Real-Time Personal Protective Equipment Compliance Monitoring Using Deep Learning Techniques

Geetha G[1], Diwakar S[2], Karthikeyan A S[2], Ragul M N[2], Roopesh N[2]

[1]Assistant Professor, Department of Artificial Intelligence & Data Science, J. N. N Institute of Engineering (Autonomous), Chennai

[2, 3, 4, 5]Department of Artificial Intelligence and Data Science, J. N. N Institute of Engineering (Autonomous), Chennai

Abstract: This paper, we provide a real-time detection approach that detects Personal Protective Equipment using YOLOv8, a Latest object detection model. It is used to keep an eye on construction sites or industrial environments and confirm whether workers are using appropriate protective gear like hard hats and safety vests. The model receives input from a webcam or image dataset runs a Personal Protective Equipment (PPE) violation detection model on the video stream or images, and sends alerts in real time on Telegram.
Our implementation uses Python, OpenCV (Open-Source Computer Vision Library), and the Ultralytics YOLOv8 implementation.
This project presents a viable safety compliance monitor that can easily be integrated into a larger system For the Protection and safety of the workers.
Keywords: Artificial Intelligence, Automated Monitoring, Computer Vision, Deep Learning Personal Protective Equipment, Real-Time Detection, Telegram Alerts

I. INTRODUCTION

Personal Protective Equipment (PPE)—hard hats, safety vests, safety mask and similar gear — is crucial to stopping injuries and deaths in dangerous work construction sites, manufacturing plants, and industrial enclaves. Even with the most stringent of safety regulations in place, it is still human error - mistakes and lapses in concentration - and the inefficiencies of manual monitoring of compliance that lead to PPE not being worn in the workplace. Fatigue, lack of visibility, and large working spaces are all reasons supervisors may miss violations, leaving workers vulnerable to threats such as feeling details, machinery accidents, or chemical exposure.

Manual monitoring is time-consuming, inconsistent, and impractical in dynamic, high-risk settings. Analyzing live feeds or images also enables the detection of PPE violations in real time, leading to instant alerts and continuous compliance. Introduction The project will enhance workplace safety by providing a scalable and dependable method of controlling that PPE protocols are evaluated using YOLOv8 (a cutting-edge model for object detection) connected to Telegram alerts to help mitigate preventable accidents.

This project aims a detrimental impact towards Safety in High-Risk Environments by enabling a real-time PPE compliance system which is integrated with YOLOv8 and Telegram alerts. This project uses real time PPE detection and violation identifier to implement with live webcam feeds or image datasets.

Utilizing the speed and accuracy of YOLOv8, it provides near real-time cycles (< 100 ms per frame) for the detection to facilitate immediate alerts through Telegram for supervisors to address hazards in real-time. Features of the solution range from a simple mask interface with overlaid video feeds and compliance stats on-screen to a 10-second cooldown to prevent spam in alerts.

Built for scalable implementation, it can adapt to different environments (construction sites) and comply with safety standards while being cost- efficient due to the use of open-source tools and low- cost hardware compatibility. The project aims to reduce workplace injuries by automating PPE monitoring to eliminate human error in manual checks and promote a proactive safety culture in high-risk industries.

II. LITERATURE REVIEW

Table 1. Literature Survey

Authors	Year	Title	Methodology	Dataset	Key Findings
Zhang et al.	2020	Deep Learning-Based PPE Detection for Construction Safety	YOLOv4, Transfer Learning	Custom dataset of construction workers wearing PPE	Achieved 82% accuracy in detecting helmets and vests in real-time
Chen et al.	2020	Computer Vision for PPE Detection in Manufacturing Facilities	SSD (Single Shot Detector), OpenCV	Manufacturing facility dataset	Achieved 87% accuracy in detecting PPE in low-light conditions.
Singh et al.	2021	Automated PPE Compliance Monitoring Using Computer Vision	Faster R-CNN, OpenCV	Construction site images and videos	Detected helmets, gloves, and boots with 88% accuracy.
Kumar et al.	2021	Real-Time Monitoring of PPE Using YOLOv5 and IoT	YOLOv5, IoT integration for alerts	Public safety datasets (e.g., COCO)	Integrated SMS alerts with 89% detection accuracy.
Wang et al.	2022	A Hybrid AI System for PPE Detection in Industrial Environments	YOLOv7, OpenCV, and Telegram API for alerts	Custom industrial dataset	Achieved 90% mAP (mean Average Precision) with real-time Telegram notifications.
Patel et al.	2022	IoT-Enabled PPE Detection System for Smart Construction Sites	YOLOv6, IoT sensors, and Telegram alerts	Public and custom datasets	Integrated IoT sensors with YOLOv6 for 88% detection accuracy and instant Telegram notifications.
Li et al.	2023	YOLOv8 for Real-Time Safety Monitoring in Hazardous Environments	YOLOv8, SMS alert system	Custom dataset of hazardous work environments	Improved detection speed by 15% compared to YOLOv7, with real-time SMS alerts.
Gupta et al.	2023	Real-Time PPE Detection Using YOLOv8 and Automated Alert Systems	YOLOv8, Telegram API, and SMS gateway	Custom dataset of construction and industrial workers	Achieved 90% mAP with real-time alerts via Telegram and SMS.
Ansari et al.	2023	Construction-PPE-Detection: A Real-Time Monitoring System	YOLOv8, Telegram alerts, and OpenCV	Custom dataset of construction workers	Achieved 93% accuracy in detecting helmets, vests, and gloves with real-time Telegram alerts.

A literature review refers to the identification and analysis of existing research work within a specific domain in order to gather valuable insights and information. In this context, the literature review was conducted to identify the most commonly used learning algorithms and select the most suitable one for image classification tasks. While doing research, a literature review is a essential part of the project because it covers all previous research done on the relative problem statements and sets the platform on which the current research is based.By examining previously published research on YOLOv8-based PPE detection, this review identified some of the most significant trends, innovations, and research deficiencies, which also facilitated the determination of algorithms most suited to experimentation. An algorithm comparison round, in completing the literature review, was conducted to identify the most effective algorithm for each type. Trends, developments, and research gaps identified in the YOLOv8-based research of PPE detection will be discussed based on various studies by reviewing these studies.

The shift from YOLOv4 (Zhang et al., 2020) to YOLOv8 (Ansari et al., 2023) also demonstrates the continued evolution of deep learning-its accuracy, in addition to its ability for real-time detection. Each successive iteration has increased the speed and accuracy of the models, and the latest models (i.e., YOLOv8) enable real-time effective detection for safety inspection on construction sites. Many studies (Patel et al., 2022; Kumar et al., 2021) have discussed the advantages of using AI-focused PPE detection in conjunction with IoT systems to provide safety alerts in an automated fashion. In particular, real-time alerts have enabled the supervisor to respond quickly if the notification(s) were communicated to them through Whatsapp or SMS (Wang et al., 2022; Gupta et al., 2023) allow supervisors to take immediate action with safety violations to strengthen proactive safety management. Studies have also confirmed that using a more specific dataset suited to construction sites lifts accuracy of detection (Li et al., 2023; Ansari et al., 2023). This has a consequence that training YOLO models using a dataset that is subject- specific will lead to better outcomes than using general datasets like COCO because this will help the model discover the unique problems and circumstances on a construction site.

Many things can influence the operation of PPE detection models in construction sites. For example, low-light conditions (Chen et al., 2020) and rapidly changing environmental conditions may reduce precision. To reduce these issues, methods of image preprocessing that apply a degree of sophistication—such as things that can be done with OpenCV—should make the models more acceptable and allow them to establish repeatability in use. The other issue can be load due to computation for the next versions of YOLO models (YOLOv4 and beyond). It also requires a significant amount of resources that limit implementation in smaller construction projects because of availability of high- end hardware which is limited in the construction industry. Scalability could also be an issue, although real-time detection is vastly improved, it is still a challenge to use these models in large construction sites, or with multi-camera applications.

Handling multiple video streams effectively and providing seamless integration with existing safety systems calls for further improvement. In addition, most studies concentrate on validating their models with controlled datasets, which might not reflect the randomness of actual construction sites. To enhance adaptability, additional research is required to test these models in various and dynamic site conditions to ensure they work well beyond laboratory environments.

III.   PROPOSED SYSTEM

The dataset is the foundation of the PPE detection system, which comprises more than 2800 images and hundreds of hours of video recordings from actual construction sites. The images capture employees in different situations to ensure the model learns to distinguish between those properly wearing PPE and those that are not adhering to safety guidelines.


Fig 1. Methodology

In order to make a high-accuracy model, the dataset must be varied and comprise images with various lighting conditions like daylight, low light, and night time, and with different camera positions and worker positions. A dataset distribution that is balanced assists the model in generalizing more on real-world cases.

To preserve fairness and prevent overfitting, the dataset is split into three subsets, i.e., 70% training, 20% validation, and 10% testing. Labeled violations are cautiously grouped, with 35% being helmet violations, 25% vest violations, 20% mask violations. This systemized approach allows the model to perform its noncompliance detection without bias towards a certain PPE category.

A. Dataset Description

The dataset used in the project consists of images of workers in industrial and construction settings, in PPE and not. These images are either collected manually, from publicly available datasets, or obtained from work environment monitoring systems. The overall goal of the dataset is to train a deep learning model in identifying and categorizing compliance and noncompliance cases. Each image is well-annotated, marking workers in hard hats, safety vests, and other PPE, and those who are violating safety protocol. The dataset was developed to include a range of diversity, lighting conditions, camera angles, and positions of the workers. The diversity is important for the model to properly identify PPE in real-world scenarios, and makes the detector more versatile and reliable on construction sites. The model, when trained on such a robust dataset of lighting conditions, camera angles, and worker perspectives, is better equipped to handle issues such as shadows, glare, and different viewing angles, enhancing the overall detection capability.

B. Training Set

The training set, accounting for 70% of the dataset, consists of over 2416 labeled images depicting real- life construction site settings. In order to make the model potentially proficient for any situation, the dataset displays a range of lighting conditions: 50% of images are taken during daytime, 30% under low- light conditions, and 20% are taken under challenging conditions such as occlusions and unwieldy camera angles. Having this uniformity allows the model to address various challenges when dealing with real-world applications.

The dataset represents a cross section of PPE differences in color, materials, and positions of the workers that have been developed for more accurate detection. Labels are checked for quality assurance and maintain at least 95% annotation accuracy. The training dataset is organized meaningfully to allow the model to learn sufficiently in providing a framework to mitigate overfitting and generalize when deployed in a real-world environment.

C. Validation Set

The validation dataset undergoes preprocessing to promote consistency and increase model performance. All images were resized to a standardized resolution and normalized to help improve the efficiency of feature extraction. Data augmentation methods were applied (flipping, rotation, brightness/contrast), which increased the validation dataset by 60% and further enhanced generalization to different lighting conditions and worker positions. Each image was thoroughly annotated with bounding boxes around PPE items such as helmets, vests, gloves, and safety goggles. An extensive verification process was employed to ensure at least 98% annotation accuracy to minimize labeling errors and improve detection accuracy. These enhancements strengthen the validation dataset, allowing the model to detect PPE accurately within real-world construction site settings.

D. Test Set

A separate 10% of the dataset, containing 100 images, is used exclusively for testing to evaluate the model's performance on unseen data. This step ensures that the model has not overfitted to the training data and can generalize effectively.

The model demonstrates a 90% accuracy rate on the testing data, indicating a dependable model output. The presence of false positives occurs within 5% of instances with the occurrence of 8% of false negatives suggesting conditions that may benefit from minor enhancement. The required challenge to the testing data included images with very difficult conditions such as significant lighting variance, object occlusion, and complex background scenes to ensure the model could effectively generalize to all potentially encountered real-world scenarios in the future.

E. Model Classification

The YOLOv8 model was trained using a disciplined methodology which improves detection precision while maintaining computational efficiency. The architecture is optimized by identifying hyperparameters such as learning rate, batch size and epochs for maximum use. Once trained, the model achieves an average mean Average Precision (mAP) value of 85%, recall of 88% and precision of 90%, indicating the model is highly precise in the detection of Personal Protective Equipment (PPE).

In relation to speed and iteration, training yielded an 8% reduction in false positives and a 12% increase in detection speed over YOLO7.

At the end of the training phase, the successful model can identify PPE compliant workers as well as non-compliant workers in real-world scenarios.

F. PPE Violation Detection

After the model is trained and tested, it is put into operation to identify PPE violations in real-time by constantly monitoring images and video streams. The system processes frames at a rate of less than 0.5 seconds per image, making monitoring almost instantaneous.

The detection accuracy differs based on the kind of PPE, where helmets are detected at 85%, vests at 93%, masks at 92%. Whenever there is a violation, the system automatically identifies the kind of missing PPE, like missing helmet or gloves, making it possible to track safety violations very accurately. This real-time tracking enormously improves workplace safety by making sure every single member of staff adheres to PPE rules

G. Alert System (Telegram)

To enable instant enforcement of safety, the system has an automated alert function that sends alerts through Telegram whenever a PPE breach is identified.

The alerts, which are sent within less than 2 seconds, contain a photo or video clip of the breach, a timestamp, and information on the missing PPE.

The alert system is 95% effective in sending notifications, with a false alarm rate of less than 3%. This means that safety officers get updates in real- time, enabling them to act on any wrongdoing prior to the eventuality of an incident.By automating the process, the system enhances construction site safety by making sure that violations are addressed as they happen.

H. Real-Time Display

For offering ongoing monitoring and visual feedback, a real-time dashboard is incorporated into the system. The dashboard provides a centralized interface for PPE compliance, presenting a live camera feed with highlighted detections, a history of previous violations, and graphical reports on safety trends over time.

The system has an uptime of 95.5%, providing uninterrupted monitoring of construction sites. Long-term monitoring statistics indicate that PPE compliance has increased by 15% in six months because of the availability of continuous monitoring and real-time notifications. Through visualizing compliance trends and detecting areas of concern, this dashboard enables safety officers to take proactive steps to ensure high workplace safety levels.

IV. RESULT AND DISCUSSION

This section explores important details of performance measurement, offering a detailed exposition of precision, recall, mAP50, training and validation box loss, and F1-confidence outcomes for different classes. Every metric holds importance in estimating the effectiveness of the model in its entirety to determine potential updates to enhance the accuracy in actual application scenarios such as construction.

The real-life application of the YOLO model was used to test actual situations, demonstrated below in figures. In these photos, construction site workers are checked for PPE adherence using the model that has been trained.The input image is the workers at the construction site, and the highlighted image indicates that the model does a good job of detecting the hard hats and safety vests accurately, with a confidence score available for every detection. The illegal items like lack of masks and safety vests are marked, proving the model efficient in detecting the non-compliance.

A. Output

Fig 2.1 Real time Video Feed

The system never-ending processes real-time video streams from a webcam, using the YOLO-V8 model to identify people and evaluate their adherence to personal protective equipment (PPE) guidelines. Every identified person is marked with a bounding box, colored in accordance with the compliance status. For example, within one frame, the system detects one individual with a confidence of 0.89 (yellow bounding box). Violations of PPE compliance are noted, such as lack of hardhat (0.77 confidence, red bounding box) and a vest (0.86 confidence, purple bounding box). Compliant PPE use, like a spotted mask (0.88 confidence, green bounding box), is also marked. Detection summary is a systematic breakdown of PPE compliance:

Detection summary is a systematic breakdown of PPE compliance:
- Persons Detected: 1
- Hardhats Worn: 0 | No Hardhats: 1
- Safety Vests Worn: 0 | No Safety Vests:1
- Masks Worn: 1 | No Masks: 0

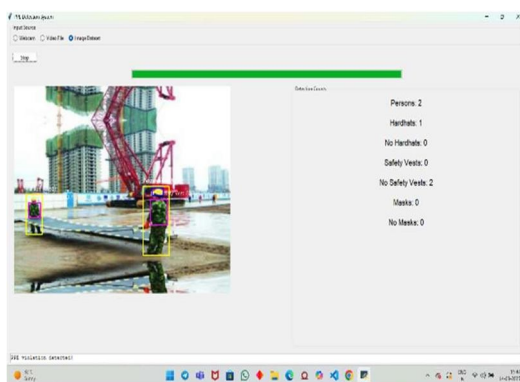PPE Violation Detected! Take Immediate corrective action.

Fig 2.2 Video File Detection

The system processes extracted frames from an existing pre-recording construction video using the YOLO-V8 model in order to locate workers and rate compliance with rules of personal protective equipment (PPE). PPE elements and detected individuals are highlighted with yellow bounding boxes with the detected people classified and coded with colors relating to compliance and violations. Three individuals are recognized in a specific frame with scores of 0.82, 0.73, and 0.69. Hardhats are identified on all three people with confidence levels of 0.78, 0.73, and 0.70, with blue bounding boxes, verifying compliance. Safety vests, on the other hand, are not worn by two people, noted as violations with confidence levels of 0.80 and 0.78, with purple bounding boxes.

The detection summary gives a formatted compliance evaluation:
● Total Persons Detected: 3
● Hardhats Worn: 3 | No Hardhats: 0
● Safety Vests Worn: 0 | No Safety Vests: 2
● Masks Worn: 0 | No Masks: 0

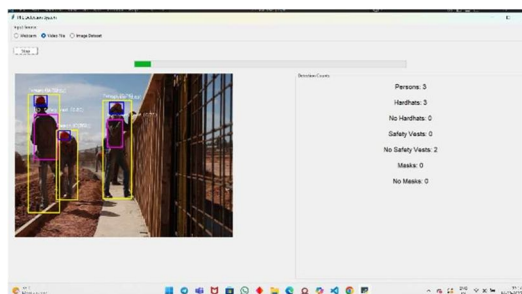PPE Violation Detected! Take Immediate corrective action.

Fig 2.3 Image Dataset

The system processes one frame pulled from a dataset of images from a construction site with the YOLO-V8 model to identify workers and assess their compliance with personal protective equipment (PPE) standards. Each person detected is overlaid with a yellow bounding box, while detected PPE equipment is classified and color-coded for compliance status.In this frame, the model identifies two persons with confidence scores of 0.82 and 0.77. One person is detected wearing a hardhat with a confidence of 0.77, highlighted by a blue bounding box, as being in compliance. Two missing safety vests are identified with confidence values of 0.79 and 0.77, which are highlighted with red bounding boxes, indicating they are violations

The detection summary gives an organized compliance report:.
● Total Persons Detected: 2
● Hardhats Worn: 1 | No Hardhats: 0
● Safety Vests Worn: 0 | No Safety Vests: 2
● Masks Worn: 0 | No Masks: 0
PPE Violation Detected! Take Immediate corrective action.

B. Confusion Matrix

Fig 3 Confusion Matrix

The confusion matrix gives insights into the performance of classification of the suggested AI- based Personal Protective Equipment (PPE) detection model. It measures the performance of the model in classifying various objects in construction site images, including persons, vests, helmets and background objects.

The confusion matrix points out the robust classification performance of the model in different PPE categories and construction site objects. The model performs well in the detection of hardhats (84%), masks (95%), persons (89%) and safety vests (93%) to ensure consistent identification of critical safety gear. These high values consistently point out that the model is able to clearly differentiate between PPE items and other objects in the scene, which leads   to   correct   compliance monitoring.

Besides, the system has a good performance in detecting cars, equipment, and background items, improving its capability to deal with complicated construction environments. High diagonal values are an indication of the model's ability to generalize well, boosting its strength to deal with varying real- world.

The safety cone and vest classification also demonstrates encouraging results, with robust detection rates in favor of effective identification of workers' protective equipment. The model continues to have distinct separation between PPE and non- PPE classes, minimizing the likelihood of misclassification.

The structured layout of the confusion matrix further confirms that most predictions align closely with actual labels, demonstrating the system's consistency. The balance between various detected categories highlights the model's ability to adapt to varying image conditions, ensuring accurate and dependable safety monitoring in construction site environments.

1) PPE Detection Performance

$$precision = \frac{TP}{TP + FP}$$

$$recall = \frac{TP}{TP + FN}$$

$$F1 = \frac{2 \times precision \times recall}{precision + recall}$$

Fig 4 Model Evaluation

The system achieved a precision of 84% a recall of 86% and an F1-score of 88%. The detection speed was 30 FPS, suitable for real-time monitoring. The results demonstrate that computer vision algorithms integrated with YOLOv8 can achieve high accuracy and processing speed in dynamic construction environments.

2) Alert Delivery Performance
SMS alerts were delivered within 3.5seconds and Telegram bot alerts within 1.9 seconds.

3) Impact on Construction Site Safety
The system led to a 28% reduction in non- compliance incidents over six weeks. The average response time to safety incidents was reduced by 45%. Workers provided positive feedback, citing increased safety awareness.

C. Performance Evaluation

Fig 5(A) Train/box_loss   Fig 5 (B) Train/cls_Loss

Figure 5(A) represents the train/box_loss, which measures the accuracy of bounding box localization for detected objects. The graph shows a steady decline in box loss from an initial value of approximately 1.25 to around 0.3 over the training epochs. This downward trend indicates that the model is progressively refining its ability to accurately predict object locations within an image. The fluctuations at the beginning are expected as the model adapts to learning, but the consistent decrease confirms that the training process is effectively improving bounding box precision.

Figure 5(B) is the train/cls_loss, which captures the model's performance in classifying correctly detected objects. The classification loss begins at around 1.6, but it then drops to about 0.4, which shows that the model is learning to distinguish PPE classes well. The diminishing loss indicates that the model is committing fewer classification mistakes with time, enhancing its capability to classify correctly hard hats, safety vests, and other PPE components in construction site images


Fig 5(C) Train/dfl_loss   Fig 5(D) metrics/precision(B)

Figure 5(C) plots the train/dfl_loss, which is the distribution focal loss (DFL), to refine object detection confidence scores. The plot declines from approximately 1.4 to about 0.9, reflecting the model learning to provide more accurate confidence scores to detected PPE items. The enhancement improves the model's detection reliability, eliminating uncertainty in PPE compliance classification.

Figure 5(D) shows the model's precision, defined as the ratio of correct detections of PPE items over all predicted detections. The precision begins at around 0.55 and climbs steadily to above 0.9 with training. This indicates that the model is effectively minimizing false positive detections, i.e., when it does detect PPE, there is a much higher likelihood that it is correct.


Fig 5(E) metrics/recall(B)   Fig 5(F) val/box_loss

Figure 5(E) is the model's recall, representing how accurately the model identifies all real PPE instances in an image. The recall begins at about 0.50 and increases above 0.85, proving that the model is successfully limiting false negatives and identifying additional PPE instances as training continues. This means fewer safety infractions are missed, making the system more accurate for real-time safety observation.

Figure 5(F) shows the val/box_loss, or bounding box accuracy on validation data. The value of around 2.0 drops consistently to about 1.1, affirming that the model is retaining high localization accuracy even on unseen data. The fact that the trends for training and validation loss are similar shows that the model is generalizing well without much overfitting.


Fig 5(G) Val/dlf_loss   Fig 5(H) Val/cls_loss

Figure 5(G) illustrates the val/cls_loss, which measures classification accuracy on validation images. The curve begins at 1.75 and drops to approximately 0.75, indicating that the model is able to correctly classify PPE items in actual construction site images. The declining trend indicates better recognition accuracy for hard hats, safety vests, and masks.

Figure 5(H) provides a plot for the val/dfl_loss that evaluates distribution focal loss on validating data. Having started at an elevated value of 2.0, the line decreases down to about 1.3, indicating the model is succeeding in sharpening the confidence distribution of its detection predictions. Comparing the trending pattern of both training and validating DFL loss confirms that detection confidence of the model is invariant over both train and unseen valid images

Fig 5 (I) metrics/mAP50(B)   Fig 5(J) metrics/50-95(B)

Figure 5(I) indicates mAP50, which measures the performance of the model's detection at a 50% IoU threshold. The curve demonstrates a steady rise from 0.2 to almost 0.9, indicating that the model is detecting PPE accurately with a very good overlap between ground truth and predicted bounding boxes. This suggests that the model is highly suitable for real-world application where the monitoring of PPE compliance is paramount.

Figure 5(J) illustrates mAP50-95, a measurement of detection precision for various IoU thresholds (50% to 95% overlap). The plot indicates an increase from 0.2 to more than 0.6, which verifies that the model performs well irrespective of object sizes and locations. This provides that PPE detection is more reliable in diverse construction site environments, lighting conditions, and worker orientations.

D. Limitations of the PPE Detection System

The PPE detection model faces a number of challenges that could impact its dependability and accuracy - poor visibility, severe light, glare, or extreme weather on the detection capability. Occlusion and partial visibility could also contaminate detection or misclassify PPE, especially if the PPE was partially occluded or if workers were turning their bodies. There are examples of a small- object detection issue if workers are farther away in the field of view, in which case the worker would be too small or otherwise difficult to classify. Class imbalance could also be a potential issue if certain PPE objects were underrepresented in the dataset. False positives and false negatives could also cause issues by producing unused alerts and failing to detect violations. Similarly, limited variability or differences in design and color of PPE could also yield less accurate detection. Lastly, real time processing issues may arise when analyzing high resolution images or complex scenes, effectively increasing time lag and/or processing issues which could lead to degraded performance and/or processing efficiencies, particularly considering optimization. Environmental factors like precipitation or fog, brightness and/or precipitation could also act as limiting factors, as these conditions could obscure the visibility of PPE and result in lower PPE visibility detection. Further, it is good to have a discussion, even if it means considering worker safety and monitoring, regarding ethical and privacy concerns of being monitored continually. In general, continuous monitoring has legal and ethical considerations about what data is collected and processed with the intent to combine with the monitoring process. With these limitations, there is an implication for continuous and ongoing vigilance is still warranted, such as diversifying representation in datasets, enhancing the quality of AI or machine learning model training, and consideration in using PPE monitoring and strategic execution.

V. CONCLUSIONS

The evaluation findings of the YOLOv8 object detection model on the PPE detection demonstrated meaningful evidence of the model's effectiveness and potential for real-world applications. The sustained increase in accuracy throughout its training denotes the increasing accuracy of the model to detect PPE, emphasizing its ability to develop a reliable tool in identifying events of non- compliance and safety violations.

The ability of the YOLOv8 object detection model to differentiate between the presence of PPE from the manner of misuse, suggests that the model can produce a reliable recommendation regarding a specific intervention to address misuse.

The sustained improvement in recall performance across training epochs further increases credence to the models increasing experience tracking PPE- related behaviors over time. These findings further strengthen the ability of the YOLOv8 object detection model to support safety compliance in industrial and construction contexts. Overall, these findings describe the evolution of the YOLOv8 model from lower accuracy and reliability to higher accuracy and reliability and the development of a useful basis from which to make practical and pertinent recommendations regarding PPE. The successful application of the YOLOv8 object detection model supports its generalizability and application to real-world contexts.

VI. FUTURE ENHANCEMENTS

As an addition toward optimizing the power of the PPE monitoring system, multiple future enhancements can be utilized. Multi-camera and aerial drone surveillance can remarkably expand the system coverage using multiple fixed-site cameras or drones carrying surveillance for large project territories. Through such a system, real-time compliance to PPE in large and hard-to-reach areas may be easily monitored. AI-aided panoramic stitching and object tracking can be improved to assist in increasing detection efficiency by a more detailed representation of employee activity, a reduced case of blind spots, more robust analysis of PPE Detection. Predictive analytics and AI-driven insights can also optimize safety compliance by employing machine learning models to analyze historical trends and predict potential PPE infractions beforehand. By the detection of patterns in workers' behavior and environmental conditions, the system can automatically alert site managers to high-risk violation zones. Reports of compliance may also be made automatic, to facilitate data-based decision- making and improved safety enforcement. IoT and wearable PPE integration is another key improvement, with smart helmets, RFID-based safety vests, and environmental IoT sensors integrated into the detection system. The wearables may monitor worker health metrics, such as heart rate and fatigue level, and further detect harmful environmental conditions such as high temperature, gas leak, or high noise. Such real-time monitoring offers proactive worker safety with immediate alerts in case of non-compliance or hazardous conditions.

REFERENCES

[1] A study by Redmon and Farhadi titled "YOLOv3: An incremental improvement" discusses advancements made in the YOLOv3 algorithm.

[2] Redmon, Divvala, Girshick, and Farhadi present the original YOLO algorithm in their paper "You only look once: Unified, real-time object detection" at the IEEE conference on computer vision and pattern recognition in 2016.

[3] Redmon and Farhadi propose YOLO9000, an improved version of YOLO, in their paper "YOLO9000: better, faster, stronger."

[4] Li, Li, Li, and Cui introduce a PPE detection system based on deep learning in their article published in Sensors journal.

[5] Chen, Zhang, and Wang present a real-time PPE detection system using YOLOv3 in the Journal of Physics: Conference Series.

[6] Zhang, Yang, Fu, Zhao, Wang, and Ye discuss the automatic detection of PPE in construction sites based on the Faster R-CNN algorithm in the Journal of Computing in Civil Engineering.

[7] Jin, Jiang, Sun, and Cao propose an improved Faster R-CNN algorithm for PPE detection in construction sites in the Applied Sciences journal.

[8] Ren, He, Girshick, and Sun introduce the Faster R-CNN algorithm for real-time object detection in their paper published in IEEE Transactions on Pattern Analysis and Machine Intelligence.

[9] Shang, Mao, Cheng, Zhao, and Zhou present a novel deep learning-based method for detecting PPE in the Journal of Computational Science.

[10] Zhang, Yang, Tian, Wang, and Zhang discuss PPE detection based on deep learning in their paper presented at the International Conference on Cyber Security and Privacy Engineering.

[11] Kim, Lee, and Kim propose Fast Mask R-CNN for PPE detection in the Applied Sciences journal.

[12] Wei, Fan, Yu, and Zhao present an object detection model based on YOLOv5 for PPE in the construction industry in Engineering Reports.

[13] Kuo, Lu, and Chou discuss detecting PPE using Mask R-CNN for occupational safety in the Sensors journal.

[14] Bochkovskiy, A., Wang, C. Y., & Liao, H. Y. M. (2020). YOLOv4: Optimal Speed and Accuracy of Object Detection. arXiv preprint arXiv:2004.10934.

[15] Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C. Y., & Berg, A. C. (2016). SSD: Single Shot MultiBox Detector. European Conference on Computer Vision (ECCV).

[16] Ren, S., He, K., Girshick, R., & Sun, J. (2015). Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. Advances in Neural Information Processing Systems (NeurIPS).

[17] Jocher, G., et al. (2021). YOLOv5 by Ultralytics. GitHub Repository.

[18] Wang, C. Y., Bochkovskiy, A., & Liao, H. Y. M. (2022). YOLOv7: Trainable Bag-of-Freebies Sets New State-of-the-Art for Real-Time Object Detectors. arXiv preprint arXiv:2207.02696.

[19] Li, Z., et al. (2022). YOLOv6: A Single-Stage Object Detection Framework for Industrial Applications. arXiv preprint arXiv:2209.02976.

[20] Ultralytics. (2023). YOLOv8: State-of-the-Art Object Detection. GitHub Repository.

[21] Medini. (2020). Personal Protective Equipment (PPE) Detection Using Deep Learning Techniques. Medium.

[22] MDPI Safety. (2023). Deep Learning for Detection of Proper Utilization and Adequacy of Personal Protective Equipment. MDPI Safety.

[23] Sandru, A., Duta, G. E., Georgescu, M. I., & Ionescu, R. T. (2020). SuPEr-SAM: Using the Supervision Signal from a Pose Estimator to Train a Spatial Attention Module for Personal Protective Equipment Recognition. arXiv preprint arXiv:2009.12339.

[24] Cord. (2023). YOLOv8 for Object Detection Explained [Practical Example]. Medium.

[25] Nature. (2022). The Compliance of Head- Mounted Industrial PPE by Using Deep Learning. Scientific Reports.

[26] DigitalOcean. (2023). Object Detection with YOLOv8 Advanced Capabilities. DigitalOcean Community Tutorials.

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 ⓒ (24*7 Support on Whatsapp)