JURNAL ELTEK
Vol. 24, No. 2, October 2026, pp. 70~79
p-ISSN: 1693 – 4024 | e-ISSN: 2355-0740
DOI: 10.33795/eltek.v24i2.10286

Edge-based multi-object PPE compliance monitoring: a critical evaluation of YOLOv8n structural reliability on resource-constrained platforms

Wahyu Nur Hidayat[1], Moechammad Sarosa[1], Sapto Wibowo[1], Erdiyan Ariyanda[1]
e-mail: wahyunurhidayat1211@gmail.com, msarosa@polinema.ac.id, sapto.wibowo@polinema.ac.id, er.erdiyanda@gmail.com
[1]Department of Electrical Engineering, Politeknik Negeri Malang, Malang, Indonesia

Article Information

Article History

Received June 29, 2026
Revised August 9, 2026
Published October 30, 2026

Keywords:

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

ABSTRACT

Occupational Safety and Health (OSH) requires continuous monitoring of Personal Protective Equipment (PPE) compliance, while manual supervision is limited by coverage, continuity, and human observation. This study evaluates an edge-based PPE monitoring prototype using YOLOv8n on a Raspberry Pi 4 for detecting six PPE classes: helmet, safety glasses, mask, safety vest, gloves, and safety shoes. The system uses a 2K QHD autofocus USB camera mounted on a static tripod and integrates object detection with a tiered SOP decision module and physical warning output. Model performance was evaluated using precision, recall, F1-score, Average Precision (AP), mean Average Precision (mAP), FPS, and inference time. Field evaluation included six initial test frames, distance tests at 3-5 m with frontal and side orientations, additional testing at 2 m, a dedicated safety-glasses test at approximately 1 m, and a qualitative indoor test. The model achieved an overall mAP@0.5 of 0.737. Helmet, safety vest, and safety shoes showed the strongest AP@0.5 values, whereas smaller or less visually distinct PPE classes remained more difficult to detect. Edge deployment achieved 8.4-11.6 FPS with inference times of 86.1-119.3 ms. Across the six initial test frames, five SOP decisions were correct, resulting in 83.3% decision accuracy and a 16.7% false alarm rate. The results demonstrate the feasibility of limited real-time PPE monitoring on Raspberry Pi 4 while showing that reliability remains sensitive to object size, distance, viewpoint, and occlusion.

Corresponding Author:
Wahyu Nur Hidayat,
Department of Electrical Engineering,
Politeknik Negeri Malang
Jl. Soekarno Hatta No. 9, Malang, East Java, Indonesia
Email: wahyunurhidayat1211@gmail.com

1. INTRODUCTION

Maintaining Occupational Safety and Health (OSH) standards requires consistent use of Personal Protective Equipment (PPE), including helmets, safety vests, protective footwear, gloves, masks, and eye protection [1]-[4]. Manual compliance monitoring remains labor-intensive, dependent on observer coverage, and difficult to sustain continuously in large or dynamic work areas [3], [4].

Computer vision and deep learning have therefore been adopted for automated PPE monitoring. YOLO-family detectors are widely used because they combine object localization with real-time inference capability [5]-[8], [10], [17], [18]. However, deployment on constrained hardware requires careful balancing of detector complexity, latency, and operational reliability.

Edge computing addresses this constraint by processing video close to the camera and reducing dependence on continuous cloud transmission. Previous PPE and lightweight-object-detection studies have demonstrated embedded or edge deployment on low-cost platforms [11], [12], [19], [20], [22], [24].

Sarosa and collaborators have also demonstrated YOLO and deep-learning vision pipelines in disaster detection, automatic target scoring, visual quality inspection, mobile agricultural object counting, and image recognition [23], [25]-[29].

A remaining gap is that many PPE studies emphasize object-level precision, recall, or mAP [9], [13]-[16], [21], whereas an operational safety system must translate individual detections into a final SOP decision. A missed required PPE item can therefore produce a system-level error even when aggregate detector metrics appear acceptable.

The novelty of this study is the evaluation of how object-level detection errors propagate into system-level SOP decisions on a Raspberry Pi 4, rather than a modification of the YOLOv8n architecture. Six PPE classes are evaluated at the detection level, while the decision module distinguishes Safe, Non-Critical Violation, and Critical Violation states. The evaluation links detector performance with decision reliability under distance, viewpoint, illumination, and partial occlusion.

The main contributions are:

1. A lightweight six-class YOLOv8n PPE monitoring implementation on Raspberry Pi 4.
2. A two-level evaluation linking object-level PPE detection with tiered SOP compliance decisions.
3. An analysis of failure propagation under distance, side-view orientation, and partial occlusion.

2. RESEARCH METHOD

2.1 System Overview

The proposed framework is organized into three stages: input, edge-based processing, and output. As shown in Figure 1, the system acquires worker images through a static USB camera, processes PPE detection and SOP decision logic on Raspberry Pi 4, and provides monitoring information together with a physical warning output.

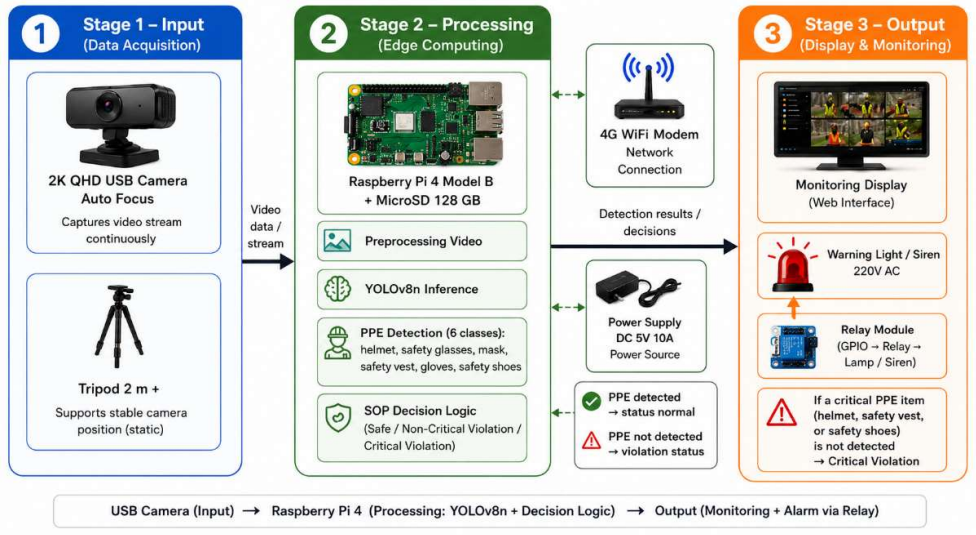


Figure 1. System block diagram

A 2K QHD autofocus USB camera mounted directly on a 2 m tripod captures images or video and remains static during testing. The frames are processed locally by a Raspberry Pi 4 Model B using YOLOv8n to detect six PPE classes. A 128 GB microSD card provides system storage, a 4G WiFi modem supports connectivity, and a 5 V, 10 A, 50 W DC supply powers the prototype. A relay interfaces the Raspberry Pi with the warning light/siren. In accordance with the thesis decision logic, the physical warning is activated for a Critical Violation when the helmet, safety vest, or safety shoes are not detected.

2.2 Data Acquisition and Experimental Conditions

The public-source PPE dataset contains six classes: helmet, goggles, mask, safety vest, gloves, and safety shoes. Each object is represented by a class label and bounding box in the source annotation. The available experiment archive preserves 6,393 annotated instances (997 helmets, 440 goggles, 254 masks, 1,176 safety vests, 1,716 gloves, and 1,810 safety shoes), but not the original image-level train/validation counts or labeling-software metadata; these values are therefore not reconstructed. System-level evaluation used independent field-acquired images and documented 14 operating conditions in addition to six initial diagnostic frames.

.

- Distance: nine main conditions at 3, 4, and 5 m (three orientations per distance), three additional conditions at 2 m, and one dedicated goggles condition at approximately 1 m.
- Viewpoint: frontal, left-facing, and right-facing orientations were evaluated at 2-5 m.
- Illumination: normal and bright ambient lighting were used. Illuminance was not measured with a lux meter and is therefore reported qualitatively.
- Environment: natural partial occlusion and background variation were retained; one additional qualitative indoor condition was evaluated with the camera kept static.

2.3 YOLOv8n-Based PPE Detection

The detector uses Ultralytics YOLOv8 nano (YOLOv8n), selected for its lightweight architecture and suitability for edge inference [7], [17], [18]. Training was performed in Google Colab on an NVIDIA Tesla T4 GPU using pretrained weights. The configuration used up to 150 epochs, batch size 16, 512 x 512 input resolution, early-stopping patience 100, initial learning rate 0.01, final learning-rate factor 0.01, momentum 0.937, and weight decay 0.0005. Augmentation included HSV variation, translation, scaling, horizontal flipping, and mosaic. The best checkpoint (best.pt) was deployed on Raspberry Pi 4.

For every input frame, the embedded YOLOv8n model generates candidate bounding boxes and confidence scores for the six PPE classes. The candidate boxes are then filtered using Non-Maximum Suppression (NMS) to obtain the final object locations.



Figure 2. Hardware configuration

Detection performance was evaluated using precision, recall, F1-score, class-wise AP@0.5, overall mAP@0.5, and the training mAP@0.5:0.95 curve.

2.4 SOP Compliance Decision Logic

After detection, a deterministic rule-based module converts the multi-class outputs into an SOP status. All six PPE classes remain part of object-level evaluation, but goggles are excluded from the final compliance status because their field detection was not sufficiently stable across monitoring distances and viewpoints.

Let D_H, D_V, D_S, D_M, and D_G denote helmet, safety vest, safety shoes, mask, and gloves, respectively (1 = detected; 0 = not detected). Safe is defined as C_safe = D_H AND D_V AND D_S AND D_M AND D_G (1). A Non-Critical Violation is assigned when D_H = D_V = D_S = 1 while the mask and/or gloves are not detected (2).

A Critical Violation occurs when at least one primary safety item is not detected: C_critical = NOT D_H OR NOT D_V OR NOT D_S (3). In this condition, the warning output is activated. Missing mask or gloves changes the displayed status to Non-Critical Violation but does not activate the physical warning.

This tiered decision strategy limits unnecessary warning activation from less stable small-PPE detections while retaining immediate warnings for missing helmet, safety vest, or safety shoes.

2.5 Performance Evaluation Metrics

Evaluation covered three levels: (1) object detection using precision, recall, F1-score, AP, and mAP; (2) SOP decision reliability using SOP Accuracy = (N_correct / N_total) x 100% (4) and FAR = (N_false_alarm / N_total) x 100% (5); and (3) edge performance using inference time and FPS on Raspberry Pi 4.

3. RESULTS AND DISCUSSION

3.1 Object Detection Model Performance and Visual Evaluation

YOLOv8n was evaluated across helmet, safety vest, gloves, safety shoes, goggles, and mask. Figure 3 shows the baseline input, while Figures 4-7 show representative frontal, profile, and partially occluded inference outputs. Visually prominent PPE was more stable in frontal views, whereas smaller PPE and footwear were more vulnerable to missed detection under profile views and occlusion.



Figure 3. Baseline image before YOLOv8n detection processing

Figure 4 shows Test Data 1, where the model localizes the helmet, vest, gloves, mask, and shoes under frontal acquisition.



Figure 4. YOLOv8n detection result for Test Data 1 under frontal acquisition

Figure 5 shows a second frontal acquisition with a similar bounding-box pattern under relatively stable distance, viewpoint, and illumination.

.



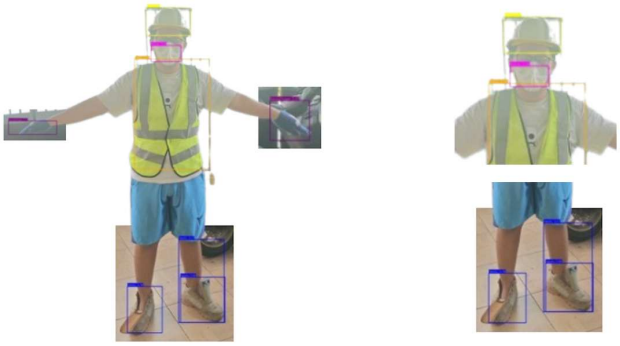
Figure 5. YOLOv8n detection result for Test Data 2 under repeated frontal acquisition

Figure 6 compares frontal and profile views. The profile case reduces lower-body visibility and makes safety-shoe localization less reliable.



Figure 6. Detection results under frontal and profile viewpoints: (a) Test Data 3 and (b) Test Data 4

Figure 7 illustrates partial occlusion and reduced small-PPE visibility using an overall output and magnified detail view.



(a) (b)

Figure 7. Detection under partial occlusion: (a) overall Test Data 5 output and (b) magnified small-PPE regions.

3.2 Quantitative Detection Performance Metrics

Object-level performance was assessed using normalized and absolute confusion matrices, confidence-based precision/recall curves, F1-confidence curves, and training/validation curves. In Figure 8, the horizontal axis is the true class and the vertical axis is the predicted class. Helmet, safety vest, and safety shoes show stronger recognition, while goggles, mask, and gloves exhibit more missed detections or class-dependent errors.

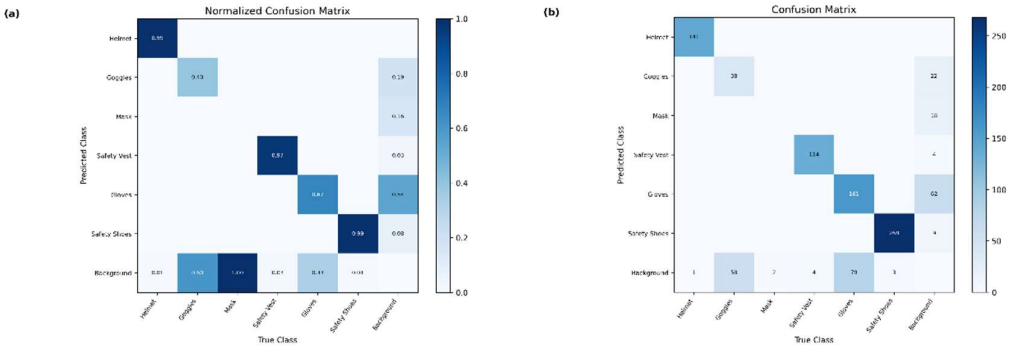


Figure 8. YOLOv8n confusion matrices: (a) normalized confusion matrix and (b) absolute confusion matrix. The horizontal axis represents the true class and the vertical axis represents the predicted class.

Table 1. Class-wise quantitative performance of YOLOv8n

PPE Class	Precision	Recall	F1-score	AP@0.5
Helmet	1.000	0.993	0.996	0.995
Goggles	0.633	0.396	0.487	0.667
Mask	0.000	0.000	0.000	0.017
Safety Vest	0.971	0.971	0.971	0.991
Gloves	0.722	0.671	0.696	0.760
Safety Shoes	0.968	0.989	0.978	0.993

Table 1 shows clear class imbalance in detection reliability. Helmet, safety vest, and safety shoes reached AP@0.5 values of 0.995, 0.991, and 0.993, respectively; gloves and goggles were weaker, and mask was the weakest class. Overall mAP@0.5 was 0.737.

.

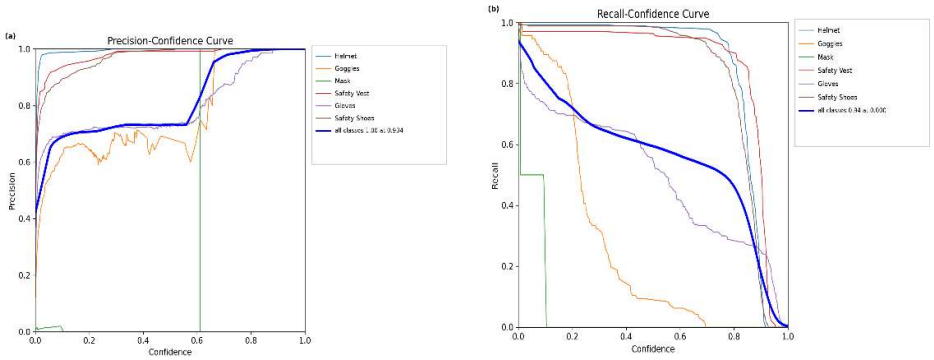


Figure 9. Confidence-based performance curves of YOLOv8n for six PPE classes: (a) precision-confidence curve and (b) recall-confidence curve.

Figure 9 shows precision and recall behavior across confidence thresholds. Helmet, safety vest, and safety shoes maintain stronger profiles, whereas goggles and especially mask are less stable. Figures 10 and 11 summarize the F1-confidence balance and training/validation convergence.

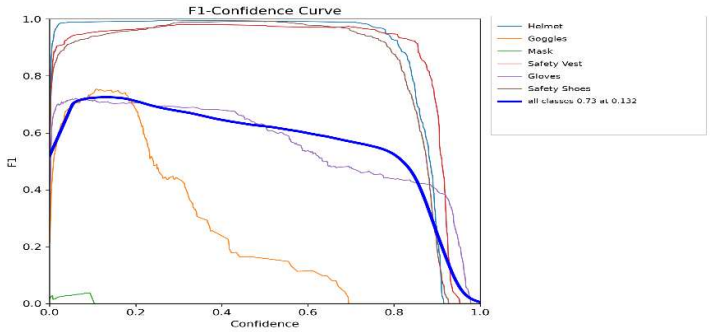


Figure 10. F1-score confidence curves for each PPE class. The horizontal axis represents the confidence threshold, and the vertical axis represents the F1-score.

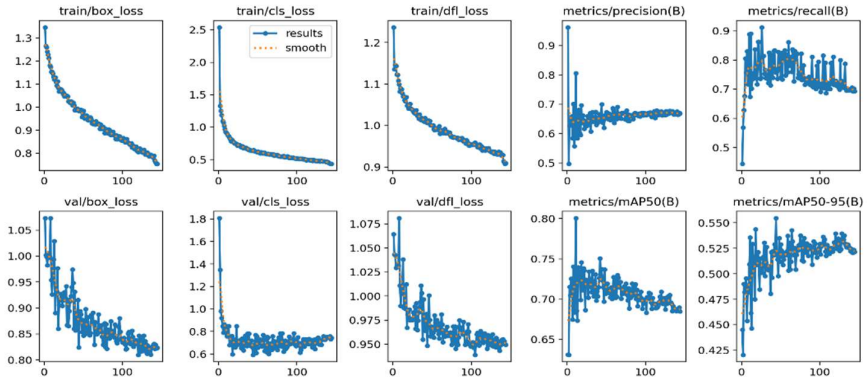


Figure 11. Training and validation performance curves of YOLOv8n. The horizontal axis represents training epochs, while the vertical axis represents loss or evaluation metric values, including precision, recall, mAP@0.5, and mAP@0.5:0.95.

3.3 SOP Compliance Decision Results

The tiered SOP rules in Section 2.4 produced five correct decisions among the six initial diagnostic frames. In Test Data 4, safety shoes were present but not detected, producing an unnecessary Critical Violation. SOP decision accuracy was therefore 83.3% (5/6), while the false alarm rate was 16.7% (1/6).

Table 2. SOP decision evaluation for six initial test frames

Test Data	System Decision	Ground Truth
1	Safe	Safe
2	Safe	Safe
3	Safe	Safe
4	Critical Violation	Safe
5	Safe	Safe
6	Safe	Safe

3.4 Real-Time Performance Results

Across the six initial frames, Raspberry Pi 4 achieved 8.4-11.6 FPS with inference times of 86.1-119.3 ms per frame (Table 3), indicating limited real-time capability.

Table 3. Real-time system performance results

Test Data	FPS	Inference Time (ms)
1	10.4	96.4
2	10.6	94
3	11.6	86.1
4	9.7	103.2
5	8.4	119.3
6	9.1	110.1

The prototype activated the physical warning when a Critical Violation was generated, confirming integration of visual inference, SOP decision logic, relay control, and the warning output. CPU utilization, memory occupancy, processor temperature, and actual power consumption were not continuously logged during the original experiments; these resource indicators are therefore reported as a limitation rather than estimated retrospectively.

3.5 Robustness Test Results

Broader field evaluation covered 14 documented operating conditions across approximately 1-5 m, three body orientations, partial occlusion, and one qualitative indoor condition, in addition to the six initial diagnostic frames (Table 4).

Table 4. Summary of robustness evaluation under different test conditions

Test Condition	Detection Result	Performance Analysis
Lighting	Normal and bright ambient lighting; low light/backlight not tested.	Illuminance not measured in lux; interpretation is qualitative.
Near distance (1-2 m)	Goggles improved at ~1 m; helmet/vest relatively stable at 2 m.	Closer range helps small objects but can reduce full-body field of view.
Main distance (3-5 m)	At 3 m helmet/vest/shoes were detectable; at 4-5 m only vest was consistent.	Smaller PPE loses pixel area as distance increases.
Viewpoint / occlusion	Side views and partial occlusion affected goggles, gloves, and shoes.	Reduced visible area increases false-negative risk and may alter SOP decisions.
Indoor / motion	One qualitative indoor test; main tests used static camera and slow/static subjects.	Broader background and fast-motion testing is still required.
Operational boundary	Stable six-class recognition was not maintained across all conditions.	Deployment should prioritize camera placement, small-object data, and temporal validation.

3.6 Discussion

Class-wise and field results show that reliability is strongly affected by object size, distance, viewpoint, and occlusion. At 3 m, helmet, safety vest, and safety shoes remained recognizable; at 4-5 m, only the vest was consistently detected. Side views reduced visible face, hand, and lower-body regions, increasing false-negative risk for small PPE.

These object-level errors directly affect SOP decisions when they involve a critical class. Test Data 4 demonstrates this propagation: a missed safety shoe produced the only false alarm, yielding 83.3% SOP accuracy and 16.7% FAR. The tiered rule limits this effect by excluding goggles from final status and treating missing mask/gloves as Non-Critical rather than physical-alarm events.

A direct same-dataset comparison with YOLOv5n, YOLOv8s, or other lightweight detectors was not performed. Published PPE studies using YOLOv8 and other lightweight variants report strong detector-level performance [18], [20], [22], but cross-study values are not directly comparable because datasets, class counts, hardware platforms, and test scenarios differ. The contribution of the present study is therefore the integrated six-class edge deployment and system-level SOP reliability evaluation rather than a claim of universally superior detector accuracy.

Accordingly, the main deployment constraint is not only throughput but the robustness of multi-object detection and decision logic under non-ideal field conditions. The absence of logged CPU, memory, temperature, and actual power measurements prevents a complete resource-efficiency comparison and is addressed as a study limitation.

3.7 Limitations and Future Work

Limitations include unstable small-PPE detection, a limited number of subjects and environments, one qualitative indoor condition, qualitative rather than lux-based illumination reporting, unavailable image-level train/validation counts in the archived dataset, no same-dataset detector comparison, and no logged CPU/RAM/temperature/actual-power profile. The results should therefore be interpreted as prototype-level evidence rather than full industrial validation.

.

Future work should use a larger balanced dataset, preserve complete split/annotation metadata, add low-light and motion tests, compare lightweight models on the same dataset and hardware, profile Raspberry Pi resources, and evaluate temporal or multi-camera decision validation to reduce short-term false alarms.

4. CONCLUSION

This study evaluated a six-class YOLOv8n PPE monitoring prototype on Raspberry Pi 4 using three SOP decision levels: Safe, Non-Critical Violation, and Critical Violation. The model achieved an overall mAP@0.5 of 0.737 and 8.4-11.6 FPS with 86.1-119.3 ms inference time. Helmet, safety vest, and safety shoes showed the strongest object-level performance, while smaller PPE remained more sensitive to distance, side-view orientation, and occlusion.

Across six initial diagnostic frames, five SOP decisions were correct (83.3%) and one false alarm occurred (16.7%) because safety shoes were missed. Field testing further showed that at 3 m the helmet, safety vest, and safety shoes remained detectable, whereas at 4-5 m only the safety vest was consistently detected. These findings confirm that edge-based PPE monitoring depends on both detector quality and the rule that translates individual detections into operational safety decisions.

5. ACKNOWLEDGMENTS

The authors thank the Department of Electrical Engineering, Politeknik Negeri Malang, for technical and laboratory support during system development and testing.

6. REFERENCES

[1] I. Yankson, N. Nsiah-Achampong, P. Okyere, F. Afukaar, E. Otupiri, P. Donkor, and E. Owusu-Dabo, "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, vol. 21, no. 1, 2021.

[2] M. Ahmed, L. Saraireh, A. Rahman, S. Al-Qarawi, A. Mhran, J. Al-Jalaoud, and M. Gollapalli, "Personal protective equipment detection: A deep-learning-based sustainable approach," Sustainability, vol. 15, no. 18, p. 13990, 2023.

[3] U. Anjum, "A descriptive study to assess the knowledge and practice on personal protective equipment among student nurses," International Journal of Nursing & Midwifery Research, vol. 7, no. 4, pp. 8-12, 2020.

[4] M. Isror, F. Nugroho, and H. Medriosa, "Evaluation of implementation of the use of personal protective equipment and construction K3 signs," Journal HMAPs, vol. 2, no. 1, pp. 1-16, 2024.

[5] O. Önal and E. Dandıl, "Object detection for safe working environments using YOLOv4 deep learning model," European Journal of Science and Technology, 2021.

[6] A. Arip, N. Sazali, K. Kadirgama, A. Jamaludin, F. Turan, and N. Razak, "Object detection for safety attire using YOLO (You Only Look Once)," Journal of Advanced Research in Applied Mechanics, vol. 113, no. 1, pp. 37-51, 2024.

[7] Ultralytics, "YOLOv8: New state-of-the-art real-time object detection and segmentation," 2023. [Online]. Available: https://github.com/ultralytics/ultralytics.

[8] D. Bao, L. Xiang, D. Yu, and K. Pan, "MARA-YOLO: An efficient method for multiclass personal protective equipment detection," IEEE Access, vol. 12, pp. 24866-24878, 2024.

[9] V. Isailovic, A. Peulic, M. Djapan, M. Savkovic, and A. M. Vukicevic, "The compliance of head-mounted industrial PPE by using deep learning object detectors," Scientific Reports, vol. 12, art. no. 16347, 2022, doi: 10.1038/s41598-022-20282-9.

[10] B. Tong, G. Li, X. Bu, Y. Wang, and X. Yu, "A deep learning-based algorithm for the detection of personal protective equipment," PLOS ONE, vol. 20, no. 5, art. no. e0322115, 2025, doi: 10.1371/journal.pone.0322115.

[11] G. Gallo, F. Rienzo, F. Garzelli, P. Ducange, and C. Vallati, "A smart system for personal protective equipment detection in industrial environments based on deep learning at the edge," IEEE Access, vol. 10, pp. 110862-110878, 2022, doi: 10.1109/ACCESS.2022.3215148.

[12] L. Alashrafi, R. Badawood, H. Almagrabi, M. Alrige, F. Alharbi, and O. Almatrafi, "Benchmarking lightweight YOLO object detectors for real-time hygiene compliance monitoring," Sensors, vol. 25, no. 19, p. 6140, 2025.

[13] Z. Wang, Z. Cai, and Y. Wu, "An improved YOLOX approach for low-light and small object detection: PPE on tunnel construction sites," Journal of Computational Design and Engineering, vol. 10, no. 3, pp. 1158-1175, 2023, doi: 10.1093/jcde/qwad042.

[14] M. Ferdous and S. M. M. Ahsan, "PPE detector: A YOLO-based architecture to detect personal protective equipment (PPE) for construction sites," PeerJ Computer Science, vol. 8, art. no. e999, 2022, doi: 10.7717/peerj-cs.999.

[15] N. Ngoc-Thoan, D.-Q. T. Bui, C. N. N. Tran, and D.-H. Tran, "Improved detection network model based on YOLOv5 for warning safety in construction sites," International Journal of Construction Management, vol. 24, no. 9, pp. 1007-1017, 2024, doi: 10.1080/15623599.2023.2171836.

[16] A. M. Vukicevic, M. Djapan, V. Isailovic, D. Milasinovic, M. Savkovic, and P. D. Milosevic, "Generic compliance of industrial PPE by using deep learning techniques," Safety Science, vol. 148, art. no. 105646, 2022, doi: 10.1016/j.ssci.2021.105646.

[17] C. Shi, D. Zhu, J. Shen, Y. Zheng, and C. Zhou, "GBSG-YOLOv8n: A model for enhanced personal protective equipment detection in industrial environments," Electronics, vol. 12, no. 22, art. no. 4628, 2023, doi: 10.3390/electronics12224628.

[18] A. Barlybayev et al., "Personal protective equipment detection using YOLOv8 architecture on object detection benchmark datasets: a comparative study," Cogent Engineering, vol. 11, 2024, doi: 10.1080/23311916.2024.2333209.

[19] S. Al-Azani, H. Luqman, M. Alfarraj, A. A. Sidig, A. H. Khan, and D. Al-Hamed, "Real-Time Monitoring of Personal Protective Equipment Compliance in Surveillance Cameras," IEEE Access, vol. 12, pp. 121882-121895, 2024, doi: 10.1109/access.2024.3451117.

[20] S. A. Dion and A. Yulianto H., "Performance evaluation of lightweight object detection models for real-time personal protective equipment detection in construction sites," JSiI (Jurnal Sistem Informasi), vol. 12, no. 1, 2025, doi: 10.30656/jsii.v12i1.9896.

[21] J.-H. Lo, L.-K. Lin, and C.-C. Hung, "Real-Time Personal Protective Equipment Compliance Detection Based on Deep Learning Algorithm," Sustainability, vol. 15, no. 1, art. no. 391, 2023, doi: 10.3390/su15010391.

[22] N. Tahilramani, P. Ahir, S. Saxena, V. P. Talreja, and P. Charanarur, "Edge-based AI solution for enhancing urban safety: helmet compliance monitoring with YOLOv9 on Raspberry Pi," Discover Internet of Things, vol. 5, 2025, doi: 10.1007/s43926-025-00113-9.

[23] M. Sarosa and N. Muna, "Implementasi algoritma You Only Look Once (YOLO) untuk deteksi korban bencana alam," Jurnal Teknologi Informasi dan Ilmu Komputer, vol. 8, no. 4, pp. 787-792, 2021, doi: 10.25126/jtiik.2021844407.

[24] D. G. Lema, R. Usamentiaga, and D. F. Garcia, "Low-cost system for real-time verification of personal protective equipment in industrial facilities using edge computing devices," Journal of Real-Time Image Processing, vol. 20, art. no. 111, 2023, doi: 10.1007/s11554-023-01368-7.

[25] E. Ariyanda, M. Sarosa, S. Wibowo, W. N. Hidayat, Y. M. Azis, and I. Mahfudi, "Performance Evaluation of YOLOv8 for Automatic Bullet Impact Detection and Scoring on Shooting Targets," in Proc. International Conference on Cybernetics, Robotics, Automation and Intelligent Digital Systems (ICCRAIDS), 2026, doi: 10.1109/ICCRAIDS67816.2026.11519736.

[26] I. Mahfudi, M. Kusumawardania, M. Sarosa, C. Setiadi, G. P. Riatma, F. A. Soelistianto, N. I. Muslimah, and N. Y. Izati, "Pengenalan kualitas tempe berbasis YOLOv8 untuk deteksi dini kegagalan fermentasi," JASIEK, vol. 7, no. 2, pp. 176-187, 2025, doi: 10.26905/jasiek.v7i2.16091.

[27] A. Evan, M. Sarosa, R. A. Asmara, M. Kusumawardani, D. F. Al Riza, and Y. M. Azis, "Detection and Counting of Grape Leaves Using YOLOv8 via TFLite on Mobile Applications," in Proc. 2024 International Conference on Electrical and Information Technology (IEIT), Malang, Indonesia, 2024, pp. 246-251, doi: 10.1109/IEIT64341.2024.10763328.

[28] M. Sarosa, P. N. Ma'rifah, M. Kusumawardani, and D. F. Al Riza, "Vitis vinera L. leaf detection using Faster R-CNN," BIO Web of Conferences, vol. 117, p. 01021, 2024, doi: 10.1051/bioconf/202411701021.

[29] A. Badriyah, M. Sarosa, R. A. Asmara, M. K. Wardani, and D. F. Al Riza, "Grapevine disease identification using ResNet-50," BIO Web of Conferences, vol. 117, p. 01046, 2024, doi: 10.1051/bioconf/202411701046.