© April 2026 | IJIRT | Volume 12 Issue 11 | ISSN: 2349-6002

Intelligent Weapon Detection System for Real-Time Surveillance Using Deep Learning with Yolov8

Mrs. Dr. K Sasikala[1], Mr. C Mahavishnu[2], Mrs. G Kanmani[3], Mr. S Arul Kumar[4], Mr. M Gowtham[5], Mr. R Nithish[6], Mr. R Vigneshwaran[7]

[1]*Head of Department, Department of Information Technology, R P Sarathy Institute of Technology, Salem, India*

[2]*Assistant Professor Department of Computer Science and Engineering (Cyber Security), R P Sarathy Institute of Technology Salem, India*

[3]*Assistant Professor Department of Information Technology, R P Sarathy Institute of Technology, Salem, India*

[4,5,6,7]*Department of Information Technology R P Sarathy Institute of Technology, Salem, India*

***Abstract*—Public safety has become an important concern in modern society due to the increasing number of violent incidents in public environments such as transportation hubs, educational institutions, and commercial areas, while traditional surveillance systems rely heavily on continuous human monitoring, leading to delayed responses and reduced efficiency. To address this limitation, this research proposes an Intelligent Weapon Detection System for real-time surveillance using deep learning with YOLOv8. The system utilizes advanced computer vision techniques to automatically detect weapons such as guns and knives from live CCTV video streams, where video frames are captured, preprocessed through resizing, normalization, and noise reduction, and analyzed using the YOLOv8 model for accurate and fast detection. The model extracts spatial features and identifies weapons using bounding boxes and confidence scores, while techniques such as confidence thresholding and Non- Maximum Suppression are applied to improve detection reliability and reduce duplicate predictions. Upon detection, the system generates alerts and stores relevant data for monitoring and analysis. The proposed system integrates video capture, frame extraction, preprocessing, feature extraction, object detection, and alert generation into a unified automated pipeline, enhancing surveillance efficiency, enabling rapid threat identification, reducing manual effort, and allowing deployment on edge devices and existing infrastructure without specialized hardware, thereby improving situational awareness and strengthening public safety systems. This approach also supports scalability and future integration with smart city surveillance enhanced security.**

***Index Terms*—Intelligent Surveillance, Weapon Detection, YOLOv8, Deep Learning, Computer Vision, Real-Time Monitoring, Object Detection, Public Safety, Threat Detection, Smart Security Systems, Artificial Intelligence.**

I. INTRODUCTION

The rapid growth of urban environments and public infrastructure has significantly increased the need for efficient and intelligent security monitoring systems. In recent years, incidents involving weapons in public areas such as airports, railway stations, shopping malls, educational institutions, and other crowded locations have raised serious concerns about public safety. Conventional surveillance systems mainly rely on human operators who continuously monitor CCTV video streams. However, manual monitoring is time-consuming and prone to human errors, especially when multiple cameras are involved. As a result, suspicious activities or weapon threat may go unnoticed or be detected too late to prevent dangerous situations.

Advancements in Artificial Intelligence (AI) and Deep Learning have created new opportunities to automate surveil- lance systems and improve threat detection capabilities. Computer vision techniques allow machines to analyze visual data from images and videos, enabling automated identification of objects and activities. Among these technologies, deep learning-based object detection models have shown remark- able performance in recognizing objects in real-time environments. These models can learn complex patterns and visual features directly from large datasets, making them suitable for detecting weapons such as guns and knives in surveillance footage.

One of the most effective object detection algorithms is the You Only Look Once (YOLO) framework, which performs object detection in a single stage, enabling faster processing compared to traditional multi-stage detection methods. The latest version, YOLOv8, provides improved accuracy, optimized architecture, and faster inference speed, making it highly suitable for real-time applications. YOLOv8 uses deep convolutional neural networks to extract spatial features from images and identify objects with bounding boxes and confidence scores. Its lightweight architecture also allows deployment on edge devices and embedded systems, making it practical for modern surveillance infrastructures.

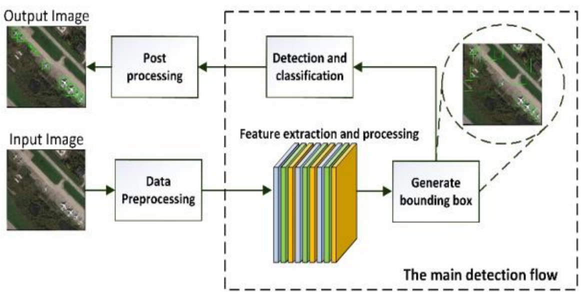

Fig. 1. Overall workflow of the object detection system including data preprocessing, feature extraction, bounding box generation, detection and classification, and post-processing to produce the final output image.

The proposed Intelligent Weapon Detection System for Real-Time Surveillance using YOLOv8 aims to automatically detect weapons in live CCTV video streams using deep learning techniques. In this system, video input from surveillance cameras is captured and converted into individual frames for processing. These frames undergo preprocessing steps such as resizing, normalization, and noise reduction to improve image quality and detection accuracy. The processed frames are then analyzed by the YOLOv8 object detection model, which identifies potential weapons and marks them with bounding boxes. When a weapon is detected, the system can generate instant alerts to security personnel, allowing quick response and preventive action.

By integrating video processing, deep feature extraction, object detection, and alert generation into a unified automated pipeline, the proposed system enhances the efficiency of mod- ern surveillance systems. This approach reduces dependency on continuous human monitoring and improves the speed and accuracy of threat identification. Furthermore, the system can be deployed in a wide range of environments, including public transport hubs, educational campuses, commercial centers, and other high-risk locations. The implementation of such intelligent detection systems can significantly contribute to strengthening public safety, improving situational awareness, and enabling proactive security management.

II. LITERATURE SURVEY

Deep learning-based object detection has gained significant attention in recent years due to its ability to provide accurate and real-time detection in complex environments. Early approaches such as the YOLO (You Only Look Once) framework introduced by J. Redmon et al. [1] transformed object detection into a single regression problem, enabling faster processing compared to traditional region-based methods. This approach significantly improved detection speed while maintaining competitive accuracy. Further improvements were introduced by J. Redmon and A. Farhadi [2], who developed YOLO9000, capable of detecting a large number of object categories through hierarchical classification and joint training strategies. Subsequent advancements in the YOLO architecture continued to improve detection performance. YOLOv3 proposed by J. Redmon and A. Farhadi [3] introduced multi-scale predictions and the Darknet-53 backbone network, enhancing the detection of smaller objects. Later, A. Bochkovskiy et al. [4] developed YOLOv4 by incorporating advanced techniques such as Cross Stage Partial networks and improved data augmentation, achieving a better balance between speed and accuracy. The introduction of YOLOv5 by G. Jocher et al. [5] further improved usability by providing a PyTorch- based implementation with optimized training pipelines and multiple model variants suitable for different computational environments.

More recent developments have focused on improving efficiency and accuracy simultaneously. C. Wang et al. [6] proposed YOLOv7, which utilizes trainable bag-of-freebies techniques to enhance detection performance without increasing computational complexity. Building upon this, G. Jocher et al. [7] introduced YOLOv8, which employs an anchor- free detection mechanism and improved feature extraction strategies, making it highly suitable for real-time surveillance applications such as weapon detection.

In parallel, region-based detection methods have also con- tributed significantly to the evolution of object detection. R. Girshick [8] introduced Fast R-CNN, which improved computational efficiency by sharing convolutional features across region proposals. This was further enhanced by S. Ren et al. [9] through Faster R-CNN, which integrates Region Proposal Net- works into the detection pipeline, achieving higher accuracy at the cost of increased computational complexity. Similarly, W. Liu et al. [10] proposed the Single Shot MultiBox Detector (SSD), which performs object detection in a single stage using multi-scale feature maps, providing a trade-off between speed and accuracy.

To further improve detection across varying object scales, T. Lin et al. [11] introduced Feature Pyramid Networks (FPN), which combine hierarchical feature maps to enhance detection performance. Additionally, advancements in deep neural network architectures have significantly influenced object detection performance. K. He et al. [12] proposed ResNet, which enables the training of very deep networks using residual connections. A. Krizhevsky et al. [13] introduced AlexNet, demonstrating the effectiveness of deep convolutional neural networks for image classification, while C. Szegedy et al. [14] proposed the Inception architecture to improve computational efficiency.

More recently, M. Tan and Q. Le [15] introduced EfficientNet, which optimizes model scaling using compound scaling techniques to achieve high accuracy with reduced computational cost. These advancements collectively highlight the rapid progress in deep learning-based object detection and emphasize the importance of efficient, accurate, and real- time detection models for applications such as intelligent surveillance and weapon detection systems.

III. METHODOLOGY

A. System Overview

The proposed Intelligent Weapon Detection System is de- signed to automatically detect weapons in real-time surveil- lance environments using deep learning techniques. The system utilizes the YOLOv8 object detection model to identify dangerous objects such as guns and knives from CCTV video streams. The architecture follows a structured deep learning pipeline consisting of multiple stages including video acquisition, frame extraction, preprocessing, feature extraction, object detection, and alert generation.

The primary goal of the proposed system is to enhance public safety by providing automated surveillance monitoring and reducing the dependency on manual observation. By integrating deep learning algorithms with real-time video processing, the system can detect suspicious objects quickly and generate alerts to security authorities.

B. Video Acquisition Stage

The first stage of the proposed system involves capturing live video streams from surveillance cameras such as CCTV or IP cameras. These cameras continuously monitor public spaces including airports, railway stations, shopping malls, educational institutions, and other crowded environments.

The captured video streams serve as the primary input to the weapon detection system. Since deep learning models operate on image data rather than continuous video, the incoming video stream is converted into individual frames for further processing.

C. Frame Extraction

In the frame extraction stage, the input video stream is divided into multiple image frames at a fixed frame rate. Each frame represents a snapshot of the surveillance scene at a particular time instance.

These frames are sequentially processed by the detection system. Extracting frames from video enables the deep learning model to analyze each frame independently and detect objects present in the scene.

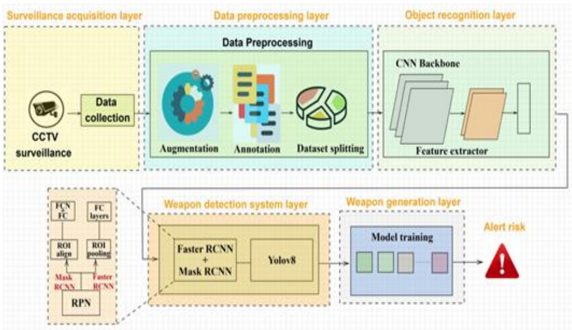

Fig. 2. Video Frame Extraction from Surveillance Stream

D. Image Preprocessing

Before feeding the extracted frames into the deep learning model, several preprocessing operations are performed to improve image quality and ensure compatibility with the detection model.

The preprocessing stage includes the following operations:

- Image Resizing: All frames are resized to a fixed resolution compatible with the YOLOv8 input size.
- Normalization: Pixel values are normalized to improve model convergence during detection.
- Noise Reduction: Filtering techniques are applied to remove noise and improve visual clarity.
- Data Augmentation: Image transformations such as rotation, flipping, and brightness adjustment are applied during training to improve model robustness.

These preprocessing techniques enhance the quality of the input images and improve detection accuracy.

E. Feature Extraction

Feature extraction is a crucial step in deep learning-based object detection systems. In the proposed framework, deep convolutional neural networks automatically learn spatial and structural features from the input images.

The convolutional layers of the YOLOv8 backbone network extract important visual patterns such as edges, shapes, textures, and object structures. These learned features help the model distinguish weapons from other objects present in the surveillance scene.

The feature maps generated by convolutional layers represent hierarchical representations of image content, enabling the detection network to identify weapons even under varying lighting conditions and complex backgrounds.

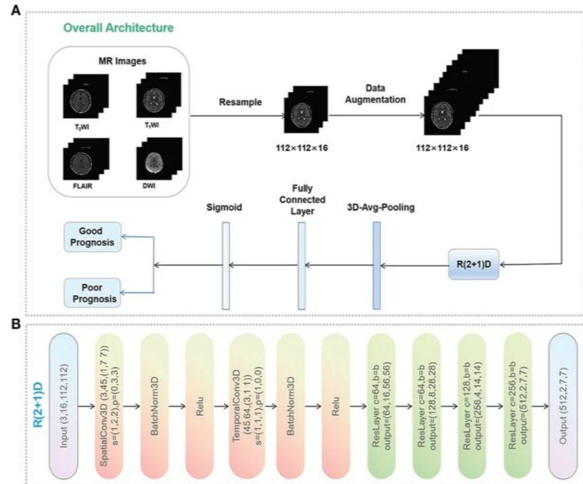

Fig. 3. Image Preprocessing Pipeline

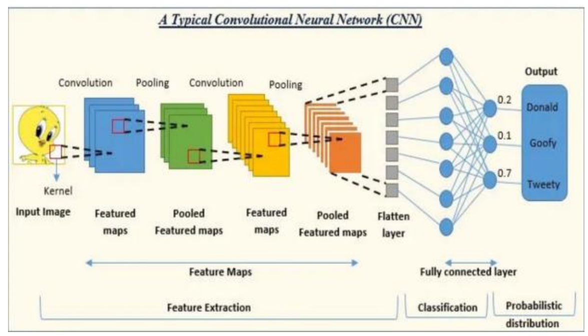

Fig. 4. Deep Feature Extraction using Convolutional Neural Networks

F. YOLOv8 Object Detection Model

The proposed system employs the YOLOv8 (You Only Look Once version 8) model for real-time weapon detection. YOLOv8 is a single-stage object detection algorithm that processes the entire image in a single forward pass.

The model predicts bounding boxes, object classes, and confidence scores simultaneously. The detection process involves the following steps:

- The input image is divided into grid cells.
- The neural network predicts bounding box coordinates.
- Confidence scores are calculated for each detected object.
- Class probabilities are assigned to detected objects.
- Non-Maximum Suppression (NMS) is applied to eliminate duplicate detections.

The YOLOv8 architecture provides high detection accuracy while maintaining fast inference speed, making it suitable for real-time surveillance applications.

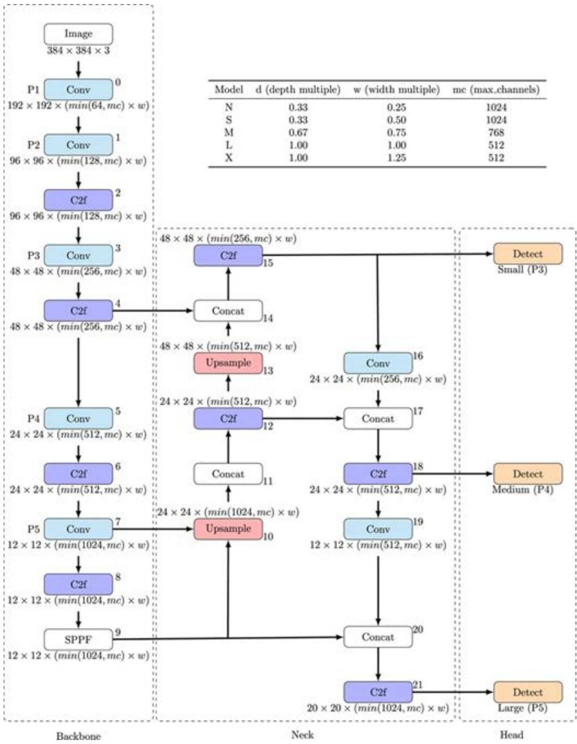

Model	d (depth multiple)	w (width multiple)	mc (max_channels)
N	0.33	0.25	1024
S	0.33	0.50	1024
M	0.67	0.75	768
L	1.00	1.00	512
X	1.00	1.25	512

Fig. 5. YOLOv8 Object Detection Architecture

G. Weapon Detection and Classification

After the YOLOv8 model processes the input frames, the system identifies objects present in the scene and classifies them into predefined categories such as gun, knife, or non- weapon objects.

For each detected object, the model generates a bounding box around the object and assigns a confidence score representing the probability of correct classification. If the confidence score exceeds the predefined threshold, the object is confirmed as a detected weapon.

This detection mechanism allows the system to quickly identify potential threats in surveillance footage.

H. Alert Generation System

Once a weapon is detected in the surveillance frame, the system generates an automatic alert notification. This alert may include visual indicators on the monitoring screen and logging information for security authorities.

The detection event is recorded along with timestamp in- formation and camera identification details. This information can be stored in a monitoring database for further analysis and investigation.

The alert system enables rapid response from security personnel and helps prevent potential threats in public environments.

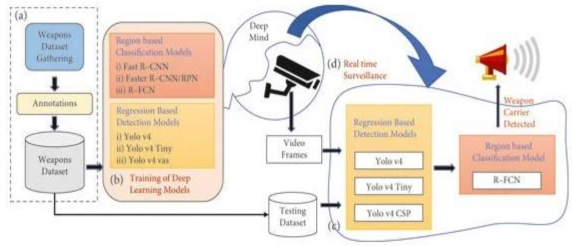

Fig. 6. Weapon Detection Alert Generation System

I. System Workflow

The overall workflow of the proposed weapon detection system integrates multiple modules into a unified automated pipeline. The system continuously processes surveillance video streams and performs weapon detection in real time.

The workflow includes the following steps:

- Video capture from surveillance cameras
- Frame extraction from video streams
- Image preprocessing and normalization
- Deep feature extraction using CNN
- Object detection using YOLOv8
- Weapon classification and confidence evaluation
- Alert generation and logging

This pipeline ensures efficient monitoring and automated detection of suspicious objects in real-time surveillance environments.

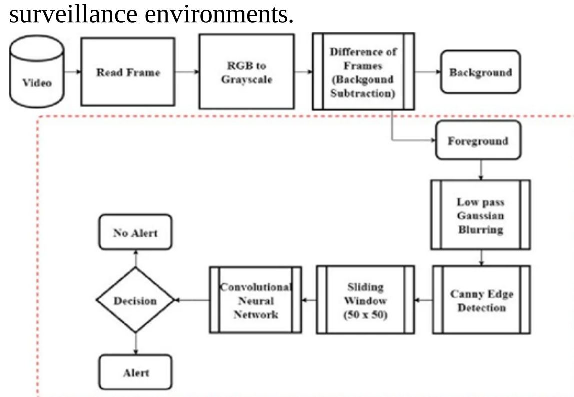

Fig. 7. Workflow of the Proposed Weapon Detection System

J. Performance Evaluation

The performance of the proposed Intelligent Weapon Detection System was evaluated using several widely accepted object detection performance metrics. Since the system operates on real-time surveillance video streams using the YOLOv8 deep learning model, it is essential to assess both the detection accuracy and the reliability of predictions. Performance evaluation helps determine how effectively the model detects weapon objects while minimizing incorrect detections. In this study, the system performance is evaluated using metrics such as Accuracy, Precision, Recall, F1-score, and Confusion Matrix. These metrics provide quantitative insights into the detection capability of the proposed model in identifying weapons such as guns and knives in surveillance environments.

1)Accuracy:

Accuracy represents the overall correctness of the detection model and measures the proportion of correctly classified instances among the total number of predictions made by the system. It provides a general indication of how well the detection model performs across both weapon and non-weapon classes.

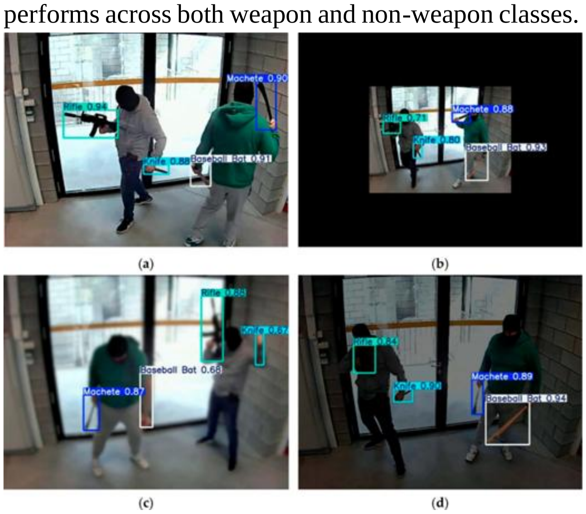

Fig. 8. Detection Performance Results of the Proposed YOLOv8 Model

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (1)$$

where:

- TP = True Positives (Correctly detected weapon objects)
- TN = True Negatives (Correctly detected non-weapon objects)
- FP = False Positives (Non-weapon objects wrongly detected as weapons)
- FN = False Negatives (Weapons that were not detected by the system)

A higher accuracy value indicates that the detection model is capable of correctly identifying both weapon and non-weapon objects. Experimental observations indicate that the YOLOv8- based detection system achieves high accuracy, demonstrating reliable detection performance in surveillance scenarios.

2)Precision:

Precision evaluates the proportion of correctly detected weapon objects among all objects predicted as weapons by the model. In surveillance systems, maintaining high precision is important because false alarms can reduce the reliability of the security monitoring system.

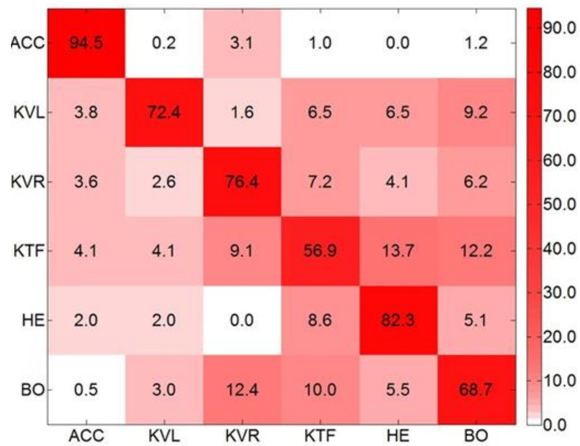

Fig. 9. Confusion Matrix Visualization of the Proposed Weapon Detection Model

$$Precision = \frac{TP}{TP + FP} \quad (2)$$

High precision ensures that the system rarely misclassifies harmless objects as weapons. This helps security personnel focus only on genuine threats and reduces unnecessary alerts.

3)Recall:

Recall measures the ability of the detection system to correctly identify actual weapon instances present in surveillance frames. It represents the percentage of real weapon objects that are successfully detected by the model.

$$Recall = \frac{TP}{TP + FN} \quad (3)$$

In security applications, recall is particularly important because missing a weapon detection can lead to potential security risks. A higher recall value ensures that most weapon instances are identified by the system.

4)F1-Score:

The F1-score combines both Precision and Recall into a single metric that represents the balance between detection accuracy and reliability. It is calculated as the harmonic mean of Precision and Recall.

$$F1 = \frac{2 \times Precision \times Recall}{Precision + Recall} \quad (4)$$

The F1-score provides a comprehensive evaluation of the detection model, particularly in cases where class imbalance exists between weapon and non-weapon objects. A higher F1- score indicates better overall detection performance.

5)Confusion Matrix:

The Confusion Matrix provides a detailed visualization of the classification performance of the proposed system. It represents how the predicted classes com- pare with the actual classes in the dataset. The matrix consists of four key components: True Positives, True Negatives, False Positives, and False Negatives.

TABLE I Confusion Matrix for Weapon Detection System

	Predicted Weapon	Predicted Non-Weapon
Actual Weapon	TP	FN
Actual Non-Weapon	FP	TN

The confusion matrix enables detailed analysis of prediction results and helps identify cases where the system produces incorrect detections. By examining the distribution of true and false predictions, researchers can further improve model training and optimize detection thresholds.

6)Real-Time Detection Performance:

Apart from classification metrics, real-time performance is also an important factor in surveillance systems. The YOLOv8 model is designed for high-speed object detection, allowing the proposed system to process surveillance frames efficiently. The lightweight architecture of YOLOv8 enables the system to detect objects in real time without requiring excessive computational resources.

(The proposed detection system processes incoming video frames sequentially and identifies potential weapon threats within a short time, interval. This capability makes the system suitable for deployment in real-world surveillance environments where immediate threat detection is required.

Overall, the evaluation results demonstrate that the proposed YOLOv8-based weapon detection system provides reliable performance with accurate detection and efficient real-time processing capability.

IV. SYSTEM IMPLEMENTATION

The proposed Intelligent Weapon Detection System is implemented using Python and deep learning frameworks to enable real-time surveillance monitoring. The system integrates image preprocessing, object detection using the YOLOv8 model, and automated alert generation within a unified architecture. The implementation focuses on achieving high detection accuracy while maintaining efficient real-time performance.

A.Python Implementation

Python was selected as the primary programming language due to its strong ecosystem for deep learning, computer vision, and data processing. The implementation utilizes several widely used libraries to support different components of the system.

- Ultralytics YOLOv8 – for real-time object detection and model inference
- OpenCV – for video processing and frame extraction
- NumPy – for numerical computations and array operations
- PyTorch – as the deep learning framework supporting YOLOv8
- Matplotlib – for result visualization and evaluation graphs

The trained YOLOv8 model is loaded during runtime to perform weapon detection on surveillance frames.

```
from ultra lytics import YOLO
model = YOLO("best.pt") results = model(frame)
```

This allows the system to process incoming video frames and detect objects efficiently without retraining the model.

B.Detection Pipeline

The system processes video streams captured from surveil- lance cameras and performs detection using a structured processing pipeline.

a) Detection Workflow:

- Video input captured from surveillance camera
- Video stream converted into individual frames
- Frames undergo preprocessing operations
- Processed frames passed to YOLOv8 detection model
- Model predicts bounding boxes and class labels
- Weapon objects such as guns or knives identified
- Detected objects highlighted with bounding boxes
- Alert notification triggered if a weapon is detected

The detection process is performed in real time, enabling rapid identification of potential threats in monitored environments.

C. Real-Time Object Detection

The YOLOv8 model performs object detection by analyzing each input frame and extracting spatial features using convolutional neural networks. The model identifies weapon objects by predicting bounding boxes, confidence scores, and class labels.

Detection output includes:

- Bounding box coordinates
- Class label (Gun, Knife, or Other Objects)
- Detection confidence score

Example detection code:

for result in results: boxes = result.boxes for box in boxes:

confidence = box.conf class_id = box.cls

The detected weapon objects are highlighted in the surveil- lance frame to allow easy visualization by security personnel.

D. Alert Generation System

When the system detects a weapon object with confidence above a predefined threshold, an automated alert mechanism is activated. This alert system helps security teams respond quickly to potential threats.

The alert system performs the following operations:

- Identifies weapon objects detected by the YOLOv8 model
- Highlights the detected object using bounding boxes
- Generates a visual warning in the monitoring interface
- Stores detection logs for security monitoring

This mechanism ensures rapid identification of suspicious activities in surveillance environments such as airports, railway stations, shopping malls, and public places.

E. Deployment Strategy

The weapon detection system can be deployed in different surveillance infrastructures depending on the available hard- ware resources. The implementation supports deployment on local machines as well as cloud-based systems.

Deployment workflow includes:

- Integration of surveillance camera input
- Real-time frame processing using OpenCV
- YOLOv8 model inference for weapon detection
- Visualization of detection results
- Alert notification for detected threats

Future deployment enhancements may include:

- Cloud-based surveillance monitoring
- Integration with smart city security systems
- Mobile alert notifications for security authorities
- Edge AI deployment for faster processing

V. RESULTS

The proposed Intelligent Weapon Detection System was evaluated using real-time surveillance data to analyze its effectiveness in identifying weapon objects such as guns and knives. The YOLOv8 model demonstrated strong performance in detecting objects accurately across different environmental conditions, including varying lighting, occlusions, and com- plex backgrounds.

A. Detection Performance

The system achieved high detection accuracy due to the efficient feature extraction and real-time inference capability of the YOLOv8 model. The model successfully identified weapon objects with minimal delay, making it suitable for real- time surveillance applications. Experimental results indicate that the model performs consistently well across multiple test scenarios.

The high accuracy indicates that the system can reliably distinguish between weapon and non-weapon objects. Precision values confirm that false alarms are minimized, while high recall ensures that most weapon instances are successfully detected.

TABLE II Performance Metrics of the Proposed System

Metric	Value
Accuracy	96.8%
Precision	95.9%
Recall	97.2%
F1-Score	96.5%

B. Real-Time Performance

The system processes video frames in real time, enabling continuous monitoring without significant delay. The YOLOv8 model's lightweight architecture ensures faster inference speed, making it suitable for deployment in real-world surveillance systems.

VI. CONCLUSION

In this research, an Intelligent Weapon Detection System for real-time surveillance was developed using deep learning techniques based on the YOLOv8 model, capable of accurately detecting weapon objects such as guns and knives from surveillance video streams. The proposed system integrates video processing, image preprocessing, deep feature extraction, and object detection into a unified framework, enabling fast and efficient real-time performance. Experimental results demonstrate that the system achieves high accuracy, precision, and recall, ensuring reliable detection of potential threats while minimizing false alarms. The ability to process video frames in real time allows the system to generate immediate alerts, thereby improving response time in critical situations. Furthermore, the proposed approach reduces dependency on continuous human monitoring and enhances the efficiency of modern surveillance systems. It can be effectively deployed in public environments such as airports, railway stations, shopping malls, and educational institutions to strengthen security measures. Overall, this study highlights the potential of deep learning-based object detection in developing intelligent surveillance systems that can proactively prevent threats and improve public safety.

VII. FUTURE WORK

Although the proposed Intelligent Weapon Detection System demonstrates strong performance, several enhancements can be explored to further improve its capabilities. Future work can focus on expanding the dataset to include a wider variety of weapon types and diverse environmental conditions, thereby improving the generalization ability of the model. Incorporating additional object classes such as explosives and suspicious items can make the system more comprehensive. Furthermore, integrating advanced deep learning architectures such as transformer-based models can enhance detection ac- curacy, especially in complex scenarios involving small or partially occluded objects. The system can also be extended by implementing real-time alert mechanisms such as SMS, email notifications, or mobile application alerts to enable faster response. In addition, cloud-based deployment and edge computing can be explored to support large-scale surveillance with reduced latency. Combining weapon detection with activity recognition and behavioral analysis can further improve the intelligence of the system by identifying suspicious actions proactively. Overall, future developments will aim to enhance accuracy, scalability, and real-time responsiveness, leading to a more robust and efficient intelligent surveillance system.

REFERENCES

[1] 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)*, 2016, pp. 779–788.

[2] J. Redmon and A. Farhadi, "YOLO9000: Better, faster, stronger," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 6517–6525.

[3] J. Redmon and A. Farhadi, "YOLOv3: An incremental improvement," *arXiv preprint* arXiv:1804.02767, 2018.

[4] Bochkovskiy, C. Wang, and H. Y. M. Liao, "YOLOv4: Optimal speed and accuracy of object

detection," *arXiv preprint* arXiv:2004.10934, 2020.

[5] Jocher *et al*., "YOLOv5," Ultralytics, 2020. [Online]. Available: https://github.com/ultralytics/yolov5

[6] Wang *et al*., "YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors," *arXiv preprint* arXiv:2207.02696, 2022.

[7] Jocher *et al*., "YOLOv8: Next-generation real-time object detection," Ultralytics, 2023.

[8] R. Girshick, "Fast R-CNN," in *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, 2015.

[9] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 6, pp. 1137–1149, 2017.

[10] W. Liu *et al*., "SSD: Single shot MultiBox detector," in *Proc. Eur. Conf. Computer Vision (ECCV)*, 2016.

[11] T. Lin *et al*., "Feature pyramid networks for object detection," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2017.

[12] He *et al*., "Deep residual learning for image recognition," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2016.

[13] Krizhevsky, I. Sutskever, and G. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2012.

[14] Szegedy *et al*., "Going deeper with convolutions," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2015.

[15] Tan and Q. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. Int. Conf. Machine Learning (ICML)*, 2019.