

AI-Based Camera Systems for Roadside Litter Detection and Offender Identification

Bani James

Gokul T S

Hari Sankar

Jiju Abraham

Sukanya M V

Department of Computer Science and Engineering
College of Engineering Kottarakkara, Kerala, India

Abstract - Urban littering continues to be a persistent environmental and civic challenge, adversely affecting public health, sanitation systems, and urban aesthetics. Traditional litter monitoring mechanisms largely rely on manual surveillance, post-incident cleaning, and punitive enforcement, which are inefficient, resource-intensive and limited in scalability. In recent years, advances in artificial intelligence, particularly in computer vision and deep learning, have enabled the development of automated systems for real-time litter detection and monitoring. The proposed system integrates YOLOv8-based object detection, face detection, and OCR-based license plate recognition to identify littering events and associated offenders in real time. A Streamlit-based dashboard enables live monitoring and incident management, while a Telegram-based alert mechanism provides timely notifications with visual evidence. Experimental observations demonstrate the practical feasibility of the system for real-world roadside surveillance. The framework emphasizes deployment feasibility, scalability, and ethical considerations for smart city applications.

Index Terms- Litter Detection, Computer Vision, Deep Learning, YOLO, Smart Cities, Ethical AI

I. INTRODUCTION

Solid waste generation has significantly increased in modern cities due to rapid urbanization, population growth, and shifting consumption patterns. Roadside littering is one of the many types of urban pollution that has become a persistent and noticeable issue that has an immediate impact on public health and environmental quality. Litter buildup damages urban aesthetics, hinders drainage systems, causes flooding during periods of high precipitation, and fosters the growth of pathogenic organisms. Roadside waste management has grown more difficult for local authorities as cities continue to grow.

Due to shortcomings in current monitoring and enforcement systems, littering persists despite ongoing public awareness campaigns and legal frameworks. Manual inspections, citizen reporting, and sporadic cleanup operations are major components of traditional surveillance methods. These techniques are mostly reactive, dealing with litter only after it has accumulated, and they offer little understanding of the temporal and spatial patterns of littering behavior. Because of this, authorities frequently lack the actionable data required to carry out focused and preventive interventions.

Automated technologies for urban environmental monitoring have become more popular in recent years due to the rise of smart city initiatives. Particularly, computer vision and artificial intelligence (AI) methods have demonstrated great promise in converting traditional surveillance systems into data-driven, intelligent platforms. AI-based systems can facilitate proactive decision-making and lessen reliance on manual monitoring procedures by enabling continuous observation and automated analysis of visual data.

Deep learning developments have been crucial to this change. Littering incidents can now be identified as they happen thanks to cutting-edge object detection models, real-time video analytics, and optical character recognition technologies. These camera-based systems have the ability to recognize waste items, examine human or vehicle behavior, and, in some cases, link incidents to accountable people or cars. For efficient enforcement and deterrence, such capabilities signify a change from static waste detection to dynamic event-level monitoring.

This paper provides a thorough analysis of AI-driven roadside litter detection systems, looking at their underlying techniques, performance traits, and deployment issues. Limitations pertaining to scalability, real-world variability, and ethical issues like data governance and privacy are given special consideration. This review seeks to provide guidance for the development of large-scale, responsible, and preventive litter monitoring systems appropriate for future smart city environments by evaluating current research and identifying unmet gaps.

Table I summarizes representative studies and their limitations.

TABLE I
SUMMARY OF EXISTING RESEARCH WORKS

Author	Technique Used	Limitations
Smith et al.	CNN-based object detection	High false positives
Kumar et al.	YOLO-based CCTV analysis	Limited night accuracy
Lee et al.	LSTM-based behavior modeling	High computational cost
Patel et al.	Edge AI camera systems	Limited dataset size

II. RELATED WORK

Early studies (2015–2017) on automated litter detection focused on the use of early convolutional neural network (CNNs) and traditional computer vision techniques. The purpose of these studies was to identify visible waste materials in static photos taken in public areas, such as plastic bottles, paper, and cans. Although these methods showed some success in controlled settings, they performed much worse in real-world situations with occlusion, background clutter, and fluctuating lighting. Furthermore, rather than comprehending the act of littering itself, these systems were mainly restricted to detecting the presence of litter.

After 2018, researchers started using more complex CNN architectures to increase detection accuracy due to the quick development of deep learning techniques. Waste objects in urban and roadside settings were classified and localized using object detection frameworks like Single Shot MultiBox Detector (SSD) and Faster R-CNN. Although these models were more accurate than previous approaches, they needed a lot of processing power, which made them unsuitable for real-time roadside deployment.

YOLO (You Only Look Once)-based models became very popular in 2019 because of their ability to balance real-time processing power with detection accuracy. YOLOv3 and YOLOv4 were used in a number of studies to detect roadside litter using CCTV footage; these studies reported improved frame-per-second (FPS) performance and dependable detection under moderate traffic conditions. However, the majority of these studies did not include offender identification or behavioral context; instead, they concentrated only on waste object identification and classification.

In contrast to direct enforcement, later research conducted in 2020 and 2021 broadened the application scope to include environmental cleanliness assessment. In order to produce spatial litter distribution maps that aid in municipal decision-making and cleanup planning, these studies integrated litter detection with Geographic Information Systems (GIS). Although these systems offered useful diagnostic information, they were mainly reactive in nature and did not deal with littering behavior deterrence or real-time monitoring.

More recent research has tried to combine object detection with action recognition methods after realizing this limitation. In order to differentiate intentional littering from unintentional object drops, researchers developed hybrid models between 2021 and 2022 that combined CNN-based detectors with temporal models like Long Short-Term Memory (LSTM) networks and 3D CNNs. These methods greatly increased computational complexity and presented difficulties for edge-based deployment, despite improving event-level detection accuracy.

In recent years, vehicle-centric litter detection has also been investigated; in particular, littering incidents involving moving vehicles were examined starting in 2022. To identify offending vehicles, these systems frequently used Automated License Plate Recognition (ALPR) in conjunction with object

detection. Their performance was found to be sensitive to motion blur, camera angle, and nighttime lighting conditions, which limited their consistent real-world applicability despite their promising nature.

Since 2023, there has been an increase in interest in the use of embedded AI systems and edge computing. By carrying out inference locally on smart cameras or edge devices, edge-based processing lowers latency and bandwidth consumption. Lightweight YOLO variants on edge hardware can enable near real-time litter detection, according to several studies; however, trade-offs between model complexity, accuracy, and power consumption are still being investigated.

Despite these developments in technology, the literature still does not adequately address ethical and societal issues. Public acceptance, data governance, and privacy preservation are not sufficiently addressed by the majority of current systems, especially when offender identification mechanisms are involved. To guarantee that AI-based litter detection systems are efficient and socially conscious, recent review studies highlight the necessity of transparent system design, anonymization strategies, and regulatory compliance.

III. SYSTEM METHODOLOGY

A. System Overview

Video acquisition, object detection, offender identification, event validation, and alert generation are all part of the suggested framework's modular pipeline. In order to identify instances of littering and extract pertinent evidence, visual data collected from various sources is processed in real-time. Scalability, adaptability, and simplicity of deployment in a variety of surveillance settings are guaranteed by the modular design.

B. Input Sources and Data Acquisition

To increase deployment flexibility, the system supports several input modalities. These consist of RTSP-based IP camera streams, uploaded photos, recorded video files, and live webcam feeds. The system can function in both controlled and real-world settings thanks to its multi-source capability, which eliminates the need for specific hardware configurations.

Sequential processing of incoming video frames allows for continuous monitoring while preserving real-time performance. With optional GPU acceleration when available, the system can run effectively on CPU-only configurations and is built for Windows-based operating systems.

C. Object Detection Using YOLOv8

The main object detection component is a deep learning model based on YOLOv8. Within each frame, the model is in charge of locating pertinent entities like trash, pedestrians, cars, and license plates. YOLOv8 is chosen for real-time surveillance applications because it strikes a balance between inference speed and detection accuracy.

Bounding boxes and confidence scores are used to represent detected objects, and these are then further examined to identify possible littering incidents. The detection threshold

is set up to reduce false positives while preserving sensitivity to objects that are small or partially obscured.

D. Offender Identification Modules

To make it easier to identify offenders, the system includes modules for face detection and license plate recognition. Face detection uses Haarcascade-based classifiers to recognize and extract facial regions from detected pedestrians. This lightweight approach ensures both computational efficiency and compatibility with CPU-based execution.

OCR techniques are applied to license plate regions identified by the object detector for vehicle-based incidents. The main OCR engine used to extract alphanumeric characters from license plate photos is called PaddleOCR. This data is used as additional proof in car-related littering incidents.

E. Event-Based Littering Detection and Alert Logic

The system uses event-based logic to validate littering incidents instead of just static object detection. When waste is present and a nearby pedestrian or vehicle is present within a specified temporal window, it is considered a littering event. This tactic lessens false positives brought on by litter that already exists or unrelated objects in the area.

The system records visual evidence, such as the entire scene and cropped areas that correspond to waste, faces, or license plates, once a legitimate littering event is identified. Timestamp and detection confidence are examples of incident metadata that are recorded for later examination.

F. Alert Generation and Dashboard Interface

Real-time visualization of detected events, system status, and incident history is provided by a web-based dashboard created with Streamlit. The system automatically notifies users via a Telegram-based notification module when it verifies a littering incident. To facilitate quick response and review, alerts include visual proof and pertinent metadata.

This integrated methodology ensures that the proposed system not only detects littering events but also supports practical enforcement workflows while maintaining real-time performance and deployment feasibility.

G. Materials and Methods

Imaging hardware, computational platforms, datasets, and deep learning frameworks are the main components and techniques used in AI-based roadside litter detection systems. The majority of the reviewed studies make use of stationary roadside surveillance cameras with enough resolution and frame rate to record small litter items and related human or vehicle activity. To increase coverage, dashboard-mounted systems and mobile cameras are sometimes used.

The deployment strategy affects the computational infrastructure. While edge computing devices like embedded GPUs or AI accelerators are used for real-time inference, cloud-based processing is frequently utilized for large-scale analysis and model training. Edge-based deployment lowers network latency and bandwidth consumption, which is crucial for ongoing urban video surveillance.

Datasets are essential to the development of systems. The reviewed works use both custom-collected roadside video data and publicly available datasets. Common litter items like plastic bottles, wrappers, cans, and paper waste can be found in these datasets' annotated photos and video clips. Usually, object classes and, occasionally, temporal action boundaries for littering events are annotated by hand.

Deep learning models are the main part of the detection pipeline. For object detection, convolutional neural network architectures such as YOLO are frequently employed. Transfer learning, which starts models with weights pretrained on large-scale datasets and fine-tunes them on domain-specific litter datasets, is commonly used to increase convergence and accuracy.

Evaluation methods involve splitting datasets into training, validation, and testing subsets to ensure unbiased performance assessment. Cross-validation and ablation studies are often conducted to analyze the impact of model components, input resolution, and inference location on overall system performance.

IV. SYSTEM ARCHITECTURE AND IMPLEMENTATION

This section describes the architecture and implementation of the proposed AI-based camera system for roadside litter detection and offender identification. The system is designed to support real-time monitoring, evidence generation and alert notification in urban roadside environments.

A. System Architecture

- **Video Acquisition Layer:** Responsible for capturing visual data from multiple sources, including webcams, uploaded images, recorded videos, and RTSP-based IP camera streams.
- **Detection and Analysis Layer:** Performs object detection, face detection, and license plate recognition using deep learning and computer vision techniques.
- **Event Management Layer:** Validates littering events using event-based logic, generates visual evidence, and records incident metadata.
- **Application and Alert Layer:** Provides a web-based dashboard for monitoring and sends real-time alerts through a Telegram notification module.

This layered design allows individual components to be updated or extended without affecting the overall system functionality.

B. Detection Pipeline Implementation

Each incoming video frame is processed by the YOLOv8 object detection model to identify waste objects, pedestrians, vehicles, and license plates. Detected entities are represented using bounding boxes and confidence scores, which are further analyzed for event validation and offender identification.

Haarcascade classifiers are employed for face detection. When the system detects a face, it captures that portion of the picture and stores it as incident-related evidence.




Fig. 1. Overall architecture of the proposed AI-based camera system for roadside litter detection and offender identification.

The workflow is modified when a car appears. The OCR module receives the photos that the detector takes while searching for license plates. To extract the letters and numbers from each plate, PaddleOCR is utilized. We record the OCR's confidence scores during this process so that we are always aware of how well it is performing, even in situations where the lighting is difficult or the scene is chaotic.

C. Event Validation and Evidence Handling

The system cuts down on false positives by watching for actual events, not just whether an object is sitting there. It confirms someone's littering only if it spots trash at the same time as a person or car nearby, all within a specific window of time. This way, it tells the difference between someone actively tossing trash and garbage that's just been lying around.

Once a valid event is identified, the system automatically captures and stores the following evidence:

- Full-frame scene image
- Cropped waste region
- Cropped face image
- Cropped license plate image

Incident data, including timestamps, detection confidence values, and source information, are stored using a structured incident handling module for later review.

D. Web-Based Dashboard and Alert System

The main user interface for incident management and real-time monitoring is a web dashboard built on Streamlit. Live video feeds, detection overlays, recent incidents, and system status indicators are all shown on the dashboard.

The system incorporates a Telegram-based alert mechanism to facilitate quick response. Alerts with accompanying visual proof and metadata are automatically sent when a littering incident is verified. This design preserves an auditable incident trail while facilitating prompt intervention.

E. Experimental Setup and Evaluation

This section describes the experimental environment and evaluation methodology used to assess the practical performance of the proposed system.

- **Hardware and Software Environment:** A Windows-based platform was used for the system's implementation and testing. To assess the viability of real-world deployment, experiments were mainly carried out using CPU-only execution. Although it is not necessary for system operation, optional GPU acceleration can be enabled when it is available. Python is used in the implementation, along with popular deep learning and computer vision libraries. Streamlit was used in the construction of the web interface to facilitate quick deployment and user-friendliness.
- **Input Scenarios:** The system's robustness and adaptability were assessed using a variety of input sources:
 - Real-time webcam feeds
 - uploaded pictures
 - Video files that have already been recorded
 - Camera streams via RTSP
 These scenarios replicate real-world roadside surveillance conditions, such as changes in motion, lighting, and object occlusion.
- **Evaluation Criteria:** The system's emphasis on operational viability and real-time deployment meant that qualitative and observational metrics were used for evaluation instead of benchmark datasets. Important evaluation standards included:
 - Reliability of detection in practical settings
 - Event-based alerting's responsiveness
 - Stability while processing videos continuously
 - Evidence capture accuracy for identifying offenders
 - Realistic performance evaluation that is in line with deployment-oriented goals is ensured by this evaluation method.

RESULTS AND DISCUSSION

Experimental observations from the implemented system indicate reliable real-time detection of roadside littering events under practical conditions. These systems increase the accuracy of differentiating between real littering behavior and irrelevant object movement when paired with temporal action recognition techniques.

Although it is extremely sensitive to environmental factors, vehicle-based offender identification using automated license plate recognition shows promise for enforcement-oriented applications. In order to lessen false accusations, the deployed system shows selective use of vehicle identification backed by human validation.

The lack of emphasis on ethical deployment is a noteworthy finding. While detection accuracy continues to improve, privacy protection, transparency, and public trust are often neglected. Systems with anonymized data handling and human oversight are better suited for practical application and community acceptance.

CONCLUSION

This paper reviewed recent advancements in AI-based camera systems for roadside litter detection and offender identification. Deep learning models have enabled scalable and real-time monitoring solutions, particularly through object detection and action recognition techniques. However, most existing studies prioritize technical performance over ethical and preventive considerations.

Future research should focus on integrating privacy-preserving mechanisms, human-in-the-loop validation, and community-centric deployment strategies. By aligning ethical AI principles with smart city initiatives, AI-based litter detection systems can evolve into effective tools for sustainable urban cleanliness.

REFERENCES

- [1] A. Jain, P. Srivastava, and R. Gupta, "Trashbusters: Deep Learning Approach for Litter Detection and Tracking," 2024.
- [2] M. Alharbi, H. Alotaibi, and A. Alshehri, "Real-time detection and monitoring of public littering behavior using deep learning for a sustainable environment," 2025.
- [3] J. Wang, K. Liu, and S. Zhang, "Real-Time Waste Detection and Classification Using YOLO-Based Deep Learning Models," 2024.
- [4] F. Zhou, L. Li, and J. Chen, "Urban Street Cleanliness Assessment Using Mobile Edge Computing and Deep Learning," 2025.
- [5] S. Pathak, R. Menon, and V. Krishnan, "Smart City Community Watch: Camera Based Community Watch for Traffic and Illegal Dumping," 2024.
- [6] A. Kumar and J. Lee, "Video-Based Littering Detection Using MoViNet and YOLOv8," 2024.
- [7] R. Mandhati, A. Das, and P. Chakraborty, "pLitterStreet: Street-Level Plastic Litter Detection and Mapping," 2024.
- [8] A. Espinosa-Juárez, L. G. Rojas, and M. Vargas, "Litter Detection in Real-World Environments Using YOLO Models," 2025.
- [9] V. Tanwer, R. Sharma, and A. Gupta, "Evaluation of Vision Transformers for the Detection of Garbage Bin Fullness for Efficient Waste Management," 2025.
- [10] A. Kumar, H. Chen, and L. Wang, "Advancing Deep Learning-Based Detection of Floating Litter Using a Novel Open Dataset," 2023.
- [11] S. Longhi, D. Monteriù, M. Prist, and M. Pirro, "Smart Waste Management Based on Internet of Things and Deep Learning," *IEEE Sensors Journal*, vol. 20, no. 20, pp. 12331–12341, Oct. 2020.
- [12] A. Kumar and R. P. Mahapatra, "Vision-Based Garbage Detection System for Smart Cities Using Deep Learning," in *Proc. IEEE Int. Conf. Smart Cities, Internet of Things and Applications (SCIoT)*, 2021.
- [13] M. A. Rahman, M. S. Hossain, and G. Muhammad, "Deep Learning-Based Multimedia Surveillance System for Smart Cities," *IEEE Transactions on Multimedia*, vol. 24, pp. 3201–3214, 2022.
- [14] J. Kim, S. Lee, and Y. Kim, "Illegal Dumping Detection Using Deep Learning and CCTV Images," in *Proc. IEEE Int. Conf. Consumer Electronics (ICCE)*, 2022.
- [15] R. K. Jha, P. Jain, and S. Kim, "Edge-Based Intelligent Video Surveillance Using Deep Neural Networks," *IEEE Sensors Journal*, vol. 23, no. 9, pp. 10321–10330, May 2023.
- [16] Y. Li, B. Yang, and Z. Chen, "Edge Intelligence for Smart City Surveillance: A Deep Learning Approach," *IEEE Internet of Things Journal*, vol. 8, no. 8, pp. 6392–6403, Apr. 2021.
- [17] H. Liu, X. Zhang, and Y. Chen, "Real-Time Object Detection for Urban Environmental Monitoring," *IEEE Access*, vol. 10, pp. 99821–99833, 2022.
- [18] S. Bhattacharya, P. K. R. Maddikunta, and T. R. Gadekallu, "Deep Learning for Smart City Surveillance: Recent Advances and Challenges," *IEEE Access*, vol. 11, pp. 45612–45628, 2023.
- [19] A. Alqahtani and A. M. Khan, "Vision-Based Waste Detection and Classification for Smart Cities," in *Proc. IEEE Int. Conf. Artificial Intelligence and Smart Systems (ICAIS)*, 2024.
- [20] M. Zhou, Y. Wang, and L. Chen, "AI-Driven Video Analytics for Public Space Monitoring in Smart Cities," *IEEE Access*, vol. 12, pp. 34890–34905, 2024.