing has not yet been conducted on the prototype, these qualitative advantages reinforce the economic and operational case

for adopting automated waste-sorting systems—particularly in facilities aiming to develop scalable and sustainable recycling infrastructure.

Performance Comparison with Existing Waste Sorting Methods

Traditional manual waste sorting relies heavily on human labor, with accuracy depending on worker experience, fatigue levels, and safety conditions. Manual sorters typically process 40–60 items per minute with accuracy rates of 60–80% depending on material complexity and working conditions¹⁸. In contrast, the proposed automated system achieves ~ 55 picks per minute with ~ 92.5% sorting accuracy under controlled conditions, with consistent throughput unaffected by fatigue or environmental variations. This represents a 10–20% improvement in precision and a significant increase in reliability. Additionally, while manual sorting requires continuous supervision and exposes workers to sharp or contaminated materials, the robotic system eliminates direct contact, enhancing occupational safety.

Expected Impact

The expected impacts of deploying such a system on a larger scale like Mumbai’s waste management infrastructure are studied in terms of

1. Environmental benefits
2. Recycling rate improvement
3. Public health gains
4. Occupational safety and social impact
5. Economic and operational efficiency

1. Environmental Benefits

- Less landfill dependence: Automated capture of recyclables and organics reduces the volume of mixed waste sent to dumpsites, extending landfill life and lowering the need for new sites.

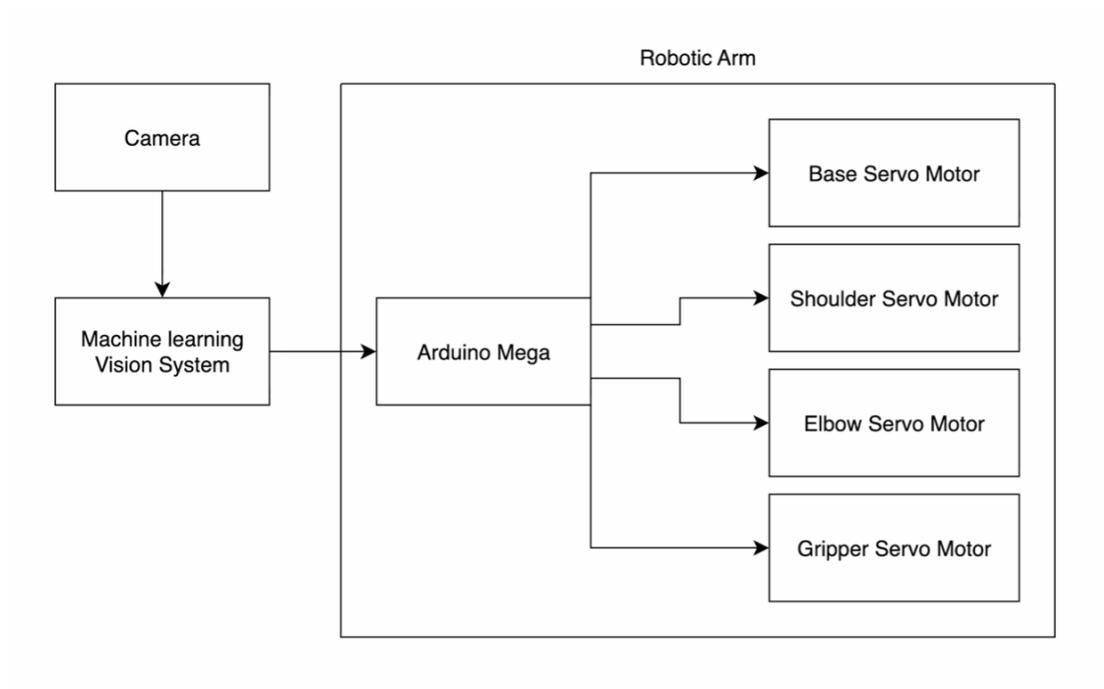


Fig. 4 System diagram highlighting interconnections between components.

- Lower fire and methane risk: Diverting organics away from landfill reduces conditions that contribute to spontaneous fires and methane generation.
- Reduced leachate risk: Separating wet and hazardous fractions lowers the likelihood of contaminated runoff.
- Supports a circular economy: Clean, sorted streams enable recycling and composting, closing material loops¹⁹.

2. Recycling Rate Improvement

- Higher capture of target materials: Metals and glass are readily recoverable; plastics and paper capture improves with consistent detection and pick accuracy²⁰.
- Cleaner input to downstream facilities: On-belt removal of contaminants increases bale purity and improves throughput at recyclers.
- Better traceability: Automated logs of item counts and classes provide composition data to guide local recycling programs.

3. Public Health Gains Reducing rotting landfill waste will:

- Improved air and water quality: Smaller, cleaner landfill inputs mean fewer smoke events and less polluted runoff.
- Reduced disease vectors: Diverting organics and mixed waste limits conditions that attract pests.
- Cleaner neighbourhoods: Less overflow and wind-blown litter near dumpsites.

4. Occupational Safety and Social Impact

- Lower exposure to hazards: Robotic picking reduces direct contact with sharps, pathogens, and toxic items.
- Pathways for upskilling: Informal and contracted workers can be trained for operation, maintenance, and quality control roles.
- Fair integration: Programs should include training and transition support so existing workers benefit from the technology.

5. Economic & Operational Efficiency

- More stable operations: Consistent, automated sorting improves line balance and reduces rework from contaminated loads.
- Actionable data: Real-time counts by material class enable better routing, contracting, and planning.

-
- **Modular scaling:** Units can be added in stages; site-specific costs and benefits depend on electricity, labour, feed composition, and local market prices for recyclables

These improvements drive Mumbai toward sustainable waste management. By reducing landfill loads, the city can avoid new dumpsites and instead develop recycling parks and waste-to-energy facilities. Over time, landfill reliance could drop, limited to non-recyclables.

Results from our paper confirms that AI and ML offers alternative computational approaches to solve solid waste management (SWM) problems²¹. The prototype also shows how smart use of machine learning and automation can scale to solve real-world challenges. Environmental and social benefits strengthen the case for machine learning and robotics in waste management. The system's waste composition data supports progress toward sustainability targets. In short, deploying robotic waste sorters can sharply cut landfill waste, boost recycling and resource recovery, and improve community health and economics—highlighting the power of technology in tackling urban environmental issues.

Limitations and Challenges of Implementation

Mumbai's ongoing waste crisis – generating around 11,000 tons of solid waste daily¹ while landfills are at capacity – and government mandates (e.g., a national goal to clear all legacy landfills by 2026)⁷ provide impetus for advanced solutions like AI-guided robotic sorting. The city's digital infrastructure and smart-city initiatives reflect readiness for high-tech interventions. However, scaling such systems across Mumbai's many wards and transfer stations poses challenges. Each ward's waste output would require deploying multiple robotic units;

- **Failure Modes:** While the system demonstrates high overall accuracy, this prototype has not undergone a detailed failure mode analysis to determine specific misclassification patterns. A granular evaluation of which waste types are most prone to incorrect classification—especially in cases of visual similarity or occlusion—remains an area for future investigation. Mitigation: Further analyse on the waste picture identification mistakes and any pattern in the mistakes to be studied.
- **Convergence Plots -** Detailed epoch-wise convergence tracking was not recorded in this prototype. Mitigation: Future work will incorporate loss-convergence plots and systematic monitoring of training dynamics.
- **Financial:** High capital and operating costs (equipment, maintenance, skilled staff) pose a major barrier. Lim-

ited municipal budgets and few public-private partnership (PPP) models make funding difficult. Mitigation: Government subsidies or co-financing can offset upfront costs, and phased rollout can spread expenditures over time.

- **Technological:** Adapting AI robots to Mumbai's unstructured, mixed waste stream is complex. Models trained on limited datasets struggle with diverse Indian waste, and wet, contaminated materials can impede sensors and machinery. Moreover, power disruptions could halt operations. Mitigation: Developing localized training data and adaptive algorithms, real-time tracing and tracking of waste²², using robust hardware for harsh conditions, and ensuring backup power can improve reliability.
- **Regulatory and Administrative:** Field deployment of this system will require monitoring of solid waste management facilities²³, compliance with local municipal waste handling, electrical safety, and e-waste disposal regulations, which were not explored in this prototype. Fragmented governance and rigid contracts hinder implementation. Waste management involves multiple contractors, making city-wide integration difficult. Mitigation: In the future, pilot testing can be conducted directly at an active dumping site using real mixed waste in real world setting and complying with regulations. Stronger coordination²⁴ and more flexible contracting (to accommodate new technologies) are needed, alongside workforce development via training programs or industry-academic partnerships to build technical skills.

Code, Model Availability, and Reproducibility Statement

This study outlines a replicable workflow for developing and testing a prototype “Robotic waste sorting system using machine learning classification techniques”. Key elements of the project, including dataset preparation (class definitions, image collection, labeling methodology, and the 80/10/10 train-validation-test split), are fully reproducible using standard tools. Model training steps—such as input resolution, data augmentation techniques, and evaluation metrics—are also described to support replication using equivalent platforms.

The computer vision model was developed using Google Teachable Machine, a cloud-based tool that does not offer public access to model weights or training code. As such, the exact model instance cannot be distributed. Additionally, certain firmware components and control APIs of the commercial 3-DOF robotic arm used in this prototype are proprietary and not open-source.

To support transparency, all dataset summaries, model configuration details, and performance metrics are provided in this paper. Preprocessing scripts and evaluation code used for

analysis can be shared upon reasonable request to the author. This ensures the scientific methodology is reproducible, while acknowledging limitations due to platform-dependent components.

Conclusion

This study developed and evaluated a robotic waste sorting system using machine learning classification techniques to automate waste segregation in a controlled environment. The goal was to reduce human exposure to hazardous materials and improve the recovery of recyclable materials. The system integrates a vision model (trained on 10,000 labeled images, achieving over 90% classification accuracy) with a 3-degree-of-freedom robotic arm to autonomously sort four categories of waste.

Performance tests demonstrated the prototype's ability to process approximately 55 items per minute with a pick-and-place success rate of around 92.5%. The expected benefits include enhanced automation of waste segregation, improved recycling efficiency, reduced landfill burden, and increased worker safety. Although the current system operates in a lab-controlled environment, the prototype lays a foundation for future deployment in real-world conditions.

In the next phase, the system can be adapted for pilot trials at active waste processing sites. These trials will provide valuable operational insights for refining the model's robustness, speed, and ability to handle greater variety in waste.

In summary, the proposed system presents a practical and scalable approach to sustainable waste management. By combining accessible machine learning tools with robotics, it offers a replicable model for cities aiming to improve material recovery and reduce environmental impact.

Supplementary Materials

To support transparency and reproducibility, the following supplementary materials will be provided with the submission. No personally identifying information is included in these materials.

- `classes.txt` — List of class labels: plastic, paper, metal, glass
- `Dataset split_10000_images.csv` — 80/10/10 split of images into train-validation-test.
- `Eval_accuracy_1000_images.csv` — Vision agent accuracy on the evaluation subset.
- `Mechanical_latency_200_picks` — Mechanical latency and pick success

- `timing_log.csv` — Per-pick timings captured during bench tests
- `images_sample.zip` — A small sample of labeled images (Plastic, Paper, Metal, Glass) to illustrate dataset format.

Note: The full dataset of 10,000 labeled images is available upon reasonable request due to file size limitations.

References

- 1 A. Kanoi, *An analysis of technologies, in order to move to a sustainable waste management system in the context of Mumbai*, NHSJS, 2023, <https://nhsjs.com/2023/an-analysis-of-technologies-in-order-to-move-to-a-sustainable-waste-management-system-in-the-context-of-mumbai/>.
- 2 F. Lopes, *Why methane emissions from landfills are concerning*, IndiaSpends, 2023, <https://www.indiaspend.com/earthcheck/why-methane-emissions-from-landfills-are-concerning-848571>.
- 3 S. Mor, K. Kaur and K. Ravindra, *Frontiers in Sustainable Cities*, 2024.
- 4 R. Jakhar and P. Sachar, *International Journal of Recent Technology and Engineering*, 2023, **12**, 110–120.
- 5 G. Vinti, V. Bauza, T. Clasen, K. Medlicott, T. Tudor, C. Zurbrügg and M. Vaccari, *International Journal of Environmental Research and Public Health*, 2021, **18**, 4331.
- 6 S. K. Singh, P. S. Salve and M. Godfrey, *Clinical Epidemiology and Global Health*, 2021, **9**, 34–40.
- 7 D. Kannan, S. Khademolqorani, N. Janatyan and S. Alavi, *Waste Management*, 2024, **174**, 1–14.
- 8 J. Lahoti, J. Sn, M. V. Krishna, M. Prasad, R. BS, N. Mysore and J. S. Nayak, *PeerJ Computer Science*, 2024, **10**, e1957.
- 9 B. Fang, J. Yu, Z. Chen, A. I. Osman, M. Farghali, I. Ihara, E. H. Hamza, D. W. Rooney and P. S. Yap, *Environmental Chemistry Letters*, 2023, **21**, 1959–1989.
- 10 T. Cheng, D. Kojima, H. Hu, H. Onoda and A. H. Pandiyaswargo, *Sustainability*, 2024, **16**, 10155.
- 11 W. Xiao, J. Yang, H. Fang, J. Zhuang, Y. Ku and X. Zhang, *Clean Technologies and Environmental Policy*, 2020, **22**, 1829–1841.
- 12 R. Sarc, A. Curtis, L. Kandlbauer, K. Khodier, K. E. Lorber and R. Pomberger, *Waste Management*, 2019, **95**, 476–492.
- 13 M. M. Islam, S. M. M. Hasan, M. R. Hossain, M. P. Uddin and M. A. Mamun, *PLOS ONE*, 2025, **20**, e0324294.
- 14 A. G. Satav, S. Kubade, C. Amrutkar, G. Arya and A. Pawar, *International Journal on Interactive Design and Manufacturing*, 2023, **17**, 2789–2806.
- 15 G. Ahmad, F. M. Aleem, T. Alyas, Q. Abbas, W. Nawaz, T. M. Ghazal, A. Aziz, S. Aleem, N. Tabassum and A. M. Ibrahim, *Scientific Reports*, 2025, **15**, 27078.
- 16 D. B. Olawade, O. Fapohunda, O. Z. Wada, O. S. Usman, A. O. Ige, O. Ajisafe and B. I. Oladapo, *Waste Management Bulletin*, 2024, **2**, 244–263.
- 17 S. Shahab, M. Anjum and M. S. Umar, *International Journal of Advanced Computer Science and Applications*, 2022, **13**, 381–395.
- 18 Evreka, *AI-Powered Waste Sorting: The Future of Recycling Efficiency*, Evreka Blog, 2025, <https://evreka.co/blog/ai-powered-waste-sorting-the-future-of-recycling-efficiency/>.
- 19 H. Wilts, B. R. Garcia, R. G. Garlito, L. S. Gomez and E. G. Prieto, *Resources*, 2021, **10**, 28.
- 20 M. A. A. Bin Daej and P. Alexandridis, *Recycling*, 2024, **9**, 59.

-
- 21 M. Abdallah, M. A. Talib, S. Feroz, Q. Nasir, H. Abdalla and B. Mahfood, *Waste Management*, 2020, **109**, 231–246.
 - 22 S. K. Singh, A. Chhabra and J. Arora, *Archives of Computational Methods in Engineering*, 2024, **31**, 1255–1276.
 - 23 CPCB, *Annual report 2020-21 on implementation of solid waste management rules*, Central pollution control board technical report, 2021.
 - 24 A. Lakhout, *Results in Engineering*, 2025, **25**, 104018.