Automated PPE Compliance Monitoring at Construction Sites Using Deep Learning

Ketan Kanjiya[1], Piyush Sonani[2], Upendrasinh Zala[3]

[1]Chief Research Officer, Information Technology Department, Kshatrainfotech Pvt Ltd
[2]Chief Technology Officer, Information Technology Department, Kshatrainfotech Pvt Ltd
[3]Chief Executive Officer, Information Technology Department, Kshatrainfotech Pvt Ltd

Abstract
Construction sites are among the most hazardous working environments, where failure to use Personal Protective Equipment (PPE) can lead to serious injuries and safety violations. Manual monitoring of PPE compliance is often labor intensive, subjective, and difficult to maintain across large and dynamic construction environments. Recent advances in deep learning and computer vision provide effective solutions for automatically detecting safety equipment from visual data. In this study, a deep learning based object detection approach using the YOLO26 architecture is investigated for detecting PPE and related safety violations in construction environments.
The model is trained and evaluated on the Construction Site Safety dataset, which contains annotated images representing ten classes such as Hardhat, Mask, Safety Vest, Person, and non-compliance categories including NO-Hardhat and NO-Mask. Targeted data augmentation and mixed-precision training are employed to improve model robustness and training efficiency. Experimental results demonstrate strong detection performance, achieving a mean Average Precision (mAP@50) of 85.39%, precision of 86.69%, and recall of 78.68%. The results indicate that the proposed approach provides an efficient and scalable solution for automated PPE compliance monitoring in construction environments.

Keywords: PPE Compliance Monitoring, Construction Site Safety, Deep Learning, Object Detection, Computer Vision, Real-Time Safety Compliance

1. Introduction

Construction sites are among the most hazardous working environments, where workers are frequently exposed to risks such as falling objects, heavy machinery operation, and unsafe working conditions. Ensuring the proper use of Personal Protective Equipment (PPE), including hardhats, safety vests, and face masks, is critical for minimizing workplace injuries and maintaining safety compliance. Regulatory bodies such as the Occupational Safety and Health Administration emphasize strict adherence to PPE guidelines to reduce occupational hazards. However, manual monitoring of PPE compliance on construction sites is challenging due to the large scale of operations, dynamic environments, and reliance on human supervision. Consequently, automated and intelligent monitoring systems are increasingly being explored to enhance safety management.

Recent advances in computer vision and deep learning have enabled the development of automated systems capable of detecting and analyzing objects within complex visual scenes. Object detection models can identify workers and determine whether they are wearing the required protective equipment in real time. Traditional monitoring approaches often rely on manual inspection or rule-based image processing techniques, which are limited in scalability and accuracy. Deep learning based object detection methods, particularly those based on convolutional neural networks, have demonstrated superior performance in visual recognition tasks. Among these approaches, the You Only Look Once (YOLO) family of models has gained significant attention for its ability to perform high speed and accurate object detection in real-world applications.

In this study, we propose an automated PPE compliance monitoring framework for construction sites using a modern deep learning based object detection architecture, specifically YOLO26. The proposed approach aims to detect multiple safety related classes, including hardhats, safety vests, face masks, and potential violations such as missing protective equipment. By leveraging a robust training pipeline, targeted data augmentation strategies, and comprehensive dataset analysis, the system is designed to improve detection accuracy while maintaining real-time performance suitable for practical deployment.

The model is trained and evaluated using the publicly available construction safety dataset from Roboflow, which contains annotated images representing various PPE categories commonly observed in construction environments. Extensive experiments are conducted to analyze class distribution, address dataset imbalance, and evaluate model performance using standard object detection metrics such as mean Average Precision (mAP), precision, and recall.

The main objective of this study is to investigate an efficient and reliable vision based approach for automatically monitoring PPE compliance in construction environments. By leveraging advanced deep learning based object detection techniques, the proposed approach seeks to improve the reliability of safety monitoring while supporting more efficient and intelligent construction site management.

2. Literature Review

Construction remains one of the most hazardous industrial sectors globally, consistently reporting high rates of occupational injuries and fatalities [1]. In this environment, Personal Protective Equipment (PPE), such as helmets, safety vests, and gloves, serves as a critical line of defense to mitigate risks like falls, electrocution, and impacts from moving objects [2-4]. Despite stringent safety regulations, non-compliance remains a significant challenge due to worker discomfort, lack of safety awareness, or the perception that PPE hinders productivity [5-6].

Historically, monitoring PPE compliance relied on manual inspections and physical patrols, which are inherently inefficient, subjective, labor intensive, and prone to human error [7]. Early attempts to automate this process introduced sensor based systems using technologies like RFID, GPS, and IoT [8]. Although these systems remain stable under varying weather conditions, they require workers to wear additional electronic devices, which can be cumbersome and expensive to maintain across large workforces [9].

The advent of Computer Vision and Deep Learning has prompted a shift toward fully non-invasive solutions for real-time safety monitoring [10]. Deep learning based object detection algorithms have become the state-of-the-art approach for PPE compliance monitoring due to their ability to automatically extract complex features and generalize across diverse environments [11]. These models are generally categorized into two-stage and one-stage detectors [12-13]. Two-stage detectors, such as Faster R-CNN, are known for high accuracy and effective feature extraction but often lack the inference speed required for real-time site applications [14]. Conversely, one-stage detectors like the YOLO (You Only Look Once) family and SSD (Single Shot Detector) prioritize speed and computational efficiency, making them well

suited for deployment on edge devices and surveillance systems [15-16].
Recent research has focused on enhancing model robustness through attention mechanisms and transformer based architectures [17-18]. Modules like Global Attention Mechanism and BiFPN are integrated into YOLO backbones to improve the detection of small objects, such as safety goggles or gloves, which are frequently missed by traditional models [19]. Furthermore, human pose estimation has emerged as a vital tool to ensure that PPE is not just present in a frame but properly worn on the body [20]. For instance, coupling pose estimators with PPE detectors allows systems to verify if a helmet is actually positioned on a worker's head, thereby reducing false positives caused by workers merely holding their equipment [21].
Despite these advancements, several critical research gaps remain. Most existing studies primarily focus on detecting helmets and safety vests, while other essential PPE items such as masks and safety goggles remain relatively underexplored [22]. Additionally, many models struggle with complex site conditions, including high worker density, significant occlusions, and variable lighting [23]. There is also a scarcity of large scale, multi-class open source datasets that accurately reflect the dynamic and cluttered nature of real construction sites [24-25]. To address these issues, recent research trends emphasize the development of lightweight and adaptive frameworks that utilize domain adaptation and edge computing to provide reliable, real-time safety monitoring while preserving worker privacy and reducing infrastructural costs [26-27]. In this work, we address these limitations by developing a deep learning based PPE compliance monitoring framework using the YOLO26 object detection architecture, designed to detect multiple PPE categories under complex construction site conditions.

3. Dataset

Figure 1. Sample images from the Construction Site Safety dataset

The experiments in this study utilize the Construction Site Safety dataset [28] available on Roboflow Universe, which contains annotated images of construction environments designed for Personal Protective Equipment (PPE) detection and safety monitoring tasks. The dataset includes a diverse set of real-world construction scenes featuring workers, machinery, and safety equipment under varying lighting conditions and site complexities.
The dataset consists of 10 object classes relevant to construction safety monitoring: Hardhat, Mask, NO-Hardhat, NO-Mask, NO-Safety Vest, Person, Safety Cone, Safety Vest, machinery, and vehicle. This annotation scheme is particularly useful because it explicitly distinguishes between compliant and non-compliant safety behavior, such as the presence of Hardhat versus NO-Hardhat. Such labeling enables automated systems to detect safety violations in addition to identifying PPE usage simultaneously.
All annotations are provided in YOLO object detection format, where each label file contains normalized bounding box coordinates and class identifiers. Sample images from the dataset are illustrated in Figure 1, demonstrating typical construction site scenarios and variations in PPE usage used for model

training and evaluation.

4. Methodology

This study proposes a deep learning based framework for automated Personal Protective Equipment (PPE) compliance monitoring at construction sites using the YOLO26 object detection architecture. The overall methodology consists of dataset preparation, exploratory data analysis, data augmentation, model training, and performance evaluation.

4.1 Data Preparation and Augmentation

The Construction Site Safety dataset was organized into training, validation, and test subsets following an 80:15:5 split ratio. Each image is associated with an annotation file in YOLO format containing normalized bounding box coordinates and class identifiers. Prior to training, a dataset integrity check was performed to ensure proper alignment between images and label files, eliminating missing or mismatched annotations. Additionally, exploratory analysis was conducted to examine the distribution of object classes across dataset splits.

To improve model robustness and mitigate class imbalance, targeted data augmentation was applied to selected minority classes in the training set. Augmentation techniques included horizontal flipping, 90-degree rotation, and brightness adjustment, increasing dataset diversity and enhancing model generalization under varying real-world construction site conditions.

4.2 Object Detection Model

The proposed system utilizes the YOLO26 object detection model, a recent generation of the YOLO family designed for efficient and accurate real-time detection tasks. YOLO26 employs a single-stage detection pipeline that predicts bounding boxes and class probabilities directly from input images, eliminating the need for region proposal stages used in two-stage detectors. The architecture integrates optimized convolutional feature extraction with multi-scale feature aggregation, allowing the model to effectively detect objects of varying sizes commonly found in construction site environments. Additionally, YOLO26 employs an NMS-free detection strategy and an efficient optimization scheme that improve inference speed and training stability. The architecture of the YOLO26 model used in this study is illustrated in Figure 2.

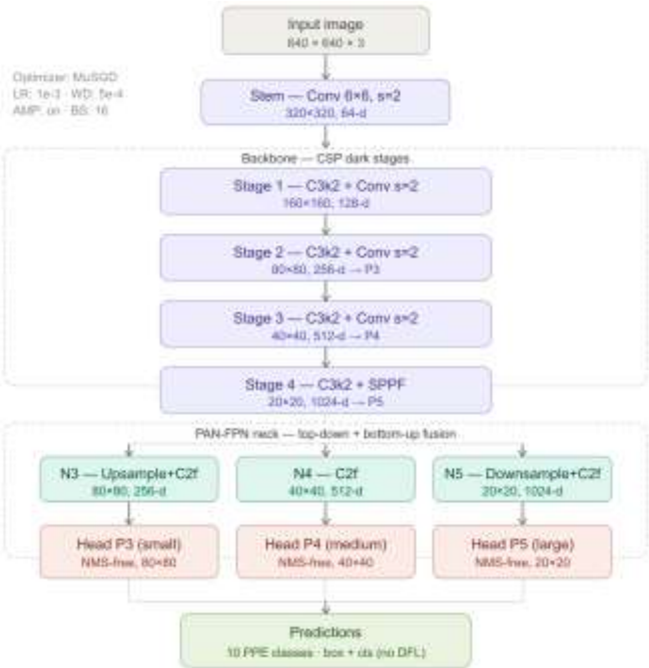

Figure 2. Overview of the YOLO26 object detection architecture used for PPE compliance monitoring.

In this study, the YOLO26-Large variant is used as the base model. The model accepts images resized to 640 × 640 pixels and processes them through a deep convolutional backbone and feature aggregation layers to extract multi-scale representations. The architecture is optimized for real-time inference and enhanced detection of small objects commonly found in PPE monitoring scenarios.

4.3 Model Training

The model was trained for 50 epochs using a batch size of 16 and an input image resolution of 640 × 640 pixels. Training was performed using the MuSGD optimizer, which is specifically designed for YOLO26 to improve convergence and stability during optimization. A learning rate of 0.001 and weight decay of 5×10^{-4} were used to regulate the learning process and prevent overfitting. Mixed precision training was also enabled to accelerate training and reduce GPU memory usage.

Early stopping with a patience of 20 epochs was implemented to prevent unnecessary training once performance stabilized. During training, both classification loss and bounding box regression loss were monitored to track model convergence.

4.4 Evaluation Metrics

The performance of the proposed PPE detection framework was evaluated using standard object detection metrics, including mean Average Precision at IoU threshold 0.5 (mAP@50), mean Average Precision across IoU thresholds from 0.5 to 0.95 (mAP@50-95), precision, and recall. These metrics provide a comprehensive evaluation of detection accuracy and localization quality. In addition to quantitative metrics, qualitative evaluation was performed by visualizing predicted bounding boxes on test images.

5. Results

5.1 Training Convergence Analysis

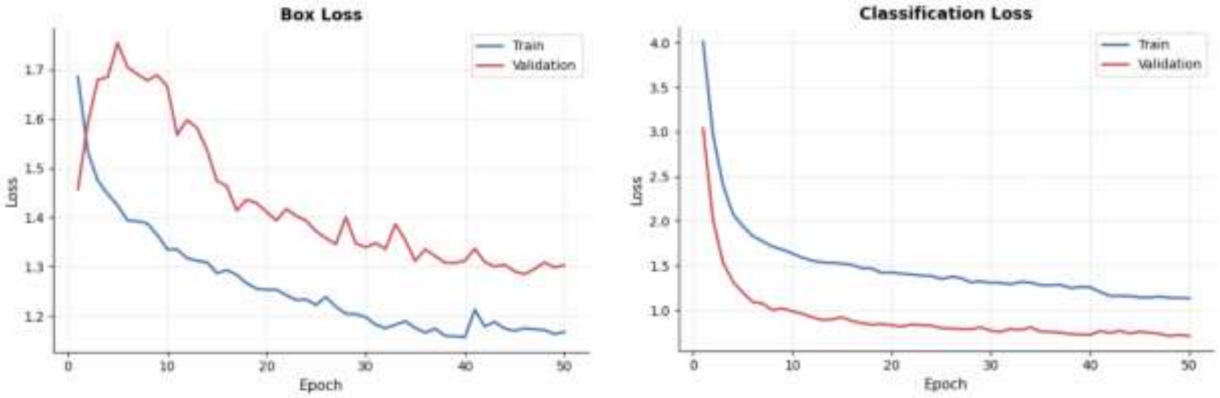

(a) **(b)**

Figure 3. Training and validation loss curves of the YOLO26 model: (a) bounding box loss and (b) classification loss across training epochs.

The training behavior of the proposed model is illustrated using the training and validation loss curves. Figure 3(a) shows the variation of the bounding box regression loss, while Figure 3(b) presents the classification loss across training epochs. Both losses decrease steadily during training, indicating stable learning and effective optimization of the model parameters. The close alignment between training and validation curves suggests good generalization with minimal overfitting.

5.2 Detection Performance Across Epochs

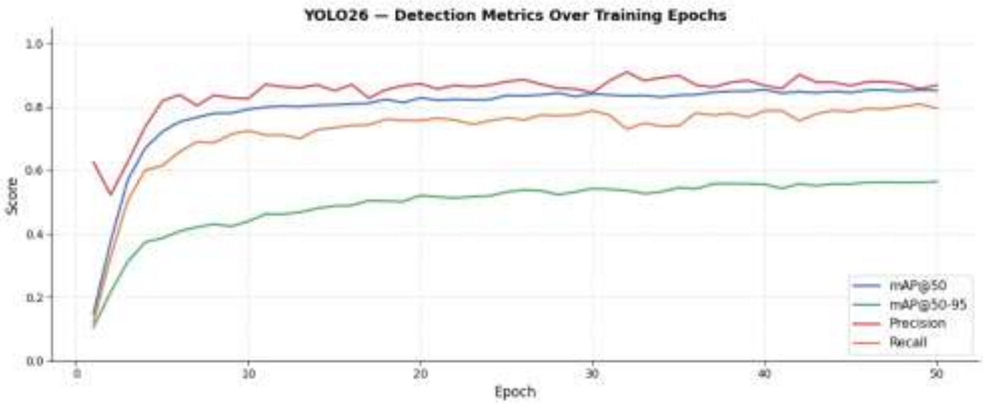

Figure 4. Progression of key detection metrics during training, including mAP@50, mAP@50-95, precision, and recall.

Figure 4 presents the evolution of key detection metrics, including mAP@50, mAP@50-95, precision, and recall, during training. The metrics improve consistently as training progresses and gradually stabilize toward the final epochs, indicating that the model successfully learns discriminative features for PPE detection and reaches stable performance.

5.3 Overall Detection Performance

Table 1. Overall detection performance

Metric	Value (%)
mAP@50	85.39
mAP@50–95	55.57
Precision	86.69
Recall	78.68

The overall detection performance of the proposed system is summarized in Table 1. The model achieves strong detection accuracy, with mAP@50 exceeding 0.85, along with high precision and recall values. These results indicate that the model can reliably detect PPE objects across different construction site scenarios.

5.4 Per-Class Detection Performance

Table 2. Per-class detection performance

Class	AP@50 (%)	AP@50–95 (%)
Hardhat	92.81	63.79
Mask	95.50	67.83
NO-Hardhat	78.57	42.40
NO-Mask	72.36	35.78
NO-Safety Vest	82.15	49.45
Person	87.52	63.89
Safety Cone	88.76	51.51
Safety Vest	88.47	62.28
Machinery	96.37	77.73
Vehicle	68.54	49.47

Table 2 presents the per-class average precision results. High detection performance is observed for PPE classes such as Hardhat, Mask, and Machinery, while relatively lower performance is observed for NO-Mask and Vehicle, likely due to occlusion, smaller object sizes, or class imbalance. Overall, the model demonstrates consistent performance across all ten classes.

5.5 Qualitative Results

Figure 5 shows sample detection results on test images. The model successfully identifies multiple PPE items and safety violations, including missing helmets and safety vests. These results demonstrate the effectiveness of the proposed system for automated PPE compliance monitoring in construction environments.

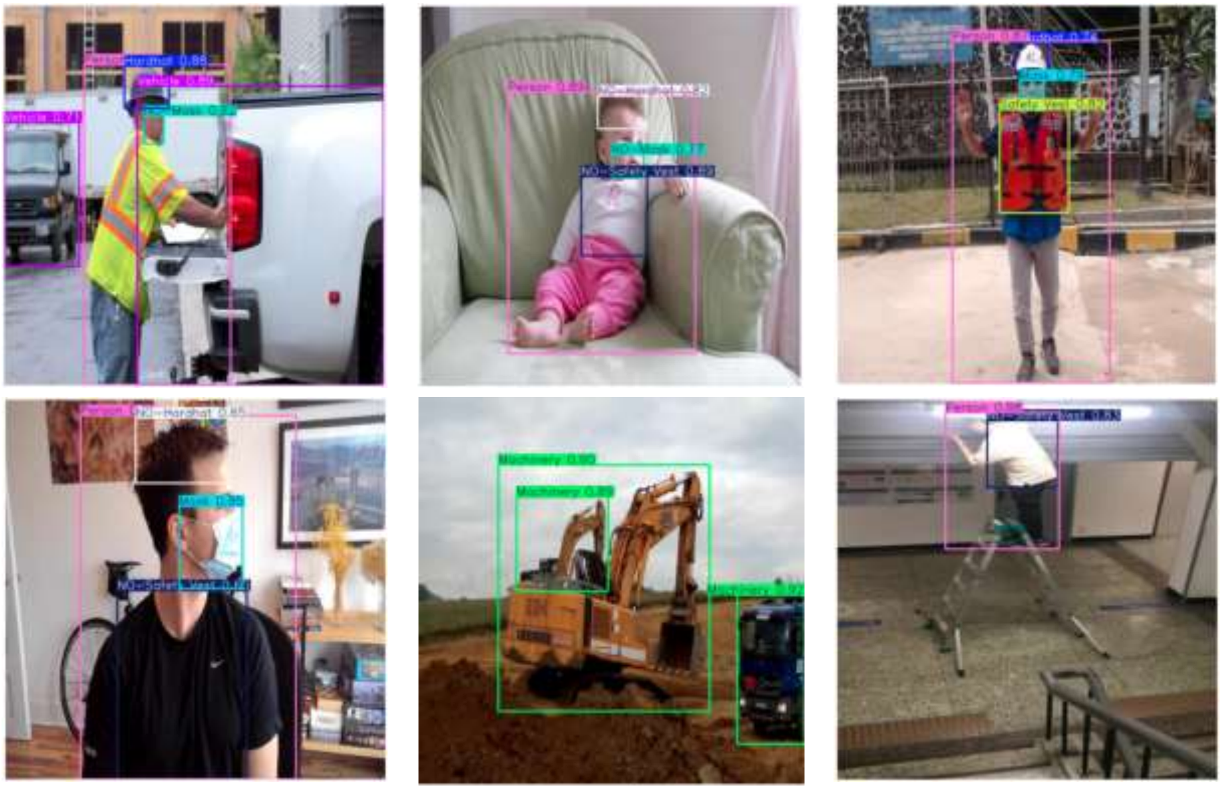

Figure 5. Sample detection results of the trained YOLO26 model on test images demonstrating PPE detection and safety violation identification.

Conclusion

This study presented an automated framework for monitoring Personal Protective Equipment (PPE) compliance at construction sites using a deep learning based object detection approach. The proposed system employs the YOLO26 architecture to detect multiple safety related objects and identify potential safety violations directly from construction site images. The model was trained and evaluated on the Construction Site Safety dataset containing ten classes representing both PPE items and non-compliance scenarios.

Experimental results demonstrate strong detection performance with high precision and recall, indicating the effectiveness of the model in identifying PPE usage across diverse construction environments. Training analysis shows stable convergence and consistent learning behavior, while per-class evaluation confirms reliable detection performance across most PPE categories.

Overall, the proposed approach provides an efficient and scalable solution for automated PPE compliance monitoring in construction environments. Future work may focus on integrating the system with real-time surveillance systems, expanding datasets with more diverse construction scenarios, and incorporating advanced techniques such as pose estimation to further improve PPE compliance verification.

References

1. Ferdous M., Ahsan S.M.M., “PPE Detector: A YOLO-Based Architecture to Detect Personal Protective Equipment (PPE) for Construction Sites”, PeerJ Computer Science, 2022, 8, e999. https://doi.org/10.7717/peerj-cs.999
2. López L., Suárez-Ramírez J., Alemán-Flores M., Monzón N., “Automated PPE Compliance Moni

toring in Industrial Environments Using Deep Learning-Based Detection and Pose Estimation", Automation in Construction, 2025, 176, 106231. https://doi.org/10.1016/j.autcon.2025.106231
3. Mubasher M., Chen W.T., Merrett H.C., Marpaung B., "Bridging the Safety Gap: A Systematic Review of Traditional and Smart Personal Protective Equipment in Construction Safety", KSCE Journal of Civil Engineering, 2026, 100510. https://doi.org/10.1016/j.kscej.2026.100510
4. Balakreshnan B., Richards G., Nanda G., Mao H., Athinarayanan R., Zaccaria J., "PPE Compliance Detection Using Artificial Intelligence in Learning Factories", Procedia Manufacturing, 2020, 45, 277–282. https://doi.org/10.1016/j.promfg.2020.04.017
5. Pisu A., Elia N., Pompianu L., Barchi F., Acquaviva A., Carta S., "Enhancing Workplace Safety: A Flexible Approach for Personal Protective Equipment Monitoring", Expert Systems with Applications, 2024, 238, 122285. https://doi.org/10.1016/j.eswa.2023.122285
6. Nazli N.A.N.M., Sabri N., Aminuddin R., Ibrahim S., Yusof S., Nasir S.D.N.M., "A Real-Time System for Detecting Personal Protective Equipment Compliance Using Deep Learning Model YOLOv5", Procedia Computer Science, 2024, 245, 647–656. https://doi.org/10.1016/j.procs.2024.10.291
7. Al-Azani S., Luqman H., Alfarraj M., Sidig A.A.I., Khan A.H., Al-Hamed D., "Real-Time Monitoring of Personal Protective Equipment Compliance in Surveillance Cameras", IEEE Access, 2024, 12, 121882–121895. https://doi.org/10.1109/ACCESS.2024.3451117
8. Ji X., Gong F., Yuan X., Wang N., "A High-Performance Framework for Personal Protective Equipment Detection on the Offshore Drilling Platform", Complex & Intelligent Systems, October 2023, 9 (5), 5637–5652. https://doi.org/10.1007/s40747-023-01028-0
9. Lo J.H., Lin L.K., Hung C.C., "Real-Time Personal Protective Equipment Compliance Detection Based on Deep Learning Algorithm", Sustainability, 2023, 15 (1), 391. https://doi.org/10.3390/su15010391
10. Ahmad H.M., Rahimi A., "SH17: A Dataset for Human Safety and Personal Protective Equipment Detection in Manufacturing Industry", Journal of Safety Science and Resilience, 2025, 6 (2), 175–185. https://doi.org/10.1016/j.jnlssr.2024.09.002
11. Shi C., Zhu D., Shen J., Zheng Y., Zhou C., "GBSG-YOLOv8n: A Model for Enhanced Personal Protective Equipment Detection in Industrial Environments", Electronics, 2023, 12 (22), 4628. https://doi.org/10.3390/electronics12224628
12. Di B., Xiang L., Yang D., Pang K., "MARA-YOLO: An Efficient Method for Multiclass Personal Protective Equipment Detection", IEEE Access, 2024, 12, 24866–24878. https://doi.org/10.1109/ACCESS.2024.3365504
13. Ahmed M.I.B., Saraireh L., Rahman A., Al-Qarawi S., Mhran A., Al-Jalaoud J., Al-Mudaifer D., Al-Haidar F., AlKhulaifi D., Youldash M., Gollapalli M., "Personal Protective Equipment Detection: A Deep-Learning-Based Sustainable Approach", Sustainability, 2023, 15 (18), 13990. https://doi.org/10.3390/su151813990
14. Wang S., "Automated Non-PPE Detection on Construction Sites Using YOLOv10 and Transformer Architectures for Surveillance and Body Worn Cameras with Benchmark Datasets", Scientific Reports, July 2025, 15 (1), 27043. https://doi.org/10.1038/s41598-025-12468-8
15. Yang W., Kim M.S., Sippel G.J., Mun A.H., McCarthy K.H., Park B., Sarcevic A., Linguraru M.G., Marsic I., Burd R.S., "An Image Dataset for Surveillance of Personal Protective Equipment Adher-

ence in Healthcare", Scientific Data, January 2025, 12 (1), 96. https://doi.org/10.1038/s41597-024-04355-0
16. Rahman A., Alatallah F.A., Almubarak A.J., Alkhazal H.A., Alzayer H.A., Shaaban Y.Z., Min-Allah N., Bakry A., Aloup K., "Deep Learning-Based Automated Inspection of Generic Personal Protective Equipment", Computers, Materials & Continua, 2025, 85 (2), 3507–3525. http://www.techscience.com/cmc/v85n2/63829
17. Sujit K., Indhumathi K., Mohan G., Jose N.N., Ramudu K., Abdennaji T.S., Cheepurupalli N.R., Saidana T., Fissha Y., "Real-Time Construction Safety Monitoring Using a Drone Based Deep Hybrid Attention Model", Scientific Reports, December 2025, 16 (1), 1812. https://doi.org/10.1038/s41598-025-31392-5
18. Huang K., Abisado M.B., "Lightweight Construction Safety Behavior Detection Model Based on Improved YOLOv8", Discover Applied Sciences, April 2025, 7 (4), 326. https://doi.org/10.1007/s42452-025-06766-z
19. Liao H., Hu X., Wang F., Liu Z., Wu M., "Lightweight Intelligent Detection System for Personal Protective Equipment in Laboratories of Higher Education Institutions", IEEE Access, 2025, 13, 202745–202757. https://doi.org/10.1109/ACCESS.2025.3637181
20. Isailovic V., Peulic A., Djapan M., Savkovic M., Vukicevic A.M., "The Compliance of Head-Mounted Industrial PPE by Using Deep Learning Object Detectors", Scientific Reports, September 2022, 12 (1), 16347. https://doi.org/10.1038/s41598-022-20282-9
21. Imam M., Baïna K., Tabii Y., Ressami E.M., Adlaoui Y., Benzakour I., Bourzeix F., Abdelwahed E.H., "Ensuring Miners' Safety in Underground Mines Through Edge Computing: Real-Time PPE Compliance Analysis Based on Pose Estimation", IEEE Access, 2024, 12, 145721–145739. https://doi.org/10.1109/ACCESS.2024.3470558
22. Vukicevic A.M., Petrovic M., Milosevic P., Peulic A., Jovanovic K., Novakovic A., "A Systematic Review of Computer Vision-Based Personal Protective Equipment Compliance in Industry Practice: Advancements, Challenges and Future Directions", Artificial Intelligence Review, October 2024, 57 (12), 319. https://doi.org/10.1007/s10462-024-10978-x
23. Liu Y., Wang J., "Personal Protective Equipment Detection for Construction Workers: A Novel Dataset and Enhanced YOLOv5 Approach", IEEE Access, 2024, 12, 47338–47358. https://doi.org/10.1109/ACCESS.2024.3382817
24. Alnahas J., "Faster-PPENet: Advancing Logistic Intelligence for PPE Recognition at Construction Sites", IEEE Access, 2025, 13, 107936–107949. https://doi.org/10.1109/ACCESS.2025.3578110
25. Wang Z., Wu Y., Yang L., Thirunavukarasu A., Evison C., Zhao Y., "Fast Personal Protective Equipment Detection for Real Construction Sites Using Deep Learning Approaches", Sensors, 2021, 21 (10), 3478. https://doi.org/10.3390/s21103478
26. Gallo G., Di Rienzo F., Garzelli F., Ducange P., Vallati C., "A Smart System for Personal Protective Equipment Detection in Industrial Environments Based on Deep Learning at the Edge", IEEE Access, 2022, 10, 110862–110878. https://doi.org/10.1109/ACCESS.2022.3215148
27. Wang S., "Domain-Adaptive Faster R-CNN for Non-PPE Identification on Construction Sites from Body-Worn and General Images", Scientific Reports, January 2026, 16 (1), 4793. https://doi.org/10.1038/s41598-026-35148-7

28. Roboflow Universe Projects, "Construction Site Safety Dataset", Open Source Dataset, Roboflow Universe, Roboflow, February 2023. https://universe.roboflow.com/roboflow-universe-projects/construction-site-safety