IJIREEICE ISSN (O) 2321-2004, ISSN (P) 2321-5526

International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering

Impact Factor 8.414 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 3, March 2026

DOI: 10.17148/IJIREEICE.2026.14389

Computer Vision–Driven PPE Compliance and Safety Violation Detection

Vanitha A[1], Rithish B[2], Safna M[3]

Assistant Professor, Department of Computer Science,

Rathinam College of Arts and Science, Coimbatore, Tamil Nadu, India[1]

Final Year AIML, Department of Computer Science,

Rathinam College of Arts and Science, Coimbatore, Tamil Nadu, India[2,3]

Abstract: If you've ever spent even a few minutes at an active construction site, you will know it's rarely quiet or predictable.There's constant movement materials being shifted, machines running, people coordinating tasks, and often someone working several feet above the ground. In that kind of setting, even a minor lapse in attention can lead to a serious incident. Most accidents don't happen because the work itself is impossible. More often, they occur because basic precautions like wearing Personal Protective Equipment (PPE) — are ignored or treated casually.

On most sites, safety is mainly handled by supervisors who walk around and keep an eye on things, or sit and watch CCTV footage. It definitely helps, but it's not flawless. After all, no person can keep track of several workers and multiple camera screens for hours without feeling tired or losing focus. Over time, small violations can slip through simply because human attention isn't unlimited. As construction projects grow larger and more complex, this gap becomes more noticeable.

This project looks at the problem from a slightly different angle. Instead of entirely relying on manual oversight, it introduces an automated safety monitoring system based on computer vision and deep learning. The system analyses live video streams from site cameras to detect workers and verify whether required PPE is being worn. Once a worker is identified, an object tracking module continues to follow that individual across frames. In other words, the system doesn't just detect once and move on it keeps observing, which helps avoid missing unsafe actions during movement.

Seeing a worker in the frame isn't enough to judge their behaviour. So, pose estimation helps interpret their body position and movement during work.For example,if someone climbs without stable support or works at height without a harness, the system can interpret that posture as potentially unsafe. A rule-based alert mechanism further checks predefined safety conditions, such as entering restricted areas or remaining without protective equipment beyond an acceptable duration. When a violation is identified, the system generates a visible warning and notifies the supervisor. The intention is not to replace human decision- making, but to strengthen it and ensure quicker responses when needed.

By combining PPE detection, worker tracking, behaviour analysis, and real-time alerts within a single framework, the system improves overall site monitoring. It eases the burden on supervisors, supports early risk detection, and encourages more consistent compliance with safety standards. Ultimately, the goal is straightforward: reduce preventable accidents and create a work environment where safety is actively supported rather than assumed.

Keywords: Computer vision, deep learning, personal protective equipment (PPE), safety monitoring, YOLO, CNN, OpenCV, object tracking, pose estimation, real-time alerting.

I. INTRODUCTION

Construction sites are chaotic almost by default. There's constant motion cranes lifting heavy loads overhead, trucks squeezing into narrow spaces, workers balancing on scaffold- ing several floors up. It doesn't really slow down. And the truth is accidents don't usually come from some dramatic, movie-style failure. Most of the time, it's the small decisions. Someone skips their helmet because it's "just for a minute." A harness isn't clipped properly because the task seems quick. Those tiny shortcuts are what end up causing real damage. PPE helmets, vests, gloves, harnesses isn't paperwork or corporate formality. It's basic protection. The challenge isn't writing safety rules; it's making sure people actually follow them when the site is loud, busy, and pushing against deadlines.

Supervisors do what they can. They walk around, check on crews, or monitor CCTV feeds from a control room. And yes, that helps. But let's be honest watching multiple screens for hours is exhausting. Attention drifts. Even experienced supervisors can miss small violations, and sometimes that's all it takes. As projects get bigger and operations become more layered, depending entirely on human observation starts to feel unrealistic.

This is exactly the gap the system is designed to fill. Rather than depending entirely on manual supervision, it relies on computer vision to keep an eye on PPE usage and unsafe actions through live video feeds. When a worker shows up on camera, the system identifies them and continues following their movement frame by frame using multi-object tracking.

IJIREEICE ISSN (O) 2321-2004, ISSN (P) 2321-5526

International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering

Impact Factor 8.414 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 3, March 2026

DOI: 10.17148/IJIREEICE.2026.14389

So even if they walk across the site or move into another camera's view, the system doesn't lose track or start from scratch — it keeps the full picture intact.

But spotting someone in the frame isn't enough. The system also looks at posture and movement through pose estimation. For example, if a worker is unstable at height or climbing without proper support, that body position can be flagged as risky. At the same time, a rule-based engine evaluates specific conditions like entering restricted zones or working without required protective gear for too long. If something doesn't meet the safety criteria, an alert is generated immediately so supervisors can respond before the situation gets worse.

The point isn't to push supervisors out of the picture they're still a key part of site safety. What this system really does is back them up. It keeps watching in the background, consistently and without fatigue, so fewer things slip through unnoticed. By bringing together worker detection, tracking, posture monitoring, and instant alerts, it reduces the pressure of having to watch everything all the time and helps flag risks sooner rather than later. At the end of the day, improving safety isn't about piling on more policies or stricter language. It's about making sure safety practices are actually followed in real time, while the job is still in progress.

II. PROBLEM STATEMENT

A. Need for Automated PPE Compliance and Safety Monitoring

Construction sites are busy, fast-moving places, and with that pace comes risk. People are lifting heavy materials, operating machinery, climbing scaffolding sometimes all within the same few minutes. Safety rules are clearly written, and PPE like helmets, vests, gloves, and harnesses are mandatory for a reason. Still, accidents continue to happen. And more often than not, it's because someone skipped a basic precaution. Maybe they forgot their helmet. Maybe they stepped into a restricted area thinking it would only take a second. Small lapses like that can have serious consequences, both for the worker and for the project as a whole.

Right now, most sites rely heavily on manual supervision. Supervisors walk around checking compliance, or they sit in control rooms watching CCTV footage. Cameras definitely help with coverage, but they don't actually understand what's happening. Someone still has to monitor those screens constantly. And let's be honest watching multiple video feeds for hours is exhausting. Attention fades. People miss things. On large sites, it's simply unrealistic to expect every unsafe action to be caught in real time. By the time a violation is noticed, the damage may already be done.

That's why there's a growing need for something smarter and more consistent. An automated system that can continuously monitor worker activity without getting tired makes a lot of sense. Ideally, it should be able to check whether PPE is being worn properly, recognise unsafe behaviours, and flag potential hazards immediately. And it has to work in real- world conditions changing lighting, crowded scenes, different camera angles not just in controlled environments.

The core issue this project addresses is the absence of a fully integrated, real-time safety monitoring solution on construction sites. What's needed isn't just more cameras or stricter policies, but a system that actively interprets what it sees. By developing a computer vision–based framework that detects PPE non-compliance and risky behaviour as they happen, the goal is to move from reactive safety management to something more proactive. In simple terms, it's about catching problems early and supporting safer decision-making before an accident occurs.

III. LITERATURE REVIEW

A. Traditional Construction Safety Monitoring

For a long time, construction site safety has mostly de- pended on people keeping an eye on other people. Safety officers walk around the site, observe ongoing work, and step in when they notice something wrong. In many projects, CCTV cameras are also installed, and someone monitors those screens from a control room. This approach isn't outdated or useless it's actually the backbone of site safety in many places. But it has its limits.

The main issue is that everything depends on human focus. On a busy site, dozens of tasks happen at once. Workers move between zones, machines operate simultaneously, and activities overlap. Expecting one or two supervisors to catch every missing helmet or unsafe action in real time is, frankly, a bit unrealistic. Even the most experienced safety officer can't be everywhere at once.

There's also the human factor. Watching multiple screens for hours is tiring. Walking across a large site under harsh weather conditions is exhausting. Over time, attention naturally drops. Small violations the ones that may seem minor at first can easily go unnoticed. And as projects become larger and more complex, relying only on manual supervision starts to feel stretched thin.

So yes, traditional monitoring does play an important role, but it isn't always steady or easy to scale as projects grow. It gets the job done just not flawlessly. And in a place as risky as a construction site, being "mostly careful" simply doesn't cut it

IJIREEICE ISSN (O) 2321-2004, ISSN (P) 2321-5526

International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering

Impact Factor 8.414 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 3, March 2026

DOI: 10.17148/IJIREEICE.2026.14389

B. Personal Protective Equipment Detection Systems

Over the past few years, computer vision has moved from being mostly experimental to something that's actually practical on job sites. One clear example is automated PPE detection. Instead of relying only on supervisors to spot missing helmets or vests, deep learning models can now be trained to recognise safety gear directly from images and live video. Helmets, reflective jackets, gloves, even harnesses the system can identify whether they're present or not within seconds.

What makes this useful is the consistency. Unlike people, these models don't get tired or distracted. They scan frame after frame and apply the same logic every time. In that sense, they can be faster and, in some situations, more reliable than traditional inspection routines. It's not about replacing safety officers, but about giving them another set of eyes that never blinks.

That said, a lot of existing PPE detection systems stop at simple verification. They check whether a helmet or vest is visible, and that's it. But safety on a construction site isn't just about wearing the right gear. A worker can be wearing a helmet and still be putting themselves at risk. The same goes for a harness just having it on doesn't automatically mean it's being used correctly. If a system only checks whether the gear is visible and doesn't look at how the person is moving or where they're positioned, it's only seeing half the situation.

So yes, PPE detection is a solid improvement, no doubt about that. But on its own, it's not enough. If we really want to make sites safer, the system has to look beyond just the equipment and understand behaviour and surroundings too. That's where the real difference starts to happen Otherwise, it becomes just another checklist automated, yes, but still limited.

C. Real-Time Object Detection Models

When it comes to monitoring busy environments like construction sites, speed really matters. You can't afford a system that takes several seconds just to understand what's in a single frame. This is where real-time detection models really start to make sense. Most of them are built using convolutional neural networks, which are basically very good at spotting patterns in images. In simple terms, they can scan a live video feed and quickly pick out workers, helmets, vests, or even machinery all without slowing the system down or interrupting the workflow.

Models such as YOLO (You Only Look Once) are especially popular for this kind of task. The name itself hints at the idea the model processes the image in one go and predicts multiple objects at once. That kind of speed really matters on construction sites, because nothing stays the same for long. On a construction site, nothing really stays still for long. Workers move from one spot to another, trucks roll in and out, and the lighting shifts as the day progresses. In that kind of environment, the system has to respond quickly and still catch the important details, without slowing down or overlooking something that could matter.

Another big plus is that these models can recognise multiple things at the same time. In just one frame, there could be several workers, different tools or machines, and a lot happening in the background. The model can sort through all of that and identify each part in real time. Compared to older, slower methods that struggled with busy scenes, this is a noticeable upgrade.

That being said, object detection isn't perfect. It can point out what's in the frame, but it doesn't always explain what's actually going on For example, it might detect a helmet, but it won't automatically know if the worker is using it correctly or standing in a hazardous position. It also doesn't naturally track how long a violation continues. So while real-time detection is a powerful foundation, it usually needs to be combined with tracking, posture analysis, or rule-based logic to fully understand unsafe behaviour.

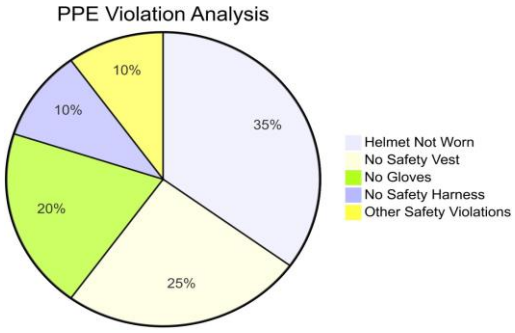

Fig. 1. Distribution of PPE Compliance and Violations

IJIREEICE ISSN (O) 2321-2004, ISSN (P) 2321-5526
International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering
Impact Factor 8.414 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 3, March 2026
DOI: 10.17148/IJIREEICE.2026.14389

D. Worker Tracking in Video Surveillance

Worker tracking helps the system keep an eye on the same person across multiple video frames instead of treating every appearance as someone new. As workers move around the site or pass between cameras, tracking algorithms assign them a consistent identity. This makes it easier to notice repeated violations, longer periods without PPE, or entry into restricted zones. Still, tracking by itself isn't enough it needs to work together with behaviour analysis to give a complete picture of site safety.

E. Pose Estimation for Behaviour Analysis

Pose estimation methods analyse human body key points to understand posture and movement patterns. These techniques assist in identifying unsafe activities such as improper lifting, climbing without support, or working at height without protective equipment. By examining body posture, systems can detect risks beyond simple PPE absence. Pose-based analysis enhances the intelligence of monitoring systems. However, many existing implementations apply pose estimation independently rather than integrating it with tracking and PPE verification.

F. Rule-Based Safety Evaluation

Detecting something is only half the job. What really matters is what the system *does* with that detection. That's where the rule-based part steps in. The model might spot a person, a helmet, or mark out a restricted area, but the rules decide if the situation is acceptable or not.

Say someone walks into a restricted zone. The system doesn't panic immediately it checks the condition. Are they authorised? If not, that's a violation. Same with PPE. If a helmet is compulsory and someone shows up without one, the system immediately treats it as a violation. There's no confusion or second guessing the rule is simple, and it applies the same way every time. And the moment that rule is broken, the system doesn't just sit there. It sends out an alert right away, so the issue can be addressed without delay. No manual checking, no waiting around. The same standards apply every single time, which honestly makes enforcement much more reliable.

That said, not every system handles it perfectly. Some just record the violation and stop there. They log it, maybe generate a report later, but they don't always escalate it immediately. And in safety-critical environments, that delay can be a real problem. When it comes to safety, quick action isn't optional it's necessary.

G. Real-Time Alert and Notification Mechanisms

In most modern safety systems, the response is immediate. If a violation is detected, an alert is triggered right away. That quick notification gives supervisors a chance to act before the situation gets worse. Sometimes it shows up on a live dashboard, other times it's sent as an automated message either way, the goal is the same: don't waste time.

In our system, if something unsafe happens at the construction site, a message is instantly sent to the site owner and supervisor. It's a simple step, but that speed can make a real difference in preventing accidents. Research Gap and Need for Integration.

Many existing studies focus on one part of the problem PPE detection, worker tracking, pose estimation, or alerts but often these pieces work separately. That separation makes monitoring less effective than it could be. Construction sites are complex, and they need everything connected. A single framework that combines detection, tracking, behaviour analysis, and instant notifications would make safety monitoring more practical and reliable, instead of depending heavily on manual supervision.

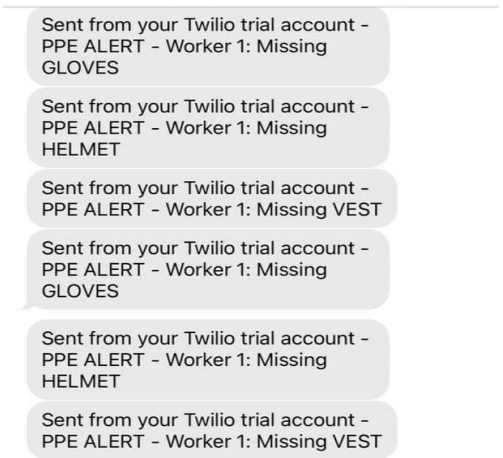

Fig. 2. SMS-Based PPE Safety Violation Alerts Generated by the System

IJIREEICE ISSN (O) 2321-2004, ISSN (P) 2321-5526
International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering
Impact Factor 8.414 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 3, March 2026
DOI: 10.17148/IJIREEICE.2026.14389

H. Research Gap and Need for an Integrated Framework

A lot of existing research handles things like PPE detection, worker tracking, or alert systems separately. And honestly, that's part of the problem. When each piece works on its own, the overall monitoring just isn't as effective as it could be. Construction sites don't need scattered solutions they need everything connected. Bringing detection, tracking, behavior analysis, and instant alerts into one real-time system would make safety management smoother and more practical, instead of relying so much on manual supervision.

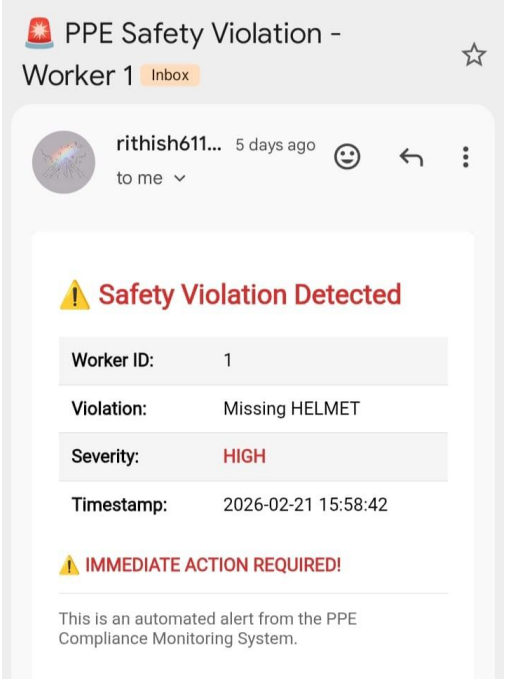

Fig. 3. Automated PPE Safety Violation Alert Notification Generated by the System

IV. METHODOLOGY

Most existing studies focus on one piece at a time PPE detection, tracking, alerts, things like that. But when these parts work separately, the system feels incomplete. It can detect issues, sure, but it's not as efficient as it should be.

Construction sites are dynamic, and everything needs to work together. A single system that combines detection, tracking, behaviour analysis, and instant alerts would honestly make safety monitoring much more practical. It cuts down on the need for someone to constantly keep an eye on everything and makes it easier to handle risks before they turn into real problems.

A. Data Collection and Annotation

We collect images and video footage directly from construction site CCTV cameras and recorded site activities. After collecting the footage, every frame is annotated by identifying workers and the PPE they're wearing helmets, safety vests, gloves, harnesses, and so on.

Once that's finished, we divide the dataset into training, validation, and testing sets. This helps us figure out whether the model is actually learning something meaningful, rather than simply memorising the data it has already seen.

B. Model Development

We use a YOLO-based object detection model to identify workers and different types of PPE in video frames. The idea is simple — the model should be able to spot multiple objects at once, even in messy, real construction site conditions.Since lighting and surroundings can change a lot on-site, the model is trained to handle those variations too. It's built to stay reliable whether it's bright daylight or a slightly dim indoor setup.

C. Pose Estimation

We also use pose estimation to understand how workers are moving and positioning their bodies. It's not just about spotting a person it's about seeing what they're actually doing.

This helps the system catch unsafe actions, like lifting something the wrong way or working at heights without proper

DOI: 10.17148/IJIREEICE.2026.14389

support. Small details in posture can say a lot, and that's exactly what this step focuses on.

D. Real-Time Monitoring

During operation, the system processes live CCTV footage frame by frame. The trained model checks for workers and confirms whether they're wearing the required PPE all in real time. At the same time, a tracking system assigns a unique ID to each worker. This way, the system can follow each worker smoothly from one frame to the next, instead of repeatedly recognising the same person as a different individual.

E. Safety Rule Evaluation

The system follows a set of clearly defined safety rules. It compares each worker's actions with those standards and flags anything that doesn't align. If there's even a small safety lapse, the system picks it up on its own no manual checking needed.

F. Alert Generation and Notification

If the system detects a violation, it immediately raises an alert and records the details of what happened. At the same time, a notification is sent to the construction owner and site supervisor so they can step in and handle it without delay.

G. Continuous Safety Enhancement

This integrated setup keeps the site under constant watch, almost like having an extra pair of eyes that never gets tired. It reduces the need for someone to manually monitor everything all the time, which honestly isn't practical on a busy construction site.

More importantly, it doesn't just record problems it spots risks early. If someone forgets safety gear or moves into a restricted area, the system reacts immediately. Catching the issue early means a minor slip doesn't get the chance to grow into something much more serious. In the long run, it makes the whole site feel more controlled and a lot safer.

V. SYSTEM ARCHITECTURE

The system follows a modular design, meaning each part has its own job but still connects smoothly with the others. It's structured in a clean, practical way that supports real-time monitoring. Since everything is organised like this, managing or expanding the system later becomes much easier. Each component does its part, and together they keep the site under steady safety supervision.

A. Data Acquisition Layer

This layer collects video and image data from CCTV cameras installed across the construction site. It continuously streams real-time footage to the processing modules. High-resolution cameras ensure that even small safety violations can be detected accurately.

B. Pre processing Layer

Before we pass the raw video to the model, we tidy it up a little so things run more smoothly. We resize each frame to a fixed size — say 224×224 — to keep everything uniform. After that, we normalise the pixel values and apply a mild smoothing filter to cut down background noise. Nothing too heavy, just enough to make a difference. We also filter out duplicate or almost identical frames. Before we pass the raw video to the model, we tidy it up a little so things run more smoothly.

C. Object Detection and PPE Verification Module Here, a YOLO-based model scans the video to detect workers and verify whether they have the necessary PPE on helmets, vests, gloves, harnesses, the usual safety gear. It basically scans each frame and matches what it sees against the safety rules.

If something's missing, it doesn't ignore it. The system flags it right away so it can be reviewed. Simple idea, but honestly, it makes safety monitoring a lot more reliable.

D. Pose Estimation and Activity Analysis Module

Pose estimation algorithms analyse the body posture of workers in each frame. This module detects unsafe activities, such as improper lifting techniques, working at heights without support, or unsafe positioning of limbs. The analysis helps in understanding not just the presence of PPE but also safe behaviour compliance.

E. Tracking and Worker Identification Module

In this module, every worker is assigned a separate ID so the system can recognise and follow the same person as the video continues. It doesn't treat each frame like a new situation it actually keeps track of who's who Because of that, it

IJIREEICE ISSN (O) 2321-2004, ISSN (P) 2321-5526
International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering
Impact Factor 8.414 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 3, March 2026
DOI: 10.17148/IJIREEICE.2026.14389

becomes easier to notice if someone keeps repeating the same safety mistake or showing unusual behaviour. It might seem like a small feature, but keeping track of workers this way makes the monitoring much more reliable over time.

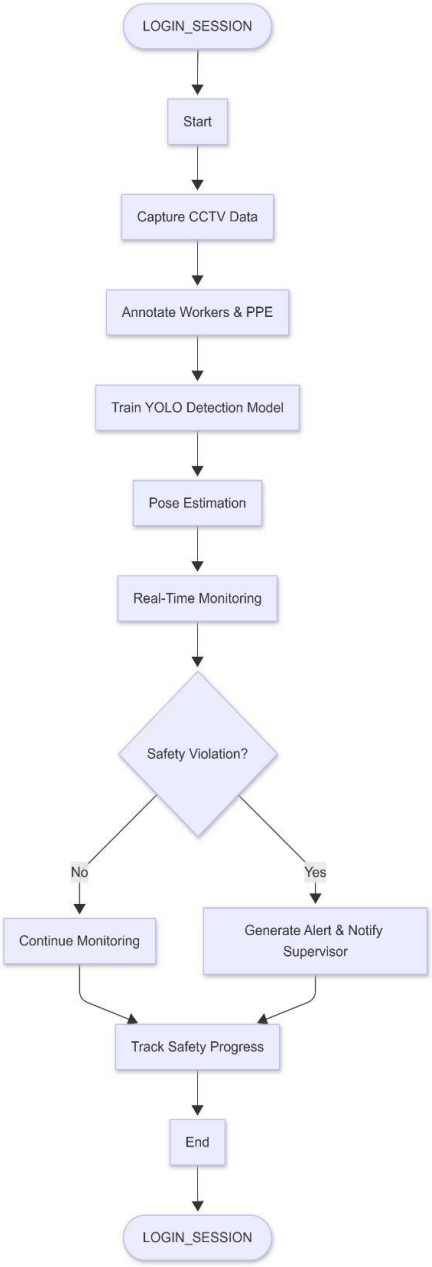

Fig. 4. Proposed Construction Site Safety Monitoring System

F. Safety Rule Evaluation and Alert Generation Module At this point, the system compares the detected activities with a set of safety guidelines. If a worker's activity doesn't match the safety guidelines, the system treats it as a possible safety concern When that happens, the incident is saved in the system and a quick alert is sent to the supervisor or site manager. This way, they can spot the problem quickly and take action before it turns into something more serious.

G. Visualisation and Reporting Layer

The system displays the results using simple dashboards and reports. These show useful information such as the number of workers on site, whether PPE is being worn properly, and any safety violations that were detected. Since the data is presented clearly, supervisors can easily understand the situation and take action whenever something needs attention.

IJIREEICE ISSN (O) 2321-2004, ISSN (P) 2321-5526

International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering

Impact Factor 8.414 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 3, March 2026

DOI: 10.17148/IJIREEICE.2026.14389

This multi-layered architecture ensures the system operates efficiently, while also providing automated safety monitoring. By combining real-time detection, behaviour tracking, and quick alerts, it supports better safety management across the construction site.

VI. RESULTS AND DISCUSSION

To see how well the system actually works, we tested it using a dataset of construction site images and video clips. The data included different situations workers wearing various PPE, people working at heights, and some common unsafe actions.Then we looked at how accurately the system could detect PPE, track workers across frames, analyse body posture, and generate alerts when a safety rule was violated. The idea was simply to check how reliable the system is in real site conditions.

A. PPE Detection Performance

The YOLO model proved capable of spotting workers and their safety equipment in video footage, all in real-time. Helmets and safety vests were detected very reliably in most cases, with the system identifying them quite accurately.Gloves and harnesses were a little harder for the system to spot because they are smaller and sometimes not clearly visible in the camera frame. Even so, the model still achieved an overall PPE detection accuracy of about 92, showing that it performs quite reliably in real construction site environments.

B. Worker Tracking and Identification

This method enables continuous monitoring of the same individual, rather than treating each frame as a separate entity.As a result, the system can accurately record repeated safety violations by the same worker. Most of the time, the system continues tracking workers even if they move frequently or disappear from the camera for a short moment, which is especially useful in the fast-moving environment of a construction site.

C. Pose Estimation and Unsafe Activity Detection

The pose estimation module facilitated the identification of hazardous actions, including lifting objects with improper posture, bending without adequate support, and working at elevated heights without the use of a safety harness. It demonstrated a capacity to distinguish between safe and unsafe postures, achieving an accuracy rate approaching 89. Consequently, supervisors are given the opportunity to detect unsafe behaviours at an earlier stage, thereby enabling timely intervention to mitigate potential problems.

D. Alert Generation and Response

Whenever the system detects a safety violation, it automatically sends an alert. In most cases, the notification is generated in less than two seconds after the issue is detected, which allows for quick response on site.The system also keeps a record of these incidents. This stored information can later be used for safety reviews, reports, or even worker training.

E. Discussion

The results show that a computer vision-based system can really help improve safety on construction sites. By continuously checking PPE usage and worker behaviour, the system can catch potential risks early. The detection works quite well overall, though smaller items or partially hidden objects can still be tricky to spot. In the future, using multiple cameras, thermal sensors, or edge AI devices could make the system even faster and more reliable on site.

VII. FUTURE WORK

The proposed system shows a practical way to automatically monitor safety on construction sites. It helps keep track of worker behaviour and safety compliance without constant manual supervision.That said, there's still room to improve. With some improvements, the system could work more reliably, handle larger environments more easily, and respond better to real situations on site.

A. Multi-Camera Integration

In later versions, the system could be expanded by placing cameras at different spots across the site. This would help capture areas that a single camera might overlook and improve how accurately the system detects activity. With coverage from several angles, the system could observe larger parts of the construction site and follow workers more steadily, even when the area becomes busy or crowded.

B. Edge Computing and On-Site Processing

Right now, the system mainly depends on central servers to process the video data. In the future, adding edge AI devices or local processing units could help handle some of that work directly on site. This would reduce delays, make real-time

DOI: 10.17148/IJIREEICE.2026.14389

responses faster, and also allow the system to run in remote areas where internet connectivity is limited.

C. Advanced Pose and Activity Analysis

The system could also be improved to identify more complicated safety hazards, like workers standing too near active machinery or not using scaffolding properly. With more advanced activity analysis that looks at movement across video frames, the system could better understand how workers behave on site. This would help make safety monitoring more accurate and practical.

D. Predictive Safety Analytics

Over time, the system can learn from past safety data instead of only reacting to problems. By studying earlier incidents, it may start identifying areas that are more risky on the site.This could help predict possible accidents and suggest preventive steps before a violation or unsafe situation actually happens.

E. Integration with IoT Sensors

One possible improvement is linking the system with wearable IoT devices used by workers, such as smart helmets or harness sensors. These devices can provide additional safety information along with the video feed. When both types of data are used together, the system can understand the situation on the site more clearly. As a result, it can react more quickly if a risky situation starts to develop.

F. User-Friendly Dashboard and Reporting

In later versions, the system could offer an interactive dash- board where supervisors can view safety updates as they occur. Simple alerts and easy-to-read reports would help them notice any problems quickly. When all the information is displayed in one place, it becomes easier for site managers to understand the situation and respond at the right time.

VIII. CONCLUSION

This system shows that computer vision and deep learning can be useful for improving safety on construction sites. It can identify workers, check if they are wearing the necessary PPE, and notice unsafe actions in real time. Since these checks run continuously, supervisors do not have to depend only on manual observation. The system can quickly highlight possible risks and allow timely action. Overall, it provides a practical and reliable way to support safer construction environments and reduce the chances of accidents.

REFERENCES

[1]. P. Teicholz, “Labor productivity declines in the construction industry: Causes and remedies,” 2013. [Online]. Available: https://doi.org/10.1061/9780784413517.001

[2]. C. M. Eastman, R. Sacks, G. Lee, and P. Teicholz, *BIM Handbook: A Guide to Building Information Modeling*, 3rd ed., 2018. [Online]. Available: https://doi.org/10.1002/9781119287568

[3]. R. Sacks, L. Koskela, B. A. Dave, and R. Owen, “Interaction of lean and building information modeling in construction,” 2010. [Online].Available: https://doi.org/10.1061/(ASCE)CO.1943-7862.0000203

[4]. C. T. Formoso, L. Soibelman, C. De Cesare, and E. L. Isatto, “Material waste in building industry,” 2002. [Online]. Available: https://doi.org/10.1061/(ASCE)0733-9364(2002)128:4(316)

[5]. M. Abanda, L. Tah, and R. Keivani, “BIM in construction risk management,” 2017. [Online]. Available: https://doi.org/10.1108/ECAM-01- 2016-0013

[6]. H. Li, G. Q. Shen, and X. Xue, “Research on safety management in construction,” 2012. [Online]. Available: https://doi.org/10.1016/j.ssci.2011.12.021

[7]. W. Lu and H. Yuan, “Investigating waste reduction potential in the construction industry,” 2010. [Online]. Available: https://doi.org/10.1016/j.jclepro.2010.02.001

[8]. A. Borrmann, M. Ko¨nig, C. Koch, and J. Beetz, *Building Information Modeling*, 2018. [Online]. Available: https://doi.org/10.1007/978-3-319- 92862-3

[9]. S. Azhar, M. Khalfan, and T. Maqsood, “Building information modeling (BIM): Now and beyond,” 2012. [Online]. Available: https://doi.org/10.5130/ajceb.v12i4.3032

[10]. J. Zhang and N. El-Gohary, “Semantic NLP-based information extraction,” 2016. [Online]. Available: https://doi.org/10.1061/(ASCE)CP.1943-5487.0000503

[11]. X. Li, H. Yi, X. Chi, X. Wang, and A. Chan, “Virtual and augmented reality in construction,” 2018. [Online]. Available: https://doi.org/10.1016/j.autcon.2017.11.003

[12]. Y. K. Juan and P. Gao, “Decision support system for sustainable construction,” 2010. [Online].

DOI: 10.17148/IJIREEICE.2026.14389

Available: https://doi.org/10.1016/j.autcon.2009.10.002

[13]. D. Arditi and H. M. Gunaydin, "Total quality management in construction," 1997. [Online]. Available: https://doi.org/10.1016/S0263- 7863(96)00076-2

[14]. K. K. Chitkara, *Construction Project Management*, 2014. [Online] Available: https://www.mheducation.co.in

[15]. T. H. Kim, J. S. Chi, and S. H. Kim, "IoT-based smart construction site management," 2020. [Online]. Available: https://doi.org/10.3390/s20030780

[16]. F. Golparvar-Fard et al., "Image-based construction progress monitoring," 2009. [Online]. Available: https://doi.org/10.1016/j.autcon.2009.06.005

[17]. J. Teizer, "Vision-based sensing on construction sites," 2015. [Online]. Available: https://doi.org/10.1016/j.aei.2015.03.006

[18]. S. Lee, F. Pen˜a-Mora, and H. Park, "Web-enabled project control systems," 2006. [Online]. Available: https://doi.org/10.1061/(ASCE)0733- 9364(2006)132:7(683)

[19]. World Economic Forum, "Shaping the future of construction," 2016. [Online]. Available: https://www.weforum.org/reports/shaping- the-future-of-construction

[20]. McKinsey Global Institute, "Reinventing construction: A route to higher productivity," 2017. [Online]. Available: https://www.mckinsey.com/industries/capital-projects-and-infrastructure/our-insights/reinventing-construction-through-a- productivity-revolution