Artificial Intelligence (AI) for industrial safety

Dr. Thanuja D. Ambegoda

Introduction: the limits of traditional safety

Industrial environments are inherently complex and risky. From construction sites to factory floors, safety depends on the coordination of people, machines, and environmental conditions. Over the years, regulations, inspections, and training programmes have reduced many risks. Yet, accidents continue to occur, often due to factors that were either unnoticed or recognized too late.

A key limitation of traditional safety systems is that they are largely reactive. Incidents are analyzed after they happen, and improvements are introduced to prevent recurrence. While necessary, this approach does little to prevent the first occurrence. Human supervisors, no matter how experienced, cannot continuously observe every worker, every machine, and every condition across a large industrial environment.

Figure 1: Evolution of industrial safety: from reactive response to predictive prevention using AI.

Artificial Intelligence (AI) introduces a different approach. Instead of depending solely on periodic inspections, AI enables continuous monitoring and early detection of risks. More importantly, it allows organizations to anticipate unsafe situations before they develop into incidents.

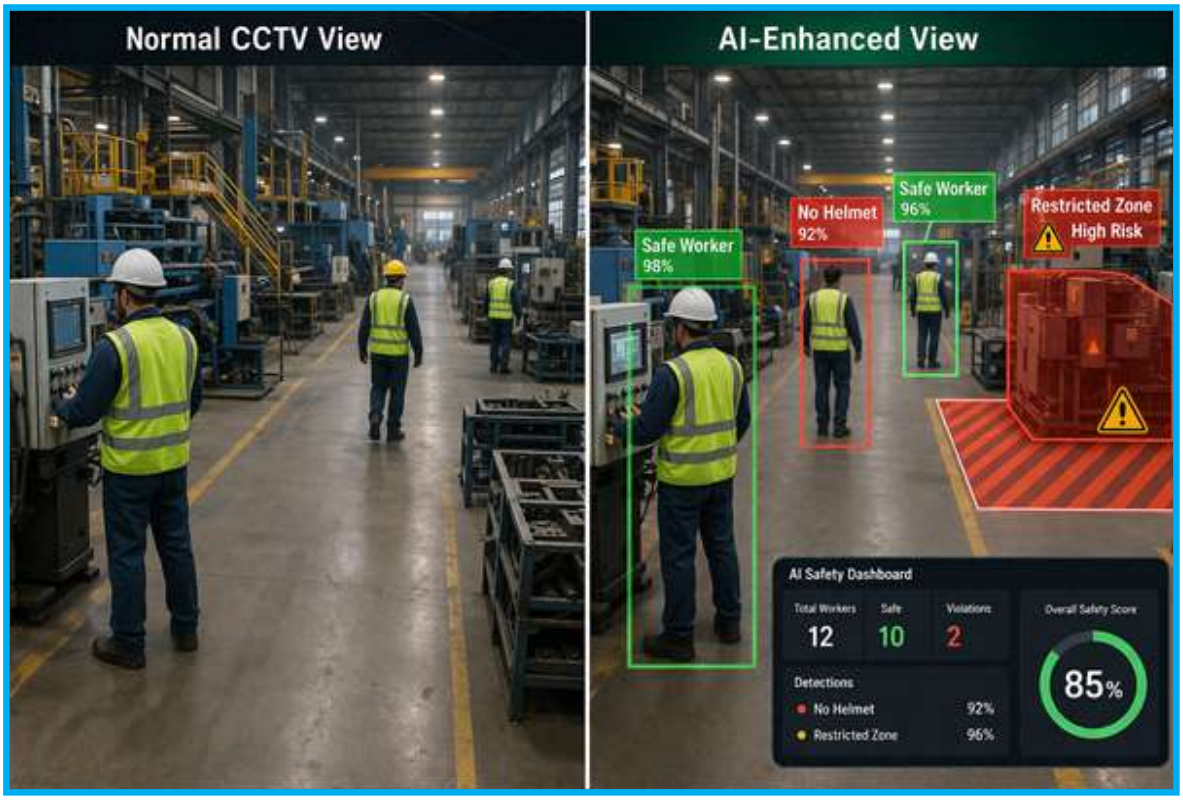

Figure 2: AI-based computer vision identifying safety violations in real time.

How AI enhances industrial safety in practice

One of the most widely adopted applications of AI for industrial safety is the computer vision. For example, in a large manufacturing plant, cameras installed above assembly lines can automatically detect whether workers are wearing helmets, gloves, or safety vests. If a worker enters a high-risk zone without proper equipment, the system can immediately trigger an alert to a supervisor. This removes the need for constant manual checking.

In construction sites, AI systems are being used to detect the fall risks. Cameras can identify whether workers are operating at height without harnesses or whether scaffolding areas are overcrowded. In some deployments, near-miss events such as a worker stepping into the path of a moving crane are automatically logged and analyzed. These insights help site managers identify patterns, such as the specific times of day or the locations where risks are higher.

Another example comes from warehouses and logistics centres. In such environments, accidents often occur due to interactions between workers and forklifts. AI-based systems can track both workers and vehicles in real time. If a worker comes too close to a moving forklift, the system can issue a warning to both the driver and the worker, sometimes even slowing down the vehicle automatically. These examples illustrate a key point: AI does not merely observe. It can be used to intervene at the right moment.

From detection to prediction: moving one step ahead

Beyond real-time monitoring, AI enables a shift toward prediction. In industrial equipment, sensors continuously monitor variables such as vibration, temperature, and pressure. Consider a thermal power plant or a cement factory. A slight but consistent increase in vibration in a rotating machine may not be noticeable to operators, but AI models can detect it as an early sign of mechanical failure. If such a failure occurs suddenly, it could lead not only to production loss but also to serious safety hazards for nearby workers. By predicting the issue early, maintenance can be scheduled

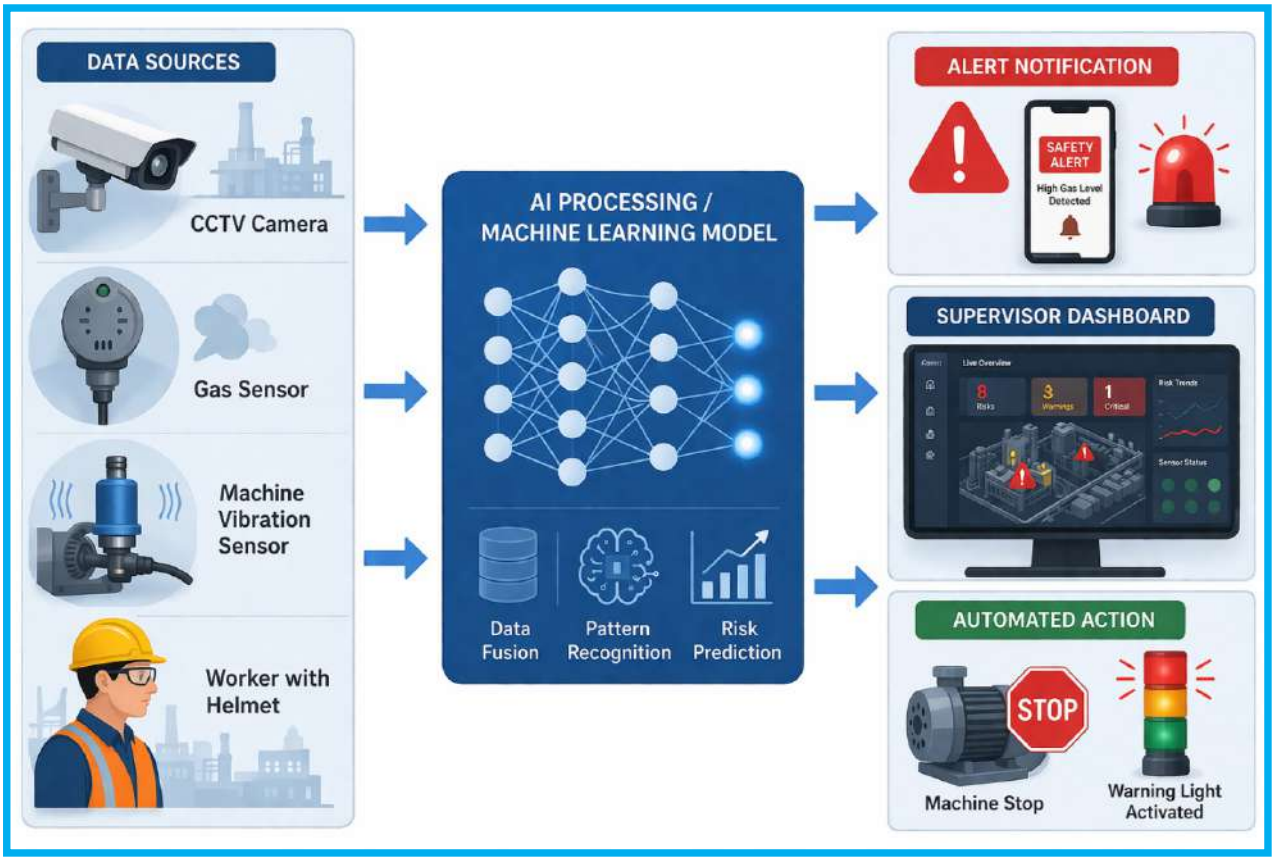

Figure 3: Typical architecture of an AI-driven industrial safety monitoring system.

before the situation becomes dangerous.

In the oil and gas sector, predictive models are used to monitor gas leaks. Instead of relying only on threshold-based alarms, AI systems analyze patterns in gas sensor readings. Subtle deviations from normal patterns can indicate early-stage leaks, allowing intervention before concentrations reach dangerous levels.

Another emerging example is the fatigue detection. In long-duration operations such as mining or transport logistics, worker fatigue is a major safety risk. AI systems using cameras inside vehicle cabins can detect signs such as eye closure patterns or head movements that indicate drowsiness. When such patterns are detected, the system can alert the operator or even trigger safety protocols. These capabilities demonstrate how safety can move from responding to hazards toward anticipating them.

Applications across industries: real-world impact

AI-based safety systems are now being used across multiple sectors, each addressing specific risks. In mining operations, autonomous haul trucks are increasingly being deployed. These vehicles navigate large mining areas without human drivers, reducing the risk of accidents caused by fatigue or human error. At the same time, AI systems monitor the surroundings to ensure that human workers are not within unsafe proximity.

In manufacturing, the robotic arms are often used alongside human workers. Modern AI-enabled systems ensure safe collaboration by continuously monitoring the distance between humans and machines. If a worker moves too close to a robotic arm, the system can slow down or stop the machine instantly.

In the construction industry, drones equipped with AI are being used to inspect sites. Instead of sending

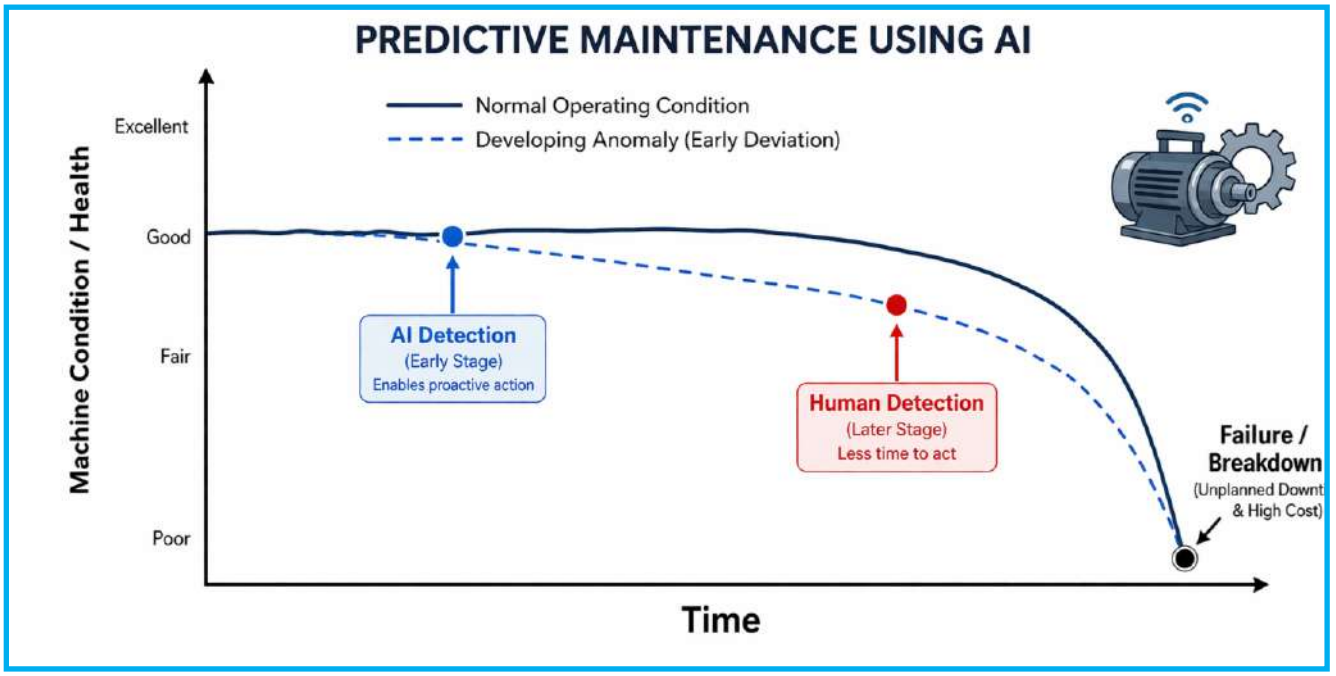

Figure 4: AI detecting early signs of equipment failure before visible breakdown occurs

workers into unstable or hazardous areas, drones can capture images and identify risks such as structural instability, unsafe material storage, or blocked emergency pathways.

Closer to the Sri Lankan context, similar approaches can be applied in industries such as apparel manufacturing, where worker density is high, or in the port operations, where heavy machinery and human activity coexist. For instance, AI systems can monitor container handling operations to detect unsafe positioning of workers relative to cranes and vehicles.

Across all these cases, the impact is not only in preventing accidents but also in improving awareness of risk patterns that were previously difficult to observe.

Challenges and ethical considerations

Despite its potential, the adoption of AI in industrial safety is not without challenges. One of the main technical issues is variability. A system trained to detect safety helmets in one environment may struggle in another with different lighting or camera angles. Ensuring consistent performance across sites requires careful design and continuous refinement.

Integration is another challenge. Many industrial facilities rely on older systems that were not designed for modern data-driven technologies. Incorporating AI into such environments requires both technical expertise and organizational commitment.

Equally important are the human aspects. Continuous monitoring can raise concerns about privacy and surveillance. If workers feel that AI systems are being used to monitor them too closely, rather than ensure their safety, resistance is likely.

To address this, organizations must be transparent about how systems are used. In many implementations, systems are designed to detect actions rather than identify individuals. Data can also be processed locally without long-term storage, reducing privacy risks. Ultimately, AI should be positioned as a tool for protection, not control.

Looking ahead: safer systems, and smarter decisions

The future of industrial safety lies in the integration of intelligence

across systems. AI is moving toward combining multiple data sources such as video, sensors, operational logs etc., into unified platforms that provide a comprehensive understanding of safety conditions. Instead of isolated alerts, systems will increasingly provide context-aware insights. For example, a combination of high temperature, increased machine vibration, and worker proximity could trigger a higher-level risk alert.

Another likely development is the emergence of real-time safety assistants. These systems could guide supervisors and workers by providing recommendations based on current conditions. For instance, suggesting temporary shutdowns, rerouting workflows, or adjusting staffing in high-risk areas. However, the success of these systems depend not only on technology but also on how they are integrated into practice. Safety remains a human-centred responsibility, and AI should enhance, not replace, human judgement.

Conclusion

Artificial Intelligence is transforming industrial safety by enabling a shift from observation to prevention. Through continuous monitoring, pattern recognition, and predictive capabilities, AI allows risks to be identified and addressed before they lead to incidents.

The real value of this transformation lies not in efficiency or automation, but in its impact on human lives. Every avoided accident, every prevented injury, and every safer workplace represents the true success of these systems.

As industries continue to evolve, AI offers an opportunity to move beyond reactive safety practices and build environments where safety is embedded into every moment of operation.

Dr Thanuja D. Ambegoda
Senior Lecturer,
Department of Computer Science and Engineering
University of Moratuwa