

Vigil AI: Workplace Safety Monitoring System

Aaliya Khan

Email: aaliyakhan4352@gmail.com

Sejal Ubale

Email: Sejalubale32@gmail.com

Rimjhim Chauhan

Email: rimjhim.chauhan.st@gmail.com

Under the guidance of

Dr. Ranjana Kumari

Email: ranjana.thakur@ltce.in

Dr. Nandini nag

Email: nandini.nag@ltce.in

Department of Computer Science Engineering (Data Science), Lokmanya Tilak College of Engineering, An Autonomous Institute Affiliated to University of Mumbai

ABSTRACT

Vigil AI is an AI-powered workplace safety monitoring system that continuously tracks workers in real-time, automatically verifying whether they are wearing required Personal Protective Equipment such as helmets and safety vests, eliminating the need for manual surveillance.

The system uses a YOLOv11 detection engine to instantly identify PPE violations through live camera feeds, triggering immediate alerts via sirens and voice warnings. Administrators can define custom restricted zones on the interface, while a centralized dashboard displays live footage, detection overlays, and compliance status in one unified view. Automated shift reports and safety scores help organizations track performance and drive continuous improvement.

What sets Vigil AI apart is its predictive capability - it anticipates danger by monitoring worker movement and issuing warnings before a restricted zone is even entered. Built using Next.js, Flask-SocketIO, and OpenCV, the system delivers reliable, low-latency performance across industrial environments. With plans to expand into IoT sensor integration, fatigue detection, and safety heatmaps, Vigil AI stands as a proactive and scalable solution for modern workplace safety.

Keywords: AI-powered workplace safety monitoring system, OpenCV, YOLOv11 detection, centralized dashboard, Next.js, Flask-SocketIO.

1. INTRODUCTION

In modern industrial and construction settings, maintaining worker safety is one of the most demanding responsibilities organisations face. Conventional monitoring relies heavily on human observation, which is often inconsistent, error-prone, and unable to cover all of a large worksite. This highlights the urgent need for an intelligent and automated system that can effectively fill this gap.

Vigil AI is an AI-based workplace safety platform that uses computer vision and live video processing to continuously monitor workers in hazardous environments. Unlike traditional CCTV systems that merely store footage for later review, Vigil AI analyses live feeds in real-time, detects safety violations instantly, and responds without delay. Automatically check

compliance with, identifies unauthorised zone entries, and notifies supervisors immediately through voice- and siren-based alerts in simple, easy-to-understand language.

Rather than reacting after incidents occur, Vigil AI anticipates risks before they escalate, transforming workplace safety into a proactive practice. With its intuitive dashboard, flexible zone configuration, and automated reporting tools, it empowers organisations to build a safer and more responsible work environment.

Motivation

We developed Vigil AI to address a deeply concerning reality in industrial and construction workplaces - safety violations frequently go unnoticed due to the limitations of manual monitoring. Supervisors cannot realistically watch every worker on an entire site at once, and existing CCTV systems merely record incidents rather than prevent them, leaving workers consistently exposed to avoidable risks.

We aimed to build a system that not just observes, but actively protects it. We wanted a solution capable of immediately detecting violations and predicting danger before it unfolds, ensuring that all workers on site remain protected without solely relying on human vigilance.

The motivation behind Vigil AI was to make workplace safety smarter, faster, and more reliable. By combining real-time PPE detection, predictive zone monitoring, instant alerts, and automated reporting, we sought to replace the burden of constant manual supervision with an intelligent system that works around the clock - because no worker should ever face harm that could have been prevented.

Problem Statement

In high-risk workplaces, manual safety monitoring is inconsistent and unreliable as supervisors cannot watch every worker at the same time. Existing systems are purely reactive, alerting only after a violation has already occurred, by which time an accident may already be in motion. There is a need for a smart automated solution that monitors workers in real-time and predicts danger before it happens.

2. LITERATURE SURVEY

2.1 Real-Time Personal Protective Equipment (PPE) Detection using YOLOv8 and Computer Vision for Industrial Safety Compliances

In August 2025, this research focused on automating workplace safety compliance in industries such as construction, manufacturing, and chemical processing. The core problem addressed was the inconsistency of manual safety monitoring, where supervisors are unable to continuously track every worker on large worksites. The system processes live video streams and instantly identifies whether workers are wearing essential safety equipment such as helmets, safety vests, and face masks, eliminating the need for manual checks entirely.

The system was built using the YOLOv8 object detection model, trained on a diverse annotated industrial dataset with data augmentation techniques including rotation, scaling, flipping, and brightness adjustments to improve real-world adaptability. Developed using Python and OpenCV for live video streaming, the system achieved a mAP@50 of 96%, maintaining smooth performance at 25–30 FPS on GPU and 12–15 FPS on CPU, confirming its practicality as a scalable and cost-effective alternative to traditional manual surveillance.

2.2 A Real-Time Advisory Tool for Supporting the Use of Helmets in Construction Sites

In September 2025, this research focused on building a real-time safety tool designed to monitor whether workers on construction sites were wearing helmets. The system went beyond simply detecting violations — it also provided instant safety advice to supervisors based on what was observed through live camera feeds. The goal was to automate a process that was

previously dependent on manual supervision, making workplace safety monitoring faster, smarter, and more reliable.

The system was trained using a large collection of construction site images to accurately distinguish between workers wearing helmets and those without helmets. Multiple detection models were tested and compared, with the best achieving very high accuracy results. These models were then connected to an AI language model that could generate clear, easy-to-understand safety recommendations in real time. The entire tool was made accessible through a simple web-based interface, allowing users to upload images, analyze recorded videos, or monitor live camera feeds, with automatic safety feedback delivered within seconds.

2.3 Artificial intelligence-driven systems for enhancing worker safety monitoring in construction environments: A comprehensive review (August 2025)

This study carried out an in-depth review of AI-based safety monitoring systems in the construction sector, drawing insights from 88 peer-reviewed articles published between 2017 and 2025. It established that conventional safety practices fall short, offering only 50-60% accuracy and taking several days to detect risks, highlighting an urgent demand for more intelligent and automated monitoring solutions.

The findings covered a broad spectrum of AI-driven technologies, including live video-based safety gear detection with an accuracy exceeding 91%, automated aerial inspections that cut the assessment time from multiple days to under two hours, sensor-equipped wearables for continuous health tracking, and data-driven models capable of predicting workplace accidents with nearly 89% accuracy. However, the study also highlighted significant obstacles to broader implementation, including concerns about data privacy, biased decision-making in AI models, steep setup costs, and notable gaps in technology adoption between wealthier and developing nations.

2.4 Why Does an Indian Construction Worker Fail to Wear Personal Protective Equipment (PPE) at Workplace?

In 2022, this research investigated why construction workers in India consistently avoid wearing safety helmets on site, despite being fully aware of their importance. Rather than focusing on technology, the study took a human-centered approach by directly engaging workers through on-site observations and a structured survey. The goal was to understand the real reasons behind non-compliance and identify what changes could make helmets more acceptable and comfortable for everyday use.

The study surveyed 36 active construction workers in different job roles and found that physical discomfort was the main reason workers refused to wear helmets. The main complaints included poor visibility, excessive heat, sweating, itching, and problems with the straps that hold the helmet in place. Based on these findings, the study recommended redesigning existing helmets to be lighter, better ventilated, and more comfortable to wear for long hours in hot conditions. It concluded that workers are more likely to follow safety rules when the equipment itself is practical and suitable for their working environment.

2.5 Existing System & its Drawbacks

Problem with the current scenario:

Real-Time Personal Protective Equipment (PPE) Detection using YOLOv8 and Computer Vision for Industrial Safety Compliances (August 2025):-

- The system is limited to detecting only helmets, safety vests, and face masks, making it insufficient for workites that require broader coverage of safety gear.
- Performance drops in CPU-only systems with frame rates as low as 12–15 FPS, making it

less reliable in environments without high-end hardware.

- The system lacks predictive capabilities and can only react to violations after they occur, rather than preventing them in advance.
- No alert mechanism is integrated to notify supervisors in real time when a violation is detected.

A Real-Time Advisory Tool for Supporting the Use of Helmets in Construction Sites (September 2025):-

- The system focuses exclusively on helmet detection and does not cover other essential safety equipment, such as vests, gloves, or goggles.
- The tool depends on an internet connection to communicate with the AI language model, making it unreliable in areas with poor connectivity.
- Safety advice is generated based only on helmet counts and lacks deeper contextual understanding of the actual risk situation on the ground.
- The system was tested only as a proof of concept and has not been validated for large-scale real-world deployment

Why Does an Indian Construction Worker Fail to Wear Personal Protective Equipment (PPE) at Workplace? (2022):-

- The study is based on a small sample of only 36 workers from a single region in northeastern India, limiting its applicability to broader populations.
- It focuses purely on identifying the problem through surveys rather than proposing or implementing any technological solution to address non-compliance.
- The research does not account for workers from different regions, climates, or construction roles, reducing the generalizability of its findings.
- -No automated or AI-based monitoring system is suggested, leaving the responsibility of enforcement entirely on human supervisors and management.

Therefore, there is a clear need for a system like Vigil AI that addresses the short-comings of existing approaches by providing real-time, AI-driven workplace safety monitoring. Rather than just recording violations, Vigil AI proactively detects missing safety equipment, anticipates potential hazards, and immediately notifies supervisors, making worker protection smarter and less dependent on manual oversight.

3. PROPOSED METHODOLOGY

3.1 Methodology

VigilAI is an AI-driven intelligent safety monitoring system built to detect PPE compliance and predict workplace hazards in real time across construction and industrial environments. The system keeps a continuous watch on every worker through live camera feeds and automatically verifies whether each worker is equipped with the necessary safety gear, including helmets and vests. Whenever a violation is identified, supervisors are immediately notified through audio and visual alerts. The key highlight of VigilAI is its predictive capability, where it tracks worker movement patterns and issues warnings before a worker steps into a restricted or hazardous zone, making safety monitoring proactive rather than reactive.

The detection pipeline is powered by YOLOv11, which analyses live video to identify workers and assess their equipment status in real time. OpenCV manages all video and image processing tasks, while Shapely enables precise geometry-based zone mapping, giving admins the flexibility to define custom restricted areas on the site. Missing PPE triggers an immediate gear violation log, unauthorised zone entry generates an intrusion alert, and movement toward a danger area activates a predictive warning ahead of time. A persistence confirmation step

ensures that brief detection errors are filtered out before an evidence screenshot is captured and the alert is dispatched to the concerned supervisors.

The system is built on a modern and reliable technology stack to deliver speed, accuracy, and ease of use. Next.js and Tailwind CSS power the front-end dashboard, displaying live video feeds, detection overlays, and restricted zone boundaries in a clean interface. Flask paired with Flask-SocketIO forms the back end, ensuring fast and seamless real-time data exchange between the AI engine and the dashboard. Shift-end safety reports are automatically generated using fpdf2. Together, these components form a scalable and efficient architecture that enables uninterrupted monitoring, instant alerting, and live safety score updates throughout the entire worksite.

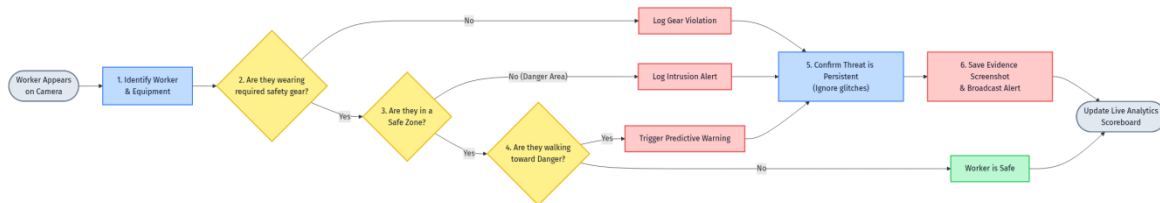


Figure 3.1: Methodology

As shown in fig.3.1, the methodology operates as a step-by-step decision pipeline. When a worker is detected on camera, YOLOv11 identifies the person and checks their PPE compliance. Any missing safety gear is immediately flagged as a violation. For compliant workers, the system then determines their zone — those in restricted areas generate an intrusion alert, while those in safe zones are tracked for movement toward danger, triggering a predictive warning if required. Every alert goes through a persistence check across frames to filter out false detections before an evidence screenshot is saved and supervisors are notified. Workers with no identified risk are logged as safe, and all results are reflected in real time on the live analytics scoreboard.

As shown in fig.3.2, the system works in four simple stages. First, a live camera feed is taken as input and passed to the Visual Detection Engine, which spots workers and checks if they are wearing the required safety gear. These findings are then sent to the Rule Engine, which compares them against set safety rules and zone boundaries to decide if anything is wrong. Depending on what is found, three things happen at the same time — the Live Dashboard updates for real-time monitoring, the Alert System sends out instant warnings on mobile and web, and daily safety score reports are automatically generated.

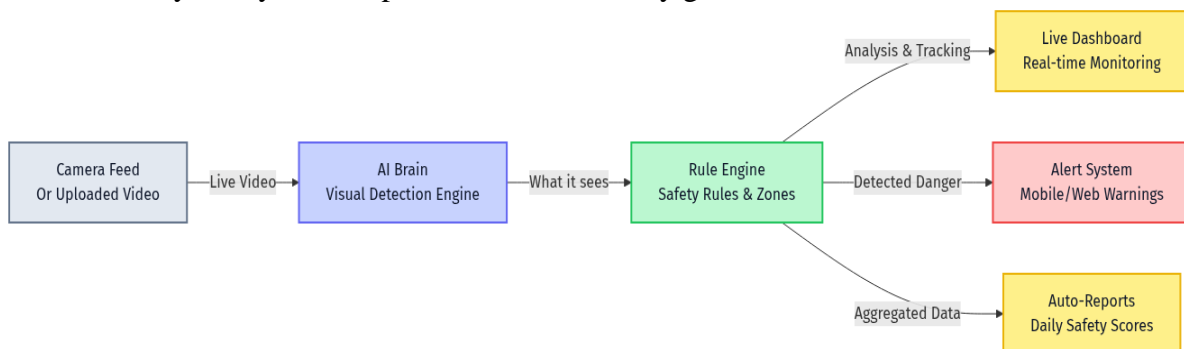


Figure 3.2: Architecture

IMPLEMENTATION Working Model

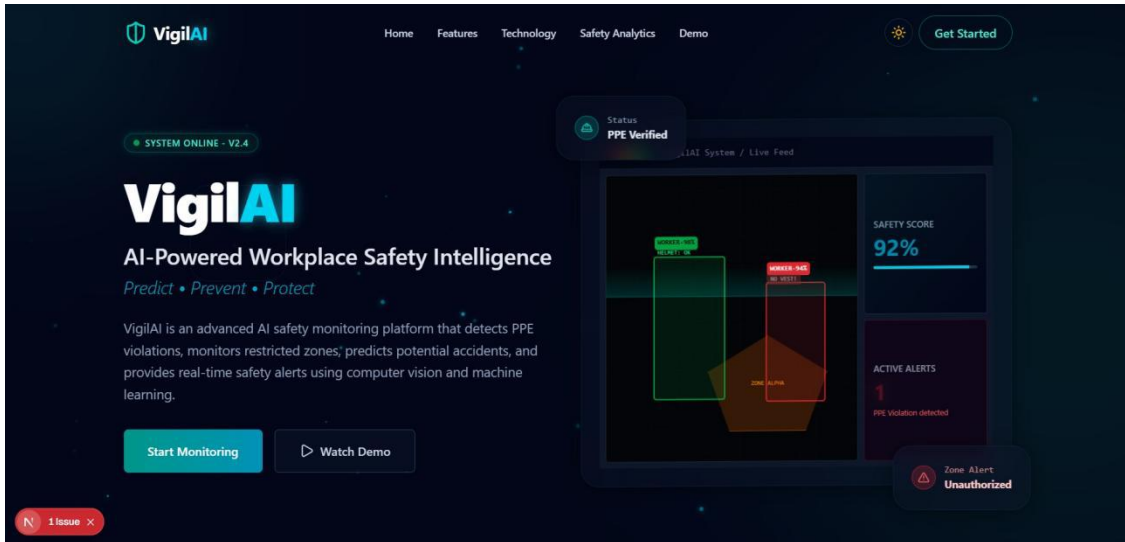


Figure 4.1: Landing page

As shown in fig.4.1, this page serves as the welcoming interface for Vigil AI, an AI-powered workplace safety monitoring platform. Its sleek dark-themed design captures attention with the bold branding of "VigilAI", the tagline "AI-Powered Workplace Safety Intelligence", and the motto "Predict • Prevent • Protect". The header features a clean navigation bar with links to Home, Features, Technology, Safety Analytics, and Demo, making it easy for users to explore the platform.

The homepage briefly describes the system's ability to detect PPE violations, monitor restricted zones, predict potential accidents, and deliver real-time safety alerts. A live feed panel on the right shows how the system is in action, displaying worker detection with bounding boxes, a safety score of 92percent, an active PPE violation alert, and an unauthorised zone alert. Two call-to-action buttons - 'Start Monitoring' and 'Watch Demo' - invite users to immediately engage with the platform.

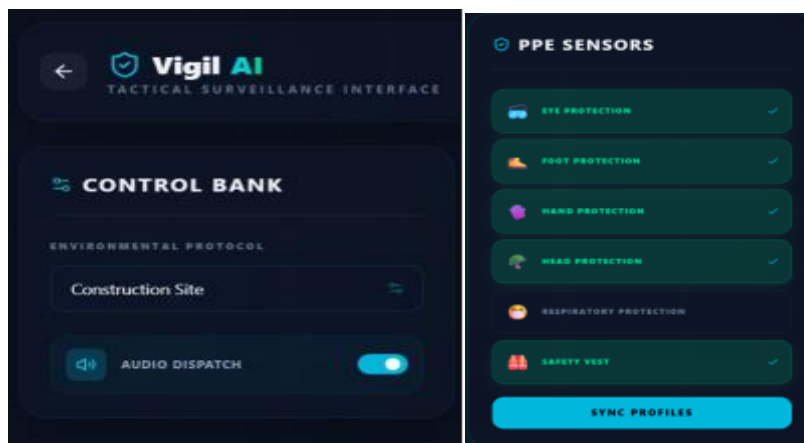


Figure 4.2: feature 1

As shown in fig.4.3, the PPE Sensors panel and the Control Bank interface are key components of the Vigil AI Tactical Surveillance Interface, built using HTML, CSS, JavaScript, and Flask-SocketIO.

The Control Bank panel serves as the configuration interface for the surveillance system. It allows administrators to set the Environmental Protocol by selecting the workites type, such as a Construction Site, ensuring the correct safety rules are applied. The Audio Dispatch toggle, when enabled, activates real-time voice and siren alerts whenever a violation is detected, ensuring that supervisors are notified immediately without any manual intervention.

The PPE Sensors panel displays a real-time checklist of all required safety equipment being monitored, including Eye Protection, Foot Protection, Hand Protection, Head Protection, Respiratory Protection, and Safety Vest. Each item is marked with a checkmark upon detection, giving supervisors a quick overview of worker compliance. The Sync Profiles button updates and saves the detected PPE data across the system instantly.

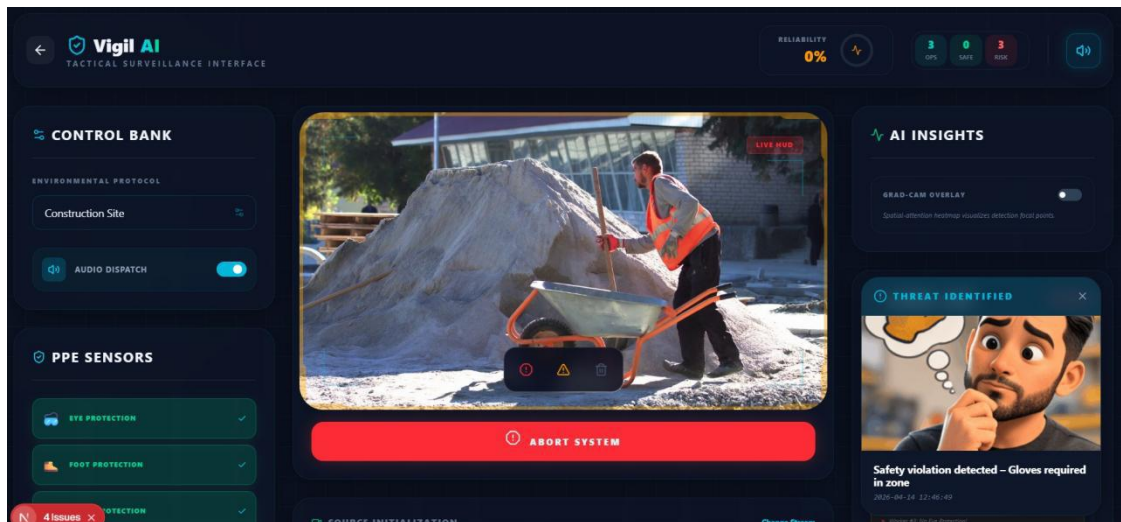


Figure 4.3: login page

As shown in the figure, this page represents the "Vigil AI Tactical Surveillance Interface", the core monitoring dashboard built using HTML, CSS, JavaScript, Flask, and OpenCV.

The dashboard is organised into three sections. The left panel contains the 'Control Bank', where administrators can configure the worksite type and enable 'Audio Dispatch' for instant voice and siren alerts, along with the 'PPE Sensors' checklist that tracks detected safety equipment in real time. The centre displays a 'Live HUD' feed showing active worker monitoring through the camera stream, with an 'Abort System' button to halt surveillance instantly when needed. The right panel features 'AI Insights' with a 'Grab-CAM Overlay' that highlights areas the AI focused on, alongside a 'Threat Identified' alert showing timestamped violation details. The header in the top summarises key metrics such as total operations, safe workers, and risk count, giving supervisors a complete worksite safety overview at a glance.

Vigil AI Shift Report

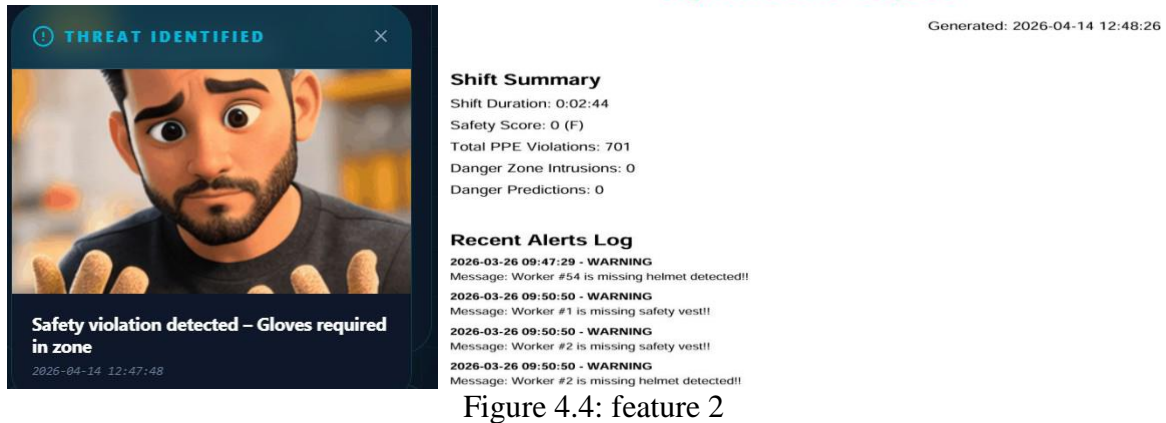


Figure 4.4: feature 2

As shown in fig.4.4, the 'Threat Identified' alert and 'Vigil AI Shift Report' are two essential output features of the system, developed using JavaScript, Flask-SocketIO, pyttx3, and fpdf2. The 'Threat Identified' pop-up is triggered automatically on the dashboard the moment a safety breach is identified, presenting a clear violation message alongside a precise timestamp, for example, 'Safety violation detected – Gloves required in zone'. A relevant animated GIF sticker accompanies each alert, offering a visual representation of the specific violation so that all workers and staff can grasp the warning without any technical knowledge.

The 'Vigil AI Shift Report' is a PDF document automatically produced at the conclusion of every monitoring session. Provides a structured overview of the shift, including duration, safety score, total PPE breaches, restricted zone intrusions, and hazard predictions. A timestamped 'Recent Alerts Log' further details each individual incident, such as workers found without helmets or safety vests, offering safety supervisors a well-organised and comprehensive record to support ongoing safety improvements.

Software and Hardware Used

Hardware

- Processor: Intel i5 or higher for smooth real-time video processing and AI detection.
- RAM: Minimum 8 GB to support simultaneous video processing, detection, and alert operations.
- Storage: 100 GB SSD or higher to store trained model weights, violation screenshots, logs, and PDF shift reports.
- Camera: HD Camera or CCTV feed to capture live video streams from the worksite.
- Internet Connection: High-speed internet for real-time communication between the frontend dashboard and the backend server.

Software

Operating System: Windows 10/11

Device:

Acer Aspire 3 Intel Core i7-1255U

Lenovo IdeaPad Slim 3 Intel Core i3 12th Gen

HP Victus Gaming Laptop, 12th Gen Intel Core i5-12450H

VS Code as IDE: We have used Visual Studio Code (VS Code) as our Integrated Development Environment (IDE) to write, edit, and manage your project's code. It is a popular choice because of its powerful features, extensions, and ease of use for development.

Programming Language:

FRONTEND:

HTML5- structures the overall dashboard layout and user interface.

CSS3- handles all styling, design, and visual presentation of the platform.

Vanilla JavaScript- manages all user interactions and dynamic content on the dashboard.

Canvas API- enables drawing and defining restricted danger zones directly on the live video feed.

WebSocket- receives real-time updates from the backend without page refresh.

Flask-SocketIO- handles all live event communication between the frontend and backend.

BACKEND:

Flask- serves as the main backend framework, handling REST API routes and server-side logic.

OpenCV- captures live frames from the connected camera for real-time video processing.

Shapely- checks and verifies whether workers are present inside defined restricted zones.

pygame- plays siren alert sounds whenever a safety violation is detected.

pyttsx3- generates text-to-speech voice alerts to notify supervisors instantly.

Grad-CAM- visually explains AI decisions by highlighting the areas the model focused on during detection.

STORAGE (File System / JSON):

Vigil AI relies on a well-organized local storage system to manage all monitoring data. Zone polygon data for restricted areas is saved in JSON format, ensuring custom danger zones remain intact across sessions. Violation screen-shots are automatically captured and stored by the video processing module, while shift-end safety reports are generated as PDF files summarizing all incidents. The trained model weights are stored locally and loaded during runtime to enable continuous detection. All alert events are also logged with timestamps on the local file system, maintaining a reliable record for future auditing and review.

MODEL TRAINING (Google Colab T4):

The core detection engine of Vigil AI is powered by YOLOv11, trained to identify helmets and safety vests through live camera feeds. The Ultralytics framework was used to streamline the training process, with PyTorch serving as the underlying deep learning backbone. A carefully prepared dataset of 10,875 annotated workplace images from Roboflow was used to ensure accurate and reliable detection. Training was performed on a Google Colab T4 GPU over 50 epochs to achieve strong model performance. Grad-CAM was also integrated to provide visual explainability, helping users understand which parts of an image influenced the AI's detection decisions.

RESULT AND DISCUSSION

The VigilAI system effectively monitors workplace safety using a YOLOv11-powered engine to detect personnel and PPE compliance, such as helmets and safety vests. It features an interactive dashboard for live video feeds, real-time bounding boxes, and restricted zone mapping via polygon overlays. By integrating Next.js with a Flask-SocketIO backend, the system ensures low-latency communication and automated alerts for safety breaches. Testing results show that VigilAI significantly reduces human error and enhances operational security, providing a robust, AI-driven solution for industrial and construction site safety.

CONCLUSIONS

VigilAI demonstrates the power of AI in automating industrial safety monitoring. Using YOLOv11 for real-time PPE detection (helmets, vests) and restricted zone mapping, it provides a reliable alternative to manual CCTV surveillance. The system ensures continuous monitoring, immediate alerts, and significantly reduces human error, making workplaces safer and more compliant.

Future Scope

Expand detection to include safety gloves, boots, goggles, and masks, making it adaptable for the mining, chemical, and food industries. We plan to integrate IoT gas and pressure sensors for early explosion detection and environmental monitoring. Using AI-driven spatial analysis, the system will generate dynamic safety heatmaps and use predictive analytics to identify worker fatigue or high-risk behaviors, offering an all-in-one proactive safety solution.

References

- Makandar A, Naik R. "Real-Time Personal Protective Equipment (PPE) Detection using YOLOv8 and Computer Vision for Industrial Safety Compliances." IJARCCCE – International Journal of Advanced Research in Computer and Communication Engineering, Vol. 14, Issue 8, August 2025, pp. 303-306. DOI: 10.17148/IJARCCCE.2025.14832.
- Jayakody J.A.U.S., Thilakarathne P.R.H.N.G. "Artificial intelligence driven systems for enhancing worker safety monitoring in construction environments: A comprehensive review." Proceedings of the 13th World Construction Symposium, 15-16 August 2025, Sri Lanka, pp. 154-166. DOI: 10.31705/WCS.2025.12.
- Vigneshkumar C, Salve UR. "Why Does an Indian Construction Worker Fail to Wear Personal Protective Equipment (PPE) at Workplace?" In: Chakrabarti D et al. (eds.), Ergonomics for Design and Innovation, Lecture Notes in Networks and Systems, Vol. 391, Springer Nature Switzerland, 2022, pp. 271–277. DOI: 10.1007/978-3-030-94277-9-24.
- Isikdag U, Cemrek H.A, Sonmez S, Aydin Y, Bekdas G, Geem Z.W. "A Real-Time Advisory Tool for Supporting the Use of Helmets in Construction Sites." Information, Vol. 16, Issue 10, September 2025, Article 824, pp. 1-17