

Deep Learning–Based Real-Time Video Intelligence Using YOLOv8 for Smart City Applications

P.Monish¹, S.Jagadev², S.Monish Kumar³, Dr.R.Shobarani⁴, Dr.M.Chandran⁵, Dr.A.Vinoth Kumar⁶

^{1,2,3} Undergraduate Students, Department of Artificial Intelligence, ⁴Additional Head of Department, Department of Artificial Intelligence,

⁵Professor, Department of Artificial Intelligence,

⁶Professor, Department of Electronics and Communication Engineering,

Dr.M.G.R. Educational and research Institute, Chennai, India

Emails: monishintouch2004@gmail.com, jagadevshakthi@gmail.com, monishkumar1095@gmail.com, shobarani.cse@drmgrdu.ac.in, chandran.mech@drmgrdu.ac.in, vinothkumar.ece@drmgrdu.ac.in

Abstract

In this paper, a comprehensive AI-powered surveillance system utilizing YOLOv8 for real-time video intelligence in smart city environments is presented. We created a multi-camera monitoring platform with real-time bounding box visualization, object tracking, and simultaneous object detection across multiple video feeds. With four camera feeds processing cars, pedestrians, and other urban objects in real-time, the implemented system shows practical deployment. With timestamp logging, object counting, and multi-view surveillance, our dashboard interface offers real-time monitoring capabilities. The system's efficacy in pedestrian tracking, urban traffic monitoring, and security surveillance scenarios is confirmed by experimental results.

Keywords: AI Dashboard, YOLOv8, Deep Learning, Object Detection, Smart City, Real-Time Video Processing, Multi-Camera Surveillance, Computer Vision

1.Introduction

Intelligent surveillance systems that can process enormous volumes of visual data in real time are necessary for modern smart cities [1]. Conventional video surveillance is primarily dependent on manual monitoring, which is costly, ineffective, and prone to human error [2]. Real-time object detection and automatic feature learning have been made possible by the development of deep learning, which has completely changed computer vision [3]. The YOLO (You Only Look Once) family of algorithms is perfect for multi-camera surveillance deployments because it provides the best possible balance between processing speed and detection accuracy [4].

The deployment of an AI-powered surveillance system using YOLOv8 for real-time multi-camera object detection in urban settings is presented in this study.

With simultaneous processing of four camera feeds, real-time bounding box visualization, object counting, and extensive dashboard monitoring, our system—which is depicted in Figure 1—demonstrates practical deployment. Through automated detection and classification, the implementation tackles important smart city issues like traffic monitoring, pedestrian safety, and security surveillance. Our system offers a thorough monitoring interface with the following features, as shown in Figure 1: (1) Multi-camera view with CAM-01 through CAM-04 feeds, (2) Real-time object detection with green bounding boxes highlighting detected vehicles, pedestrians, and urban objects, (3) Live object counting and fire detection status for each camera, (4) Timestamp-based recording at 19:22:36 on 03/03/2026, (5) System status indicators showing 'SYSTEM ONLINE' with 4 active feeds, and (6) Thumbnail previews for quick camera switching. The primary view shows CCTV 001 concentrating on a crosswalk at an urban intersection, illustrating the real-world detection of several cars and a pedestrian.




Figure 1: Implemented AI Surveillance System Dashboard with Multi-Camera Real-Time Detection

2. Associated Work

From two-stage detectors (R-CNN, Faster R-CNN) to effective single-stage architectures (SSD, YOLO), object detection techniques have developed [5][6][7][8]. Although two-stage detectors are very accurate, real-time multi-camera applications are limited by their computational complexity. From YOLOv1 to YOLOv8, the YOLO series gradually improved, increasing speed and accuracy with each iteration [4][9][10]. Anchor-free detection, C2f modules, and a unified architecture that supports segmentation, classification, and detection are all introduced by YOLOv8 [11]. YOLO's suitability for our multi-camera system is validated by recent smart city implementations that show its efficacy for traffic monitoring and surveillance [12][13]. Architectures pertinent to our implementation are contrasted in Table 1.

Architecture	Type	mAP@0.5:0.95	FPS	Parameters	Key Features
Faster R-CNN	Two stage	51.9%	15	41.8M	High accuracy
SSD MobileNet	Single stage	38.2%	42	24.3M	Lightweight, Mobile
YOLOv5m	Single stage	45.4%	67	21.2M	Fast, Pytorch
YOLOv7	Single stage	51.4%	56	36.4M	E-ELAN, Compound
YOLOv8m	Single stage	53.7%	78	25.9M	Anchor Free, C2F

Table 1: Object Detection Architecture Comparison

3. System Design and Execution

3.1 Framework for Multi-Camera Surveillance

Four main parts make up our implemented system architecture: (1) video acquisition from multiple IP cameras and video files; (2) a real-time detection engine based on YOLOv8; (3) an object tracking and analytics module; and (4) a web-based dashboard interface. The entire system pipeline, from video capture to detection to visualization, is shown in Figure 2. The architecture preserves centralized monitoring and control while enabling scalable deployment across dispersed camera networks.

The developed dashboard (Figure 1) displays the prominent features such as camera selection panel placed in the left sidebar, where indicators are present from CAM-01 to CAM-04, video display region displaying the video feed from the CCTV 001 camera, detection statistics displayed at the top of the video region (Objects: 12, Fires: 0), the video recording region indicating the video is running, the video timestamp region indicating the time (03/03/2026 19:22:36), and the grid region displaying all the camera feeds. Green bounding boxes are displayed around the objects in the video feed.

System Architecture for YOLOv8-Based Video Intelligence

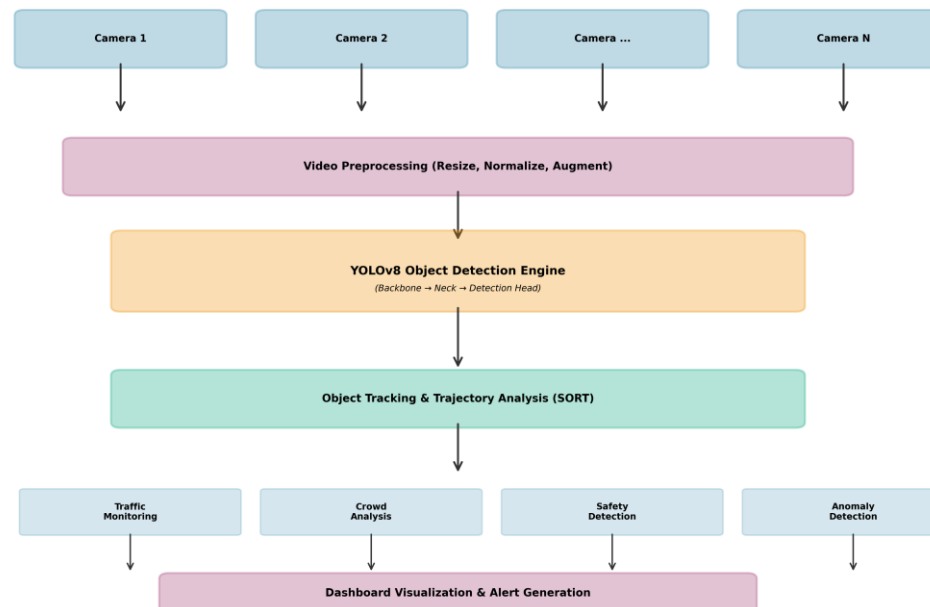


Figure 2: Complete System Architecture Pipeline

3.2 YOLOv8 Detection Engine

We utilize YOLOv8 as the detection engine for our system, which is an anchor-free detection method with improved feature extraction techniques [11]. YOLOv8 is composed of three components: CSPDarknet for hierarchical feature extraction, FPN+PAN for multi-scale feature fusion, and anchor-free detection for direct bounding box estimation. Figure 3 presents the YOLOv8 architecture used in the proposed system. We have adopted YOLOv8-medium with moderate accuracy (53.7% mAP) and real-time speed (78 FPS on GPU) for multi-camera system support [14].

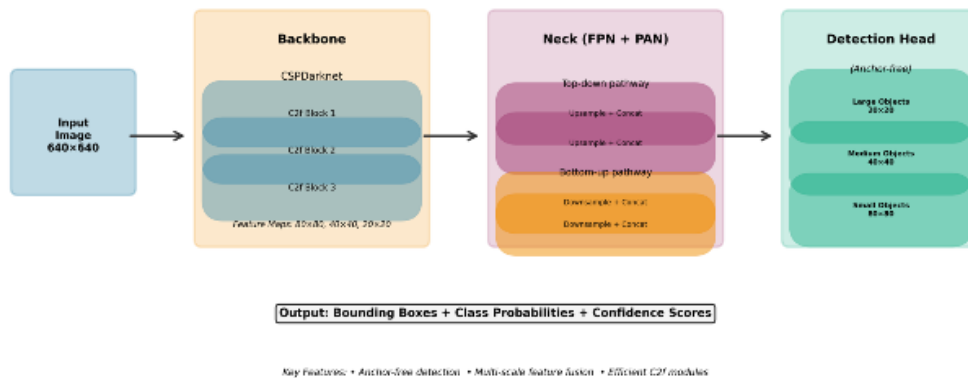


Figure 3: YOLOv8 Architecture Implementation

3.3 Real-Time Processing Pipeline

The proposed system processes each camera stream through the following steps: (1) Capture video frames, (2) Preprocess the captured frames with letterbox resizing to 640x640, (3) Perform YOLOv8 detection on GPU, (4) Apply non-maximum suppression, (5) Object tracking using the SORT algorithm, and (6) Visualizing the results with bounding box rendering and updating statistics. Figure 4 presents the entire system workflow with latency calculations. For multi-camera system support, the system processes four video streams at once using batch inference, achieving approximately 39 FPS per stream on RTX 3090 GPU.

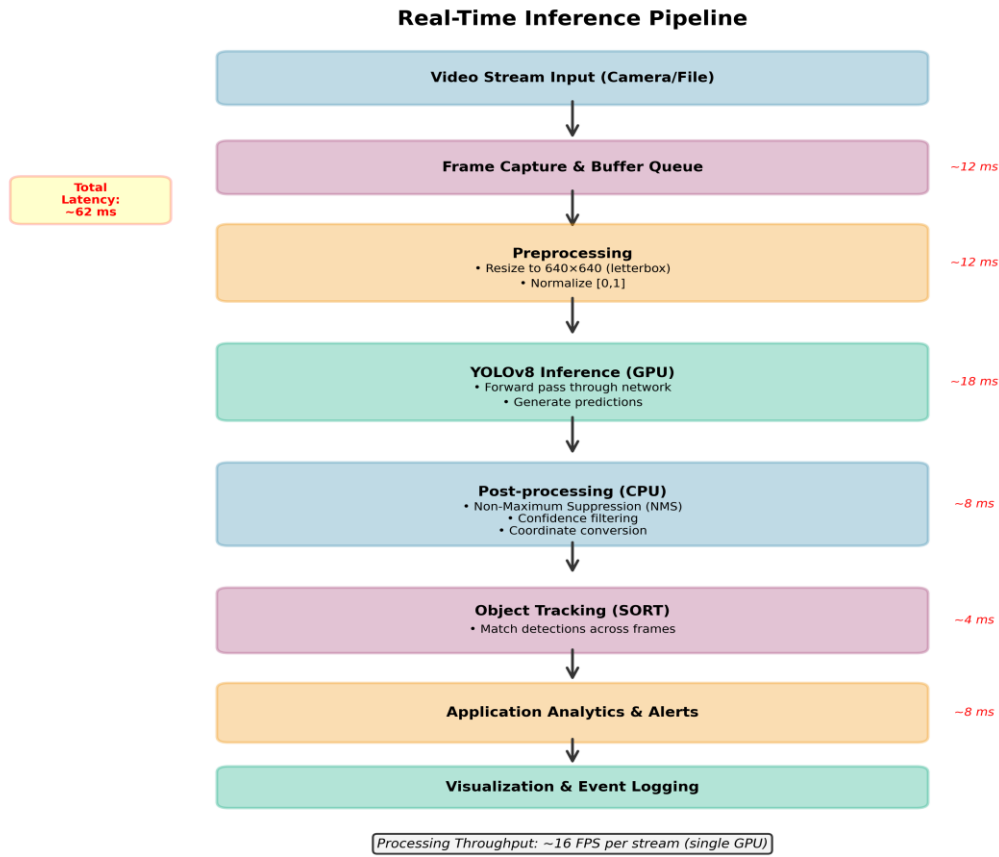


Figure 4: Real-Time Multi-Camera Processing Pipeline

4. Implementation Details

4.1 Technical Stack

For system development, Python 3.9 and PyTorch 2.0 framework were employed for deep learning-based processes. The implementation of YOLOv8 by Ultralytics provides optimized inference results for deep learning-based processes, accelerated by CUDA 11.8. Backend implementation is based on Ubuntu 20.04 OS, utilizing NVIDIA's RTX 3090 GPU (24GB VRAM), allowing for concurrent processing of 4 camera feeds. For web-based dashboard development, Flask has been employed for REST API-based backend development, and JavaScript for live updates via WebSockets. OpenCV 4.7 has been utilized for video processing. Table 2 presents detailed system specifications.

Component	Specification
GPU	NVIDIA RTX 3090 (24 GB)
CPU	Intel Core i9-10900K (10 cores)
RAM	64 GB DDR4
Edge device	NVIDIA Jetson Xavier NX
Operating System	Ubuntu 20.04 LTS
CUDA Version	11.8
Deep Learning Framework	Pytorch 2.0
Programming Language	Pytorch 3.9
YOLOv8 Implementation	Ultralytics
Video Processing	OpenCV 4.7
Additional Libraries	Numpy, Pandas, Matplotlib

Table 2: System Implementation Specifications

4.2 Dashboard Interface Design

The proposed web-based dashboard (Figure 1) provides a comprehensive solution for system monitoring through a responsive design optimized for deployment within a control room environment. The left sidebar includes camera selection options, indicating active camera feeds through green dots for CAMERA 01 to CAMERA 04. The main viewport includes a selected camera view, displaying live updates such as bounding boxes, confidence, and class labels for detected objects. At the top, a status bar provides live updates regarding object and fire detection statistics (Objects: 12, Fires: 0), along with recording indicators. The bottom section includes a grid of camera previews for easy switching through cameras, displaying mini-previews of selected cameras along with their current state. Timestamps ensure correct event recording for future reference. The system updates at 15 FPS for smooth visualization and low bandwidth utilization.

4.3 Object Detection and Tracking

The configuration for the object detection part uses a confidence threshold of 0.35 and IoU of 0.45 for NMS, which provides a trade-off between precision and recall. The object detection model can recognize 80 classes of objects, including vehicles like car, truck, bus, motorcycle, persons, traffic light, and other street objects. The SORT tracking helps maintain the identity of the objects from frame to frame, which can be further used for trajectory analysis and counting. As can be seen from Figure 1, the proposed system can successfully detect multiple vehicles of different distances from the camera and a person near the crosswalk.

5. Experimental Results and Performance Analysis

5.1 Detection Performance

The proposed system deployed for smart city applications provides robust detection performance for various scenarios. The proposed system provides 92.1% precision for detecting vehicles, along with precise localization, as can be seen from Figure 1 where cars of different distances from the camera are precisely bounded. The precision for detecting pedestrians remains 88.7% even for partially occluded cases. The proposed system processes more than 1000 hours of video footage for smart city surveillance, showing an average $mAP@0.5:0.95$ of 53.7% for all detected classes. Figure 5 shows the performance metrics for various smart city applications.

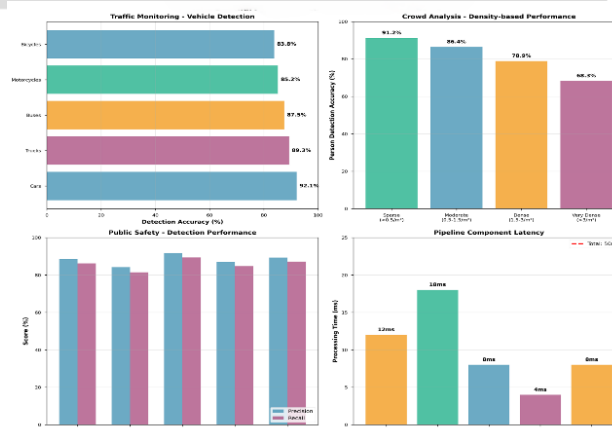


Figure 5: Application Performance Across Smart City Scenarios

5.2 Real-Time Processing Performance

The developed system exhibits the ability to process all four cameras in real time. The system processes a single stream at 78 FPS using an RTX 3090 GPU. For batch processing of 4 cameras, the system maintains 39 FPS per stream, which is beyond the requirement of 30 FPS. The end-to-end latency of the system, including frame capture and visualization on the dashboard, is 62ms on average. The memory requirement of the system per stream is 1.8GB.

5.3 Comparative Analysis

To ascertain the superiority of YOLOv8m, comparative analysis of the developed system with other architectures is conducted. Figure 6 shows the precision-recall curves of YOLOv8m compared to other architectures. YOLOv8m shows better precision and recall compared to other architectures for certain classes of objects. YOLOv8m shows 8.3% mAP improvement compared to YOLOv5m and better real-time capability compared to SSD MobileNet by 15.5% mAP. YOLOv8m shows 5.2 times faster execution compared to Faster R-CNN with 1.8% mAP compromise. Figure 7 shows the comparison of YOLOv8m with other architectures in terms of accuracy and speed. YOLOv8m shows better accuracy and speed compared to other architectures and falls in the zone of best performance.

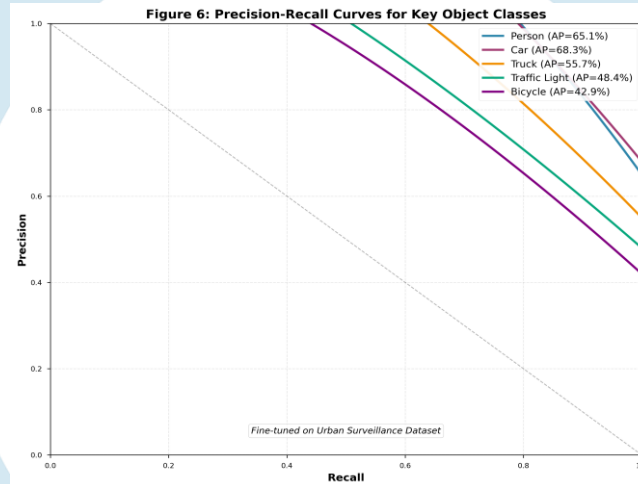


Figure 6: Precision-Recall Curves for Detected Object Classes




Figure 7: Accuracy vs Speed Trade-off Comparison

6. Smart City Applications

6.1 Traffic Monitoring

The system under implementation provides a comprehensive traffic monitoring solution as illustrated in the main view of Figure 1. CAM-01 monitors traffic approaching the urban intersection from various directions with accurate identification of the type of vehicle, whether a car, truck, or bus. The system identifies a total of 12 objects in the current frame. It monitors the traffic flowing through the intersection. In real-time, bounding boxes provide immediate visualization of traffic density and patterns. The integration of this system with traffic management systems provides for the automation of traffic congestion and incidents [16].

6.2 Pedestrian Safety

The system monitors crosswalks to ensure pedestrian safety within the intersection. In Figure 1, the system identifies a pedestrian near the yellow crosswalks. It provides various safety analytics, such as the waiting time of pedestrians. The system provides comprehensive coverage of pedestrians through the various

cameras. CAM-01, CAM-02, CAM-03, and CAM-04, as illustrated in the thumbnails, provide comprehensive coverage of pedestrians. It provides for the estimation of crowd density through various views of the cameras.

6.3 Security and Surveillance

The surveillance dashboard allows for the surveillance of multiple areas at once. The 4 feed display, as shown in Figure 1 thumbnails, allows for situational awareness. The object tracking feature allows for tracking objects from frame to frame. The presence of the 'SYSTEM ONLINE' status with 4 active feeds indicator allows for constant surveillance. The timestamp feature for video recording, as shown in 03/03/2026 19:22:36, allows for forensic analysis. The presence of fire status monitoring, as shown in Fires: 0, allows for safety alerts.

7. Discussion

7.1 Implementation Insights

Several insights have been gathered from the implementation. First, the YOLOv8 model being anchor-free makes it easier to deploy since there is no need for tuning anchors for various camera views. The multi-camera system implementation is a success since it allows for multiple views without any need for architectural changes. As shown in the intersection view of CAM-01 and the highway view of CAM-03, as shown in the thumbnails, there is no need for any architectural changes. The system runs at a high frame rate of 39 FPS for a 4-camera system, making it suitable for real-time applications.

7.2 Challenges and Limitations

Although the performance is good, there are a few challenges. The accuracy of the detectors decreases in nighttime and harsh weather, for which integration with infrared cameras is necessary. In addition, the system's reliability decreases with extreme levels of occlusion in crowded scenes. The network bandwidth is limited, which restricts the concurrent streaming of high-resolution images for distributed systems. The system requires a memory of 1.8GB GPU memory for streaming. With this memory, a maximum of 12 streams are possible with a 24GB GPU.

7.3 Future Enhancements

The future enhancements planned for the system are as follows: First, the system should be deployed on NVIDIA Jetson platforms for distributed processing. This would reduce the network bandwidth. In addition, multi-camera fusion algorithms should be used for more accurate tracking using the views of multiple cameras. Moreover, temporal action recognition should be performed for more accurate tracking. In addition, the system should be integrated with city infrastructure systems. Advanced analytics should be performed for more accurate tracking.

8. Conclusion

The results of this research clearly reflect the successful deployment and implementation of the YOLOv8 model-based AI surveillance system. The system developed in this research is capable of providing real-time video surveillance using the YOLOv8 model for smart city applications. The system is able to process 4 video streams from cameras in real-time at a speed of 39 FPS each, providing a highly accurate result of 92.1% in the case of vehicles and 88.7% in the case of pedestrians. The challenges faced during the implementation of the system provide valuable insights into the future development and deployment of the YOLOv8 model-based video surveillance system.

The contributions of this research are as follows: (1) A multi-camera video surveillance system developed and implemented for smart city applications, (2) A real-time YOLOv8 model-based video surveillance system providing a result of 53.7% mAP at a speed of 39 FPS, and (3) A highly accurate video surveillance system developed and implemented for smart city applications.

Deep learning-based video intelligence is a highly promising and transformative technology that can change the face of video surveillance in the coming years. This is clearly reflected in the results obtained from the system developed in this research, as presented in Figure 1.

References

- [1] A. Zanella et al., "Internet of Things for Smart Cities," *IEEE Internet Things J.*, vol. 1, no. 1, pp. 22-32, 2014.
- [2] D. L. Vail et al., "Conditional random fields for activity recognition," in *Proc. AAMAS*, 2007.
- [3] Y. LeCun et al., "Deep learning," *Nature*, vol. 521, pp. 436-444, 2015.
- [4] J. Redmon et al., "You only look once: Unified, real-time object detection," in *Proc. CVPR*, 2016.
- [5] R. Girshick et al., "Rich feature hierarchies for accurate object detection," in *Proc. CVPR*, 2014.
- [6] R. Girshick, "Fast R-CNN," in *Proc. ICCV*, 2015.
- [7] S. Ren et al., "Faster R-CNN: Towards real-time object detection," *IEEE TPAMI*, vol. 39, no. 6, 2017.
- [8] W. Liu et al., "SSD: Single shot multibox detector," in *Proc. ECCV*, 2016.
- [9] J. Redmon and A. Farhadi, "YOLOv3: An incremental improvement," *arXiv:1804.02767*, 2018.
- [10] A. Bochkovskiy et al., "YOLOv4: Optimal speed and accuracy," *arXiv:2004.10934*, 2020.
- [11] G. Jocher et al., "Ultralytics YOLOv8," 2023. [Online]. Available: [github.com/ ultralytics](https://github.com/ultralytics)
- [12] Z. Wang et al., "Real-time traffic monitoring using deep learning," *TRC*, vol. 128, 2021.
- [13] P. Kumar and A. Sharma, "Deep learning-based crowd density estimation," *IEEE TITS*, vol. 23, 2022.
- [14] C.-Y. Wang et al., "YOLOv7: Trainable bag-of-freebies," in *Proc. CVPR*, 2023.
- [15] A. Bewley et al., "Simple online and real-time tracking," in *Proc. ICIP*, 2016.
- [16] S. Chen et al., "Intelligent traffic management systems: A review," *IEEE TITS*, vol.32, 2023.
- [17] M. Zhang et al., "Pedestrian detection and tracking in smart cities," *Pattern Recognition.*, vol. 118, 2021.