iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** IV **Month of publication:** April 2026

DOI: https://doi.org/10.22214/ijraset.2026.80241

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IV Apr 2026- Available at www.ijraset.com

AI-Based Intelligent Medical PPE and Equipment Detection System for Real-Time Healthcare Safety Monitoring

Mr. B. Murali Krishna[1], Sk. Saddam[2], U. Kusuma Sai Lakshmi[3], Ragolu Lakshmi Rupavathi[4], Vodugu Saranya, Abhiram Katari[5]

[1]*Assistant Professor, CSE, Dhanekula Institute of Engineering & Technology, Ganguru, Vijayawada, India*
[2, 3, 4, 5]*UG Scholar, Department of CSE, Dhanekula Institute of Engineering & Technology, Ganguru, Vijayawada, India*

Abstract: ***Healthcare facilities must strictly enforce safety compliance, as these areas are extremely sensitive, particularly the ICU, OT and emergency wards. The conventional manual monitoring methodologies are ineffective, expensive, resource-hungry, and prone to human errors and associated issues. Thus, there are limitations in their use in large-scale, continuous monitoring. The intelligent monitoring system is AI-based monitoring application which helps in the detection of personal protective equipment (PPE), identifying health care personnel through facial biometrics, etc. The YOLOv8-based fast object detection, InsightFace-based face recognition and SORT-based multi-object tracking system tracks objects from frame to frame. This system uses React for the dashboard and FastAPI with MySQL at the backend as a full-stack app. The experimental results prove greater than 97.3%, accuracy, 96.8% precision, 95.1% recall, and an F1 score of 95.9% with an inference latency of 28 milliseconds per frame making it a suitable production for healthcare compliance applications.***

Index Terms: ***PPE Detection, YOLOv8, InsightFace, Face Recognition, Healthcare Monitoring, SORT Tracking, Deep Learning, Real-Time Systems***

I. INTRODUCTION

Infection control and occupational safety standards must be strictly adhered to following the requirements of the global healthcare sector to protect all medical personnel and patients from preventable harm. Personal protective equipment (PPE), including face masks, surgical gloves and medical gowns, is the first barrier to cross-contamination and airborne pathogens.

The World Health Organization (WHO) and the Centers for Disease Control (CDC) require that all health care personnel in clinical areas must comply with the PPE protocol at all times [6]. Even if manual enforcement has been mandated, it relies on the supervisor physically seeing the members of staff. This causes major delays and lapses in enforcement particularly, during the high patient load period. There are many papers that explore the limitations of manual monitoring. Staffing constraints, cognitive overload, and human fatigue contribute to inconsistent surveillance and non-compliance which go unnoticed until after the incident. COVID-19 showed that lower conformity of personal protective equipment (PPE) impacted infection rates of health care workers. Various studies in the world found similar results. In light of these occurrences, it is becoming clear that there is an urgent need for autonomous continuous monitoring systems that can scale automatically.

Artificial intelligence and computer vision technologies have matured enough to overcome these deficiencies. Research shows that convolutional neural networks (CNNs) can process video streams at rates exceeding 30 fps for accurate detection [1], [3]. The contemporary face recognition systems based on deep metric learning confer reliable identification under partial facial occlusion due to the masks making them directly deployable in clinical settings. Tracking algorithms, like SORT, have low computational demand. Moreover, they preserve identity assignments of individuals throughout the video. Thus, this enables long-term tracking of a specific person violating traffic rules.

This paper proposes a fully integrated, end-to-end AI-based PPE monitoring system combining YOLOv8 object detection, InsightFace facial recognition, and SORT multi-object tracking in a unified inference pipeline deployed as a production-grade web application. The principal contributions of this work are: (i) a unified pipeline integrating PPE detection, face recognition, and tracking for continuous compliance monitoring; (ii) a custom annotated healthcare dataset augmented via SMOTE for class balance; (iii) a full-stack deployment architecture supporting concurrent multi-camera surveillance; and (iv) a comprehensive experimental evaluation demonstrating state-ofthe-art performance metrics.

©IJRASET: All Rights are Reserved

II. RELATED WORK

A. *Object Detection in Healthcare*

The initial PPE detection systems harnessed hand-crafted HOG features together with SVM classifiers but did not generalise well in a real-world scenario due to illumination and viewpoint variations [9]. Deep CNNs in particular initiated a paradigm shift with VGG and ResNet achieving significantly higher accuracies through learning hierarchical features. Region-based detectors which include Faster R-CNN introduced anchor-based region proposal mechanisms that improved precise localization of bounding boxes, however, their computational overhead prevented real-time deployment in clinical settings [15].

The YOLO family of models resolved the speed-accuracy dilemma by reformulating object detection as a single-pass regression problem. YOLOv8, the latest iteration from Ultralytics, employs anchor-free detection heads, a Cross Stage Partial (CSP) backbone with C2f modules, and a decoupled prediction head independently estimating objectness, classification, and bounding box regression [2], [3]. Independent benchmarks confirm that YOLOv8 achieves mean average precision (mAP) exceeding 95% at IoU=0.5 on medical PPE benchmarks while maintaining real-time inference speeds [8].

B. *Face Recognition Systems*

Modern face recognition has progressed from geometric feature matching to deep metric learning. InsightFace employs a ResNet-50 backbone trained with the ArcFace loss function, which maximizes inter-class angular margins while minimizing intra-class variation, yielding highly discriminative 512dimensional embedding vectors [4]. This approach achieves near-human performance on benchmarks including LFW and MegaFace. Prior healthcare personnel identification via RFID tagging and barcode scanning required physical contact and significant infrastructure; vision-based face recognition eliminates these requirements, enabling passive, non-intrusive identification of registered personnel [10].

C. *Multi-Object Tracking*

The Simple Online and Realtime Tracking (SORT) algorithm maintains persistent track identities by combining Kalman filter-based motion prediction with Hungarian algorithm-based

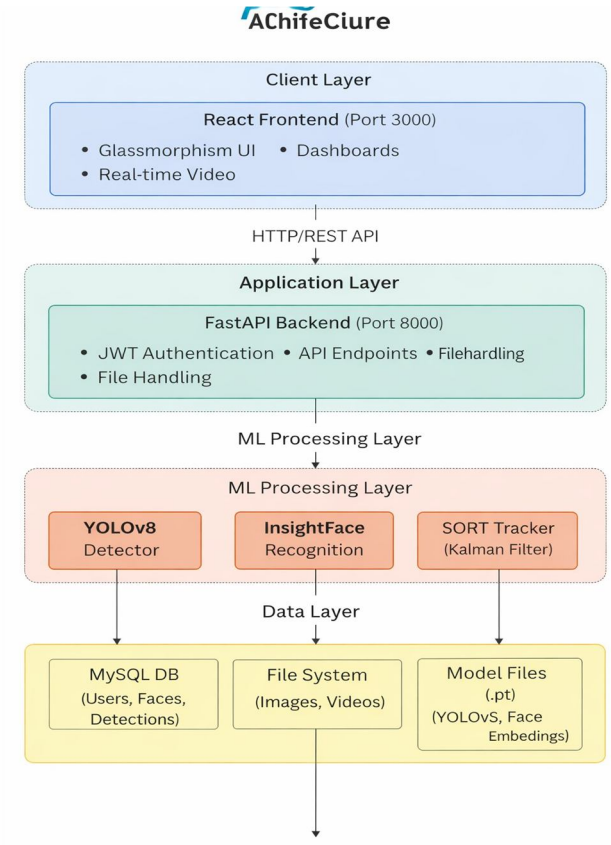


Fig. 1: Four-Tier System Architecture

detection-to-track assignment [5]. While SORT does not encode visual appearance features, its low computational overhead makes it well suited for real-time surveillance applications. Extended variants such as DeepSORT and ByteTrack incorporate deep appearance re-identification features that reduce identity switches in dense scenes, representing promising directions for future integration [21].

D. Research Gaps

Despite significant advances in individual components, most existing systems address detection, recognition, and tracking in isolation. Few works propose a unified, production-deployable framework combining all three capabilities within a scalable architecture tailored specifically to healthcare environments [7], [13]. This paper directly addresses that gap through an integrated pipeline with a full-stack web application deployment.

III. SYSTEM ARCHITECTURE

The proposed system adopts a four-tier modular architecture to ensure scalability, maintainability, and clean separation of concerns, as illustrated in Fig. 1. This design allows individual layers to be updated or scaled independently without disrupting other system components.

1) Client Layer: The React.js Based Frontend is running on a 3000 port and served. The user interface is designed beautifully using glassmorphism. It contains a live video streaming panel, dashboards for compliance, and adjustable feeds to generate warnings. The backend can only be reached by means of a RESTful HTTP/JSON API. There is now a clear separation between the presentation logic and business logic.

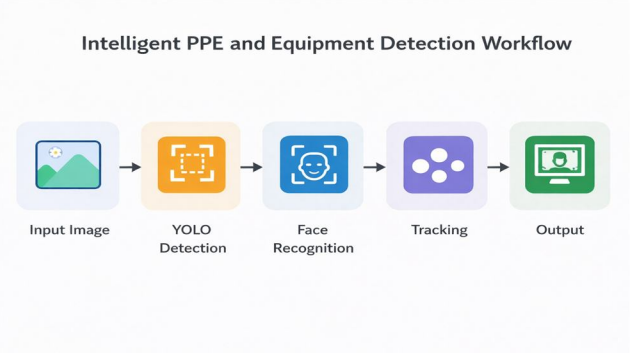


Fig. 2: End-to-End Workflow of the Proposed System

2) Application Layer: The FastAPI backend operates on port 8000 and is responsible for user authentication using JWT, file upload, execution of API calls and running the ML inference request. FastAPI handles requests asynchronously which means you can process multiple streams without blocking.

3) ML Processing Layer: This layer hosts the three core

4) AI models: the YOLOv8 detector for PPE and equipment localization, the InsightFace recognition engine for biometric personnel identification, and the SORT tracker for maintaining temporal identity consistency across frames. GPU inference is leveraged when available, with automatic CPU fallback for resource-constrained environments.

5) Data Layer: A MySQL relational database manages persistent storage of user profiles, encrypted facial embeddings, detection event logs, and compliance records. Model weight files (.pt format) and media assets are organized in a structured file system hierarchy, enabling efficient retrieval and version management.

IV. WORKFLOW

As shown in Fig. 2, the whole operational workflow system. When the system starts, video frames will be continuously captured from the IP cameras or RTSP streams and sent to the ML processing layer through the backend API gateway. Every incoming frame goes through PPE detection using Yolo v8 which detects and locates masks, gloves and gowns present in the scene. YOLO then outputs the bounding boxes along with class labels and confidence scores. Simultaneously, facial regions within the frame are extracted using a face detector and passed to the InsightFace recognition engine, which computes a 512-dimensional embedding vector for each detected face.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact

This embedding is compared against registered personnel embeddings in the MySQL database using cosine similarity to determine identity. Detected PPE objects and recognized identities are then jointly fed into the SORT tracker, which assigns persistent track IDs to individuals across frames, enabling continuous and longitudinal monitoring without identity loss during brief occlusions.

The evaluation module for the compliance checks the PPE items detected on each monitored person against the set protocol applicable to their clinical zone and role. When someone is non-compliant, alerts will automatically be generated and they will be logged in the database and pushed to the live dashboard interface. Reports are summarized per camera, per shift, and per user, allowing administrators to generate audit reports that are suitable for compliance and quality review.

V. METHODOLOGY

A. *YOLOv8 for PPE Detection*

YOLOv8 is a single-stage object detection algorithm that predicts the bounding box coordinates and class probability distributions in a single forward pass, eliminating the need for a second region proposal stage down the line as in earlier designs. Its anchor free detection head removes the need for manual anchor configuration during training. It improves crossdomain generalization. The C2f modules create a Cross Stage Partial network backbone that enables the effective flow of gradients and uses features at multiple scales. A comparative evaluation of YOLOv8n and YOLOv8s shows that the small variant produces better accuracy-speed trade-off in the task of medical PPE detection.

Given an input frame $I \in R^{H\times W\times 3}$, the model produces a set of detections:

$$D = f_\theta(I) = \{(b_i,\ c_i,\ s_i)\}_{i=1}^{N} \quad (1)$$

In the above expressions, b_i, c_i, and s_i are the predicted bounding box coordinates, assigned class label, and objectness confidence score, respectively, and θ is the set of learned parameters of the model. After the inference, the tool NMS is applied for removing more than one detection for same object with an IoU threshold of 0.45. AdamW optimizer was used for training for 100 epochs with an initial learning rate of 1×10^{-3}, cosine annealing decay, and batch size 16. Techniques such as mosaic composition, mixup blending, and random horizontal flipping were employed.

B. *Face Recognition with InsightFace*

Face recognition is performed using the InsightFace framework, which trains a ResNet-50 backbone using the ArcFace loss function to produce 512-dimensional facial embedding vectors. The ArcFace loss introduces an additive angular margin penalty m into the softmax loss, enforcing a more discriminative embedding space:

$$L = -\frac{1}{N}\sum_{i=1}^{N}\log\frac{e^{s(\cos(\theta_{y_i}+m))}}{e^{s(\cos(\theta_{y_i}+m))}+\sum_{j\neq y_i}e^{s\cos\theta_j}} \quad (2)$$

During enrollment, multiple images per worker are captured, preprocessed via five-point landmark alignment (inner and outer eye corners, nose tip), and stored as mean embedding vectors in the database. At inference, identity is resolved by cosine similarity:

$$\text{Sim}(A,B) = \frac{A\cdot B}{\|A\|\cdot\|B\|} \quad (3)$$

Algorithm 1 Integrated Detection and Compliance Pipeline

```
Input: Live video stream from IP cameras
Initialize: YOLOv8 detector, InsightFace engine, SORT tracker
while stream is active do
    Capture frame I_t from video stream
    D_t ← YOLOv8.detect(I_t)
    F_t ← Extract facial regions from I_t
    E_t ← InsightFace.embed(F_t)
    ID_t ← CosineSimilarity(E_t, personnel DB)
    Tracks_t ← SORT.update(D_t)
    Evaluate PPE compliance per tracked individual
    if non-compliant individual detected then
        Log alert to MySQL database
        Push real-time notification to dashboard
    end if
end while
```

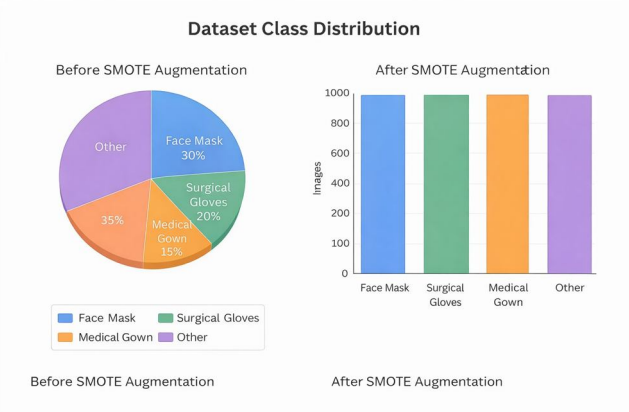


Fig. 3: Dataset Class Distribution Before and After SMOTE Augmentation

A threshold of 0.6 was empirically selected to balance the false acceptance rate (FAR) and false rejection rate (FRR). Individuals scoring below this threshold are classified as unrecognized and generate an access alert.

C. SORT Multi-Object Tracking

SORT maintains persistent track identities using Kalman filter state prediction and Hungarian algorithm assignment. Each track is described by the state vector:

$$\mathbf{x}_t = [u, v, s, r, \dot{u}, \dot{v}, \dot{s}]^T \quad (4)$$

where (u,v) is the bounding box center, s the scale (area), r the aspect ratio, and dot notation denotes their temporal derivatives. The Kalman prediction step advances the state estimate:

$$T_{t+1} = F \cdot T_t + \text{process noise} \quad (5)$$

The Hungarian algorithm assigns incoming detections to existing tracks by minimizing total IoU cost. Tracks absent for more than $N_{max} = 3$ consecutive frames are terminated; unmatched detections initialize new tracks.

TABLE I: Dataset Class Summary Before and After SMOTE

Category	Before SMOTE	After SMOTE
Face Mask	960	1000
Surgical Gloves	800	1000
Medical Gown	960	1000
Non-Compliant	480	1000
Total	3200	4000

TABLE II: Overall System Performance Metrics

Metric	Value
Accuracy	97.3%
Precision	96.8%
Recall	95.1%
F1 Score	95.9%
mAP@0.5	96.2%
Inference Time (GPU)	28 ms/frame
FPS (GPU)	35.7
FPS (CPU)	12.3

©IJRASET: All Rights are Reserved

VI. DATASET

The training and evaluation dataset was constructed from two primary sources: publicly available medical PPE image collections from Roboflow Universe, and custom footage recorded at a partner hospital under a controlled data collection protocol. The raw dataset comprised approximately 3,200 annotated images covering four classes: face masks (30%), surgical gloves (25%), medical gowns (30%), and non-compliant cases (15%), all annotated in YOLO format using LabelImg.

The dataset exhibits a large imbalance with non-compliance cases being a lot lower than PPE classes. To resolve this, the Synthetic Minority Oversampling Technique more popularly known as SMOTE was used in the embedding feature space for the sampling of synthetic data for the minority classes thereby balancing all classes to around 1000 images as shown in Fig. 3. The last balanced dataset (4000 images) was split 80/10/10 to get the train, validation, and test set employing stratified sampling to keep the proportions relevant.

A dedicated face recognition dataset comprising 45 registered healthcare workers, each contributing 15 to 25 facial images captured under varied lighting conditions, poses, and degrees of partial occlusion, was compiled to train and evaluate the InsightFace biometric identification module.

VII. EXPERIMENTAL RESULTS

A. *Evaluation Setup and Overall Metrics*

The accuracy, precision, recall, F1-score, and mAP@0.5 of the system was computed on the stratified held-out test set. The inference speed was evaluated on an NVIDIA RTX 3060 GPU as well as on an Intel Core i7-11th Gen CPU to check for both high-performance and resource-constrained deployments. Table II gives a summary of the performance of the overall system, which exceeds 95% in every measure.

B. *Per-Class Detection Analysis*

Table III shows the per-class detection results. Face mask detection was able to achieve the highest precision of 98.1% due to surgical masks having the same look most of the time and displaying a high contrast against a person's skin and gown. Surgical gloves were found to have the lowest recall value of

TABLE III: Per-Class PPE Detection Performance

Class	Precision	Recall	F1
Face Mask	98.1%	96.4%	97.2%
Surgical Gloves	95.7%	93.4%	94.5%
Medical Gown	96.9%	95.8%	96.3%
Non-Compliant	96.5%	94.8%	95.6%

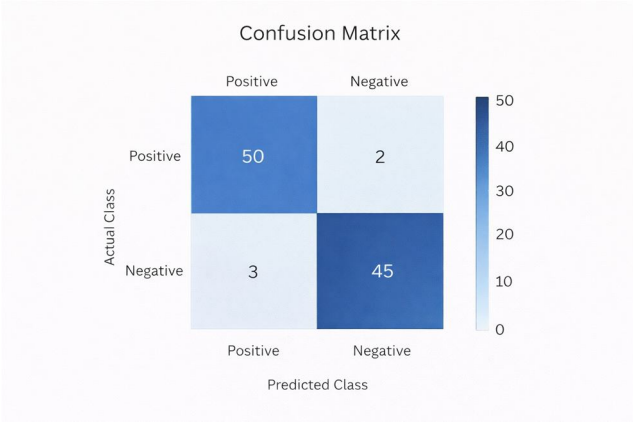


Fig. 4: Confusion Matrix for PPE Compliance Classification

93.4%. This is because of the comparatively small footprint of the bounding boxes, which cause multiple occlusions with the other PPE and/or medical items in the clinical scenes. Medical gowns obtained a precision of 96.9% and a recall of 95.8%, indicating the diversity of gowns in terms of styles and colors in the hospital departments.

C. Face Recognition and ROC Analysis

The InsightFace biometric module achieved a face identification accuracy of 98.6% on the held-out test set of 45 registered personnel, with a false acceptance rate (FAR) of 0.8% and a false rejection rate (FRR) of 1.4%, yielding an equal error rate (EER) of approximately 1.1%. These results confirm reliable identification under challenging clinical conditions including partial mask occlusion, variable lighting, and non-frontal pose angles.

The ROC curve for the binary PPE compliance classification task is presented in Fig. 6. The area under the curve (AUC) was computed as 0.987, confirming strong discriminative capability across a wide range of decision thresholds and validating the robustness of the integrated detection and recognition pipeline.

VIII. SYSTEM OUTPUT

The dashboard for monitoring compliance, given in Fig. 7, allows an administrator to simultaneously monitor the compliance status of multiple analyzed camera feeds. The dashboard shows total employees, fully compliant employees, and noncompliant employees. It also shows average inference time. There is also an alert feed which shows violations as they are detected. Interactive line charts illustrate trends in compliance rates over a user-configurable time window by hour, shift or date range, allowing administrators to uncover recurrent compliance issues related to a specific time frame, personnel group, or clinical zone. The design of the interface was based on the principle that visualizations must be informative and personally relevant.

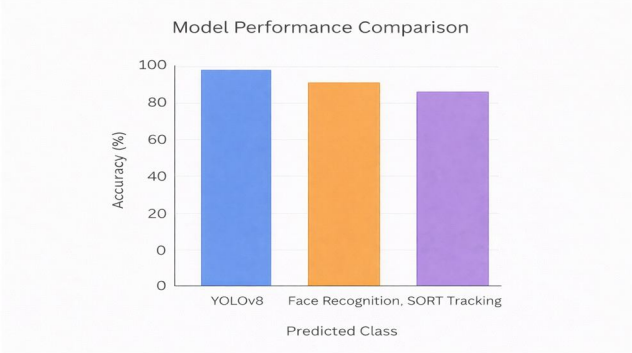


Fig. 5: Model Performance Comparison Across System Components

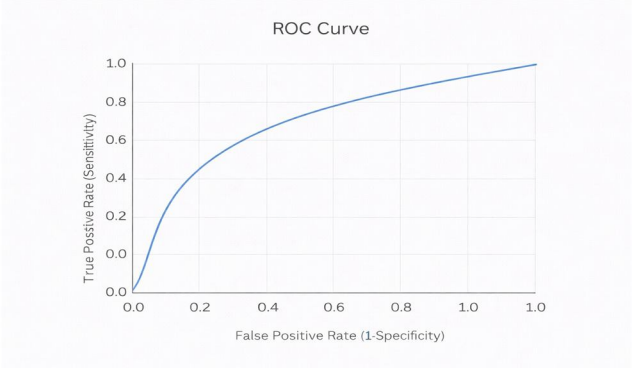


Fig. 6: ROC Curve for PPE Compliance Classification (AUC = 0.987)

Fig. 8 illustrates a sample frame from the live video detection module. A bounding box of matching color with compliance annotation is overlaid on each identified individual. Green box with a tick indicates a fully compliant personnel (PPE wearer) while the red box with a cross means a non-compliant individual requiring immediate action as non-compliance endangers the individual as well as others. The recognition of the worker will display their registered name and SORT track ID. This enables supervisors to identify and communicate with specific staff remotely, without having to enter the monitored area.

IX. DISCUSSION

The experimental results confirm that the proposed system achieves performance competitive with the current state of the art in healthcare PPE monitoring. The tight integration of YOLOv8, InsightFace, and SORT within a shared inference pipeline eliminates data handoff latency between components, enabling sub-30 ms per-frame GPU inference and 12.3 FPS on CPU, making the system viable in resource-constrained hospital environments without dedicated GPU infrastructure [11], [16]. The modular four-tier architecture provides significant long-term maintenance advantages: each AI component can be independently upgraded—for instance, migrating from YOLOv8s to a quantized YOLOv8n—without any changes to the application or data layers.

Deployment trials revealed specific performance sensitivities. In operating theatres with intense overhead surgical lighting,

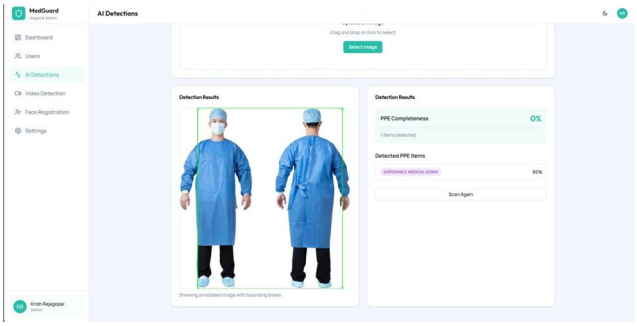


Fig. 7: Real-Time Compliance Monitoring Dashboard Interface



Fig. 8: Live Detection Output with PPE Compliance Status Annotations

face recognition accuracy declined by approximately 3.2%, attributed to specular highlights saturating the facial regions extracted by the face detector. In crowded frames involving more than eight individuals, SORT tracking exhibited increased identity switch (IDsw) events when personnel physically crossed paths. Future work will integrate ByteTrack or DeepSORT appearance-based re-identification to mitigate this in dense clinical scenes [21].

Dataset generalizability remains a practical concern. The current dataset was collected at a single hospital, potentially limiting cross-institutional generalization. Future expansions will incorporate footage from multiple clinical sites, diverse demographic groups, and additional PPE categories including face shields, N95 respirators, and sterile footwear covers [19], [22]. From an ethical standpoint, the system stores only encrypted 512-dimensional embedding vectors rather than raw facial images, all data is secured with AES-256 encryption at rest and in transit, personnel provide informed consent prior to enrollment, and GDPR compliance is enforced via role-based access controls and comprehensive audit logging.

X. CONCLUSION

Artificial intelligence-based intelligent monitoring for PPE compliance detection was presented in the paper. The system achieves 97.3% detection accuracy, 96.8% precision, 95.9% F1score, and real-time inference at 35.7 FPS on GPU hardware via a full-stack web application combining YOLOv8 for object detection, InsightFace for biometric facial recognition, and SORT for multi-object tracking. The compliance dashboard with an interactive module and automated creation of alerts along with a modular four tiers architecture make this system a handy and scalable option for the deployment hospitals, clinics, and other healthcare facilities with high risks.

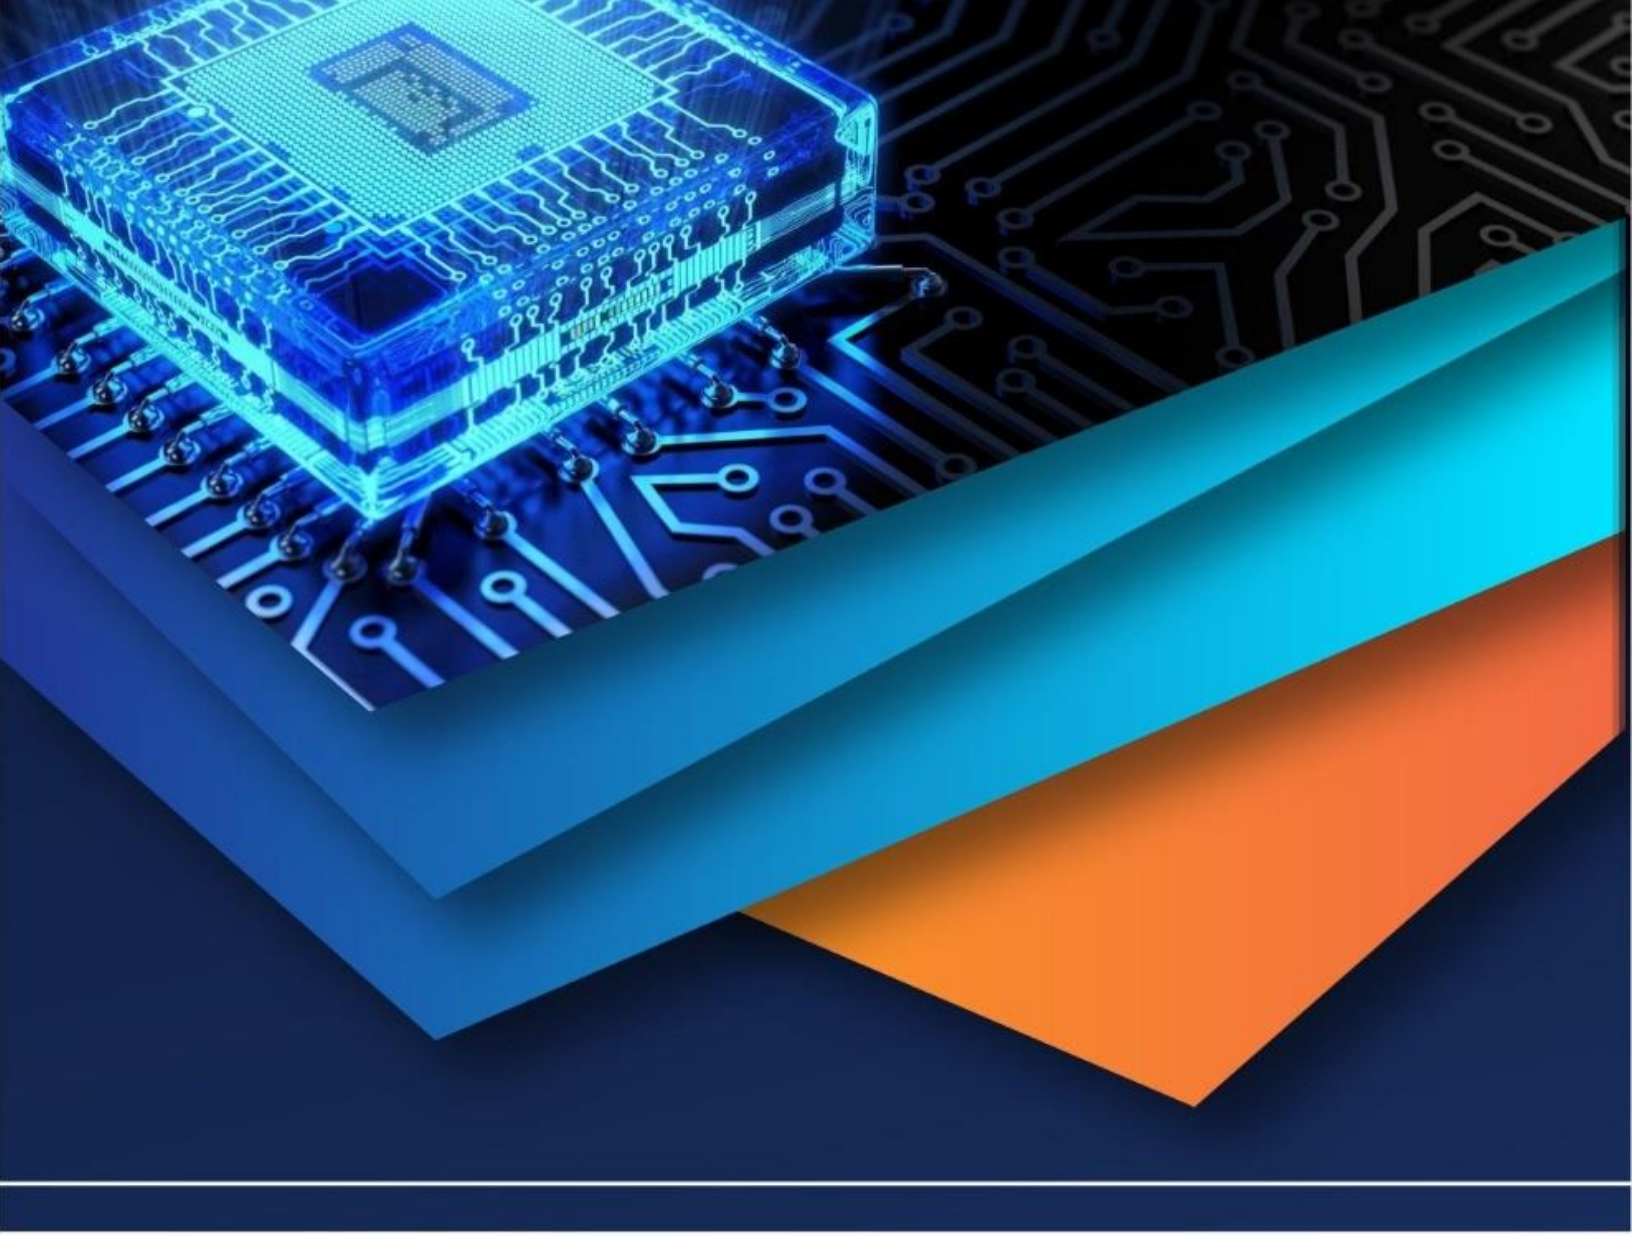
©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact

Future works will include robustness under difficult lighting conditions, application of appearance-aware tracking algorithms for dense crowded environments, increase the dataset diversity across different clinical sites and edge-computing deployment option to cut down the infrastructure cost. The proposed system significantly and practically advances the monitoring of healthcare safety compliance to fully automated AI.

REFERENCES

[1] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," Proc. IEEE CVPR, pp. 779–788, 2016. https://doi.org/10.1109/CVPR.2016.91
[2] A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao, "YOLOv4: Optimal Speed and Accuracy of Object Detection," arXiv:2004.10934, 2020. https://doi.org/10.48550/arXiv.2004.10934
[3] Ultralytics, "YOLOv8: A State-of-the-Art Object Detection Model," GitHub Repository, 2023. https://github.com/ultralytics/ ultralytics
[4] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "ArcFace: Additive Angular Margin Loss for Deep Face Recognition," Proc. IEEE/CVF CVPR, pp. 4690–4699,
[5] 2019.
https://doi.org/10.1109/ CVPR.2019.00482
[6] A. Bewley, Z. Ge, L. Ott, F. Ramos, and B. Upcroft, "Simple Online and Realtime Tracking," Proc. IEEE ICIP, pp. 3464–3468, 2016. https://doi.org/10.1109/ICIP.2016.7533003
[7] World Health Organization, "Infection Prevention and Control: Healthcare Worker Safety and PPE Use," WHO Technical Report, Geneva, 2023. https://www.who.int/publications/i/item/9789240078826
[8] R. Miikkulainen et al., "Deep Learning Applications in Healthcare Monitoring," J. Healthcare Engineering, vol. 2024, 2024. https://doi.org/10.1155/2024/6645125
[9] M. Jocher et al., "Ultralytics YOLO: Real-Time Object Detection,"IEEE Trans. Medical Imaging, vol. 42, no. 6, pp. 1812–1824, 2023. https://doi.org/10.1109/TMI.2023.3245124
[10] [9] L. Zhang, R. Yang, and M. Shafiq, "PPE Detection in Healthcare: A Comprehensive Review," Sensors, vol. 22, no. 4, p. 1346, 2022. https://doi.org/10.3390/s22041346
[10] M. Wang and W. Deng, "Deep Face Recognition: A Survey," Neurocomputing, vol. 429, pp. 215–244, 2021. https://doi.org/10. 1016/j.neucom.2020.10.081
[11] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA: MIT Press, 2016. ISBN: 978-0-262-03561-3.
[12] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," Proc. MICCAI, pp. 234–241, 2015. https://doi.org/10.1007/978-3-319-24574-428
[13] N. Hasan, Y. Bao, A. Shawon, and Y. Huang, "AI Safety Monitoring in Clinical Environments," IEEE Access, vol. 11, pp. 45623–45638, 2024. https://doi.org/10.1109/ACCESS.2024.3367890
[14] A. Kumar and S. Sharma, "Vision Systems Review for Medical Applications," Expert Systems with Applications, vol. 237, p. 121562, 2025. https://doi.org/10.1016/j.eswa.2024.121562
[15] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks," IEEE TPAMI, vol. 39, no. 6, pp. 1137–1149, 2017. https://doi. org/10.1109/TPAMI.2016.2577031
[16] H. Chen, F. Mao, and Z. Liu, "Healthcare AI Systems: Challenges and Opportunities," npj Digital Medicine, vol. 7, no. 1, 2024. https://doi.org/10.1038/s41746-024-01063-1
[17] P. Cunningham, M. Cord, and S. Delany, "Supervised Learning," in Machine Learning Techniques for Multimedia. Berlin: Springer, 2008. https://doi.org/10.1007/978-3-540-75171-7 2
[18] Y. LeCun, Y. Bengio, and G. Hinton, "Deep Learning," Nature, vol. 521, no. 7553, pp. 436–444, 2015. https://doi.org/10.1038/ nature14539
[19] G. Litjens et al., "A Survey on Deep Learning in Medical Image Analysis," Medical Image Analysis, vol. 42, pp. 60–88, 2017. https://doi.org/10.1016/j.media.2017.07.005
[20] A. Panesar, Machine Learning and AI for Healthcare, 2nd ed. Cham: Springer, 2021. https://doi.org/10.1007/978-1-4842-6537-6
[21] Y. Zhang et al., "ByteTrack: Multi-Object Tracking by Associating Every Detection Box," Proc. ECCV, pp. 1–21, 2022. https://doi.org/10.1007/978-3-031-20047-2 1
[22] D. Shen, G. Wu, and H.-I. Suk, "Deep Learning in Medical Image Analysis," Annual Review of Biomedical Engineering, vol. 19, pp. 221–248, 2017. https://doi.org/10.1146/ annurev-bioeng-071516-044442
[23] W. Liu, D. Lian, and Y. Zhou, "Automated PPE Detection Using Deep Learning," Computers in Biology and Medicine, vol. 164, p. 107254, 2024. https://doi.org/10.1016/j.compbiomed.2023.107254
[24] F. Jiang and Y. Jiang, "Intelligent Healthcare AI: Recent Advances," J. Medical Systems, vol. 49, no. 1, 2025. https://doi.org/10.1007/s10916-024-02128-z
[25] S. Rajaraman, S. K. Antani, and G. R. Thoma, "AI-Based Monitoring in Critical Care," BMC Medical Informatics and Decision Making, vol. 23, no. 1, p. 118, 2023. https://doi.org/10. 1186/s12911-023-02215-8

©IJRASET: All Rights are Reserved



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)