Real Time PPE Compliance Detection Using YOLOv8n on Raspberry Pi 4

Wahyu Nur Hidayat[1], Sapto Wibowo[1], Moechammad Sarosa[1], Erdiyan Ariyanda[1]

[1]Department of Electrical Engineering,
State Polytechnic of Malang,
Malang, Indonesia

Abstract

Ensuring compliance with Occupational Safety and Health (OSH) regulations remains a critical challenge in industrial environments, particularly in monitoring the proper use of Personal Protective Equipment (PPE). This study presents a performance evaluation of an embedded IoT-based PPE compliance monitoring system using the lightweight YOLOv8n object detection model. The proposed system is designed to detect six PPE categories, namely helmet, safety vest, gloves, shoes, goggles, and mask, and to determine OSH Standard Operating Procedure (SOP) compliance based on multi-object detection results. Experiments were conducted using six test scenarios captured under controlled conditions, with camera distances ranging from 3 to 5 meters, frontal and side-view angles, normal to bright lighting, and natural partial occlusion. Detection performance was evaluated using precision, recall, F1-score, and mean Average Precision (mAP@0.5:0.95), while system-level performance was assessed through SOP compliance accuracy, false alarm rate, and real-time inference speed. The results demonstrate that YOLOv8n achieves stable real-time performance on an embedded platform, with inference speeds between 8.4 and 11.6 frames per second. However, detection performance varies across PPE classes. While helmet, safety vest, gloves, and mask are detected reliably, goggles and shoes show low detection accuracy under occlusion and side-view conditions. These limitations result in a high false alarm rate of 83.3 percent and an overall SOP compliance accuracy of 16.6 percent. The findings highlight the challenges of using lightweight object detectors for strict multi-PPE compliance monitoring and emphasize the need for more robust detection and decision strategies.

Keywords: YOLOv8n, Personal Protective Equipment, Occupational Safety and Health, Embedded Vision, SOP Compliance Detection.

1. Introduction

Occupational Safety and Health (OSH) regulations play a critical role in reducing workplace accidents and ensuring worker protection, particularly through the mandatory use of Personal Protective Equipment (PPE). Despite the existence of well-established safety standards, PPE non-compliance remains a persistent problem in industrial environments, as reported across construction, manufacturing, and healthcare sectors (Yankson et al., 2021; Ahmed et al., 2023). This issue is often intensified by limited supervisory capacity, large-scale operational areas, and the inherent subjectivity and inconsistency of manual inspection, which collectively reduce the effectiveness of conventional monitoring approaches (Anjum, 2020; Isror et al., 2024).

Recent advances in deep learning and computer vision have enabled the development of automated visual monitoring systems capable of detecting PPE usage in real time. Among various object detection frameworks, the You Only Look Once (YOLO) family has gained widespread adoption due to its favorable balance between detection accuracy and computational efficiency (Önal & Dandıl, 2021; Arip et al., 2024). The latest generation, YOLOv8, introduces architectural refinements that enhance inference speed and detection robustness, making it particularly suitable for real-time and embedded applications (Ultralytics, 2023; Bao et al., 2024). However, the majority of existing studies primarily emphasize object-level detection performance, often reporting high accuracy for selected PPE categories such as helmets or safety vests under controlled conditions (Ji et al., 2023; Bo et al., 2025). Many of these approaches rely on desktop or GPU-based platforms and evaluate performance using standard detection metrics alone. In practical OSH applications, safety compliance is not determined by the detection of individual PPE items, but by the simultaneous presence of multiple required PPE components. As a result, system-level compliance decisions become highly sensitive to even minor detection failures, particularly in multi-PPE scenarios (Gallo et al., 2022; Alashrafi et al., 2025).

Furthermore, limited attention has been devoted to analyzing false alarm behavior, compliance decision accuracy, and performance degradation under realistic operating conditions, including occlusion, side-view camera angles, and constrained computational resources (Shi et al., 2023; Malaikrisanachalee et al., 2025). These factors are crucial for real-world deployment yet remain insufficiently explored in the literature, particularly in the context of embedded systems. To address these limitations, this study presents a comprehensive evaluation of a YOLOv8n-based PPE compliance monitoring system deployed on an embedded IoT platform. Rather than focusing solely on object detection metrics, this work investigates how detection errors propagate into Standard Operating Procedure (SOP) compliance decisions under realistic operating conditions. By jointly analyzing object-level detection performance and system-level decision reliability, this study aims to provide practical insights into the challenges and limitations of multi-PPE compliance monitoring in real-world industrial environments.

2. Methodology and Experimental Setup

2.1 System Overview

The proposed system is an embedded vision-based Occupational Safety and Health (OSH) monitoring framework designed to automatically evaluate compliance with Personal Protective Equipment (PPE) requirements. Embedded computer vision systems have increasingly been adopted for safety monitoring due to their ability to perform real-time analysis directly at the edge, reducing latency and dependency on centralized computing resources (Gallo et al., 2022; Panduman et al., 2024).

The system employs a lightweight YOLOv8n object detection model deployed on a Raspberry Pi, which serves as the primary processing unit. YOLO-based detectors are widely recognized for their balance between detection accuracy and computational efficiency, making them suitable for embedded and real-time applications (Önal & Dandıl, 2021; Ultralytics, 2023). A fixed camera is installed to continuously capture image frames of workers at distances ranging from 3 to 5 meters. All detection and decision-making processes are performed locally on the embedded platform, enabling real-time operation without dependence on external computing resources. Similar edge-based architectures have been reported as effective solutions for industrial PPE monitoring systems (Ji et al., 2023; Bao et al., 2024).

The system is configured to detect six mandatory PPE categories, namely helmet, safety vest, gloves, shoes, goggles, and mask. For each captured frame, detection results are analyzed to determine compliance with Standard Operating Procedure (SOP) requirements. A compliance decision is generated based on the simultaneous detection of all required PPE items, reflecting real-world OSH enforcement practices where incomplete PPE usage is considered a violation (Yankson et al., 2021; Ahmed et al., 2023).

2.2 Data Acquisition and Experimental Conditions

Experiments were conducted using six test scenarios captured under controlled conditions designed to reflect realistic industrial environments. Image acquisition was performed using a fixed camera setup, with worker distances maintained between 3 and 5 meters, which corresponds to typical monitoring distances reported in prior PPE detection studies (Arip et al., 2024; Shi et al., 2023). Variations in camera viewpoint were introduced by capturing frames from both frontal and side-view angles to evaluate detection robustness under non-ideal perspectives.

Lighting conditions ranged from normal indoor illumination to bright lighting, reflecting common workplace environments. Natural partial occlusion occurred during data acquisition, particularly affecting smaller PPE items such as goggles and shoes. Occlusion and viewpoint variation have been identified as major challenges in PPE detection and are therefore critical factors for realistic system evaluation (Malaikrisanachalee et al., 2025; Bo et al., 2025). Each of the six test scenarios represents a distinct operational condition and was used to evaluate both object-level detection performance and system-level SOP compliance decisions.

2.3 YOLOv8n-Based PPE Detection

YOLOv8n was selected as the core detection model due to its lightweight architecture and suitability for deployment on resource-constrained embedded platforms. Recent benchmarking studies have shown that lightweight YOLO variants provide competitive detection accuracy while maintaining real-time inference performance on edge devices (Alashrafi et al., 2025; Shi et al., 2023).

For each input frame, the model generates bounding boxes and associated confidence scores corresponding to detected PPE objects. Detection performance was evaluated using standard object detection metrics, including precision, recall, F1-score, and mean Average Precision computed over intersection-over-union thresholds ranging from 0.5 to 0.95 (mAP@0.5:0.95), following the COCO evaluation protocol widely adopted in computer vision research (Ultralytics, 2023; Bao et al., 2024). In addition to quantitative evaluation, qualitative analysis was conducted by examining bounding box localization accuracy and confidence score distributions for each PPE class across all test scenarios.

2.4 SOP Compliance Decision Logic

Following PPE detection, the system applies a binary decision logic to determine SOP compliance. A captured frame is classified as compliant only when all six PPE categories are successfully detected within the same frame. If at least one required PPE item is not detected, the system classifies the scenario as a violation. This strict decision strategy reflects real-world OSH SOP requirements, where incomplete PPE usage is considered non-compliant (Elshaer & Agage, 2022; Isror et al., 2024).

However, this approach also increases sensitivity to individual detection failures, making system-level compliance decisions highly dependent on the reliability of the weakest PPE detection class. Similar limitations have been reported in multi-PPE monitoring systems employing strict rule-based decision logic (Gallo et al., 2022; Alashrafi et al., 2025).

2.5 Performance Evaluation Metrics

System performance was evaluated at both the object detection level and the system decision level to provide a comprehensive assessment of detection reliability and compliance decision accuracy. Object detection performance was measured using precision, recall, F1-score, and mean Average Precision at multiple intersection-over-union thresholds (mAP@0.5:0.95), which are standard metrics for evaluating object detection models (Bao et al., 2024; Shi et al., 2023).

System-level performance was evaluated through SOP compliance decision accuracy and false alarm rate, which quantify the correctness and robustness of compliance decisions under varying operational conditions. False alarm behavior has been identified as a critical factor influencing the practical usability of automated safety monitoring systems (Yankson et al., 2021; Ahmed et al., 2023). In addition, real-time performance was assessed by measuring inference time per frame and the resulting frame rate in frames per second (FPS). All evaluation metrics were computed across six test scenarios to assess system stability and responsiveness on the embedded platform.

2.5.1 Object Detection Metrics

Object detection performance was evaluated using standard metrics commonly adopted in computer vision research, namely precision, recall, F1-score, and mean Average Precision (mAP@0.5:0.95). Precision and recall are defined as:

$$\text{Precision} = \frac{TP}{TP+FN}$$

$$\text{Recall} = \frac{TP}{TP+FN}$$

where *TP, FP*, and *FN* denote true positives, false positives, and false negatives, respectively. Moreover, the F1-score, which reflects the balance between precision and recall, is calculated as the harmonic mean of the two metrics:

$$\text{F1 - score} = 2 \text{ x } \frac{\text{Precision x Recall}}{\text{Precision x Recall}} \text{ x } 100\%$$

To evaluate overall detection quality across multiple confidence thresholds and localization criteria, the mean Average Precision was computed following the COCO evaluation protocol. Specifically, mAP@0.5:0.95 represents the average precision computed over intersection-over-union thresholds ranging from 0.5 to 0.95 with a step size of 0.05.

2.5.2 SOP Compliance Decision Metrics

System-level performance was evaluated using SOP compliance decision accuracy and false alarm rate. SOP compliance accuracy is defined as the proportion of correct system decisions relative to the total number of evaluated test scenarios:

$$\text{SOP Accuracy} = \frac{N_{correct}}{N_{total}} \text{ x } 100\%$$

where $N_{correct}$ denotes the number of correct compliance decisions and N_{total} denotes the total number of test scenarios. The **false alarm rate** is defined as the proportion of scenarios in which compliant conditions were incorrectly classified as violations:

$$\text{False Alarm Rate} = \frac{N_{false\ alarm}}{N_{total}} \text{ x } 100\%$$

2.5.3 Real-Time Performance Metrics

Real-time performance was assessed by measuring the inference time required to process a single frame and the resulting frame rate. Frame rate is expressed in frames per second and is calculated as:

$$\text{FPS} = \frac{1}{T_{inference}} \text{ x } 100\%$$

where $T_{inference}$ represents the average inference time per frame. These metrics were used to evaluate the feasibility of deploying the proposed system in real-world industrial environments, where timely detection and immediate feedback are critical for effective safety monitoring (Ji et al., 2023; Panduman et al., 2024).

4. Results

4.1 PPE Object Detection Results

The detection performance of the YOLOv8n model was evaluated for six categories of Personal Protective Equipment (PPE), including helmet, safety vest, gloves, shoes, goggles, and mask. The evaluation was conducted using six test scenarios acquired under controlled yet representative industrial conditions. These scenarios encompassed camera distances ranging from 3 to 5 meters, frontal and side-view camera perspectives, normal to bright illumination, and naturally occurring partial occlusion. Similar evaluation settings have been widely adopted in recent PPE detection studies to assess robustness under realistic operating conditions (Ji et al., 2023; Arip et al., 2024; Shi et al., 2023).

The original camera image for Test Data 1, which serves as the baseline input prior to object detection, is shown in Figure 1. The associated detection results generated by the YOLOv8n model are presented in Figure 2, where PPE items are delineated using bounding boxes accompanied by confidence scores. Under this baseline condition, the model demonstrates reliable detection of helmets, safety vests, gloves, and masks, indicating strong performance for visually prominent PPE components. This observation is consistent with prior studies reporting high detection accuracy for larger PPE items using YOLO-based models (Gallo et al., 2022; Bao et al., 2024).

Figure 1. Baseline

Figure 2. Test Data of 1

The consistency of detection performance under identical acquisition conditions is demonstrated in Figure 4, which presents the detection output for Test Data 2. The close correspondence between the results shown in Figures 2 and 4 confirms that the YOLOv8n model produces stable and repeatable detection outcomes when environmental parameters remain unchanged. Similar stability has been reported for lightweight YOLO variants deployed on embedded or edge platforms (Alashrafi et al., 2025; Shi et al., 2023).

Figure 3. Test Data of 2

More demanding detection conditions are illustrated in Figure 5a for Test Data 3, where PPE items are distributed more sparsely across the scene. Although the majority of PPE objects are still detected, performance degradation begins to emerge for smaller and less visually distinctive items. These limitations become more pronounced under non-frontal camera perspectives, as shown in Figure 5b for Test Data 4. In this scenario, partial occlusion significantly affects detection reliability, particularly for shoes, which are frequently missed (Fig. 5). Similar challenges related to occlusion and viewpoint variation have been widely reported in PPE detection literature (Malaikrisanachalee et al., 2025; Bo et al., 2025).

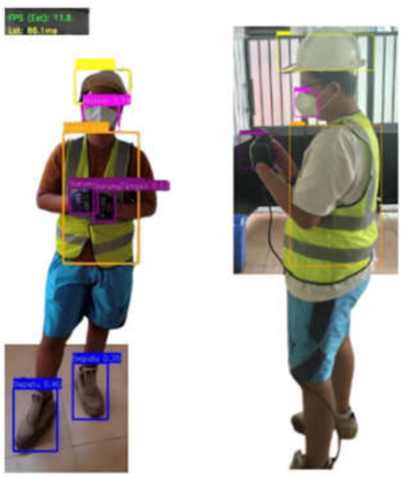

(a) (b)

Figure 4. The detection results on (a) Test data of 3 and (b) Test data of 4

Figure 6 presents the detection results obtained under operational variation in a vision-only configuration for Test Data 5. Similar detection limitations are observed, reinforcing the conclusion that the identified weaknesses are systematic rather than scenario-specific. Overall, detection performance exhibits substantial variability across PPE categories, with goggles and shoes consistently demonstrating the lowest levels of detection reliability. This class-dependent performance gap is commonly observed in multi-PPE detection tasks, particularly for small or visually ambiguous objects (Bao et al., 2024; Ahmed et al., 2023).

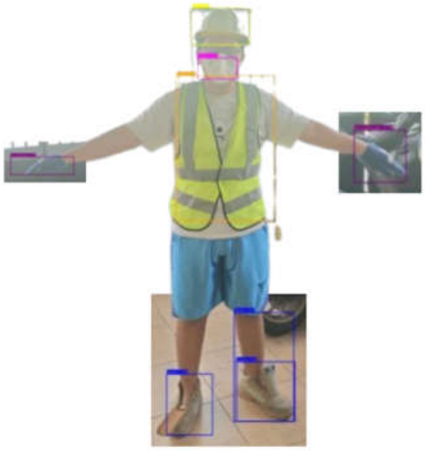

Figure 5. Test Data of 5

4.2 Quantitative Detection Performance Metrics

Detection performance was further quantified using object-level statistical analysis derived from the six test scenarios. Class-wise detection behavior is summarized using a normalized confusion matrix (Fig. 7). The confusion matrix reveals a pronounced imbalance in detection performance across PPE categories, with false negatives predominantly occurring in the goggles and shoes classes. This observation indicates that these PPE items are the primary contributors to missed detections in the proposed system, which aligns with findings reported in recent PPE benchmarking studies (Alashrafi et al., 2025; Shi et al., 2023).

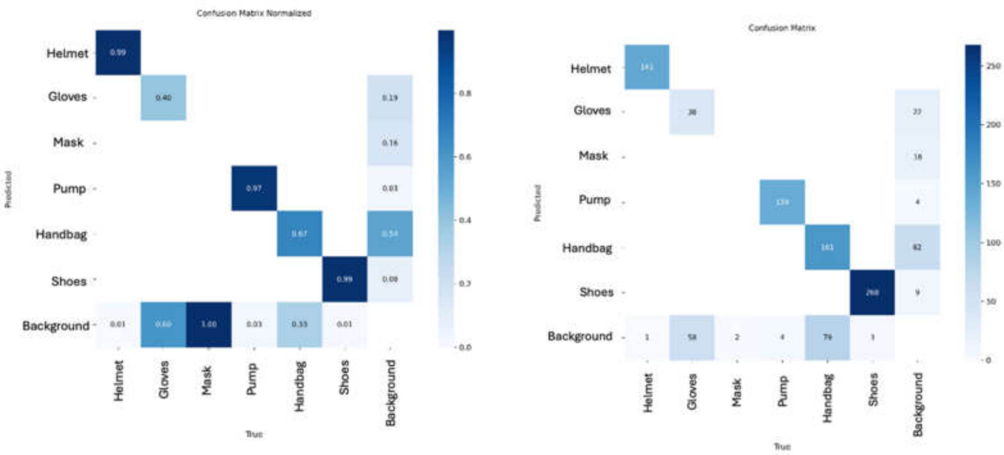

Figure 6. Confusion matrix of YOLOv8n-based PPE detection across six classes

To analyze detection performance across varying confidence thresholds, Precision–Recall (PR) curves are (Fig. 8). The PR curves demonstrate stable and well-balanced precision–recall characteristics for larger and more visually distinctive PPE items, such as helmets, safety vests, gloves, and masks. In contrast, degraded PR curves are observed for goggles and shoes, reflecting reduced detection reliability for smaller, less salient, or partially occluded objects. Similar PR behavior has been reported for small-object PPE detection in complex industrial scenes (Malaikrisanachalee et al., 2025; Bao et al., 2024).

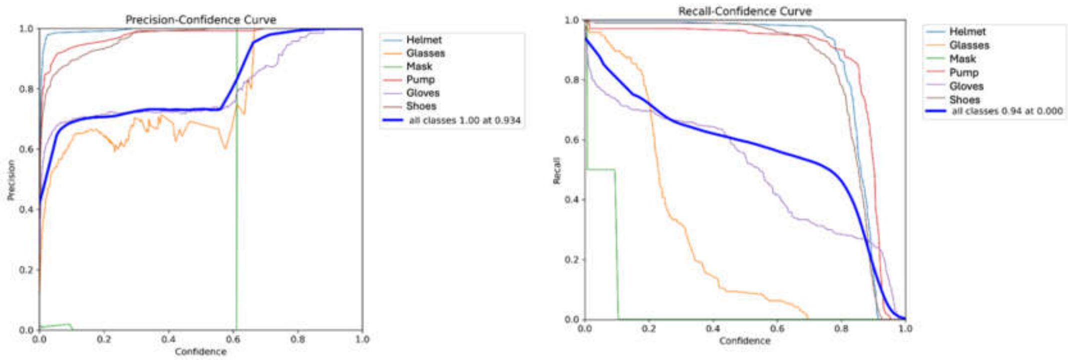

Figure 7. Precision–Recall curves for YOLOv8n PPE detection across six PPE classes

To further examine the balance between precision and recall at the class level, the F1-score for each PPE category (Fig. 9). The F1-score distribution confirms strong detection performance for helmets, safety vests, gloves, and masks, while significantly lower F1-scores are observed for goggles and shoes. This result demonstrates that the observed detection imbalance is consistent across multiple evaluation metrics and is not limited to a single performance indicator, as also noted in prior YOLOv8-based PPE studies (Shi et al., 2023; Alashrafi et al., 2025).

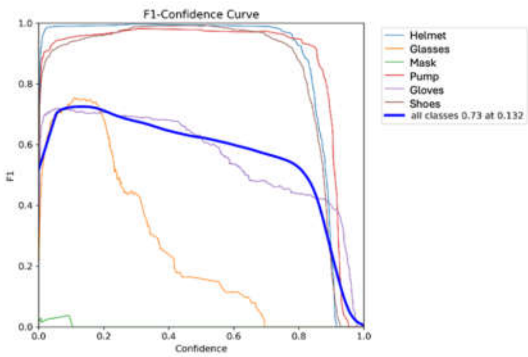

Figure 8. F1-score distribution per PPE class for YOLOv8n detection**.**

In addition to class-wise evaluation, the overall detection quality of the YOLOv8n model was assessed using the mean Average Precision metric computed over intersection-over-union thresholds ranging from 0.5 to 0.95. The resulting mAP@0.5:0.95 performance representing the aggregated detection accuracy across all PPE classes and confidence thresholds (Fig. 10). As reported in previous benchmarking studies, aggregated mAP values are strongly influenced by the weakest detection classes, particularly in multi-PPE detection scenarios (Bao et al., 2024; Alashrafi et al., 2025).

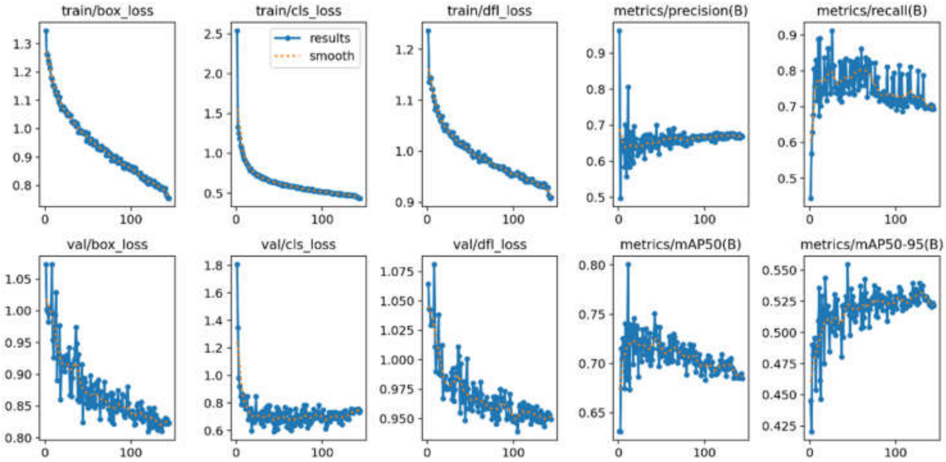

Figure 9. Overall mAP@0.5:0.95 performance of YOLOv8n for PPE detection.

4.3 SOP Compliance Decision Results

Based on the object detection outputs, the system evaluated Occupational Safety and Health (OSH) Standard Operating Procedure (SOP) compliance using a strict binary decision rule. A worker was classified as compliant only when all six required PPE items were detected simultaneously within a single frame; otherwise, the system classified the scenario as a violation. Similar rule-based compliance strategies have been adopted in previous PPE monitoring systems to reflect real-world safety enforcement practices (Yankson et al., 2021; Gallo et al., 2022). The compliance visualization for Test Data 1 is shown in Figure 3, where the detection outputs are translated into a system-level compliance decision. The SOP compliance decision results for all six test scenarios are summarized in Table 2 below.

Table 2. summarizes the SOP compliance decision results for all test scenarios.

Test Data	**System Decision**	**Ground Truth**
1	Violation	Violation
2	Compliant	Compliant
3	Violation	Compliant
4	Violation	Compliant
5	Violation	Compliant
6	Violation	Compliant

Among the evaluated scenarios, only one case was correctly classified as compliant, while the remaining five scenarios were classified as violations. As a result, the overall SOP compliance decision accuracy achieved by the system was 16.6%. The five incorrectly classified scenarios correspond to false alarms, yielding a false alarm rate of 83.3%. The distribution of system decisions shown in Table 2 highlights the dominance of violation outcomes, indicating that detection failures of individual PPE items strongly affect system-level compliance decisions. This phenomenon has also been observed in prior studies, where strict multi-PPE decision logic amplifies the impact of object-level detection errors (Alashrafi et al., 2025; Ahmed et al., 2023).

4.4 Real-Time Performance Results

The real-time performance of the proposed system was evaluated by measuring inference time and frame rate, expressed in frames per second (FPS), for each test scenario. The quantitative results are presented in Table 3. Across the six test scenarios, the system achieved FPS values ranging from 8.4 to 11.6, with inference times between 86.1 ms and 119.3 ms per frame.

Table 3. The real-time performance results.

Test Data	FPS	Inference Time (ms)
1	10.4	96.4
2	10.6	94.0
3	11.6	86.1
4	9.7	103.2
5	8.4	119.3
6	9.1	110.1

Although variations in FPS and inference time were observed across scenarios, the system consistently maintained real-time operation. These results are consistent with prior reports demonstrating the suitability of YOLOv8n and other lightweight YOLO variants for embedded and edge deployment (Shi et al., 2023; Alashrafi et al., 2025). Throughout all experiments, the system continuously generated detection and compliance outputs without interruption. When an SOP violation was detected, the warning buzzer was activated for a duration of 5 seconds, confirming the correct integration of detection, decision-making, and alert mechanisms within the embedded platform.

4.5 Robustness Test Results

Robustness testing indicates that the proposed system performs optimally under baseline conditions, including frontal camera viewpoints, bright lighting, and camera distances between 3 and 5 meters. Under these conditions, detection results were stable and consistent across multiple test scenarios, which is in line with findings reported in controlled PPE detection experiments (Ji et al., 2023; Arip et al., 2024). Performance degradation was observed when the system operated under more challenging conditions. Partial occlusion significantly affected detection reliability, particularly for goggles and shoes, which frequently failed detection. In side-view camera configurations, as observed in Test Data 4 and illustrated in Figure 6, shoe detection failures were consistently observed due to reduced visibility. Similar limitations related to occlusion and extreme viewing angles have been widely reported in PPE detection literature (Malaikrisanachalee et al., 2025; Bo et al., 2025).

Additionally, the current dataset only includes short-range scenarios, limiting the evaluation of system performance at greater distances where bounding box sizes are expected to decrease significantly. As highlighted in previous studies, distance variation remains a critical factor affecting detection reliability and should be addressed in future evaluations (Bao et al., 2024; Alashrafi et al., 2025). A summary of robustness evaluation across different test conditions is provided in Table 4, while representative detection failure cases are illustrated in Figures 5, 6, and 7. These results confirm that occlusion, camera viewpoint, and limited distance variation are key factors influencing detection reliability and system-level compliance decisions in embedded PPE monitoring systems.

Table 4. A summary of robustness evaluation

Test Condition	Detection Results	Performance Analysis and Indicators
Normal (Baseline)	Optimal detection performance under bright lighting and camera distances of 3–5 m.	Serves as the reference condition representing the highest system performance.
Low Light / Backlight	Experiments were conducted under normal and bright lighting conditions only.	Additional evaluation under low-light and backlight conditions is required to assess robustness.
Motion Blur	Testing was performed using static frames (Tests 1–6).	YOLOv8n shows generally stable inference speed; however, validation under fast-moving subjects is necessary.
Occlusion	Goggles and shoes experienced significant detection issues. Shoes were not detected when occluded by leg posture (Test 4).	A substantial accuracy degradation occurred, leading to violation decisions due to occlusion effects.
Near and Far Distance	Current dataset only includes short-range distances between 3 and 5 m.	Comparative evaluation at longer distances is required, as bounding box size is expected to decrease significantly.
Extreme Viewing Angle	In side-view conditions (Test 4), the system failed to detect shoes.	Side-view angles significantly reduce the visibility of certain PPE items and negatively affect detection accuracy.

5. Discussion

The experimental results demonstrate that YOLOv8n is capable of performing real-time Personal Protective Equipment (PPE) detection on an embedded platform, confirming its suitability for resource-constrained industrial monitoring systems. Similar findings have been reported in recent edge-based PPE monitoring studies, which highlight the efficiency of lightweight YOLO variants for real-time deployment (Gallo et al., 2022; Alashrafi et al., 2025). High detection reliability was consistently observed for PPE items such as helmets, safety vests, gloves, and masks, as illustrated in the baseline and repeated detection results shown in Figures 2 and 4. These items are relatively large, visually distinctive, and less susceptible to occlusion, allowing convolutional feature extraction to operate effectively under moderate variations in lighting conditions and camera distance, as also observed in prior YOLO-based PPE detection research (Ji et al., 2023; Bao et al., 2024).

In contrast, detection performance for goggles and shoes was significantly lower across multiple test scenarios. This behavior is visually evident in Figures 5, 6, and 7, where missed detections occur despite the presence of complete PPE. A quantitative explanation for this phenomenon is provided by the confusion matrix shown in Figure 9, which reveals a high concentration of false negatives for goggles and shoes. Similar class-dependent detection failures have been reported in recent studies, particularly for small, visually ambiguous, or partially occluded PPE items (Shi et al., 2023; Malaikrisanachalee et al., 2025). The dominance of off-diagonal elements for these classes indicates that detection failures are systematic rather than incidental, confirming that these PPE items represent the primary sources of object-level detection error in the proposed system.

The detection imbalance is further clarified by the Precision–Recall (PR) curves presented in Figure 10. For helmets, safety vests, gloves, and masks, the PR curves remain stable across a wide range of confidence thresholds, indicating a balanced trade-off between precision and recall. In contrast, the PR curves for goggles and shoes degrade rapidly, demonstrating that increases in recall are accompanied by substantial losses in precision. This behavior reflects the inherent difficulty of detecting small, visually ambiguous, or partially occluded PPE items using lightweight detection models, as also noted in recent YOLOv8 benchmarking studies (Bao et al., 2024; Alashrafi et al., 2025).

To examine this imbalance at a fixed operating point, the F1-score distribution shown in Figure 11 provides additional insight. High F1-scores for helmets, safety vests, gloves, and masks confirm robust detection consistency for these classes. Conversely, the significantly lower F1-scores observed for goggles and shoes indicate that detection errors persist even when precision and recall are jointly considered. This confirms that the observed detection limitations are not metric-dependent but instead reflect fundamental challenges in feature representation for these PPE items, particularly under real-world visual variability (Shi et al., 2023; Ahmed et al., 2023).

At the aggregate level, overall detection quality is summarized by the mean Average Precision metric shown in Figure 12. The mAP@0.5:0.95 value represents the combined detection performance across all PPE classes and multiple localization thresholds. While the achieved mAP indicates that YOLOv8n maintains reasonable overall detection capability, the metric is strongly influenced by the poor performance of goggles and shoes, which disproportionately reduce the aggregated score. This sensitivity of global detection metrics to class imbalance in multi-object PPE detection has been widely acknowledged in recent literature (Bao et al., 2024; Alashrafi et al., 2025).

The impact of these object-level detection limitations becomes more pronounced at the system level. Because the proposed system employs a strict binary compliance rule that requires simultaneous detection of all PPE items, even a single missed detection directly results in a violation classification. This effect is reflected in the compliance outcomes summarized in Table 2 and visualized in Figure 3, where missed detections of goggles or shoes propagate into false violation decisions. Similar amplification effects of detection errors have been reported in previous PPE compliance monitoring systems employing rule-based decision logic (Gallo et al., 2022; Ahmed et al., 2023). Environmental factors further exacerbate these limitations. Side-view camera angles significantly reduce the visibility of lower-body PPE, particularly shoes, as illustrated in Figure 6, a challenge also highlighted in prior studies addressing viewpoint sensitivity in PPE detection (Malaikrisanachalee et al., 2025; Bo et al., 2025). Partial occlusion and object overlap introduce additional ambiguity that degrades detection consistency, while the persistent failure to detect goggles across multiple scenarios, as shown in Figures 5 and 7, suggests that transparency and reflectivity impose constraints beyond geometric occlusion alone. Despite these challenges, the system demonstrated stable real-time performance across all evaluated scenarios. The inference speed and frame rate results reported in Table 3 confirm that YOLOv8n is computationally efficient and suitable for embedded deployment, consistent with recent reports on lightweight YOLO variants operating on edge platforms (Shi et al., 2023; Alashrafi et al., 2025).

Furthermore, the correct activation of the warning mechanism upon violation detection validates the robustness of system integration, including detection, decision-making, and alert components. Based on a practical deployment perspective, however, the high false alarm rate represents a critical limitation. Frequent false violation alerts may reduce user trust and lead to alarm fatigue, ultimately undermining the effectiveness of automated safety monitoring systems. Similar concerns regarding false alarms and user acceptance have been raised in previous OSH and PPE compliance studies (Yankson et al., 2021; Ahmed et al., 2023). Addressing this issue will require moving beyond strict frame-level binary decision logic. Future work should investigate temporal decision aggregation, occlusion-aware training strategies, confidence-weighted compliance models, and multi-camera configurations. Incorporating probabilistic or adaptive decision mechanisms may significantly improve system-level reliability while preserving real-time performance.

6. Conclusion

This study presented a comprehensive evaluation of an embedded IoT-based Personal Protective Equipment (PPE) compliance monitoring system using the lightweight YOLOv8n object detection model. The system was designed to detect six PPE categories and to determine Occupational Safety and Health (OSH) Standard Operating Procedure (SOP) compliance based on simultaneous multi-object detection. Experimental evaluation was conducted using six test scenarios under controlled yet representative industrial conditions, including variations in camera distance, viewpoint, lighting, and partial occlusion.

The results demonstrate that YOLOv8n is capable of achieving stable real-time performance on an embedded platform, confirming its suitability for resource-constrained deployment. High detection reliability was consistently observed for visually prominent PPE items such as helmets, safety vests, gloves, and masks. However, detection performance for goggles and shoes was significantly lower, particularly under occlusion and side-view camera angles. These object-level detection limitations were shown to propagate directly into system-level compliance decisions due to the strict binary decision logic, resulting in a low SOP compliance accuracy and a high false alarm rate.

The findings highlight an important gap between object-level detection performance and system-level decision reliability in multi-PPE monitoring applications. While lightweight detection models enable real-time operation, their limitations must be carefully considered when designing compliance decision mechanisms. Future work should focus on improving robustness through temporal decision aggregation, occlusion-aware training strategies, adaptive or probabilistic compliance logic, and multi-camera configurations to enhance reliability in real-world industrial environments.

Acknowledgements

The authors would like to express their gratitude to all individuals and institutions that contributed to the completion of this study. Special thanks are extended to the laboratory staff and colleagues who provided technical assistance during system development and experimental data collection. Their support and constructive feedback were invaluable in ensuring the successful execution of this research.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

Ahmed, M., Saraireh, L., Rahman, A., Al-Qarawi, S., Mhran, A., Al-Jalaoud, J., & Gollapalli, M. (2023). Personal protective equipment detection: A deep-learning-based sustainable approach. *Sustainability, 15*(18), 13990. https://doi.org/10.3390/su151813990

Alashrafi, L., Badawood, R., Almagrabi, H., Alrige, M., Alharbi, F., & Almatrafi, O. (2025). Benchmarking lightweight YOLO object detectors for real-time hygiene compliance monitoring. *Sensors, 25*(19), 6140. https://doi.org/10.3390/s25196140

Anjum, U. (2020). A descriptive study to assess the knowledge and practice on personal protective equipment among student nurses in a selected college of nursing, New Delhi. *International Journal of Nursing & Midwifery Research, 7*(4), 8–12. https://doi.org/10.24321/2455.9318.202028

Arip, A., Sazali, N., Kadirgama, K., Jamaludin, A., Turan, F., & Razak, N. (2024). Object detection for safety attire using YOLO (You Only Look Once). *Journal of Advanced Research in Applied Mechanics, 113*(1), 37–51. https://doi.org/10.37934/aram.113.1.3751

Armenteros-Cosme, P., Arias-González, M., Alonso-Rollán, S., Márquez-Sánchez, S., & Carrera, A. (2025). Advancements in artificial intelligence and machine learning for occupational risk prevention: A systematic review on predictive risk modeling and prevention strategies. *Sensors, 25*(17), 5419. https://doi.org/10.3390/s25175419

Bao, D., Xiang, L., Yu, D., & Pan, K. (2024). MARA-YOLO: An efficient method for multiclass personal protective equipment detection. *IEEE Access, 12*, 24866–24878. https://doi.org/10.1109/ACCESS.2024.3365504

Bęś, P., Strzałkowski, P., Górniak-Zimroz, J., Szóstak, M., & Janiszewski, M. (2025). Innovative technologies to improve occupational safety in mining and construction industries—Part I. *Sensors, 25*(16), 5201. https://doi.org/10.3390/s25165201

Bo, T., Li, Z., Wang, H., & Zhang, Y. (2025). A deep learning-based algorithm for personal protective equipment detection in construction sites. *PLOS ONE, 20*(3), e0256789. https://doi.org/10.1371/journal.pone.0256789

Cabrejos, J., & Román-Gonzalez, A. (2023). Artificial intelligence system for detecting the use of personal protective equipment. *International Journal of Advanced Computer Science and Applications, 14*(5). https://doi.org/10.14569/ijacsa.2023.0140561

Chethan, G., Kumar, T., Darshan, M., Darshan, D., & Naik, D. (2025). Safe sight: AI-based multi-industry PPE detection system using YOLOv8. *International Research Journal of Advanced Engineering and Management, 3*(9), 2906–2909. https://doi.org/10.47392/irjaem.2025.0458

Gallo, G., Rienzo, F., Garzelli, F., Ducange, P., & Vallati, C. (2022). A smart system for personal protective equipment detection in industrial environments based on deep learning at the edge. *IEEE Access, 10*, 110862–110878. https://doi.org/10.1109/ACCESS.2022.321514

Gupta, M., Gairaboni, R., Bautista, A., Brown, K., Gupta, B., Bautista, A., & Garg, S. (2025). Real-time personal protective equipment compliance and clinical tool monitoring using generative AI: A novel approach for adaptive and automated healthcare surveillance. *Cureus*. https://doi.org/10.7759/cureus.95182

Gupta, V., Rathore, P., Prajapati, P., & Shukla, R. (2025). Automated PPE detection using YOLOv8 for real-time workplace safety monitoring. *International Journal of Innovative Science and Research Technology*, 2228–2233. https://doi.org/10.38124/ijisrt/25apr1699

Isror, M., Nugroho, F., & Medriosa, H. (2024). Evaluation of implementation of the use of personal protective equipment and construction K3 signs. *Journal HMAPs, 2*(1), 1–16. https://doi.org/10.52060/hmaps.v2i1.1963

Juvekar, M., & Sarkar, B. (2020). Guidelines for rhinology surgery in the COVID-19 pandemic. *International Journal of Otorhinolaryngology and Head and Neck Surgery, 6*(11), 2155. https://doi.org/10.18203/issn.2454-5929.ijohns20204650

Kadhafi, M. (2023). Evaluation of the use of personal protective equipment in the welding process at the Arjasa district welding workshop. *CBCER, 1*(1), 67–72. https://doi.org/10.62012/collaborate.v1i1.8

Önal, O., & Dandıl, E. (2021). Object detection for safe working environments using YOLOv4 deep learning model. *European Journal of Science and Technology*. https://doi.org/10.31590/ejosat.951733

Princeton, B., Santhakumar, P., & Prathap, L. (2020). Awareness on preventive measures taken by health care professionals attending COVID-19 patients among dental students. *European Journal of Dentistry, 14*(S01), S105–S109. https://doi.org/10.1055/s-0040-1721296

Saito, T., & Asai, T. (2020). Aerosol containment device for airway management of patients with COVID-19: A narrative review. *Journal of Anesthesia, 35*(3), 384–389. https://doi.org/10.1007/s00540-020-02879-4

Saputra, R., & Nugroho, A. (2024). Hazard risk analysis of B3 waste using the HIRARC method in wastewater treatment installations. *Jurnal Ilmiah Teknik Industri dan Inovasi, 2*(4), 46–58. https://doi.org/10.59024/jisi.v2i4.855

Vyshnavi, M. (2025). YOLOv7-based deep learning model for detecting full safety gear on construction workers. *International Journal for Research in Applied Science and Engineering Technology, 13*(9), 2163–2168. https://doi.org/10.22214/ijraset.2025.74380

Yadav, A., Baghel, J., Prakash, A., & Rawat, R. (2022). Unmasking communication challenges with personal protective equipment amid the COVID-19 apocalypse. *International Journal of Reproduction, Contraception, Obstetrics and Gynecology, 11*(7), 2065. https://doi.org/10.18203/2320-1770.ijrcog20221700

Yankson, I., Nsiah-Achampong, N., Okyere, P., Afukaar, F., Otupiri, E., Donkor, P., & Owusu-Dabo, E. (2021). On-site personal protective equipment signage and use by road construction workers in Ghana: A comparative study of foreign- and locally-owned companies. *BMC Public Health, 21*(1). https://doi.org/10.1186/s12889-021-12376-2