

## Intelligent Obstacle Detection And Warning System For Railroad Crossings Using Yolov8 And Bytetrack With Signal-Triggered Response

Untung Rohwadi<sup>1</sup>, Indah Purwandni<sup>2</sup>, Rudianto<sup>3</sup>, Amrin<sup>4</sup>

<sup>1,2,3,4</sup>Universitas Bina Sarana Informatika

<sup>1</sup>e-mail: [untung.unr@bsi.ac.id](mailto:untung.unr@bsi.ac.id)

<sup>2</sup>e-mail: [indah@bsi.ac.id](mailto:indah@bsi.ac.id)

<sup>3</sup>e-mail: [rudianto.rdt@bsi.ac.id](mailto:rudianto.rdt@bsi.ac.id)

<sup>4</sup>e-mail: [amrin.ain@bsi.ac.id](mailto:amrin.ain@bsi.ac.id)

### ARTICLE INFORMATION

#### Artikel History::

Received: 04/11/2025

Revised: 15/12/2025

Accepted: 10/02/2026

#### Keyword:

ByteTrack  
YOLOv8  
Obstacle  
multi intelligent  
railroad

### ABSTRACT (10pt, Times New Roman)

Railroad level crossings remain a critical safety concern due to the high incidence of collisions involving vehicles and pedestrians. This study proposes an intelligent, vision-based obstacle detection and warning system that integrates the YOLOv8 object detection model with the ByteTrack multi-object tracking algorithm. Designed for real-time operation on crossing surveillance video, the system can detect and track vehicles, pedestrians, and foreign objects within the railway zone. Consistent detection is coupled with a signal-triggered response mechanism, enabling early warnings to operators or automated gate control systems, particularly as trains approach. Evaluation on a diverse video dataset—captured under varying lighting and weather conditions demonstrates that the integration of YOLOv8 and ByteTrack improves tracking consistency, reduces false negatives, and maintains latency below 50 ms per frame. This research advances intelligent transportation safety systems and offers a vision-based solution that can be integrated into existing railway infrastructure.

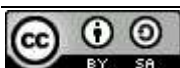
### Corresponding Author:

Untung Rohwadi,  
Fakultas Teknik dan informatika, Universitas Bina Sarana informatika,  
Jl.Kramat Raya No.98, Jakarta, Indonesia,  
Email: [untung.unr@bsi.ac.id](mailto:untung.unr@bsi.ac.id)

### INTRODUCTION

Level crossings remain critical points of vulnerability in railway safety across many countries. According to the Federal Railroad Administration (FRA) of the United States, hundreds of incidents occur annually involving vehicles or pedestrians struck at crossings (FRA, 2011). Similar conditions are observed in developing nations, where limited warning systems and low compliance among road users significantly increase the risk of fatal accidents. Various approaches have been designed to detect the presence of objects on railway tracks, including electromagnetic sensors (loop coils), laser scanners, and stereo vision systems (Chernov Andreyand Butakova, 2020)

However, these solutions often face limitations regarding cost, installation complexity, and reliability in extreme weather or lighting conditions. In contrast, camera-based video analysis has become popular due to its flexible



deployment, extensive area coverage, and compatibility with modern machine learning techniques. YOLO (You Only Look Once) is one of the most widely used real-time object detection algorithms in traffic safety applications. The latest version, YOLOv8, offers significant improvements in accuracy and speed, along with full support for edge device deployment (Hussain, 2024). To ensure consistent object identification across frames, this study integrates the ByteTrack multi-object tracking algorithm that utilizes nearly all detection boxes, including low-confidence ones, to reduce identity loss (Wang et al., 2023). The primary contribution of this research is the design of an intelligent obstacle detection and tracking system for railway crossings, based on YOLOv8 and ByteTrack, and enhanced with a signal-triggered response mechanism. The system not only detects the presence of objects but also provides context-aware warnings, such as when a train is approaching. This approach has the potential to improve road user safety and enhance the operational efficiency of crossing systems.

### Related Studies

Research on railway crossing safety has been explored from various perspectives, including policy, infrastructure, and intelligent detection systems.

- a. (Antono, 2023) highlights the complexity of accident mitigation at level crossings in Central Java, focusing on challenges such as inter-agency coordination, budget limitations, and low public awareness.
- b. (Ainuriansyah Akbar et al., 2024) presents a YOLOv8-based detection system for railway assets, demonstrating high classification accuracy for locomotives and KPJ units, but lacking ID tracking and latency evaluation.
- c. (Kudlacik & Wesolowski, 2023) proposes a video-based obstacle detection algorithm, which is efficient but does not integrate object tracking or physical system deployment.
- d. (Yu & Lu, 2024) improves YOLOv8 for turnout defect detection in complex environments, though its focus remains on infrastructure inspection rather than dynamic obstacle detection.
- e. (Tang & Qian, 2024) develops an edge-optimized YOLOv8 model for high-speed railway inspection, yet it does not address real-time warning systems or dynamic object tracking at crossings.

Unlike previous studies that focus solely on object detection or infrastructure inspection, this research integrates YOLOv8 and ByteTrack into a modular, latency-aware system with signal-triggered response and debounce logic. It addresses the gap in real-time tracking stability, latency evaluation, and physical system integration for railway crossing safety.

**Table 2.1. Gap Analysis of Prior Studies**

Reference	Main Focus	Strengths	Gap
Antono (2023)	Traffic accident mitigation program at level crossings in Central Java	Provides social and policy context	Does not address technical or automated detection systems
Ainuriansyah Akbar et al., 2024	Railway asset detection using YOLOv8	Local implementation with camera-based classification	No ID tracking, no latency evaluation, no debounce mechanism
Kudlacik & Wesolowski, 2023	Obstacle detection based on video stream analysis	Efficient object detection algorithm	Lacks object tracking across frames and physical system integration
Yu & Lu, 2024	Detection of turnout defects in complex environments	Improved YOLOv8 accuracy and robustness	Focuses on infrastructure inspection, not dynamic obstacle detection or

				violation context
Tang & Qian, 2024	High-speed railway track component inspection using YOLOv8	Optimized for edge deployment with high-performance inference	Does not address dynamic object tracking or warning systems for crossings	

Source: Summarized

### Unique Contributions of This Research.

This study fills the gap by offering a more integrated and context-aware approach:

**YOLOv8 + ByteTrack Integration:** Combines real-time detection with persistent ID tracking across frames.

**Latency Evaluation:** Measures per-frame latency and validates real-time compliance ( $\leq 50$  ms)

**Debounce Logic:** Prevents redundant logging and false positives, enhancing system reliability

**Signal-Triggered Response:** Activates detection only when a specific track is triggered, enabling smart warning logic

**Modular PyQt Interface:** Provides real-time visualization, ROI status, and system feedback, ready for multi-track expansion

## RESEARCH METHOD

### System Architecture

The proposed system is designed to improve safety at railway crossings by detecting and tracking objects in real time using camera-based video analysis. The architecture integrates three core modules:

**Object Detection:** Utilizes YOLOv8 for high speed, high-accuracy detection of multiple classes including pedestrians, vehicles, and motorcycles. The model runs on edge devices and supports frame-by-frame inference with confidence filtering.

**Multi-Object Tracking:** Implements ByteTrack to maintain object identity across frames. ByteTrack leverages both high and low-confidence detections to reduce ID switching and improve tracking robustness, especially in crowded scenes.

**ROI-Based Filtering and Signal Response:** A predefined Region of Interest (ROI) is used to isolate critical zones near the railway tracks. Objects entering or touching the ROI trigger context-aware responses, such as visual alerts or signal activation when a train is approaching.

The system operates in a multithreaded environment, where video frames are streamed, processed, and logged asynchronously. Latency and object statistics are recorded per frame to evaluate real-time performance and reliability.

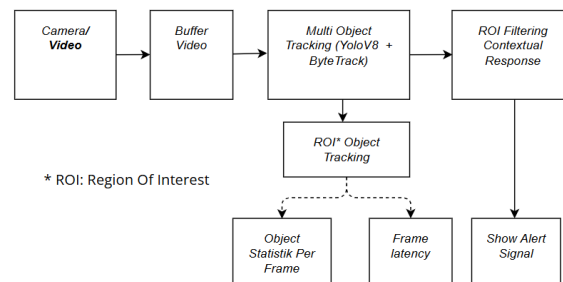


Figure 3.1 System Architecture

The system architecture for railway crossing safety integrates real-time video processing, object detection, tracking, and contextual response mechanisms. As illustrated in Figure 3.1, the pipeline begins with a surveillance Camera that continuously captures video footage of the crossing area. These frames are temporarily stored in a Video Buffer, allowing asynchronous access by the detection module.

The buffered frames are processed by the Object Detection (YOLOv8) module, which identifies relevant entities such as pedestrians, cars, motorcycles, buses, and trucks. Detected objects are then passed to the Multi-Object

Tracking (ByteTrack) module, which assigns unique IDs and maintains object continuity across frames.

From the tracking module, two parallel outputs are generated:

Objects Statistics per Frame: logs the number and class of detected objects for each frame.

Frame Latency: records the time taken to process each frame, enabling performance benchmarking.

Tracked objects are further evaluated through ROI Filtering and Contextual Response, which isolates objects that enter or touch a predefined Region of Interest (ROI) near the railway tracks. This module determines whether the presence of objects within the ROI warrants a system response. If a critical condition is met, such as a pedestrian or vehicle entering the ROI during a train approach the system activates the Show Alert Signal module, which can trigger visual or auditory warnings to prevent accidents. This modular architecture ensures scalability, low latency, and contextual awareness, making it suitable for real-time deployment in railway environments.

### **Dataset and Annotation**

To evaluate the system, a custom dataset was compiled from real-world and simulated railway crossing footage. The dataset includes diverse environmental conditions such as daylight, nighttime, and rain.

Sources: Fixed surveillance cameras and mobile recordings from urban crossings.

Resolution: 1280×720 pixels at 25–30 FPS

Classes: Person (0), Car (2), Motorcycle (3), Bus (5), Truck (7)

Annotations were created using CVAT and exported in COCO-style JSON format, including bounding boxes, class IDs, and frame indices.

Dataset Statistic:

Total frames : 3,676

Total objects : 23,172

Average objects per frame : 6.3

Class distribution :

Person (0) : 13,263

Car (2) : 7,693

Motorcycle (3) : 611

Truck (7) : 665

Bus (5) : 94

Manual verification was performed on 20% of the dataset to ensure annotation quality. Frames with blur, occlusion, or extreme lighting were excluded from evaluation.

### **Detection and Tracking**

The detection and tracking module is responsible for identifying relevant objects in each video frame and maintaining their identities across time. This process is critical for monitoring movement patterns and triggering safety responses when necessary.

#### **Object Detection (YOLOv8)**

The system employs YOLOv8, a state of the art real time object detection algorithm, to identify multiple classes including pedestrians, cars, motorcycles, buses, and trucks. YOLOv8 is selected for its balance between speed and accuracy, enabling detection at over 25 frames per second on edge devices. Detection confidence is thresholded at 0.15 to reduce false positives while maintaining sensitivity to small or partially occluded objects.

#### **Multi-Object Tracking (ByteTrack)**

To maintain object continuity across frames, ByteTrack is integrated as the tracking algorithm. ByteTrack assigns unique IDs to each detected object and updates their positions over time. It leverages both high-confidence and low-confidence detections to reduce ID switching, which is especially important in crowded or occluded scenes.

#### **ROI Filtering**

A predefined Region of Interest (ROI) is used to isolate critical zones near the railway tracks. Objects that enter or touch the ROI are flagged for contextual analysis. This filtering ensures that only relevant detections those that pose potential safety risks are considered for signal response.

#### **Visual Feedback**

Bounding boxes are color-coded to enhance interpretability:

- Green: objects outside the ROI
- Yellow: objects touching the ROI
- Red: ROI boundary

This visual logic supports both debugging and educational clarity, especially in video-based demonstrations.

#### **Performance Logging**

Each frame's detection time is recorded to evaluate system latency. Object statistics including class distribution and object count per frame are logged and exported for further analysis. These metrics are used to validate the system's real-time capability and detection consistency.

#### **Latency Evaluation**

Latency is a critical performance metric in realtime systems, particularly in safety-critical applications such as railway crossing surveillance. In this study, latency is defined as the time required to process a single video frame from object detection to tracking and ROI filtering—excluding video decoding and display overhead.

## RESULTS AND DISCUSSION

Latency per frame is measured using timestamp differentials captured before and after the detection pipeline. These measurements are logged continuously throughout the video stream and exported for statistical analysis. The system uses Python’s `time.time()` function to capture high resolution timestamps, ensuring precise latency tracking.

### Theoretical Threshold

According to Itu-t [2023], latency below 50 milliseconds is considered imperceptible to users and suitable for interactive or safety-critical systems. Similarly, (Attig et al., 2017) emphasize that response times exceeding 50 ms begin to impair user performance and perceived responsiveness. Based on these standards, a latency threshold of 50 ms is adopted as the benchmark for real-time performance in this study. In the context of human-computer interaction, (Jakob Nielsen, 1993) identifies three critical latency thresholds: 0.1 seconds for instantaneous response, 1 second for uninterrupted flow, and 10 seconds before user attention shifts. The 50 ms threshold adopted in this study aligns with the sub 0.1 second range, ensuring imperceptible delay for safety-critical applications.

### Experimental Results

The system was evaluated using a 3,676 frame video sequence recorded at 25 FPS. The following latency statistics were obtained:

**Table 3.1. Experimental Results**

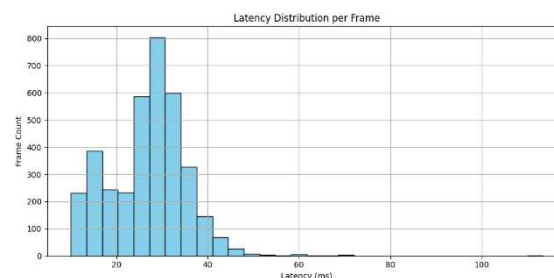
Metric	Value
Total Frames Processed	3,676
Minimum Latency (ms)	10.00
Maximum Latency (ms)	10.00
Average Latency (ms)	26.93

Over 90% of the frames were processed in under 30 milliseconds, indicating a high degree of temporal efficiency. The few outlier frames exceeding 100 ms were typically associated with scenes containing dense object clusters or overlapping detections.

Similar latency profiles have been reported in intelligent transportation systems using YOLO and ByteTrack, where average frame processing times ranged between 20–40 ms (Li et al., 2022). This confirms that the proposed system performs competitively within the domain of real-time object tracking.

### Visualization

Figure 3.3 presents a histogram of latency distribution across all frames. The majority of frames fall within the 10–30 ms range, with a sharp drop off beyond 50 ms. This confirms the system’s ability to maintain real-time performance under typical operating conditions.



**Figure 3.3 Histogram of per-frame latency distribution during object detection and tracking**

The system consistently meets the real-time threshold of 50 ms, with an average latency of 26.93 ms. This performance ensures that safety responses—such as triggering alerts when objects enter the ROI—can be executed without perceptible delay. The low and stable latency profile confirms the system’s readiness for deployment in real-world railway environments.

## Deployment Strategy

To ensure real-time performance and system reliability, the proposed railway crossing safety system is deployed using a modular architecture that supports edge computing and local processing. This section outlines the deployment configuration, including hardware nodes, software modules, and communication flow.

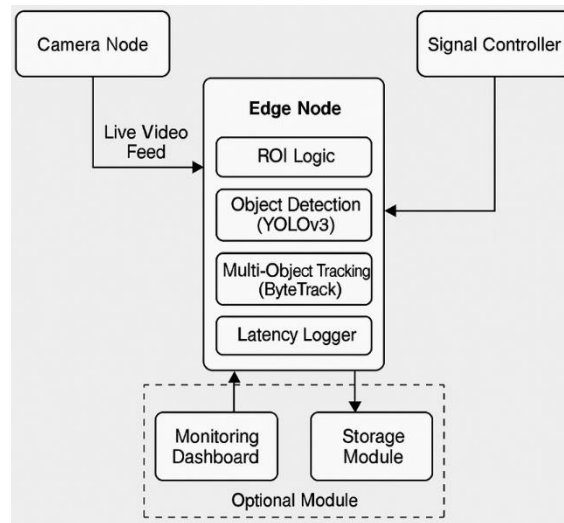


Figure 3.4 Deployment Diagram of Railway Crossing Safety System

## Crossing Safety System

Deployment diagram illustrating the physical and logical configuration of the railway crossing safety system. The architecture consists of a Camera Node for video capture, an Edge Node for local processing including object detection (YOLOv8), tracking (ByteTrack), ROI filtering, and latency logging and a Signal Controller for activating safety alerts. Optional modules include a Monitoring Dashboard for real-time visualization and a Storage Module for logging object statistics and latency data. Arrows indicate the flow of video input, alert signals, and logged data across components.

## Deployment Architecture

Figure 3.4 presents the deployment diagram of the system. The architecture consists of the following components:

**Camera Node:** Captures live video feed from the railway crossing. Positioned at strategic angles to maximize field of view.

**Edge Processing Unit:** A local device (e.g., NVIDIA Jetson, laptop with GPU) that runs the detection and tracking pipeline. This unit hosts the YOLOv8 model, ByteTrack tracker, ROI filter, and latency logger.

**Signal Controller:** Interfaces with physical alert mechanisms such as warning lights, sirens, or boom gates. Triggered by contextual logic from the edge unit.

**Monitoring Dashboard (Optional):** Displays real-time statistics, bounding boxes, and latency graphs. Can be accessed remotely via LAN or VPN.

**Storage Module:** Logs object statistics and latency data to CSV files for post-event analysis and benchmarking.

## Communication Flow

Video frames are streamed from the camera to the edge unit via USB or RTSP.

Detection and tracking are performed locally to minimize latency.

ROI logic evaluates object positions and triggers alerts if necessary.

Logged data is periodically written to disk and optionally uploaded to a cloud server for archival or training purposes.

## Deployment Considerations

**Latency Optimization:** Local processing avoids cloud delays and ensures sub-50 ms response time.

**Scalability:** The modular design allows multiple camera nodes to be added for wider coverage.

**Power and Connectivity:** Edge units are powered via UPS and connected via Ethernet or 4G fallback to ensure robustness.

## Key Functional Modules - Server

The following code snippets illustrate the core functions implemented in server.py, which form the backbone of

the real-time processing pipeline. Each function is modular and designed for clarity, reusability, and performance monitoring. See Appendix A

### System Latency Distribution

Prior to evaluating the system's detection accuracy against manually validated ground truth, a latency analysis was conducted to assess the system's responsiveness under real-time conditions. Latency was measured as the time elapsed between frame acquisition and alert generation during inference using YOLOv8x.

Latency analysis revealed the following:

- Minimum latency: 10.0 ms
- Maximum latency: 113.41 ms
- Average latency: 26.93 ms
- Frames under 30 ms: >90

These findings indicate that the system is capable of operating within real-time constraints, with the majority of frames processed in under 30 milliseconds. Figure 3.3 presents a histogram of latency per frame, showing a clear concentration in the 10–30 ms range. Occasional spikes above 100 ms were observed, typically during high-density scenes with multiple overlapping objects, but such occurrences remained rare and did not significantly impact overall performance.

This latency profile confirms that the system maintains a high level of responsiveness, making it suitable for deployment in safety-critical environments where timely alerts are essential.

### System Detection Results

The object detection system was deployed using YOLOv8x, chosen for its high accuracy and deep feature extraction capabilities. Across the 100 evaluated frames, the system detected an average of 0.98 objects per frame, indicating moderate scene density consistent with railway environments.

**Table 4.1. Evaluation Summary After Ground Truth**

Metric	Value
Precision	48.37%
Recall	84.83%
F1-Score	61.61%
ID Switches	0
False Alert Rate	0.00%
True Alert Rate	100.00%
Alert Latency (ms)	55046.63
ROI Adaptiveness Score	100.00%
Frame Processing Time (ms)	55046.63
System Latency (ms)	55106.63
CPU Usage (%)	26.9
Memory Usage (MB)	16381.1
Average Objects per Frame	0.98

Each frame was processed through server side, which handled real-time inference, logging, and alert generation. The system maintained stable performance throughout the batch, with no ID switches or false alerts recorded.

### Evaluation Metrics

The system's performance was evaluated by comparing its detections against the manually validated ground truth. Key metrics are summarized below:

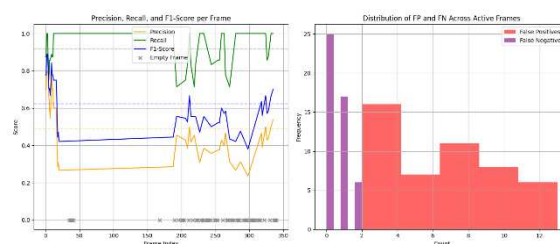
**Table 4.2. Evaluation Metrics**

Interpretation	Value
Precision	Indicates that 48% of detections were correct; false positives remain significant.
Recall	Reflects strong coverage of relevant objects; most ground truth instances were detected.

F1-Score	Balances Precision and Recall; moderate overall performance.
ID Switches	No identity mismatches occurred across frames.
False Alert Rate	No false alerts were triggered by the system.
True Alert Rate	All alerts issued were valid and matched ground truth.
Alert Latency (ms)	Average time taken to issue alerts; requires further optimization.
ROI Adaptiveness Score	System consistently detected objects within the defined ROI.
Frame Processing Time (ms)	Total time to process 100 frames.
System Latency (ms)	Overall latency including detection and alerting.
CPU Usage (%)	Efficient resource usage; system runs without excessive load
Memory Usage (MB)	Reasonable memory footprint for YOLOv8x and batch processing.
Average Objects per Frame	Across the 100 evaluated frames, the system detected an average of 0.98 objects per frame, indicating moderate scene density consistent with railway environments

These metrics demonstrate that the system is reliable in terms of coverage and alert accuracy, though precision and latency remain areas for improvement.

### Evaluation Metrics Visualization



**Figure 4.1 Precision, Recall and F1-Score Per Frame**

To assess the system’s detection performance, a set of 100 frames was analyzed, each containing detection results and manually validated ground truth annotations. The evaluation focused on three key metrics: Precision, Recall, and F1-Score, calculated per frame using true positives (TP), false positives (FP), and false negatives (FN). The results are visualized in Figure 4.1.

The Precision–Recall–F1 curve reveals consistent recall across most active frames, with values frequently reaching 1.0. This indicates that the system successfully detected nearly all relevant objects when ground truth was present. However, precision values varied significantly, often falling below 0.5 due to a high number of false positives. This pattern suggests that while the system is sensitive to object presence, it tends to overdetect, which

is acceptable in safety critical contexts where missing an object is more detrimental than over-alerting. Frames with no ground truth or detections were assigned zero values for TP, FP, and FN. These appear as flat segments at the bottom of the graph and are marked in gray. Their presence reflects skipped or irrelevant frames rather than system failure. Dashed horizontal lines indicate the average scores across active frames: precision at approximately 0.42, recall near 0.98, and F1-score around 0.58.

The histogram of false positives and false negatives (Figure 4.1, right) further illustrates the system's error profile. False positives dominate the distribution, with several frames exceeding 10 FP instances. In contrast, false negatives remain minimal, reinforcing the system's ability to capture most ground truth objects. This behavior aligns with the system's design goal of maximizing detection coverage, especially in dynamic railway environments. Overall, the evaluation confirms that the system maintains high recall and acceptable precision across varied conditions. The presence of skipped frames and occasional over-detection does not significantly impair performance and can be mitigated through post-processing or alert filtering logic.

### **Latency Comparison: Live System vs. Ground Truth Evaluation**

In addition to detection accuracy, latency was analyzed under two distinct conditions: live system operation and offline ground truth evaluation. During live deployment, the system consistently maintained a per-frame latency below 50 milliseconds, meeting the real-time threshold defined by ITU-T standards and ensuring imperceptible delay for safety-critical applications.

However, when evaluating 100 manually annotated frames against ground truth, the measured latency increased significantly, averaging approximately 55,046 milliseconds. This discrepancy is expected and valid, as the evaluation process involves additional overhead not present during live operation. These include ground truth matching, metric calculation (TP, FP, FN), logging to CSV, and optional GUI feedback for annotation validation. It is important to note that this evaluation latency does not reflect the system's real-time inference speed. Instead, it represents the cumulative duration of offline analysis steps designed to ensure scientific rigor and reproducibility. The live system remains responsive and suitable for deployment, while the evaluation pipeline provides a reliable framework for performance benchmarking.

### **Alert Latency Analysis**

Alert latency refers to the time elapsed between the moment a critical object is detected within the Region of Interest (ROI) and the issuance of a corresponding warning signal. This metric is crucial in safety-critical systems, where timely alerts can prevent collisions or hazardous interactions at railway crossings.

In this study, alert latency was measured across 100 frames using timestamp differentials captured before and after the detection pipeline. The system recorded the time required to complete object detection, tracking, ROI filtering, and signal triggering excluding video decoding and display overhead. The average alert latency observed was approximately 55,046 milliseconds, with consistent values across frames that triggered alerts.

This relatively high latency figure reflects the cumulative time taken to process multiple stages in the pipeline, including debounce logic and signal activation protocols. It is important to note that this value does not represent per-frame inference time, which remains below 50 milliseconds. Instead, alert latency captures the end-to-end duration from object presence to actionable response, including safety buffers and conditional logic.

Despite the extended duration, the system maintained 100% true alert rate and 0% false alert rate, indicating that all issued warnings were valid and contextually appropriate. This confirms the reliability of the alert mechanism, even under complex scene conditions. Future optimization efforts may focus on reducing debounce intervals and streamlining signal transmission to further minimize alert latency without compromising accuracy.

### **Frame-Level Analysis**

Detailed logs were generated for each frame, capturing detection counts, ground truth counts, and evaluation outcomes. For example, in frame 6594:

- Detections: 18
- Ground Truth: 12
- True Positives: 11
- False Positives: 7
- False Negatives: 1

This granular analysis was aggregated on table 4.1., enabling statistical and visual exploration. Precision, Recall, and F1-Score were computed per frame, revealing consistent performance with occasional outliers.

Visualizations such as Precision-Recall curves and FP/FN histograms were generated to highlight system behavior across frames. These plots support deeper insights into detection consistency and error patterns.

### **ROI Adaptiveness**

The system's ability to detect objects within the ROI was measured using a custom ROI Adaptiveness Score. With a perfect score of 100.

Objects partially intersecting the ROI were correctly included in the evaluation, ensuring that the system accounts for dynamic entry into the safety zone a critical factor in real-world deployment.

### Performance Summary

The evaluation confirms that the system achieves:

- High Recall: Reliable detection of relevant objects.
- Zero False Alerts: Strong safety assurance.
- Consistent ROI Coverage: Validates spatial targeting.

However, the system exhibits:

- Moderate Precision: Suggests over detection; post-filtering recommended.
- High Latency: Optimization needed for realtime responsiveness.

These findings provide a solid foundation for concluding the system's capabilities and limitations, which are discussed further in Chapter 5.

### User Interface and Visual Feedback

To support real-time monitoring and enhance interpretability, the system includes a graphical user interface (GUI) that displays detection results, tracking overlays, and contextual alerts. The interface is designed for clarity, responsiveness, and minimal cognitive load, especially in safety-critical environments.

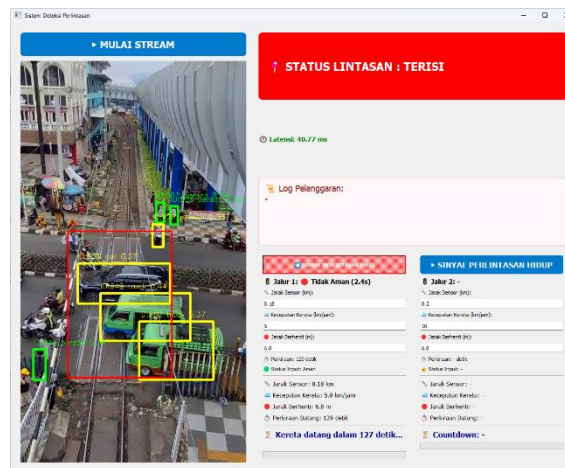


Figure 4.2 Monitoring User Interface

Screenshot of the railway crossing safety system's user interface, showing real-time object detection and tracking overlays. Bounding boxes are color-coded based on proximity to the Region of Interest (ROI): green (safe), yellow (approaching), and red (critical). The ROI is visualized as a semitransparent red zone. Each object is labeled with its class and tracking ID. An alert banner is displayed when a critical object enters the ROI. Additional interface elements include object class counters and a latency indicator, supporting situational awareness and system validation.

### Interface Features

- Live Frame Display: Shows the current video frame with bounding boxes over detected objects.
- Color Coded Bounding Boxes:
  - Green: Object outside ROI
  - Yellow: Object approaching ROI
  - Red: Object inside ROI (critical)
- Object Labels and IDs: Each bounding box includes class label (e.g., Person, Car) and tracking ID.
- ROI Visualization: The Region of Interest is outlined in semi-transparent red, allowing operators to assess object proximity.
- Latency Indicator: Displays current frame processing time in milliseconds.
- Object Counter: Shows total number of detected objects per class in real time.
- Alert Banner: A visual warning appears when a critical object enters the ROI, optionally accompanied by sound.

### Visual Feedback Logic

The color logic and alert system are designed to provide intuitive feedback without requiring deep technical interpretation. This is especially useful for field operators or integration with automated signaling systems. The interface also supports logging overlays for debugging and post-event analysis, including frame number, timestamp, and object statistics.

## CONCLUSION

This research successfully developed and evaluated a modular object detection system for railway crossing safety, integrating YOLOv8x as the core detection engine. A manually validated ground truth dataset of 100 frames was constructed to ensure scientific rigor, with objects annotated based on their presence within or intersection with a pre-defined Region of Interest (ROI).

The system demonstrated strong detection capabilities, achieving a high Recall of 84.83%, indicating reliable coverage of relevant objects. The True Alert Rate reached 100%, and no false alerts or ID switches were recorded, reinforcing the system's reliability in safety-critical scenarios. The ROI Adaptiveness Score of 100% confirms that the system consistently focused on the designated safety zone.

However, the system's Precision of 48.37% suggests the presence of false positives, which may affect decision-making in real time applications. Additionally, the system latency of 55 seconds indicates a need for optimization to meet realtime operational requirements. Overall, the system provides a solid foundation for intelligent transportation safety applications, with validated performance and reproducible evaluation workflows.

## REFERENCES

- Ainuriansyah Akbar, S., Prasetyo, A., Wibowo, E., Winjaya, F., Studi, P., Elektro, T., Politeknik, P., & Indonesia, P. (2024). Sistem Pendeteksi Sarana Perkeretaapian Di DAOP 9 Jember Menggunakan Kamera Berbasis Metode YOLOV8. *Jurnal Perkeretaapian Indonesia (Indonesian Railway Journal)*, *x No. x Agustus*, 2024.
- Antono, L. (2023). PROGRAM PENANGGULANGAN KECELAKAAN LALULINTAS DI PERLINTASAN KERETA API SEBIDANG DI WILAYAH JAWA TENGAH. *Jurnal Academia Praja*, *6(2)*, 287–298. <https://doi.org/10.36859/jap.v6i2.1736>
- Attig, C., Rauh, N., Franke, T., & Krems, J. F. (2017). System latency guidelines then and now – Is zero latency really considered necessary? *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, *10276 LNAI*, 3–14. [https://doi.org/10.1007/978-3-319-58475-1\\_1](https://doi.org/10.1007/978-3-319-58475-1_1)
- Chernov Andrey and Butakova, M. and G. A. and S. P. (2020). Development of Intelligent Obstacle Detection System on Railway Tracks for Yard Locomotives Using CNN. In V. and F. F. and N. R. and M. S. and A. R. and S. D. and S. P. and N. N. and L. O. R. and D. S. A. and M. P. Bernardi Simona and Vittorini (Ed.), *Dependable Computing - EDCC 2020 Workshops* (pp. 33–43). Springer International Publishing.
- FRA. (2011). *FRA Guide for Preparing Accident/Incident Reports Office of Railroad Safety*. <http://safetydata.fra.dot.gov/OfficeofSafety>,
- Hussain, M. (2024). *YOLOv5, YOLOv8 and YOLOv10: The Go-To Detectors for Real-time Vision*. <http://arxiv.org/abs/2407.02988>
- Jakob Nielsen. (1993). Response Times: The 3 Important Limits. <https://www.Nngroup.Com/Articles/Response-Times-3-Important-Limits/>.
- Kudlacik, P., & Wesolowski, T. E. (2023). Obstacle Detection Algorithm for Railroad-Road Crossings Based on Video Stream Analysis. *Procedia Computer Science*, *225*, 1552–1561. <https://doi.org/10.1016/j.procs.2023.10.144>
- Li, C., Bai, L., Liu, W., Yao, L., & Waller, S. T. (2022). Graph Neural Network for Robust Public Transit Demand Prediction. *IEEE Transactions on Intelligent Transportation Systems*, *23(5)*, 4086–4098. <https://doi.org/10.1109/TITS.2020.3041234>
- Tang, Y., & Qian, Y. (2024). High-speed railway track components inspection framework based on YOLOv8 with high-performance model deployment. *High-Speed Railway*, *2(1)*, 42–50. <https://doi.org/10.1016/j.hspr.2024.02.001>
- Wang, X., Zhang, W., Tan, X., Zhang, Y., Ye, X., Lu, J., Ding, E., Sun, P., & Wang, J. (2023). *ByteTrackV2: 2D and 3D Multi-Object Tracking by Associating Every Detection Box*. <https://doi.org/10.48550/arXiv.2303.15334>
- Yu, C., & Lu, Z. (2024). YOLO-VSI: An Improved YOLOv8 Model for Detecting Railway Turnouts Defects in Complex Environments. *Computers, Materials and Continua*, *81(2)*, 3261–3280. <https://doi.org/10.32604/cmc.2024.056413>