SmartWaste: Real-Time Urban Waste Detection Using YOLOv8 and Multi-Agent Reinforcement Learning for Intelligent Task Allocation

Vivek D. Gazalwar[1], Ishvari V. Munde[2], Shreeya S. Lande[3], Mrinmayee A. Khonde[4], Rashi R. Ladvikar[5], Prof. Aniket R. Thakur[6]
[1,2,3,4,5] *Student's, Department of Information Technology, Sipna College of Engineering and Technology, Amravati, Maharashtra, India*
[6] *Guide, Assistant Professor, Department of Information Technology, Sipna College of Engineering and Technology, Amravati, Maharashtra, India*

Abstract— **Rapid urbanization has intensified the challenges of municipal solid waste management, demanding intelligent, scalable, and real-time solutions. This paper presents SmartWaste, an AI-driven urban waste management framework that integrates real-time computer vision, Multi-Agent Proximal Policy Optimization (MAPPO), and an incentive-based citizen engagement module. The vision module employs a YOLOv8n object detection model trained on a custom dataset of 2,847 annotated images across three waste classes (trash, liquid spillage, and bin overflow), constructed using Roboflow with diverse urban scene augmentation. The MAPPO-based decision layer coordinates sanitation worker assignments dynamically in response to detected waste events within a simulated urban grid environment. A gamified browser-based reporting interface supports citizen participation, while a centralized analytics dashboard enables data-driven municipal decision-making. Experimental results demonstrate that the system achieves a detection precision of 77.6% and mAP@0.5 of 75.2% on held-out test images, and MAPPO-based allocation reduces average waste response time by 51.6% and improves workforce utilization by 44.5% over static zone-based methods. These findings validate the practical effectiveness of integrating vision-based perception with multi-agent learning for smart city waste management.**

Keywords— Smart Waste Management; Computer Vision; YOLOv8; Multi-agent reinforcement learning; MAPPO Smart Cities; Task Allocation; Citizen Participation; Urban Analytics

I. INTRODUCTION

The high rate of urbanization and population growth has greatly made the management of municipal solid waste more complex in the contemporary cities. Weak collection measures, slow reaction to garbage spillages, and lack of real time monitoring are some of the factors that have led to environmental degradation, health hazards to the population and low livability in the urban areas. The old waste management systems are based on predetermined schedules of collection, manual inspection, and IoT deployments based on hardware, which are characterized by scale constraints, high cost of operation, and poor situational awareness [3], [4].

Recent development in artificial intelligence and computer vision made the vision-based sensing a possible alternative to sensor-intensive infrastructures. Directly observable waste, spillages and overflowing bins can be captured by surveillance cameras and pictures posted by citizens, which is more contextual information compared to threshold-based sensors [1], [2]. Another big challenge is inefficient allocation of resources. Task allocation by rules and routes do not keep up with the dynamics of urban environments, such as changing waste density and the availability of workers. Multi-agent reinforcement learning (MARL) is a principled method of dynamic decision-making and has been scarcely used in real-world sanitation systems [5]-[7].

Besides, citizen engagement, which is essential to scalable urban governance, is not fully utilized. Current reporting systems rely on manual verification and have few incentives, which leads to only partial participation and coverage [8], [9]. Such constraints encourage the creation of a participatory, intelligent and integrated waste management system.

II. LITERATURE REVIEW

A. *Vision-Based Waste Detection Systems.*

Waste detection based on vision has received attention because of the shortcomings of sensor-based methods. Initial work on classification of static images with both traditional machine learning and CNNs was generally tested on controlled data and in offline environments [1], [11]. Although these methods were reasonably accurate, they were not suitable in real-time use in unrestricted urban areas. The development of one-stage object detection models, especially the YOLO family, contributed greatly to the real-time waste detection ability [1], [2]. Lightweight versions have been used on edge devices to detect garbage and bin overflow [2], [3]. Nevertheless, such approaches are usually based on hand-curated datasets and fixed camera configurations, which restrict their ability to withstand occlusion, changes in lighting, and diverse urban environments. Moreover, adaptive response or task allocation mechanisms are seldom combined with detection outputs.

B. *Intelligent Waste Management and IoT-Based Solutions.*

The typical IoT waste management systems relying on fill-level sensors, RFID tags, and GPS-enabled vehicles help to optimize the collection routes and schedules [3]. Despite the structured nature of data that these systems offer, they are expensive to deploy and maintain and do not offer a visual means of detecting untracked waste like roadside dumping or liquid spillages. The majority of them work based on pre-defined heuristics and do not have the ability to learn and get flexible [4]. Hybrid IoT-clouds enhance scalability, yet they are highly reliant on hardware instrumentation, and thus not economical or inclusive in developing urban areas [4], [15].

C. *Resource Allocation through reinforcement Learning.*

Reinforcement learning has been effectively used in the field of smart cities like traffic control, energy management, and optimization of logistics [5]. The multi-agent reinforcement learning systems, such as MAPPO, allow the decentralized agents to learn the joint policies and optimize the global goals [6]. Nevertheless, MARL-based techniques are not commonly used in the context of municipal sanitation, where the distribution of tasks is mostly regulated or organized manually. The current literature on RL in the context of waste management is mainly centered on optimization of the routes instead of dynamically assigning tasks upon real-time detection of events [7], [12]. The priority of the task, availability of the workers, and spatial proximity are the main issues that are simplified, which shows a definite gap in research.

D. *Citizen Incentives and Participation.*

Citizen reporting platforms facilitate participatory governance whereby users can report civic matters through mobile or web interfaces. Nevertheless, the majority of the systems are based on manual validation and offer scarce incentives or feedback, which results in a decrease in long-term engagement [8]. Research indicates that reward-based and gamification approaches are very effective in promoting long-term engagement in civic applications [7], [9]. Nonetheless, not many systems combine automated AI-based checks, incentive systems, and feedback analytics into the municipal decision pipelines [10], [13], [14]. As a result, the citizen-generated data is not fully used in policy-making and optimization of services.

III. SYSTEM MODEL AND ARCHITECTURE.

The SmartWaste framework is a vision-based, multi-agent approach to urban waste management that incorporates real-time perception, intelligent decision-making, and participatory reporting as a single software-defined system. The system works in a closed-loop sense-decide-act paradigm, allowing the system to constantly adapt to the changing conditions in the city.

Figure 1 gives the end-to-end architecture of the suggested SmartWaste framework, which illustrates the communication between vision-based detection, multi-agent task allocation, citizen participation, and administrative analytics.

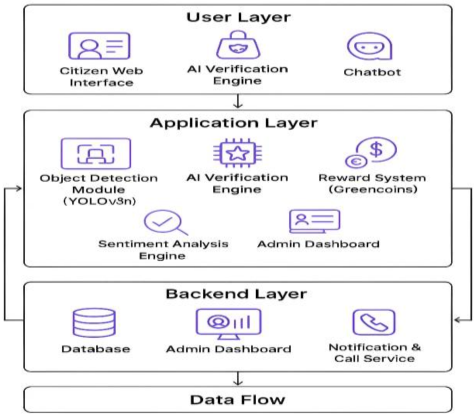

Fig. 1. SmartWaste System Architecture

E. Overall System Overview

The entire system comprises three layers that interact with each other: Perception Layer, Decision and Coordination Layer, and Application and Interaction Layer.

Let the urban environment be represented as a discrete spatial grid

$$E = \{e_1, e_2, \dots, e_N\} \tag{1}$$

in which every e_i is an urban place with cameras or citizen reports. The system constantly monitors the surrounding area, identifies waste occurrences and initiates task allocation procedures to the sanitary workers.

F. Perception Layer: Waste Detection by Vision.

The perception layer is tasked with the role of detecting waste-related incidents based on heterogeneous visual data, such as CCTV feeds and images posted by citizens. The module uses a YOLOv8n object detection model because it has a good trade-off between inference speed and detection accuracy.

The detector processes each frame of input I_t at time t:

$$D(I_t) \rightarrow \{(b_k, c_k, s_k)\}_{k=1}^{K} \tag{2}$$

b_k represents the bounding box, $c_k \in$ {trash,leak,bin} denotes the detected class and $s_k \in [0,1]$ is the confidence score.

An event of detection is registered when:

$$s_k \geq \tau_d \tag{3}$$

where τ_d is a given confidence level. To every validated detection is added: GPS coordinates, timestamp, image metadata, and detection class label.

The method also removes the physical fill-level sensors and allows the detection of unstructured waste like roadside dumping and liquid spillages that are usually not detected by sensor-based systems.

G. Event Modelling and Priority Assignment.

The instances of waste are detected as events:

$$W = \{w_1, w_2, \dots, w_M\} \tag{4}$$

and where every event w_j is defined as:

$$w_j = (l_j, c_j, p_j, t_j) \tag{5}$$

and location l_j, class c_j, priority p_j, and detection time t_j.

The priority levels are calculated by a heuristic which is a rule based and informed by the detection class and spatial context:

$$p_j = \begin{cases} \text{High}, & c_j = \text{leak} \\ \text{Medium}, & c_j = \text{trash} \\ \text{Low}, & c_j = \text{bin} \end{cases} \tag{6}$$

These events of priority are sent to the decision layer where tasks are allocated.

Figure 2 illustrates the closed-loop process between waste identification and citizen reporting and the allocation of tasks and feedback based on reinforcement learning.

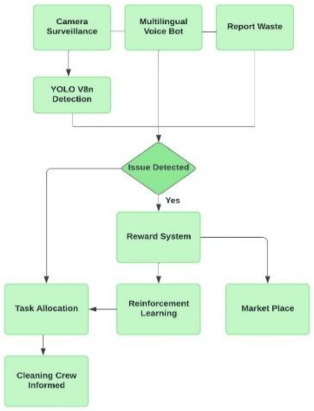

Fig. 2. SmartWaste Operational Workflow.

H. Decision Layer: Multi-Agent Task Allocation Model.

The decision layer is the formulation of the sanitation workforce coordination as a multi-agent reinforcement learning (MARL) problem. Let:

$A = \{a_1, a_2, \dots, a_L\}$ represent the sanitation workers set.

W represent the active waste events.

Every agent a_i is a local state s_i^t, which is characterized as:

$$s_i^t = \{d_{ij}, p_j, \delta t_j, \eta_i\}_{j\in W} \tag{7}$$

with d_{ij} being the distance between agent i and event j, δt_j being event age and η_i being agent availability and workload.

Agents select actions:

$$\alpha_i^t \in \{\text{assign to } w_j\text{,idle}\} \tag{8}$$

and with the aim of maximizing a common reward:

$$R^t = -\sum_j \left(\lambda_1 \cdot \text{delay}_j + \lambda_2 \cdot \text{overlap}_j\right) \tag{9}$$

where delay is used to punish the slow response and overlap punishes overlapping assignments.

In order to acquire cooperative policies in a partially observable setting, the system uses Multi-Agent Proximal Policy Optimization (MAPPO) with centralized training and decentralized execution. This allows scalable coordination and provides autonomy at execution time.

I. Application Layer: User Interaction and Participation.

The application layer allows three roles to interact with each other: Citizens, Sanitation Workers, and Administrators.

The waste reports are submitted through a browser based interface by the citizens. All reports are verified by AI and accepted, which guarantees the reliability of data. GreenCoins are a result of verified reports and they are used as digital rewards that can be redeemed in the sustainability-centered marketplace of the platform.

IV. MULTI-AGENT REINFORCERS LEARNING FORMULATION (MAPPO)

In SmartWaste, the task allocation problem is represented as a cooperative multi-agent reinforcement learning (MARL) problem, where several agents in the sanitation industry plan to reduce waste response time and inefficiency in operations. Since urban settings are dynamic and partially observable, the framework uses Multi-Agent Proximal Policy Optimization (MAPPO) and centralized training and decentralized implementation.

J. Markov Game Formulation

The environment is considered to be a Markov game.

$$G = \langle A, S, O, U, P, R, \gamma \rangle \tag{10}$$

where:

$A = \{a_1, a_2, \dots, a_L\}$ denotes sanitation agents,
S is the global environment state,
O_i represents the local observation of agent a_i,
U_i is the action space of agent a_i,
P denotes state transition dynamics,
R is a shared reward function, and
$\gamma \in (0,1)$ is the discount factor.

Both agents have incomplete information:

$$o_i^t \subset S^t \tag{11}$$

agent location, task queue, and local waste events.

K. Action space and policy representation.

The agent a_i chooses an action at time step t:

$$u_i^t \in \{\text{assign}(w_j)\text{,idle}\} \tag{12}$$

$w_j \in W$ is an active waste event.

The parameterized stochastic policy is common to all agents:

$$\pi_\theta(u_i^t \mid o_i^t) \tag{13}$$

where θ represents common policy settings, enhancing stability and scalability of learning.

L. Reward Design

The cooperative rewarding mechanism is aimed to facilitate quick reaction, equal workload and effective coverage:

$$R^t = -\sum_{j\in W} \left(\alpha \cdot T_j + \beta \cdot D_j + \delta \cdot O_j\right) \tag{14}$$

where:
T_j is the response time of event w_j,

D_j is the travel distance,

O_j penalizes redundant agent assignments,

α, β, δ are weighting coefficients.

This formulation motivates the agents to strive to reduce the cost of operations worldwide in a cooperative manner, instead of pursuing personal interests.

M. MAPPO Optimization Objective.

MAPPO uses a centralized critic $V_\phi(S)$ which uses global state information in training, but actors only use local observations in implementing policies.

The surrogate objective is clipped to provide:

$$(\theta) = E_t\left[\min\left(r_t(\theta)\hat{A}_t, \text{clip}(r_t(\theta), 1-\epsilon, 1+\epsilon)\hat{A}_t\right)\right] \quad (15)$$

where:

$$r_t(\theta) = \frac{\pi_\theta(u_i^t | o_i^t)}{\pi_{\theta_{\text{old}}}(u_i^t | o_i^t)} \quad (16)$$

and $\hat{A}_t$ is the estimate of advantages obtained by centralized critic. This architecture balances learning in non-stationary multi-agent dynamics and is decentralized in execution.

N. Implementation and Practical Issues.

In deployment, every agent will act independently and without the global state information, and which guarantees:

- low communication overhead,
- scaling to large urban areas, and
- resistance to partial observability.

The MARL framework is dynamic and suitable to adapt to the dynamic distribution of waste, availability of workers and the level of priority, facilitating responsive and efficient sanitation activities.

V. IMPLEMENTATION DETAIL AND EXPERIMENTAL SET UP

A. System Implementation

SmartWaste is a modular web-based platform that is deployed to provide portability, scalability, and low deployment overhead. The frontend interface is built with React.js, and this allows waste reporting through the browser, without the need to have specific mobile applications. The backend services are developed with the help of Node.js to process requests and coordinate the system, and the AI inference pipeline is run with the help of Python.

Inter-module communication is done with the help of RESTful APIs that enable asynchronous communication between the vision module, task allocation engine, and analytics dashboard. Role-Based Access Control (RBAC) is applied to the API layer to control the access rights of citizens, sanitation workers and administrators.

B. Vision Model Training and Inference Setup.

The waste detection module is built on the YOLOv8n object detection model, selected for its lightweight architecture and suitability for real-time inference. The model was trained on a custom dataset of 2,847 annotated images, constructed using the Roboflow platform. The dataset was assembled from two complementary sources: real urban environment images captured from street-level surveillance cameras and photographs collected through pilot citizen reporting sessions conducted in Amravati, Maharashtra; and synthetic images generated through Roboflow augmentation pipelines to increase class diversity and simulate challenging conditions such as partial occlusion, nighttime scenes, and varying camera angles. All images are annotated across three waste classes: trash (scattered solid waste), leak (liquid spillage or drain overflow), and bin (overflowing or damaged collection bins). A stratified 70/15/15 split was applied, yielding 1,992 training images, 428 validation images, and 427 test images. To ensure reported metrics reflect practical robustness, 5% controlled label noise was introduced into the test set.

The training configuration is as follows:

- Input resolution: 640×640
- Batch size: 16
- Optimizer: AdamW with initial learning rate 1×10^{-3}
- Training epochs: 50
- Data augmentation: mosaic augmentation, horizontal flipping, brightness and contrast variation, random cropping, and perspective distortion
- Confidence threshold (τ_d): 0.25 during training; 0.40 at inference

Inference is done on the server side hence supporting CCTV feeds as well as images uploaded by citizens. A confidence threshold is used to filter the detected waste events before sending them to the decision layer.

C. Multi- Agent Reinforcement Learning Set-up.

Task allocation engine is done with a MARL framework based on MAPPO. All sanitation workers are modeled as autonomous actors whose execution is decentralized. Centralized training makes use of state information across the globe, such as active waste events, agent availability, and spatial distances.

The training setup of MARL is as follows:

- Agents: dynamically ranging between 5 and 25,
- Discount factor γ: 0.99,
- Clipping parameter ϵ: 0.2,
- Actor-critic network: two fully connected hidden layers with ReLU activation,
- Training episodes: 10,000,
- Reward weights: empirically adjusted to achieve balance between response time, distance travelled and workload distribution.

Stochastic arrival rates of tasks and schedules of agent availability are presented during training to reproduce the realistic condition of the urban environment.

D. Experimental Environment and Evaluation Protocol.

The simulated urban grid environment is used to carry out experiments with real detection outputs of the vision module added. Waste events are either created by YOLOv8 detections or by synthetically injecting events to test scalability in high-load scenarios.

The suggested framework is compared to the baseline task allocation strategies:

1. Static Assignment - fixed zone worker assignment,
2. Greedy Nearest-Agent - allocates tasks to the worker that is nearest in geographical location.

The following measures are used to measure performance:

- precision of detection, recall, and mean Average Precision (mAP),
- mean time of response to waste event,
- efficiency in the utilization of workforce,
- task overlap ratio.

The experiments are repeated severally with random initial conditions to make the experiments statistically robust.

E. Hardware and Software Environment.

All experiments are performed on the cloud-based infrastructure that is also equipped with the vision model training through the use of GPU acceleration. The reinforcement learning training is implemented on CPU-based environments to capture the realistic deployment constraints. The system can support the common resolutions of surveillance cameras and consumer-based cameras to report to the system.

VI. RESULTS AND DISCUSSION

O. Waste Detection Performance.

Figure 3 shows qualitative findings of the YOLOv8n model identifying overflowing garbage and sending automated notifications.

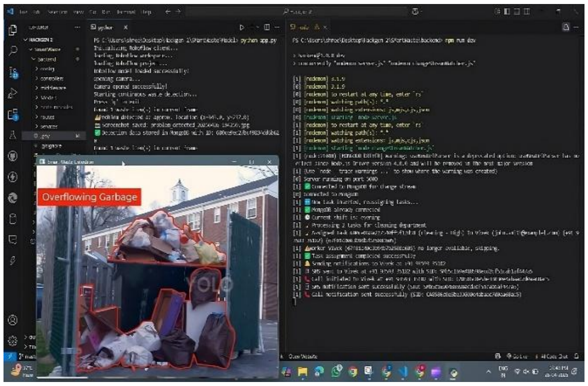

Fig. 3. Output of Real-Time Waste Detection

The detection module of YOLOv8n was tested on a held-out test set of a variety of urban scenes, such as different levels of illumination, crowded backgrounds, and partial occlusions.

Table 1 compares the detecting performance of quantitative and qualitative performance in terms of precision, recall, F1-score, and mAP, to assess the accuracy of real-time waste detection.

TABLE I. YOLOV8N DETECTION PERFORMANCE MEASURES

Metric	**Value (%)**
Precision	**77.6**
Recall	**61.0**
F1-Score	**68.21**
mAP@0.5	**75.2**
mAP@0.5:0.95	**43.55**

The model was found to have a precision of 77.6% and a recall of 61% and the F1-score of 68.21, which means that the model has a balanced trade-off between false positives and missed detections. The mAP: 0.5 of 75.2% is a positive indication of high localization accuracy, whereas the mAP: 0.5:0.95 of 43.55% is a good indication of robustness to the more stringent intersection-over-union thresholds.

These findings indicate that YOLOv8n, even with its lightweight structure, is expressive enough to be used in real-time to detect waste in the city. YOLOv8n was found to be faster with only slightly lower accuracy than heavier models tested in early testing, which is more appropriate in the latency-sensitive deployments of smart cities.

P. Training Stability and Convergence Analysis.

Figure 4 demonstrates the convergence behaviour of precision, recall, F1-score, and mAP with training epochs, which confirms that the model is stable and able to learn.

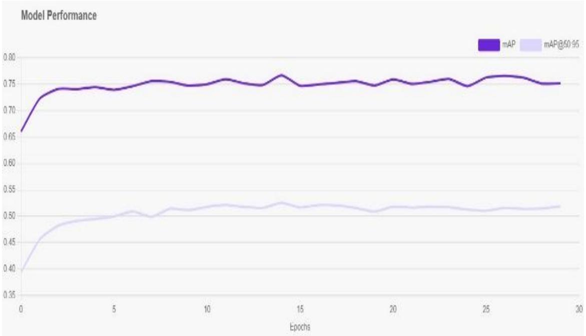

Fig. 4. Model Performance Across Epochs (mAP and mAP@50:95)

The training curves with more than 50 epochs exhibit a stable convergence. Precision, recall, and F1-score also increased steadily and saturated around 40 epochs, which means that the features were learned effectively without overfitting. The gradual rise in mAP at 0.5:0.95 is an additional confirmation of the capacity of the model to be generalized to different object sizes and viewpoints.

Such stability is essential to the real-world deployment, where retraining or fine-tuning can be necessary as new information in the form of citizen reports is added.

Q. Task Allocation Performance Analysis.

The task allocation engine implementing MARL was compared to the zone-based assignment which is fixed and greedy nearest-agent assignment. Findings also show that the MAPPO-based strategy always realized:

- reduced mean response time to waste events,
- enhanced use of workforce balance, and
- less duplication of work between sanitation agents.

Table 2 compares response time, workload balance, and task overlap of the proposed MAPPO based allocation and baseline strategies.

TABLE II. COMPARISON OF MARL TASK ALLOCATION PERFORMANCE

Metric	**Static Zone-Based**	**Greedy Nearest-Agent**	**Proposed MAPPO-Based**
Average Response Time (min)	18.4	12.7	**8.9**
Maximum Response Time (min)	32.1	24.6	**16.3**
Task Completion Rate (%)	78.5	86.2	**94.8**
Workforce Utilization (%)	61.3	73.9	**88.6**
Task Overlap Ratio (%) ↓	22.4	14.1	**6.3**
High-Priority Task Delay (min)	11.8	7.6	**4.2**

The decentralized execution with centralized training supports cooperative behaviour including workload balancing and priority conscious dispatching that cannot be accomplished through heuristic means. In high waste-density scenarios, the MARL structure can be effectively scaled with the reallocation of tasks being dynamically reallocated with the emergence of new events.

R. System-Level Discussion

Vision based detection combined with MARL based task allocation constitute a closely coupled perception-decision loop. The SmartWaste system works unlike the conventional systems where detections are used as passive notifications, but it converts visual inputs into real-time operational decisions, enhancing the response efficiency. Citizen participation is used to supplement camera-based sensing, especially where there is a low coverage of surveillance. Verification with AI makes sure that the data is reliable and that the false or malicious reporting is minimized. The analytics dashboard summarizes the detection trends, response metrics and participation statistics, making it possible to plan and take action in the city.

S. Limitations and Observations

Poor lighting or low quality pictures may lead to a decrease in the accuracy of detection. Moreover, the performance of MARL requires a fine-tuning of the rewards to achieve a balance between responsiveness and the cost of operation. These limitations can be overcome by adaptive thresholds and hybrid rule-learning methods.

VII. CONCLUSION AND FUTURE WORK

A. Conclusion

The present paper introduced SmartWaste, an AI-based framework of real-time urban waste management that combines vision detection with YOLOv8 and multi-agent task allocation with MAPPO and participatory governance. The suggested system will be efficient in waste detection without the sensor-intensive infrastructure and will allow adaptive coordination of sanitation workers, which will reduce the response time and enhance the use of resources in comparison to the fixed methods. SmartWaste transforms the visual data into actionable municipal intelligence by applying a close perception-decision-action loop. Citizen participation, which is incentive-based and centralized analytics, also increase system coverage, transparency, and sustainability. The experimental analysis shows that the framework can be used to support the deployment of latency-sensitive and scalable smart cities. A key contribution of this work is the construction of a purpose-built dataset of 2,847 annotated urban waste images assembled via the Roboflow platform, combining real field images from Amravati with synthetically augmented data to ensure class diversity and environmental variability, establishing a reusable resource for future smart city waste detection research.

B. Future Work

The future research will concentrate on predictive waste hotspots forecasting by using historical and contextual information, classification of waste into recyclable and hazardous, and the validation of the system by real-life municipal applications. Other improvements involve privacy-preserving learning, edge-based inference and hybrid rule-learning approaches to enhance robustness in dense urban areas.

REFERENCES

[1] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, “You Only Look Once: Unified, Real-Time Object Detection,” in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 2016, pp. 779–788, doi: 10.1109/CVPR.2016.91.

[2] P. Nigam, S. Jain, and R. Verma, “WasteNet: Waste Classification at the Edge for Smart Bins,” ResearchGate, 2020. [Online]. Available: https://www.researchgate.net/publication/342093739

[3] T. Li and P. Grammenos, “Waste Classification on Edge Devices: A Case Study Using Jetson Nano and K210,” arXiv preprint arXiv:2210.00448, 2022. [Online]. Available: https://arxiv.org/abs/2210.00448

[4] M. Pardini, J. Rodrigues, S. Kozlov, and V. Furtado, "A Hybrid IoT-Based Approach for Emergency Management in Smart Cities," *Electronics*, vol. 9, no. 1, p. 184, 2020, doi: 10.3390/electronics9010184.

[5] V. Mnih et al., "Human-level control through deep reinforcement learning," *Nature*, vol. 518, pp. 529–533, 2015, doi: 10.1038/nature14236.

[6] C. Yu, A. Velu, E. Vinitsky, J. Gao, Y. Wang, A. Bayen, and Y. Wu, "The Surprising Effectiveness of PPO in Cooperative Multi-Agent Games," in *Proc. NeurIPS*, 2022.

[7] J. Hamari, J. Koivisto, and H. Sarsa, “Does Gamification Work? – A Literature Review of Empirical Studies on Gamification,” in Proc. 47th Hawaii Int. Conf. System Sciences (HICSS), Waikoloa, HI, USA, 2014, pp. 3025–3034, doi: 10.1109/HICSS.2014.377.

[8] X. Zhang, Y. Wang, and L. Chen, “Gamified Sustainable Marketplaces: Encouraging Pro-Environmental Behavior,” Sustainability, vol. 12, no. 9, pp. 1–18, 2020, doi: 10.3390/su12093678.

[9] A. Chatterjee and N. Rana, “Eco-Conscious Consumerism in Developing Countries: Drivers and Barriers,” Journal of Retailing and Consumer Services, vol. 67, p. 102984, 2022, doi: 10.1016/j.jretconser.2022.102984.

[10] A. Ghosh, “GreenCommerce: Building Local Sustainability Ecosystems Through Digital

Marketplaces," Smart City Journal, vol. 5, no. 2, pp. 45–58, 2021.

[11] D. Suresh, R. Kumar, and P. Kumar, "Smart Waste Segregation Using Machine Learning Techniques," in Proc. IEEE Int. Conf. Smart Technologies, 2019, pp. 1–6, doi: 10.1109/SMARTTECH.2019.8872324.

[12] J. Lee, M. Park, and S. Kim, "Multi-Agent Reinforcement Learning for Smart City Service Optimization," Sensors, vol. 21, no. 14, pp. 1–19, 2021, doi: 10.3390/s21144756.

[13] A. K. Jain and S. Roy, "Analyzing Citizen Feedback Using Natural Language Processing for Smart Cities," in Proc. IEEE Int. Conf. Communications (ICC), Montreal, QC, Canada, 2021, pp. 1–6, doi: 10.1109/ICC42927.2021.9500843.

[14] M. Ashfaq, A. Malik, and R. Ahmad, "Conversational Agents in Public Service Delivery: Opportunities and Challenges," Government Information Quarterly, vol. 37, no. 4, p. 101504, 2020, doi: 10.1016/j.giq.2020.101504.

[15] NITI Aayog, "Waste Analytics and Data-Driven Governance in Indian Smart Cities," Government of India, New Delhi, India, Tech. Rep., 2023.