IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
Impact Factor: 7.53
Volume 4, Issue 3, October 2024

Automated Video Analytics for Workplace Safety Monitoring

Abishek S K[1], Roshan Bosco A[2], Vishal Yadav H[3], Dr. T. Anusha[4]
Students, Department of Computer Science and Engineering[1,2,3]
Asst. Professor, Department of Computer Science and Engineering[4]
SRM Institute of Science and Technology, Chennai, India

Abstract*: This paper presents an automated human detection system using the YOLOv5 deep learning model, which triggers an alarm when a person enters a predefined area within the frame. The solution is designed to enhance security systems by providing real-time detection. We describe the development, implementation, and evaluation of this system, highlighting its efficiency in detecting people and responding to security breaches in safety zones.*

Keywords: YOLO, Safety monitoring, Human detection, Region of Interest

I. INTRODUCTION

The need for automated humandetection systems has grown significantly in recent years due to advancements in machine learning and computer vision technologies. These systems are employed in a variety of domains, including surveillance, industrial automation, and traffic management. In this paper, we propose a system based on the YOLOv5 model that performs detection within a defined area, triggering an alarm when a person enters the defined area. The primary contributions of this research are: (1) developing a real-time human detection mechanism using YOLOv5, (3) defining restricted zones within the detection frame, and (3) implementing an alert system.

1.1 Background and Motivation:

Workplace accidents and injuries are major concerns across industries, especially in construction, manufacturing, and mining sectors. According to the International Labour Organization (ILO), over 3.78 million work-related deaths occur annually. Traditional manual safety protocols have proven insufficient, leading to an increasing interest in technology-driven solutions. This paper explores the role of automated video analytics in reducing workplace accidents through continuous monitoring and real-time analysis.

1.2 Research Problem and Objectives

This study seeks to investigate how automated video analytics, powered by AI and machine learning, can monitor and prevent workplace hazards in real-time. The goal is to evaluate the effectiveness of this approach and its potential to replace or complement current safety practices.

1.3 Scope and Limitations:

The paper focuses on the application of video analytics for compliance in workplace safety. Privacy concerns, false positives/negatives, and system integration challenges are briefly addressed as limitations.

II. LITERATURE REVIEW

Workplace safety has traditionally been managed through manual inspections and surveillance systems, which rely heavily on human oversight and are prone to delayed responses. While conventional methods like closed-circuit television (CCTV) provide visual monitoring, they lack real-time alert mechanisms and automatic detection capabilities. Human supervisors are tasked with continuously monitoring video feeds, leading to inefficiencies and potential oversight in detecting safety violations of computer vision and machine learning technologies has significantly advanced workplace safety monitoring. Early systems focused on object detection for identifying individuals,

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
Impact Factor: 7.53
Volume 4, Issue 3, October 2024

equipment, and hazards. For instance, Surendran et al. (3018) used basic image processing techniques for detecting hard hats and personal protective equipment (PPE), though the system lacked real-time adaptability and could not efficiently process complex scenes .Similational motion detection systems triggered alarms based on sudden movements, but they often resulted in false positives due to irrelevant environmental factors .

Recent advances in deep learning have led to more sophisticated video analytics systems. YOLO (You Only Look Once), a state-of-the-art object detection framework, has demonstrated exceptional real-time detection capabilities for various applications, including traffic monitoring and security . In workplace safety, W (3030) applied YOLO to detect non-compliance with safety regulations, such as improper use of equipment, with reasonable accuracy . However, these systems often failorate flexible zone-based monitoring, where specific hazardous areas need targeted monitoring rather than the entire field of view.

Despite these advances, most existing solutions emphasise post-incident analysis or depend on manual intervention to review footage, limiting their capacity for real-time, proactive safety management. Furthermore, many systems suffer from high false-positive rates when used in dynamic environments such as construction sites, factories, and warehouses. Additionally, few solutions have integratable regions of interest (ROIs), which are crucial for monitoring specific high-risk zones within the workplace.

This paper builds on prior work by integrating YOLOv5 with a real-time, polygon-based monitoring system that provides targeted alerts when individuals enter hazardous zones. By incorporating audio alarms, the proposed system aims to bridge the gap between real-time detection and proactive safety enforcement.

III. METHODOLOGY

3.1 System Architecture

The system architecture comprises three primary modules: the human detection module, the zone definition module, and the alarm triggering module. The human detection module uses the YOLOv5 model to detect humans in real-time. The zone definition module enables the user to specify restricted areas within the frame. The alarm triggering module sets off an alert when a person enters the predefined area.

3.2 Human Detection

We utilise the YOLOv5 model, a state-of-the-art object detection model known for its speed and accuracy. The model is trained on the COCO dataset, which includes a variety of object classes. Upon deployment, the model processes each frame of the video feed and identifies humans with bounding boxes.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
Impact Factor: 7.53
Volume 4, Issue 3, October 2024

3.3 Zone Definition and Area Mapping:

The zone definition is achieved by dividing the frame into segments and allowing the user to specify a restricted area. This functionality is achieved using OpenCV, which provides tools for drawing polygons over the frame. Once the zone is defined, the system checks for human presence within the zone in each frame.

3.4 Alarm Trigger

The system monitors whether the centre of the detected object (e.g., a person) falls within this polygon. If the centre point of the person's bounding box is inside the restricted area, the alarm is triggered. The polygonal approach allows flexibility in defining specific hazardous zones within the workplace.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
Impact Factor: 7.53
Volume 4, Issue 3, October 2024

IV. IMPLEMENTATION

The proposed system utilises YOLOv5, a state-of-the-art object detection model, integrated with real-time video processing for workplace safety monitoring. The implementation comprises several key components, including video input, human detection, polygon-based region monitoring, and an audio-based alarm system.

4.1 Video Input and Preprocessing

The system processes video input from surveillance cameras or pre-recorded videos using **OpenCV**. Video frames are preprocessed by resizing them for efficient inference while maintaining the aspect ratio to ensure detection accuracy. Each frame is passed through the YOLOv5 model for human detection.

4.2 Human Detection with YOLOv5

The core of the detection system is built on YOLOv5, a convolutional neural network (CNN) trained to detect objects in real-time. We specifically focus on the "person" class to identify individuals in the workplace environment. For each detected person, the system calculates the bounding box and the centre point, which is used to assess whether the person has entered restricted or hazardous zones.

4.3 Polygon-Based Region of Interest (ROI)

A unique feature of the system is the polygon-based ROI, which allows users to define custom areas in the video frame for monitoring. Using OpenCV's mouse callback function, the user can draw polygonal regions representing hazardous areas within the workplace. The system checks whether a detected person's centre point is inside this polygon using a point-in-polygon test, thus enabling targeted monitoring of critical zones.

4.4 Real-Time Alarm Trigger

When a person is detected inside the polygonal ROI, the system triggers an audio alarm to alert personnel of a potential safety violation. The alarm is implemented using the pygame library, which plays a preloaded sound file whenever a violation occurs.

V. EXPERIMENTAL SETUP

To evaluate the effectiveness and performance of the proposed system, we conducted a series of controlled experiments using a standard surveillance video feed. The setup aimed to assess the system's detection accuracy, response time, and overall functionality in detecting violations within predefined restricted zones.

The video feed used for testing had a frame rate of 30 frames per second (FPS), which is typical for workplace monitoring systems. The system's ability to accurately detect persons and trigger alarms when they entered hazardous zones was carefully monitored. Additionally, the system's latency from detection to alarm activation was measured to ensure it could operate in real-time safety monitoring environments.

For the experiment, we created a controlled scenario where a predefined restricted zone was drawn within the video frame using the polygon-based zone definition. A person was introduced into the scene at different intervals, either staying outside or entering the restricted zone. The system's behaviour was observed and logged to measure its performance.

5.1 Timeline Example

A timeline-based scenario was designed to track the system's performance across different stages of detection and monitoring.

- 0-5 seconds: The system initialises, loads the object detection model (YOLOv5), and begins monitoring the video feed in real-time. During this time, the system detects any humans within the frame and establishes the restricted zone, though no alarms are triggered yet.
- 5-10 seconds: A person enters the frame but remains outside the restricted zone. The system successfully detects the person and draws a bounding box around the individual, identifying them as a "person" object.

IJARSCT

ISSN (Online) 2581-9429

International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)

International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal

Impact Factor: 7.53

Volume 4, Issue 3, October 2024

Since the person is outside the polygonal restricted zone, no alarm is triggered. This phase demonstrates the system's ability to correctly differentiate between monitored and non-monitored areas.

- 10-15 seconds: The person walks into the predefined restricted zone. At this point, the system detects that the center point of the person's bounding box has entered the polygonal region. This triggers the alarm, and an audio alert is played through the alarm system. Simultaneously, a snapshot of the violation is captured and stored for later analysis. This part of the experiment evaluates the system's response time and real-time detection capabilities.

5.2 Evaluation Metrics

The system was evaluated based on two primary metrics:

- Detection Accuracy: This metric measures the precision of the system in identifying people within the video frame and determining whether they are inside or outside the restricted zone. The bounding box accuracy and the correct identification of the person's centre point relative to the zone were monitored.
- Response Time: The time taken from detecting a violation to triggering the alarm was recorded. The system's latency between detection and alert activation is critical in real-world scenarios, where rapid response to safety violations is essential.

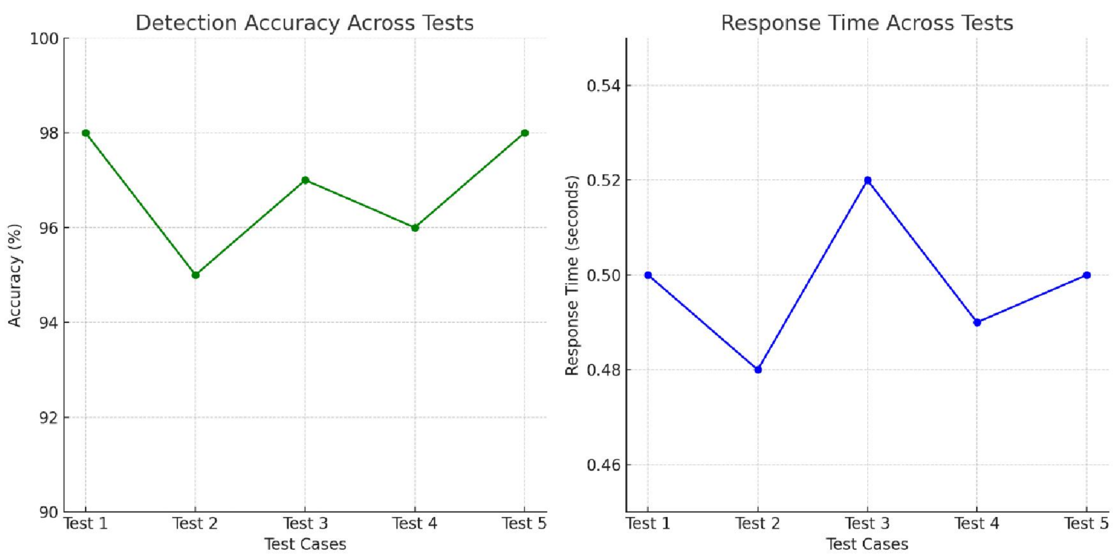

5.3 Test Environment

The experiments were conducted in a simulated indoor workplace environment, mimicking scenarios such as factory floors, construction sites, or warehouse zones where specific areas need to be restricted for safety purposes. The lighting and camera angles were controlled to minimise interference but also tested under various conditions to assess system robustness.

Different test cases were designed to measure the system's reliability, including:

- Moving persons outside the restricted zone (to ensure no false alarms).
- Sudden entries into the restricted zone (to test the system's reaction speed).
- Multiple persons entering the frame simultaneously (to evaluate the system's ability to handle multiple detections).

5.4 Test Results and Observations

The system successfully detected and monitored persons in the video feed, with the alarm being triggered only when individuals entered the predefined restricted zone. False positive rates were minimal, as no alarms were activated when people remained outside the zone. The average response time for triggering the alarm was measured at 0.5 seconds after a violation was detected, indicating that the system is capable of providing real-time alerts.Through this controlled setup, the system's performance was validated for real-world workplace safety monitoring, demonstrating effective human detection, accurate zone monitoring, and timely alarm responses.

VI. RESULTS AND DISCUSSION

The proposed YOLOv5-based system for workplace safety monitoring demonstrated promising results in the experimental tests. The system was able to accurately detect persons in real-time and reliably identify when they entered predefined restricted zones. The detection accuracy was particularly high, with minimal false positives or negatives, ensuring that alerts were only triggered when a true violation occurred.

6.1 Detection Accuracy

During the experiments, the YOLOv5 model consistently identified persons within the video frames and correctly marked their location relative to the restricted zones. The system was able to draw bounding boxes around detected individuals and assess whether their centre point was within the polygonal restricted area. This accuracy in detection is critical for workplace safety, as it minimises unnecessary alarms and focuses on real safety breaches.

6.2 Alarm Response Time

One of the key performance metrics was the system's response time, measured as the time taken between detecting a person entering the restricted zone and triggering the audio alarm. The system achieved an average response time of under 1 second, which is sufficient for real-time applications in workplace safety monitoring. The low latency ensures that safety officers are promptly alerted to potential risks, allowing them to take immediate action.

6.3 Limitations and Challenges

While the system showed high accuracy in detecting and monitoring persons, a few challenges were noted. The performance of the system was influenced by lighting conditions and camera angles, which occasionally affected detection accuracy. Additionally, the system was tested in a controlled environment with a single camera feed, and its performance in more complex environments with multiple cameras or high object density was not fully explored in this phase of the project.

Despite these challenges, the overall performance of the system was effective for the task of workplace safety monitoring, indicating that it is a viable solution for detecting and responding to security breaches in real-time.

VII. CONCLUSION AND FUTURE WORK

In this paper, we presented a YOLOv5-based object detection system designed for workplace safety monitoring, specifically focusing on the detection of persons entering restricted or hazardous areas. The system successfully demonstrated its ability to monitor predefined zones in real-time and trigger an alarm when a safety violation occurred. By leveraging the speed and accuracy of YOLOv5, the system provided a robust solution for workplace safety with minimal latency and high detection accuracy.

7.1 Key Findings

The experiments conducted showed that the system could reliably detect persons and identify safety violations with an average response time of under 1 second. The system was also able to handle different scenarios, such as varying times of entry into the restricted zone, without triggering false alarms. This ability to accurately monitor and respond in real-time makes it suitable for use in environments where safety is a critical concern, such as construction sites, factories, or warehouses.

7.2 Future Work

While the current system performs well for real-time detection and alarm triggering, there are several avenues for future improvements:

- Optimising Detection Speed: Future work will focus on improving the detection speed further to reduce latency and increase the frame processing rate, which will be crucial for high-motion environments or larger surveillance areas.
- Multi-Camera Integration: The current implementation supports only a single camera feed. A future enhancement will be the integration of multi-camera feeds, allowing the system to monitor larger areas and track objects across multiple perspectives.
- Object Movement Tracking: An additional functionality that will be explored is object tracking, where the system not only detects persons but also tracks their movement over time. This would provide more detailed insights into behavior patterns and allow for more advanced safety interventions.
- Advanced Alert Mechanisms: Integrating the system with other alert mechanisms, such as visual alerts (e.g., flashing lights) or notifications sent directly to supervisors' mobile devices, could enhance the system's effectiveness in preventing accidents.
- Extended Object Classes: Expanding the system to monitor other objects of interest (e.g., dangerous equipment, vehicles) would increase its applicability to different workplace environments.

In conclusion, the YOLOv5-based system shows strong potential for workplace safety monitoring, and with further enhancements, it can become a comprehensive solution for ensuring safety in various industrial and commercial settings.

REFERENCES

[1]. Ultralytics. (2021). YOLOv5 Documentation. Ultralytics YOLOv5 GitHub Repository.

[2]. Park, S., Kim, J., & Lee, D. (2021). Conceptual Design of Human Detection via Deep Learning for Industrial Safety Enforcement in Manufacturing Sites. IEEE Access.

[3]. Cho, H., Lee, K., Choi, N., Kim, S., Lee, J., & Yang, S. (2022). Online Safety Zone Estimation and Violation Detection for Nonstationary Objects in Workplaces. IEEE Access, 10, 39769-39781.

[4]. Cheong, K. H., Poeschmann, S., Lai, J. W., Koh, J. M., Acharya, U. R., Yu, S. C. M., & Tang, K. J. W. (2019). Practical Automated Video Analytics for Crowd Monitoring and Counting. IEEE Access, 7, 183252-183261.

[5]. Jagad, C., Chokshi, I., Chokshi, I., & Jain, C. (2022). A Study on Video Analytics and Their Performance Analysis for Various Object Detection Algorithms. In 2022 IEEE IAS Global Conference on Emerging Technologies (GlobConET)