Politecnico di Torino

Master's Degree in Environmental and Land Engineering
Academic Year 2025/2026
Graduation Session: July 2026

SmartHSE: An Integrated Framework for Health, Safety, and Environmental Management Using Drones, IoT, and Artificial Intelligence in Construction Projects

Supervisor:

Daniele Martinelli

Candidate

: Ali Nasir

Acknowledgment

This thesis marks the completion of my Master's studies in Environmental and Land Engineering at Politecnico di Torino. It has been an intense and rewarding journey that allowed me to combine my background in engineering with my interest in health, safety, and emerging technologies. I would like to express my sincere gratitude to Professor Daniele Martinelli, whose valuable advice, patience, and continuous encouragement guided me through every stage of this research. His expertise in safety engineering has been a true inspiration for me. I am also thankful to all the professionals and colleagues who kindly shared their knowledge and practical experience regarding HSE management and digital technologies in the construction sector. Their insights greatly helped me shape the practical direction of this work. Finally, I wish to thank my family and friends for their constant love and support throughout my studies and my stay in Italy. Their encouragement and belief in me made this achievement possible.

Contents

Figures

Tables

Abstract

Construction continues to record fatality and serious injury rates that other industries reduced decades ago. The persistence of this gap is not simply a matter of inadequate regulation or insufficient worker awareness. It reflects a deeper problem: the management tools in widespread use were not designed for environments where the hazard profile changes daily, where dozens of contractors share the same physical space, and where the gap between a developing risk and a serious incident can be measured in minutes.

This thesis examines whether the integration of unmanned aerial vehicles, Internet of Things sensor networks, and artificial intelligence can address that structural limitation. A systematic literature review, conducted under the PRISMA 2020 protocol and covering 73 studies selected from an initial corpus of 1,242 records across five major databases, provided the evidential foundation. The review confirmed that each technology has demonstrated measurable HSE benefits when deployed in isolation. Drones reduce inspector exposure and increase the frequency of visual monitoring; IoT sensors provide the temporal continuity that periodic inspection cannot achieve; AI converts heterogeneous data streams into risk predictions that human analysts could not produce at comparable speed or scale. What the existing literature has not produced is a coherent account of how these contributions fit together into a complete, functioning system.

The SmartHSE framework proposed in Chapter 5 addresses that gap. Structured across five layers from physical data acquisition to feedback and continuous learning, and grounded in systems theory, ISO 31000, and socio-technical design principles, the framework specifies how drone and sensor data should flow, where processing should occur, and how outputs should translate into proportionate site-level responses. A phased implementation roadmap accompanies the architecture, along with explicit treatment of the technical, legal, financial, and organisational barriers that conceptual frameworks in this domain have often left unaddressed.

The framework's primary limitation is that it has not yet been tested under real construction conditions. The quantitative impact projections cited draw on evidence from adjacent industries rather than direct field data. Field-scale validation on construction projects in developing country contexts, where the specific challenges of irregular power supply, complex multi-contractor environments, and variable workforce literacy apply most acutely, represents the most important direction for future empirical research.

Keywords: HSE management, unmanned aerial vehicles, Internet of Things, artificial intelligence, construction safety, SmartHSE, systematic literature review

Chapter 1

Introduction

1.1 Background and Significance of the Topic

Construction underpins nearly every dimension of modern life, from the hospitals where people are born to the roads that connect communities. Yet the same industry responsible for building these structures continues to impose a disproportionately heavy burden on the workers who build them. Year after year, it claims more lives than almost any other area of work. According to the International Labour Organization, over 60,000 construction workers die annually across the globe, while millions more suffer serious, life-changing injuries (ILO, 2023). In Iran, despite massive investment in infrastructure over the past two decades, the situation remains troubling. Domestic records continue to show a clear gap between official safety rules and the harsh realities workers face on-site (Ashraf, 2024; Shafei, Radzi, Algahtany, & Rahman, 2022).

The problem, however, runs deeper than individual incidents. Traditional HSE systems were designed for environments with relatively stable hazard profiles, but construction sites are defined by constant change: the workforce, subcontractor configuration, and physical conditions all shift from day to day, creating gaps that periodic inspection and static checklists cannot cover. In such a setting, periodic human inspections and paper-based checklists simply cannot keep up. This creates dangerous blind spots. Delays in spotting risks often turn small problems into serious incidents.

Recent advances in drone technology, IoT sensor networks, and AI-based analytics have created practical opportunities to address these limitations. Aerial survey by UAVs, for instance, provides coverage of large sites from angles that ground-level inspectors cannot access. IoT sensors provide continuous data on air quality, vibration, and worker vital signs. Artificial intelligence can process all this information and identify hazards before they escalate. When these three tools work together, they create something fundamentally different from traditional safety systems. In that they shift from reacting to problems after they happen to preventing them in the first place (Daniel, Oshodi, Nwankwo, Emuze, & Chinyio, 2025; Waqar et al., 2023).

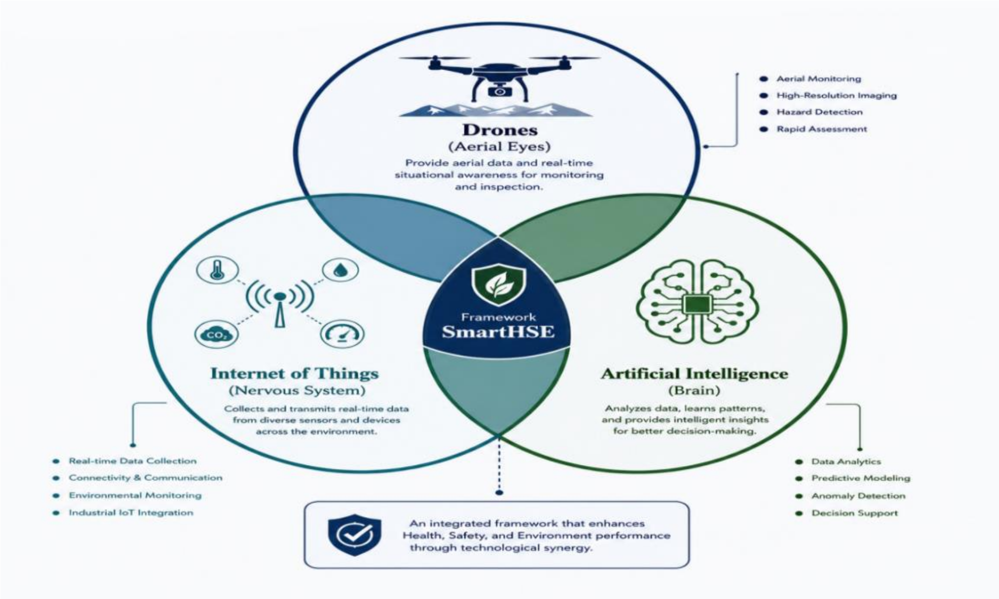

Figure 1. 1: Conceptual Model of Synergy among Drones, IoT, and AI in HSE Management

Beyond immediate safety gains, this technological convergence also supports larger goals. It aligns directly with the United Nations Sustainable Development Goals, particularly Goal 3 (Good Health and Well-being) and Goal 9 (Industry, Innovation and Infrastructure). In other words, investing in smarter HSE is not just about reducing accidents. It is about building a more responsible and sustainable construction industry (United Nations, 2023).

1.2 Problem Statement and Research Gap

The individual capabilities of drones, IoT, and AI are now well-documented. What remains surprisingly underdeveloped, however, is their genuine integration into a single, functioning HSE system. Most existing studies still treat these technologies in isolation. One paper praises drones for visual monitoring. Another highlights AI for hazard detection. Very few examine how the three can actually work together in real construction environments (Shafei et al., 2022).

This fragmentation creates a serious limitation. Drones produce excellent visual data, but without real-time context from IoT sensors, their usefulness is restricted. AI models can predict risks, but they need rich, continuous data streams to be truly effective. The result is a collection of promising tools that have not yet been combined into a coherent system (Alsamarraie, Ghazali, Hatem, & Flaih, 2022; Shah & Mishra, 2024).

The gap becomes even more pronounced in developing countries like Iran. Large infrastructure projects face unique challenges, such as dust storms, unstable power supply, remote locations, and complex multi-contractor environments. Most international research does not adequately address these realities. Moreover, critical questions about data privacy, legal responsibility for AI decisions, and integration with existing regulations remain largely unanswered (Agapiou, 2024; Waqar et al., 2023).

The gap identified across the reviewed literature is therefore not a shortage of individual tools but the absence of a unifying framework that specifies how these tools interact across the full HSE risk management cycle. Not just technical integration, but a practical system that covers the entire risk management cycle from early detection to learning and continuous improvement. This research aims to address exactly that gap by developing the SmartHSE framework.

1.3 Research Objectives

1.3.1 Main Objective

The primary purpose of this study is to design and present an integrated conceptual framework named SmartHSE that effectively combines drones, IoT sensors, and artificial intelligence to improve health, safety, and environmental management in construction projects.

1.3.2 Sub-Objectives

1. Map the current state of emerging technology application in construction HSE management through a systematic review of published literature, with particular attention to evidence of integrated multi-technology deployment.

2. Examine the capabilities, documented benefits, technical limitations, and organisational, legal, and ethical barriers associated with each technology and with their combination.

3. Specify the mechanisms through which UAV, IoT, and AI contributions interact synergistically within the full HSE risk management cycle.

4. Analyse the strategic decision-making requirements and socio-legal dimensions of implementing an integrated system in the construction industry.

5. Develop practical recommendations and a phased implementation roadmap applicable to project managers and policymakers in developing country contexts.

1.4 Research Questions

1. What is the current status of separate and integrated applications of drones, IoT, and AI in construction HSE management?

2. How can these three technologies be structured into an intelligent system to improve the complete HSE risk management cycle?

3. What are the main benefits, challenges, and ethical considerations involved in implementing such an integrated system?

4. In what ways can the proposed SmartHSE framework enhance HSE performance compared to traditional approaches?

1.5 Scope and Limitations

This research is primarily conceptual and based on a systematic literature review. No original field data were collected. Therefore, the findings and framework require practical validation through future pilot studies. The rapid evolution of technology also means some recent developments after 2025 may not be covered. The framework has been designed to remain adaptable rather than tied to specific hardware or software versions.

1.6 Significance and Innovation of the Research

This study makes several important contributions. Scientifically, it introduces one of the first comprehensive frameworks that moves beyond listing individual technology capabilities to modeling their actual synergy in a real-world HSE system. Practically, it offers construction companies, especially in developing countries, a clear pathway to move from reactive safety management to predictive, intelligent systems.

The social value is perhaps most significant. Every improvement in HSE performance means fewer injuries, fewer families left grieving, and healthier work environments. From a policy perspective, the research also provides evidence that can help shape better regulations for drone usage, AI decision-making, and data protection in construction.

Chapter 2: Health, Safety, and Environmental Management in Construction Projects

2.1 The Concept and Importance of HSE

The integration of health, safety, and environmental management into a single system emerged from practical experience rather than theoretical design. On active construction sites, separating these three domains consistently generated overlapping procedures and gaps in coverage: a chemical spill, for instance, simultaneously poses an immediate injury risk, a fire hazard, and a long-term pollution problem that no single-domain management system can address in isolation.

The foundation of this integrated approach rests on two key international standards. ISO 45001:2018 focuses on occupational health and safety. It pushes organizations to think proactively about risks, involve workers in decision-making, and maintain strong leadership commitment. ISO 14001:2015, on the other hand, deals with environmental management and requires companies to consider the full life cycle of their projects from material sourcing to final demolition. When these two standards are implemented together, they create a powerful, unified system that is both more effective and more efficient (International Organization for Standardization, 2018, 2023).

Construction companies that adopt this integrated model often report lower overall costs and better performance. What used to require two separate management systems now runs through shared processes such as audits, training, and corrective actions. For large projects and international clients, having certified HSE systems based on these standards has become almost mandatory.

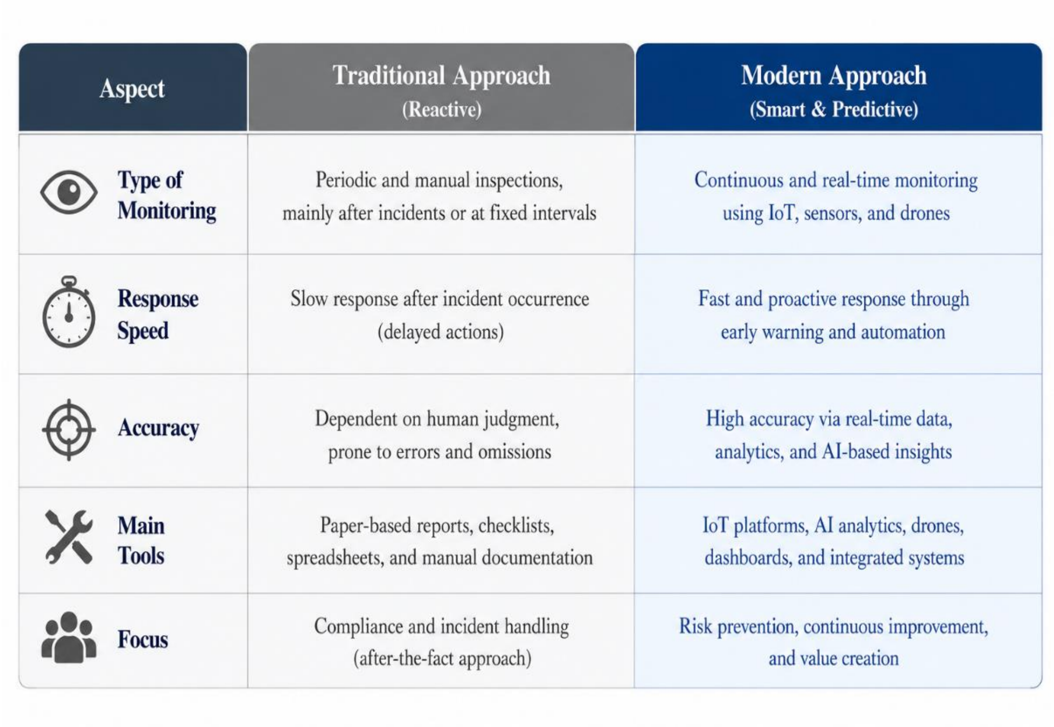

Aspect	Traditional Approach (Reactive)	Modern Approach (Smart & Predictive)
Type of Monitoring	Periodic and manual inspections, mainly after incidents or at fixed intervals	Continuous and real-time monitoring using IoT, sensors, and drones
Response Speed	Slow response after incident occurrence (delayed actions)	Fast and proactive response through early warning and automation
Accuracy	Dependent on human judgment, prone to errors and omissions	High accuracy via real-time data, analytics, and AI-based insights
Main Tools	Paper-based reports, checklists, spreadsheets, and manual documentation	IoT platforms, AI analytics, drones, dashboards, and integrated systems
Focus	Compliance and incident handling (after-the-fact approach)	Risk prevention, continuous improvement, and value creation

Figure2 .1: Comparison of Traditional and Modern HSE Management Approaches

Beyond its moral imperative, strong HSE management has a well-documented financial rationale. Studies consistently find that proactive safety investment yields returns of four to six dollars for every dollar spent, through reduced incident costs and lower insurance premiums (Waqar et al., 2023). The environmental dimension adds a third layer of accountability: decisions made during construction can contaminate soils and water sources for decades, long after the project team has moved on.

Table 2. 1: Comparison of Traditional and Modern HSE Management Approaches

Aspect	Traditional Approach (Reactive)	Modern Approach (Predictive & SmartHSE)
Focus	Analysis of past incidents	Risk prediction using AI and real-time data
Main Tools	Human inspection and PPE	Drones, IoT, AI, and intelligent analytics
Performance Indicators	Lagging indicators	Leading + Predictive indicators
Risk Response	After incident occurrence	Before incident occurrence (Prevention)

2.2 Characteristics and Inherent Risks of the Construction Industry

What makes construction sites structurally different from other industrial environments is the combination of temporariness and change. Teams assemble at the start of a project and disperse at its end, with limited institutional memory carried forward. The hazard profile on any given day reflects the current construction phase, and that profile changes as excavations deepen, structures rise, and new subcontractors bring different equipment and working practices onto the same site. Coordinating safety across this mix of organisations, each with its own procedures, priorities, and workforce, is qualitatively different from managing a fixed-facility workforce (Harvey, 2019; Shafei et al., 2022).

These characteristics produce three major types of risk. Safety hazards cause sudden and often tragic accidents. Falls from height remain the leading cause of death in construction across the world. Being struck by falling objects, caught between equipment, and electrocution complete what professionals call the “Fatal Four.” Together, they account for the majority of fatalities. Many of these incidents are predictable if someone is watching carefully, but constant human supervision is almost impossible on large, busy sites.

Health risks work more slowly but cause serious, long-term damage. Silica dust from cutting concrete leads to serious lung diseases. Noise from heavy machinery destroys hearing over time. Poor ergonomics and repetitive movements create chronic pain and disability. These problems do not appear immediately, so they often receive less attention than sudden accidents. Yet their human and economic cost is enormous(Shah & Mishra, 2024); ILO, 2023).

Environmental risks generated on construction sites extend well beyond the project perimeter. Dust from concrete cutting, chemical runoff from surface treatments, and noise from pile driving create documented health and ecological impacts on surrounding communities that persist long after construction ends. It consumes huge quantities of water and energy. Dust, chemical runoff, and carbon emissions pollute the air, soil, and rivers. Some of these impacts only become obvious years later. Modern HSE must address all three risk categories together, not in isolation.

Table 2. 2: Classification of Risks Inherent to the Construction Industry

Risk Category	Main Examples	Approx. Contribution to Accidents (%)	Proposed SmartHSE Technology
Safety	Falls, struck-by, electrocution	70–80%	Drones + AI for real-time detection
Health	Silica, noise, ergonomics	15–20%	IoT wearables + AI data analysis

Environmental	Waste, carbon emissions, pollution	Variable	IoT sensors + AI prediction

2.3 Traditional HSE Management Methods and Their Weaknesses

Tools such as the Hierarchy of Controls and Job Safety Analysis were developed for relatively stable industrial environments. Construction sites, however, present a fundamentally different operational context: the hazard profile changes daily, contractor configurations shift mid-project, and the workforce is rarely constant enough for standardised procedures to take hold.

One major problem is the reactive nature of the whole system. Traditional HSE usually measures success through lagging indicators, such as accident rates, injury statistics, and investigation reports. The system only reacts after something has already gone wrong. It offers little warning about dangers that are quietly developing right now. This backward-looking approach leaves construction sites exposed in an industry where conditions change by the hour (Pilskog Orvik, Albrechtsen, Angvik, Sæther, & Holen, 2024).

Another serious limitation is the heavy dependence on human inspectors. Even the most dedicated supervisor cannot be everywhere at once. Large sites stretch across vast areas. Walking from one point to another takes time. During those gaps, hazardous situations can develop unnoticed. When project schedules become tight, and pressure increases, the first thing that suffers is usually the quality and frequency of inspections. This is not a failure of effort. It is a structural weakness built into the traditional model (Waqar et al., 2023).

Human judgment itself adds another layer of difficulty. Experienced inspectors still miss hazards, especially when they are tired, rushed, or too familiar with the site. Cognitive biases are natural. Expectations shape what people see. These are not personal failings. They are universal limitations of human attention. Any safety system that relies almost entirely on human observation will always have blind spots.

Table 2. 3:Weaknesses of Traditional HSE Methods and SmartHSE Solutions

Traditional Weakness	**Explanation**	**Proposed SmartHSE Solution**	**Traditional Weakness**
Reactive Nature	Analysis after the incident	Risk prediction using AI and real-time data	Reactive Nature
Human Dependency	Human inspection errors	Drones + IoT for automated monitoring	Human Dependency
Poor Data Management	Scattered and unstructured data	Edge-Cloud layer + intelligent analytics	Poor Data Management
Lack of Coordination in Multi-Contractor Environments	Information gaps between contractors	Integrated decision-making system	Lack of Coordination in Multi-Contractor Environments

These limitations reinforce one another. Reactive measurement systems provide no early warning; inspectors without warning arrive at hazards after they develop; and cognitive bias reduces the reliability of what inspectors do find. The combined effect is a safety management model whose weakest moments coincide with the highest-risk conditions. The combination of drones, IoT sensors, and artificial intelligence does not just add new tools.

Chapter 3: Emerging Technologies in the Construction Industry (Drones, IoT, and Artificial Intelligence)

3.1 Drones: Capabilities, Types, and Recent Developments

Bringing a human inspector close to a site's specific hazards has always meant exposing that person to the high risk they are trying to assess. A scaffolding connection at 40 metres cannot be inspected from the ground; a tunnel face after a blast cannot be inspected without someone entering a potentially unstable space. For decades, this was simply accepted as an unavoidable feature of construction safety work, with exposure managed rather than eliminated.

Drones change this. Where traditional inspection requires physical proximity, a drone can hover over the site and transmit images without anyone leaving the ground. This is not a minor improvement on the efficiency of existing methods. It is a different operating model, one that removes the inspector from the danger zone in situations where access and risk have historically been inseparable.

The technology did not originate in construction. Military reconnaissance and topographic mapping were among its early applications, and its hardware reflected these priorities. The shift toward commercial use of drones in construction was less a result of a single breakthrough and more of a gradual convergence: sensors became smaller and cheaper, flight control software became more reliable, and onboard processors reached a point where real-time image analysis during flight became possible. By the mid-2010s, this combination had made the deployment of drones on active project sites financially realistic, and the adoption curve has rapidly climbed from there. In many large construction projects today, drones are part of the standard site toolkit, not a specialized add-on that requires separate justification (Ashraf, 2024; Olatunde, Anugwo, Awodele, Ramabodu, & Adegoke, 2026).

The relationship between drone deployment and inspector safety is more consequential than it may initially appear. A human inspector assessing the structural integrity of elevated scaffolding is simultaneously exposed to the hazard under investigation: climbing to the level of concern, navigating an unstable structure, and conducting the assessment under conditions of direct personal risk. UAVs invert this exposure relationship, allowing the hazardous location to be observed remotely, with image quality that often exceeds what an inspector could achieve while managing their own stability at height. This inversion also removes the time constraints of human-paced inspection. A multi-rotor drone can complete a visual survey of a scaffold structure that would require several hours on foot in under fifteen minutes, with the complete visual record stored for later review and legal documentation (Akinsemoyin, Awolusi, Chakraborty, Al-Bayati, & Akanmu, 2023). The practical consequence is not simply speed, but the feasibility of inspection frequencies that were previously resource-prohibitive: daily or intra-shift assessments of high-risk areas become operationally realistic in a way they are not when each inspection requires committing a trained inspector to a hazardous environment for an extended period.

The sensor payload a drone carries determines both what it can detect and what analytical role it plays in an integrated HSE system. Standard high-resolution cameras enable surface defect identification, including cracks in concrete members, loose fasteners, and progressive corrosion, at distances that would otherwise require scaffolded access. Thermal and infrared sensors extend detection capability to conditions where visible-light imaging provides inadequate information: electrical faults manifest as thermal anomalies before they produce observable damage; gas leaks can be identified through temperature differentials in conditions where visual indicators are absent; personnel can be located in smoke-obscured environments where visible-light identification fails. Gas-measurement sensors, increasingly integrated directly into UAV payloads, provide an environmental monitoring dimension that ground-level fixed sensors cannot replicate, particularly when hazardous concentrations form at altitude or vary spatially in response to wind patterns. LiDAR systems enable point-cloud generation and three-dimensional structural modelling with millimetre-level precision, supporting applications ranging from deformation monitoring to volume measurement of excavations and material stockpiles (Mustapha, Akomah, Kwaku, & Tieru, 2024). Different sensor payloads serve different analytical purposes within the same HSE system. Thermal cameras detect electrical faults before visible damage appears; gas sensors identify atmospheric hazards at altitude that fixed ground-level sensors cannot reach; LiDAR produces structural deformation data at millimetre precision. The practical implication is that the drone subsystem is not a single instrument but a configurable data acquisition platform whose specific contribution depends on the sensor configuration deployed.

3.1.1 UAV Platform Selection: A Trade-off Framework

Selecting an appropriate UAV configuration for a given construction HSE application requires understanding a design constraint that no single platform fully resolves: the tension between hovering stability and range efficiency. Multi-rotor platforms prioritise the former; fixed-wing platforms prioritise the latter; hybrid VTOL designs attempt to balance both at increased mechanical complexity and cost. Understanding where a specific project sits along this axis is a prerequisite to hardware specification rather than an afterthought.

Drones used in the construction industry are mainly divided into four main categories based on their physical design and flight mechanism.

Multi-rotor drones dominate current construction site deployments for reasons directly related to the spatial constraints of built environments. Their capacity for vertical take-off and landing, sustained hovering, and close-quarters manoeuvring makes them well-suited to the inspection tasks that dominate urban building projects. Within the multi-rotor category, the number of rotors carries direct operational implications for HSE applications. Quadcopters offer sufficient payload capacity and stability for most monitoring tasks at acquisition costs that make deployment at project scale financially viable for mid-sized contractors. Hexacopter configurations provide improved stability and higher payload thresholds, enabling the carriage of heavier sensor systems such as LiDAR units, where platform vibration directly compromises point-cloud accuracy. Octocopters represent the highest specification within the category: their redundant propulsion architecture means the platform can sustain controlled flight after losing a single motor, a fault-tolerance property that is not trivial when operating over active construction zones where an uncontrolled landing carries significant injury risk to ground personnel (Alsamarraie et al., 2022). The constraint that applies uniformly across the multi-rotor category is energy consumption: the continuous power draw required to maintain multiple counter-rotating rotors restricts flight duration to approximately 20 to 40 minutes per battery cycle, a limitation with direct implications for the design of monitoring schedules on large sites.

Fixed-wing platforms occupy a different position in the operational matrix. The aerodynamic efficiency of generating lift through forward motion rather than continuous rotor thrust translates into flight endurance measured in hours, and operational ranges extending to tens of kilometres on a single charge. For linear infrastructure projects such as road construction, pipeline routing, and railway development, where the relevant site extends beyond the realistic sortie range of a multi-rotor, fixed-wing UAVs enable comprehensive photogrammetric mapping that would otherwise require multiple sorties and significantly greater operational coordination. The fundamental limitation is the absence of hovering capability: fixed-wing platforms cannot station-keep above a specific point of interest, cannot navigate the spatial constraints of a congested urban building site, and require either a prepared runway or a dedicated launch system for deployment.

Hybrid VTOL platforms combine vertical lift capability with fixed-wing cruise efficiency through either a tilt-rotor mechanism or separate propulsion systems for vertical and horizontal flight phases. The practical result is a platform that can launch from a confined site without runway infrastructure, transition to efficient forward flight for extended-range coverage, and return to hovering mode for targeted point inspection. This capability profile matches the requirements of large-footprint projects where both area mapping and specific site inspection are operationally necessary within a single sortie. The trade-off is mechanical complexity: the additional systems required to manage the transition between flight modes increase both acquisition cost and maintenance demand relative to either pure-configuration alternative.

The fourth category comprises purpose-designed specialised platforms, which address operating conditions that fall outside the performance envelope of standard configurations. Weather-hardened drones, with sealed enclosures and integrated thermal regulation, maintain operational capability in precipitation, particulate-laden air, and sub-zero temperatures that render commercial platforms inoperable. Indoor mapping drones, equipped with SLAM navigation systems that construct environmental maps from sensor data rather than relying on GNSS signals, enable safety inspection of enclosed structures where satellite positioning is unavailable, including tunnels, underground utility corridors, and confined concrete tanks. Swarm configurations introduce a qualitatively distinct operational capability: multiple coordinated drones executing a distributed mission under shared algorithmic direction can simultaneously survey multiple work fronts, parallelise hazard searches across large areas, or maintain continuous coverage of a monitored zone while individual units cycle through battery replacement (Raj & Kos, 2026).

Choosing a UAV configuration is therefore inseparable from defining the monitoring requirements of the site. A project dominated by confined vertical structures calls for a different platform than a linear road project spanning several kilometres, and selecting hardware before specifying these requirements is likely to produce a mismatch. It requires mapping the dominant hazard profile of the project against the operational characteristics of available platform types, and this mapping exercise should precede hardware specification.

3.1.2 Technological Developments Shaping Construction Drone Capability

The integration of computer vision algorithms into drone data processing pipelines represents one of the more consequential developments for HSE applications over the period from 2018 to 2025. Where early construction drone deployments produced image records requiring subsequent human review, current systems enable real-time automated analysis of visual feeds during flight. The algorithmic families most relevant to HSE applications operate through different mechanisms suited to different detection tasks. YOLO architectures, currently at version 8 in active deployment, process entire image frames in a single network pass, achieving detection latency measured in milliseconds and making them suitable for real-time worker behaviour monitoring. Faster R-CNN architectures sacrifice some processing speed for increased detection granularity, making them preferable for structural defect identification where the relevant feature may occupy only a small fraction of the image field. Vision Transformer models demonstrate superior performance in detecting complex spatial relationships and patterns that extend across large regions of the image, a property particularly relevant for identifying developing structural deformations that manifest across multiple connected elements rather than at a single point (Akinsemoyin et al., 2023; Jha et al., 2025). Reported detection accuracy across these algorithm families, after appropriate training on construction-specific datasets, falls in the range of 93 to 97 percent, which in field conditions compares favourably with human inspection performance.

The communication infrastructure available to drone systems has undergone a parallel transformation. The data volumes generated by high-resolution cameras, LiDAR units, and multi-spectral sensors operating simultaneously exceed the bandwidth capacity of 4G networks under realistic deployment conditions. Fifth-generation cellular networks, with theoretical downlink capacity of 20 Gbps and latency specifications of one millisecond, remove this bottleneck for most current sensor configurations and enable the dense connectivity required by integrated drone-IoT systems where multiple data streams operate concurrently.

Flight endurance limitations, historically among the most significant operational constraints on drone monitoring programmes, have improved substantially through advances in energy storage technology. Recent advances in lithium-polymer and lithium-ion battery technology, combined with improvements in motor and controller efficiency, have extended the operational duration of current-generation industrial drones to more than 60 minutes. Emerging technologies such as solid-state batteries and hydrogen fuel cells also promise multi-hour flights. In addition to increasing battery capacity, the development of automated charging stations and drone-in-a-box systems, which automatically receive the drone, replace its battery, and prepare it for the next mission, has made continuous 24-hour surveillance possible (Cheng, Gheisari, & Jeelani, 2023).

The integration of drone-collected data with Digital Twin technology represents a further development that extends the HSE utility of aerial survey beyond its direct monitoring function. A digital twin, in its operational form, is a continuously updated virtual model of the physical project that reflects the actual constructed state at the current moment. In the traditional approach, static 3D models were unable to display real-time site changes, but by integrating drone data, the digital twin becomes a living model that reflects any change from physical progress to ground subsidence or scaffolding tilting in a short time. This capability enables the simulation of hazard scenarios such as fire spread and emergency evacuation, automatic detection of hazardous changes such as millimetre displacement of a retaining wall, and the training of new workers before entering the site (Olatunde et al., 2026).

Empirical studies have shown that the use of drones can reduce the time required for a comprehensive safety inspection of a construction site by up to 85 percent(Albeaino, Brophy, Jeelani, Gheisari, & Issa, 2023). For example, a manual inspection of a tall scaffold that takes several hours can be completed in less than 15 minutes with a multi-rotor drone. This significant time reduction enables the possibility of increasing the frequency of inspections to daily or multiple times per day, freeing up specialised human resources for analytical tasks, and most importantly, eliminating the inspector's exposure to the hazards of heights, confined spaces, and contaminated areas.

This efficiency gain was recorded under specific site conditions and should not be assumed to transfer uniformly across all project types and inspection scenarios.

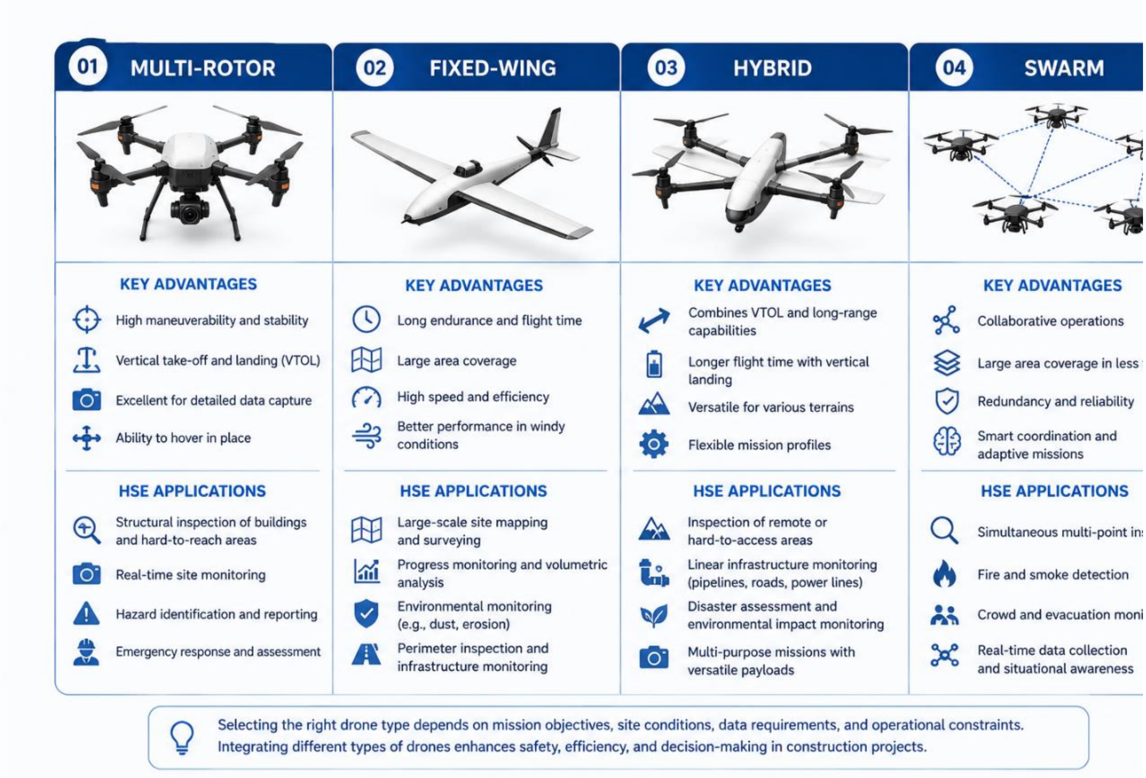

Figure 3 .1:Types of Drones and Their HSE Applications in Construction

3.1.3 Challenges

Despite the numerous and proven benefits of drones in HSE management, their widespread and systematic use in construction projects faces significant challenges. Regulatory airspace restrictions constitute the most immediate constraint in urban project environments. In the majority of jurisdictions, UAV operations above residential areas, within proximity zones of airports or heliports, or above defined altitude thresholds require specific operational authorisations that may involve significant lead time, operational restrictions, and ongoing compliance costs(Agapiou, 2024) .The regulatory landscape in this domain is also dynamic. As national aviation authorities have developed more sophisticated frameworks for commercial drone operations, the permissions required for specific deployment configurations have changed, sometimes with limited notice.

Privacy concerns represent a parallel structural challenge with dimensions that are simultaneously legal and organisational. Continuous or frequent aerial surveillance of an active construction site inevitably captures imagery of individual workers, including faces, movements, physiological states, and behaviours, in a manner that raises questions under data protection frameworks such as GDPR in European jurisdictions and equivalent instruments elsewhere. Beyond the legal dimension, worker acceptance of surveillance technology is not guaranteed, and historical experience with the introduction of monitoring systems in industrial environments suggests that resistance is a realistic risk that warrants proactive management. The Privacy by Design framework, which integrates privacy protection as a design requirement rather than a post-implementation compliance measure, offers one approach to managing these concerns.

Flight endurance, despite the improvements described in the preceding section, remains a binding constraint in certain deployment scenarios. Adverse weather compounds this constraint: wind speeds above 25 to 30 kilometres per hour, precipitation, dense construction dust, and temperatures outside the platform's rated operating range all degrade flight stability, sensor performance, and battery efficiency simultaneously (Cheng et al., 2023). In construction environments where adverse conditions are not exceptional events but regular operational contexts, the contribution of drone monitoring is correspondingly reduced during the periods when it is most needed.

Operator training requirements are frequently underestimated in implementation planning. Professional operation of a UAV in a complex construction environment requires both formal certification and accumulated practical experience that cannot be acquired quickly. In developing country contexts, where the pool of formally certified drone operators is typically small and geographically concentrated, this represents a workforce development constraint that can delay or limit implementation independently of hardware availability or regulatory clearance (Agapiou, 2024).

Organisational resistance, budget allocation limitations, and the absence of standardised data integration protocols between drone platforms and existing project management systems constitute a further category of challenges that are neither technical nor regulatory but institutional. Project managers accustomed to traditional inspection methodologies may be sceptical of automated detection systems, particularly in the early stages of deployment when algorithm performance on site-specific conditions is still being optimised and false alarm rates may exceed initial expectations (Agapiou, 2024; Cheng et al., 2023). Overcoming these challenges requires cooperation between regulatory bodies, technology companies, and project managers.

Within the SmartHSE architecture, drones serve as the primary source of visual and spatial data, capturing site conditions from angles and elevations that are physically inaccessible or hazardous for human inspectors. However, drones provide vision from above and from a distance, while lacking the ability to measure environmental parameters at ground level continuously. The full potential of drone deployment is only realised when visual data from UAVs, continuous sensor readings from the IoT network, and the pattern-recognition capacity of AI are processed within a single, closed-loop architecture, with each subsystem compensating for the inherent limitations of the other two.

3.2 Internet of Things (IoT): Architecture, Sensors, and Applications

The data collection limitation of drone systems has a temporal dimension that is worth making explicit before discussing the IoT layer. A drone-based inspection programme, however frequent, produces discrete samples of site conditions at specific moments: images captured during a morning sortie reflect conditions at that time, not the state of a gas concentration that began building three minutes after the drone landed, or the vibration pattern of a scaffold component whose connection has been progressively loosening over six hours. IoT sensor networks address this temporal gap through an architectural principle that airborne surveillance cannot replicate: continuous measurement at fixed locations, with data collection frequencies that can be calibrated from several times per second to several times per day depending on the parameter being monitored(Jha et al., 2025).

In the SmartHSE architecture, the Internet of Things plays the role of the main infrastructure; that is, without the IoT sensor network, there would be no raw and continuous data required for any intelligent analysis and real-time decision-making. The fundamental difference between IoT and drones is summarised in two characteristics: the fine-grained accuracy and temporal continuity of the data. IoT sensors measure physical parameters with very high accuracy and continuously, whereas drones collect data only at specific time intervals (Khan, Alrasheed, Waqar, Almujibah, & Benjeddou, 2024). The IoT network in smart construction sites collects three main categories of data: environmental data, including air quality, flammable gases, temperature, humidity, and dust; structural data, including strain, member deformation, scaffolding vibration, and foundation settlement; and human data, including worker presence, body temperature, and heart rate via wearable sensors (Musarat, Irfan, Alaloul, Maqsoom, & Ghufran, 2023).

3.2.1 Network Architecture

Effective implementation of IoT in construction projects requires the adoption of a network architecture that is adaptable to the difficult and dynamic conditions of construction sites. Unlike office or industrial environments with stable infrastructure, construction sites face challenges such as power outages, network fluctuations, interference from metal and concrete structures, dust and humidity, and the need to respond to safety alerts in real time. For this reason, simple, single-layer architectures are ineffective for such environments, and advanced layered or hybrid models are necessary (Kanellopoulos, Sharma, Panagiotakopoulos, & Kameas, 2023).

The layered network model used in industrial IoT separates perception, transmission, processing, storage, and application functions into distinct tiers. This separation has a practical rationale: each layer can be diagnosed and upgraded independently, which matters in construction environments where sensor hardware is replaced frequently as the site configuration changes.

For critical HSE applications where every millisecond of delay can lead to an incident, the Hybrid Edge-Cloud Architecture is the superior option (Kanellopoulos et al., 2023). In this architecture, time-sensitive data, such as a sudden gas alarm, abnormal scaffold vibration, or a critical sensor disconnection, is processed at the edge of the network, and decisions are made in real time without the need to send data to a remote cloud centre. Non-critical and large data, such as daily reports or drone images, are transferred to the cloud for in-depth analysis and training of AI models. This hybrid approach both minimises latency to less than 10 milliseconds and increases system reliability in unstable network conditions.

3.2.2 Sensor Categories

The IoT network in the SmartHSE system consists of a wide range of sensors, which are classified into five main categories based on the type of data required for safety, health, and environmental management.

The first category, Environmental Sensors, includes sensors that measure temperature, humidity, concentrations of flammable and toxic gases such as methane, carbon monoxide, and hydrogen sulphide, particulate matter, air pressure, and wind speed. These sensors provide the contextual information within which safety decisions are made and serve as the primary early warning system for atmospheric hazards invisible to unassisted human senses.

The second category, Structural Sensors, is used to monitor the health of structures and equipment and includes Strain Gauges, Accelerometers, Inclinometers, and Vibration Sensors. These devices detect tensile and compressive stress in structural members, changes in characteristic vibration frequency, and progressive angular displacement in retaining walls and slopes, providing early indicators of impending mechanical failure that precede visible structural change.

The third category, Positioning Sensors, includes Ultra-Wideband (UWB) technology to determine the location of workers and equipment with centimetre accuracy in closed environments without GPS, and RFID to track key materials and tools at the entrance and exit of the site (Hossain, Hosen, Sunny, & Tarapder, 2025). The combination of position data with environmental and structural data enables spatial risk assessment that is significantly more actionable than either data type in isolation.

The fourth category, Wearable Sensors, is installed directly on the worker's body and monitors vital signs and physiological status. These sensors include heart rate monitors, skin thermometers, fatigue detection sensors, and helmet-mounted accelerometers for impact and fall detection. Beyond their obvious clinical significance, these sensors provide information relevant to heat stress assessment, physical exertion management, and the detection of fatigue-induced cognitive impairment, a hazard type that is entirely invisible to external observation.

The fifth category is Specialised HSE Sensors, designed for specific applications such as Radiation Detectors for nuclear or medical projects, ultrasonic leak detectors for pressurised pipelines, and Flame Detectors for areas with explosion risks. Combining data from these five categories simultaneously provides a comprehensive, multidimensional picture of the site's safety, health, and environmental status that cannot be achieved by any other traditional method.

3.2.3 Communication Protocols

Choosing the right communication protocol for an IoT network on construction sites has a direct impact on the efficiency, energy consumption, and reliability of the system. The MQTT protocol is the most common option for transmitting sensor data to a central server due to its lightweight and efficient design. Its publish-subscribe architecture reduces the data overhead of each transmission to a minimum, enabling sensors with limited processing capability and power budgets to communicate reliably over unreliable networks.

The LoRaWAN protocol is suitable for applications that require a range of several kilometres and extremely low power consumption, such as environmental sensors scattered in linear and wide sites. Although it has limited bandwidth and cannot be used to transmit large data such as video, it enables environmental and structural monitoring at peripheral site areas where installing wired power and network infrastructure is not economically justifiable.

The 5G protocol has brought about a fundamental revolution in industrial IoT communications. The very high bandwidth of up to 20 Gbps allows for the simultaneous transmission of 4K videos from multiple drones, LiDAR data, and the output of hundreds of sensors. Ultra-low latency of down to 1 millisecond is critical for safety-critical applications such as remotely stopping heavy equipment or sending warning commands to workers via wearable devices. Massive IoT connectivity of up to one million devices per square kilometre also enables dense sensor networking.

The NB-IoT protocol, which runs on top of existing cellular networks, is a complementary option for situations where LoRaWAN coverage is insufficient or 5G is not available. With deep coverage in urban construction sites with tall structures and significant penetration into buildings and basements, this protocol is suitable for environments where satellite positioning and alternative protocols are inadequate. In practice, advanced HSE IoT systems use a combination of several protocols simultaneously, with MQTT for communicating sensors with the local gateway, LoRaWAN or NB-IoT for large sites with low sensor density, and 5G for drones and real-time critical applications.

3.2.4 IoT Applications in HSE

The IoT network in construction projects covers four main application categories in HSE management. First, real-time environmental monitoring includes continuous measurement of concentrations of toxic and flammable gases, particulate matter, temperature, humidity, and noise. This data allows for early detection of hazardous conditions and immediate warning to workers. Second, Worker and Equipment Tracking using UWB and RFID sensors pinpoints the exact location of personnel and instantly detects and warns people entering prohibited or high-risk areas (Hossain et al., 2025).

The third key application is Structural Health Monitoring. Strain sensors, accelerometers, and inclinometers installed on scaffolding, cranes, excavations, and concrete formwork record any deformation, abnormal vibration, or settlement in real time, providing warnings of impending collapse before an incident occurs. The fourth application is Dynamic Risk Management; unlike traditional risk assessments that are performed at the beginning of a project and remain fixed until the end, IoT enables continuous updating of risk levels based on real-time data from the environment, structure, and worker behaviour (Waqar et al., 2023).

The effectiveness of these applications in reducing accidents has been quantified by numerous empirical studies. (Hossain et al., 2025) found that deploying comprehensive IoT systems in construction projects can reduce accident rates by 25 to 68 percent, depending on the type of hazards, system maturity, and sensor coverage. (Khan et al., 2024) reports that projects deploying comprehensive IoT systems in construction achieved accident rate reductions ranging from 25 to 40 percent, with the upper end of this range associated with AI-integrated real-time warning systems.

Table 3. 1:IoT Applications in HSE

IoT Layer	Sensors / Devices	HSE Applications
Perception	Wearable sensors, Gas sensors	Fatigue detection, Toxic gas monitoring
Network	LoRaWAN, 5G	Data transmission across large construction sites
Processing	Edge Computing	Local data analysis and latency reduction

3.2.5 Challenges

Cybersecurity is the challenge that, if inadequately addressed, transforms IoT infrastructure from a safety asset into a safety liability. Every sensor connected to the network is a potential entry point for cyberattacks. Hacking a gas sensor or accelerometer can lead to false alarms, reducing user trust, or, worse, failure to issue an alarm in real dangerous situations. The threat model for industrial IoT in construction differs from office IT environments in important ways: physical access to sensors and gateways is often possible for workers, contractors, and even site visitors, and many IoT devices lack the processing capability required for strong encryption or complex authentication protocols.

The challenge of interoperability between different devices and protocols from different manufacturers remains unsolved in the construction IoT market. The absence of unified standards makes the integration of sensors, gateways, and analysis software difficult and costly, and creates long-term vendor dependency. Power consumption and power management of sensors, especially in remote sites without stable access to electricity, present another critical challenge. Designing ultra-low-power sensors, using low-power long-range protocols, and implementing sleep modes are common solutions, though balancing energy savings with frequent data updates remains a complex optimisation problem.

The high initial cost for purchasing, installing, commissioning, and training users represents a significant barrier, particularly for contractors in developing countries. A comprehensive IoT network for a medium to large construction project may include hundreds of sensors, multiple gateways, communication infrastructure, edge processing servers, and annual software and maintenance costs. However, economic studies show that the initial costs are recovered over a two- to four-year horizon through reduced incidents, reduced project downtime, and reduced insurance premiums(Chung, Tariq, Mohandes, & Zayed, 2023). Overcoming these challenges requires collaborative effort between technology manufacturers, contractors, insurers, and regulatory bodies.

3.2.6 Trends

Emerging trends in IoT for HSE applications are moving beyond simple data collection and transmission and evolving towards intelligent, self-adaptive systems. The most important of these trends is Cognitive IoT, in which an artificial intelligence layer, especially machine learning and deep learning, is injected directly into the network of sensors and edge nodes. In this approach, sensors are no longer just data collectors, but are also capable of recognising patterns, predicting events, and making local decisions in real time. For example, an accelerometer sensor equipped with Cognitive IoT can distinguish the vibration pattern at the beginning of a fall from normal patterns, such as the passage of heavy equipment, and issue a critical alert without sending data to a central server.

A key trend is the deep integration of IoT with Digital Twin technology. In this integration, a continuous stream of sensor data, including temperature, vibration, strain, and position, is injected directly into the virtual model of the project, turning it into a living twin. Any change in the physical state of the site is reflected in the digital model within seconds and can be viewed and analysed. This capability enables the simulation of what-if scenarios: for example, the impact of strong winds on the stability of a crane can be simulated on the digital twin before it occurs.

The overarching trend is the convergence and full integration of IoT, AI, and digital twin into a single integrated platform. In this advanced paradigm, IoT provides raw, continuous data from the physical world; artificial intelligence analyses this data, uncovers hidden patterns of risk, and provides short- and long-term predictions; and the digital twin displays these predictions in a visual and interactive model to project managers. The result is a system that not only reports what is happening now, but also projects what will happen in the near future and recommends the optimal solution. Together, these trends are shaping the future of smart HSE in construction projects.

3.3 Artificial Intelligence and Machine Learning: HSE-Related Techniques

In the SmartHSE architecture, drones serve as the primary aerial data source and the IoT sensor network functions as the continuous measurement infrastructure. Artificial intelligence constitutes the core processing layer of the system, the component that determines whether raw sensor readings and drone imagery translate into timely, accurate risk decisions or remain unused data. AI receives, processes, and integrates raw and heterogeneous data and transforms it into actionable knowledge (Shah & Mishra, 2024). Without this layer, data cannot be converted into actionable knowledge.

AI's position in the SmartHSE architecture is one of conditional value: its outputs are only as useful as the data it receives. Drone imagery without AI analysis remains a collection of images requiring human review; IoT readings without AI processing accumulate as raw time-series that no site manager has the capacity to interpret in real time. The architecture is designed around this dependency; data collection is specified to meet

AI processing requirements, not the other way around. This asymmetry means that the value of AI in an HSE context is entirely contingent on the data infrastructure that precedes it.

The process of transforming data into actionable knowledge is performed in four consecutive steps: integrating visual data from the drone with numerical output from IoT sensors; pattern recognition and discovery of hidden correlations with machine learning methods; predicting future events with time series models; and generating a specific recommendation, such as immediate evacuation of a designated area. Example: combining abnormal scaffold vibration, drone images, and AI prediction with an 85 percent probability of collapse produces the recommendation of immediate evacuation of the northern scaffold(Waqar et al., 2023). These systems reduce the response time to hazards from hours to minutes and increase the accident prevention rate by more than 70 percent.

Computer Vision is a key AI technique in SmartHSE that processes image data from drones and cameras. Algorithms such as CNN, YOLO versions 5 through 8, and Vision Transformers are able to automatically detect non-use of personal protective equipment, unsafe behaviours, and structural cracks with 93 to 97 percent accuracy, which in many cases exceeds the accuracy of a human inspector. Predictive Analytics with algorithms such as XGBoost and LSTM is used to model nonlinear relationships between environmental variables and predict trends such as soil subsidence or increased gas concentrations. Anomaly Detection identifies patterns of deviations from the norm without the need for labelled data, a method particularly valuable for discovering unknown and emerging hazards. Natural Language Processing extracts root causes of incidents from HSE text reports and categorises hazards. Reinforcement Learning also enables the intelligent agent to learn the optimal control strategies through interaction with the site environment. The combination of these five techniques makes the SmartHSE system a comprehensive, self-learning, and adaptable software platform.

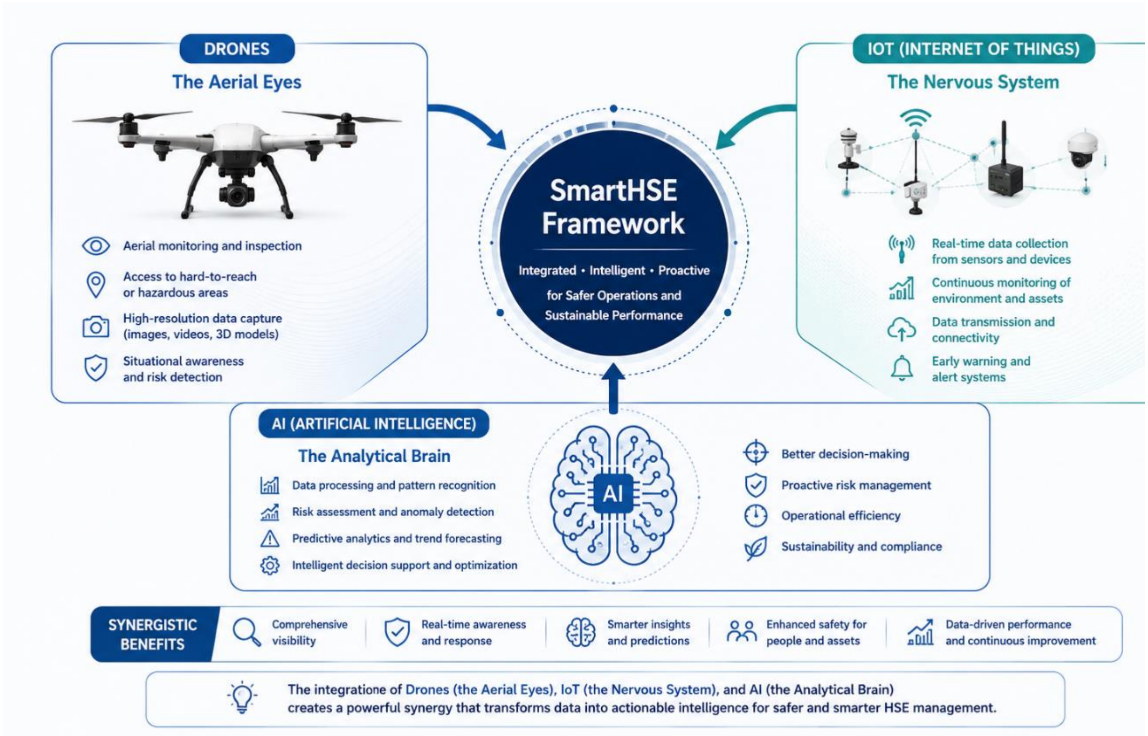

Figure 3. 2:Synergy of Drones, IoT, and AI Technologies

3.4 AI Techniques in SmartHSE

3.4.1 Computer Vision

Computer vision is the AI technique most directly enabled by drone integration and most immediately applicable to the worker safety dimension of HSE management. The fundamental task is object detection and classification within image frames: identifying the presence, absence, or condition of specific objects such as a worker, a safety helmet, a scaffolding member, or a crack in a concrete surface within the visual field captured by a drone or fixed camera.

Three algorithmic families have emerged as dominant in construction HSE computer vision applications, each with a distinct performance profile. YOLO architectures process the entire image frame through a single neural network pass, achieving detection at speeds that make real-time analysis of video feeds computationally feasible even on edge hardware with limited processing capacity. This speed advantage is decisive for applications where the relevant events are transient, such as a worker momentarily removing a helmet or a vehicle entering a restricted zone. Faster R-CNN architectures apply a two-stage detection process, first identifying candidate regions of interest and then classifying each region, which achieves higher accuracy on small targets and partially visible objects. For structural defect detection, where the relevant feature may occupy a small fraction of the image and where speed is less critical than precision, this accuracy advantage outweighs the latency cost. Vision Transformers demonstrate superior capability in detecting complex spatial patterns that extend across large portions of an image, including progressive deformation in structural members and spatial clustering of workers in hazardous configurations (Akinsemoyin et al., 2023). Detection accuracy across these architectures, measured against human-labelled construction site datasets, falls in the range of 93 to 97 percent (Waqar et al., 2023).

These figures should be interpreted with caution: they derive from controlled datasets assembled under laboratory or semi-controlled conditions, and documented accuracy on unstructured, active construction sites has been reported at lower levels, typically 75 to 88 percent, depending on lighting, occlusion, and camera angle (Akinsemoyin et al., 2023).

3.4.2 Predictive Analytics

Predictive analytics addresses a different problem from computer vision: not the detection of currently existing hazardous conditions but the anticipation of conditions that will develop over the coming minutes, hours, or days based on patterns in historical and real-time data. XGBoost is the most widely deployed algorithm for tabular predictive modelling in construction HSE applications. Its architecture, which combines multiple decision trees through a gradient boosting procedure that iteratively focuses attention on observations that previous trees misclassified, produces models that perform well on heterogeneous tabular data with complex non-linear relationships. LSTM networks, a specialised recurrent neural network architecture designed for sequence data, address the temporal dependency structure in sensor time series that tabular models cannot capture: the significance of a current vibration reading depends on whether it represents a new development or the continuation of an established pattern (Waqar et al., 2023).

3.4.3 Anomaly Detection

Anomaly detection occupies a distinct position in the AI toolkit for HSE management because it addresses a knowledge gap that predictive analytics cannot fill: the identification of hazardous conditions that have not previously been observed and therefore cannot be the subject of supervised learning models trained on labelled historical incident data. In an environment where novel hazards emerge from new materials, new construction techniques, or unusual site configurations, a system that can only identify hazards matching historical patterns will systematically fail to detect the conditions most likely to cause serious novel incidents. Unsupervised anomaly detection methods, including One-Class SVM and Deep Autoencoder architectures, learn a model of normal operating conditions from unlabelled data and then identify observations that deviate significantly from this learned normality.

3.4.4 Natural Language Processing

Natural language processing contributes to the SmartHSE analytical framework through the textual records of HSE management, including incident reports, near-miss notifications, inspection findings, and safety meeting minutes. These documents, accumulated over the life of a project and across an organisation's project history, contain information about causal patterns and risk factors that is not captured in sensor data. BERT-based classification models perform well on the task of categorising incident reports by hazard type, causal factor, and severity, enabling automated structured database entry from narrative text. GPT-based generative models enable more sophisticated applications: summarising incident patterns across a large report corpus and identifying clusters of incidents with similar root cause narratives.

3.4.5 Reinforcement Learning

Reinforcement learning represents the most advanced AI application in the SmartHSE framework and the one with the greatest gap between demonstrated research capability and mature construction deployment. The architecture involves an agent that learns, through iterated interaction with an environment, which actions in which states produce the best outcomes as defined by a reward function. The applications that reinforcement learning is best positioned to address involve dynamic resource allocation under uncertainty: scheduling drone sorties across a large site to maximise hazard detection probability given limited flight time, allocating human inspection capacity across a site with multiple concurrent high-risk activities, or sequencing the closure and reopening of work zones in response to a developing hazard while minimising project disruption (Palanisamy, 2018).

3.5 Progress 2020-2025

Several AI developments between 2020 and 2025 addressed specific deployment constraints that had limited earlier systems. Multimodal architectures, for example, emerged partly because single-modality models were failing to exploit the correlations between drone imagery and IoT sensor readings that carry the strongest predictive signal for certain hazard types.

The emergence of Multimodal AI addresses a limitation that single-modality models share: the inability to integrate information from fundamentally different data types within a single analytical framework. A computer vision model processes image arrays; a predictive analytics model processes numerical tabular data; a natural language model processes text sequences. These models, operating independently, each contribute useful outputs but cannot exploit correlations between data types that may carry the strongest predictive signal. Multimodal architectures such as Meta's ImageBind and GPT-4V are designed to process multiple data modalities simultaneously, producing integrated analyses that reflect the actual structure of complex, multi-causal construction hazard scenarios.

Explainable Artificial Intelligence (XAI) overcomes the challenge of deep models being black boxes. Techniques such as LIME, SHAP, and Grad-CAM increase transparency by displaying heat maps of the points that influence AI decisions. For example, XAI can show that a scaffold collapse alert was issued based on 70 percent abnormal vibration from sensor data, 20 percent material congestion from imagery, and 10 percent wind speed. This provides transparency, trust for human users, and allows for auditing of the model (Agapiou, 2024).

Federated learning addresses the tension between the data volume requirements of effective AI training and the privacy and commercial sensitivity constraints that limit data sharing across organisations. Unlike a centralised approach that sends raw data to a cloud server, in this method, the AI model is trained locally on each site's data, and only the updated parameters are sent to a central server. In this way, knowledge is extracted from multiple projects without any proprietary and confidential data from each project leaving its premises. This method is essential for compliance with privacy regulations such as GDPR.

3.6 Application in HSE

3.6.1 Behavioural Safety Monitoring

By analysing videos from drones and fixed cameras, the AI in SmartHSE is able to instantly detect unsafe behaviours such as walking under a crane load, removing a helmet in a prohibited area, or standing on the edge of a cliff without protection. Algorithms such as YOLOv8 and Vision Transformers detect these violations with high accuracy and can immediately issue an audible or vibration alert via wearable devices. The deterrent effect of an active monitoring system, attributable to the perception of continuous observation independent of any specific alert, has been estimated to produce up to a 30 percent reduction in unsafe behaviours (Shah & Mishra, 2024).

3.6.2 Risk Prediction and Quantified Decision Support

Risk prediction models based on LSTM and XGBoost, trained on historical data and combined with real-time streams of IoT sensors including vibration, temperature, gas concentration, wind speed, and workforce density, can calculate the probability of an accident occurring for short-term intervals at any point on the site. For example, the system might indicate that in the crane loading area, with a wind speed of 15 m/s and 12 workers in the critical radius, there is a 78 percent chance of a load falling. Such a specific prediction, unlike traditional approaches that simply identify high-risk areas in general terms, enables targeted preventive action such as evacuation or additional restrictions.

Complementing these capabilities, Dynamic Risk Heat Maps are updated on the central dashboard every few minutes with real-time data, highlighting high-risk areas in red and allowing project managers to optimise the allocation of monitoring resources at a glance without having to read multiple reports. This spatial representation of risk transforms the nature of the management decision: rather than responding to qualitative risk flags, managers receive probability-quantified predictions with identified contributing factors that enable proportionate, targeted responses.

3.6.3 Challenges

The application of AI in HSE management faces fundamental challenges that are critical to the widespread adoption of this technology in the construction industry. The black-box nature of deep learning models, especially multi-layer neural networks, is one of the most important of these obstacles. When the system issues an alert, the project manager cannot easily identify which features the alert was generated based on and with what degree of confidence. This lack of transparency reduces the trust of human users and makes relying solely on AI output in critical cases where workers' lives are at stake very questionable (Agapiou, 2024). For this reason, the development of explainable architectures is a research priority in this area.

Algorithmic bias in data and models is another challenge with serious ethical and practical implications. If the training data of an AI system is mainly collected from projects in developed countries with high safety standards, modern equipment, and standard worker behaviours, the system will perform poorly when faced with different conditions in developing countries, including less educated workforces, outdated equipment, and specific environmental conditions. The system may also unintentionally bias its diagnosis towards certain groups of workers based on clothing type, skin colour, or language.

Legal liability when AI makes a mistake and an accident occurs is still not clearly defined in most legal systems around the world. Whether the liability rests with the algorithm manufacturer, the sensor installer, or the project manager who trusted the machine's output remains an open question. Ultimately, all advanced AI capabilities depend on the availability of high-quality, labelled data that is representative of real-world construction site conditions. Many organisations lack the infrastructure and technical know-how to generate and maintain this data, and without it, AI predictions will be neither accurate nor reliable.

Table 3. 2:AI Techniques and HSE Applications

AI Technique	Main Application in HSE	Reported Accuracy	Key Reference
Computer Vision	PPE detection, unsafe behaviour recognition	93-97%	Akinsemoyin et al. (2023)
Predictive Analytics	Accident prediction, trend modelling	High	Waqar et al. (2023)
Anomaly Detection	Identification of previously unobserved hazardous conditions	Context-dependent	Shah & Mishra (2024)
NLP	Root cause extraction from incident reports	-	(Shah & Mishra, 2024)
Reinforcement Learning	Dynamic resource allocation	Emerging	(Peng, Zhang, Feng, & Liu, 2023)

Chapter 4: Systematic Literature Review on the Application of Emerging Technologies in HSE Management

4.1 Research Design and Philosophical Positioning

Research on technology-enabled management change cannot be value-neutral: the methods chosen to study it reflect assumptions about what HSE risk is, how it is best measured, and what kind of evidence is sufficient to justify a new framework. This section makes those assumptions explicit before the review protocol is described. Making these assumptions explicit serves two purposes: it clarifies the interpretive choices that underpin the review methodology, and it allows the boundaries of the resulting conceptual framework to be understood rather than merely accepted. This section sets out the philosophical and methodological positioning of the research before the systematic review protocol is described.

4.1.1 Ontological and Epistemological Stance

HSE risk has a dual nature. Its physical dimension, which includes sensor readings, deformation data, number of incidents, and is essentially measurable. Its social dimension, which is less controllable: what is considered an acceptable level of risk, which danger signals attract management attention, and how causes of incidents are attributed, all depend on the organizational culture and the interpretative habits of the people involved. Rather than committing to either a purely objectivist or a purely subjectivist position, the pragmatist stance treats these two dimensions as complementary: both are relevant to understanding how integrated technologies can change HSE management practice, and neither alone is sufficient.

The epistemological position that follows from this ontological stance is interpretivist. The knowledge generated by this study does not arise from measuring a single observable phenomenon under controlled conditions but from synthesising and interpreting findings across a heterogeneous body of literature produced in different countries, institutional settings, and technological contexts. The researcher's analytical judgement in comparing study findings, identifying patterns of evidence, and inferring the implications of research gaps is not a source of bias to be minimised but an integral part of the knowledge-building process. This does not mean the process is arbitrary: the systematic review protocol described in the following section imposes transparency and reproducibility conditions that constrain subjective interpretation. It does mean, however, that the framework produced by this research should be read as a structured interpretation of available evidence rather than an empirical discovery.

Table 4. 1:Summary of Philosophical and Methodological Positioning

Dimension	Position Adopted	Rationale in the Context of This Study
Ontology	Pragmatist	The study acknowledges that the reality of HSE risk on construction sites is multifaceted: partially measurable through sensor data and partly constructed through organisational practice and human judgement. A pragmatist stance accommodates both dimensions without forcing a strictly objectivist or subjectivist commitment.
Epistemology	Interpretivist	Knowledge about how emerging technologies improve HSE management is produced by interpreting and synthesising findings across a heterogeneous body of empirical and conceptual literature, rather than by measuring a single observable phenomenon. The researcher's analytical judgement is an integral part of the knowledge-building process.
Research Approach	Inductive	The study moves from the particular findings of 73 individual studies towards the broader generalisations embodied in the SmartHSE framework. No pre-existing integrated theory was available to test deductively; the framework emerges from the patterns identified in the literature.
Research Strategy	Systematic Literature Review and Conceptual Framework Development	An empirical case study would test the framework under specific site conditions, but could not establish the generalised reference architecture that addresses the integration gap across the field. A conceptual strategy grounded in a systematic review provides broader applicability at this early stage of the technology's adoption curve.

4.1.2 Research Approach and Strategy

The research follows an inductive approach. No pre-existing, validated, integrated theory of UAV-IoT-AI synergy in HSE management was available from which hypotheses could be derived and tested. The SmartHSE framework presented in Chapter 5, therefore, emerges from the patterns and gaps identified across the reviewed literature, moving from the particular to the general rather than from an a priori theory to its empirical confirmation.

The choice of conceptual framework development as the research strategy, rather than an empirical case study or experimental design, requires explicit justification. A case study approach would have allowed the proposed framework to be tested under specific site conditions and would have produced quantitative evidence of impact on HSE indicators. However, the primary gap identified in the literature, documented in detail in Section 4.7, is not the absence of empirical evidence about individual technologies but the absence of a reference architecture that specifies how the three technologies should be integrated across the full HSE risk management cycle. Filling this architectural gap requires a generalisable conceptual contribution rather than a site-specific empirical one. A conceptual framework grounded in a comprehensive systematic review provides the breadth of coverage necessary to establish that reference architecture, while also ensuring it is compatible with the diverse project types and geographic contexts in which it will ultimately need to operate. Empirical validation of the framework in real construction projects is accordingly identified as the primary direction for future research in Chapter 6.

4.2 Methodology of the Systematic Literature Review

Systematic literature reviews differ from conventional narrative reviews primarily in their procedural transparency: the search strategy, inclusion decisions, and quality assessment criteria are pre-specified and documented, so that another researcher could replicate the process and, in principle, arrive at the same set of included studies. Unlike traditional narrative reviews, which are often selective and subjective in their coverage, a systematic review follows a pre-specified protocol governing search strategy, inclusion and exclusion decisions, quality assessment, and synthesis. There were strong methodological reasons for adopting this approach in the present study. The field of UAV, IoT, and AI applications in construction HSE management has expanded rapidly since approximately 2018, producing a large and dispersed body of articles across multiple disciplinary databases. A systematic review is the appropriate instrument for organising this body of knowledge, assessing the quality of individual studies, and synthesising findings in a manner that is both defensible and replicable. Most importantly, the systematic review is the most reliable tool for accurately identifying research gaps, which is the central analytical objective of this chapter.

4.2.1 Protocol and Framework

The PRISMA 2020 protocol was adopted as the primary methodological framework. PRISMA, which stands for Preferred Reporting Items for Systematic Reviews and Meta-Analyses, provides a standardised checklist and flow diagram that covers the stages from initial research question definition through search, screening, eligibility assessment, and final reporting. Adherence to this protocol significantly increases the transparency and reproducibility of the review and enables other researchers to replicate or update the search as the literature continues to develop. The overall approach combined two complementary methods: a systematic review procedure for identifying and screening articles, and an in-depth narrative synthesis for comparing findings, characterising technology contributions, and drawing conclusions about research gaps. This combination preserves scientific rigour while enabling the rich content analysis required to inform framework design.

The search strings were constructed around three thematic clusters. For UAVs: ("unmanned aerial vehicle" OR "drone" OR "UAV") AND ("construction" OR "building site") AND ("safety" OR "HSE" OR "hazard"). For IoT: ("internet of things" OR "IoT" OR "sensor network" OR "wearable") AND ("construction safety" OR "occupational health"). For AI: ("artificial intelligence" OR "machine learning" OR "computer vision") AND ("construction" OR "site safety" OR "HSE"). Boolean operators AND/OR were applied across all three clusters to capture integrated studies.

Five major databases were selected for the search: Scopus, Web of Science, IEEE Xplore, ScienceDirect, and ASCE Library. The search was restricted to the period 2018 to 2025, reflecting the period during which the use of these technologies in the construction industry began growing significantly and producing peer-reviewed evidence. The initial search retrieved 1,242 articles, which, after removing duplicates across the five databases, yielded 889 unique records.

Table 4. 2:Initial Search Results by Database

Database	Initial Results	After Duplicate Removal
Scopus	487	
Web of Science	312	
IEEE Xplore	198	
ScienceDirect	156	
ASCE Library	89	
Total	1,242	889

4.2.2 Inclusion and Exclusion Criteria

The selection criteria were formulated with the explicit aim of minimising two forms of bias: over-inclusion of tangentially relevant studies that would dilute the quality of evidence, and under-inclusion through criteria so narrow that important contributions would be excluded. Eligible articles were required to examine at least one of the three target technologies in the context of HSE management in construction or closely related infrastructure projects, to be published in English, to fall within the 2018 to 2025 publication window, and to be indexed in Scopus or Web of Science as indicators of minimum publication quality. Articles not directly related to HSE, grey literature, books, letters to the editor, and purely technical studies that examined hardware or algorithm performance without addressing HSE implications were excluded from the review process. Borderline cases during the screening process were resolved through a two-step procedure. First, the full text was retrieved and assessed against all inclusion and exclusion criteria systematically. Second, where ambiguity remained, particularly for studies addressing productivity or quality management with only indirect safety relevance, the deciding criterion was whether the study generated findings directly applicable to the design of an integrated UAV-IoT-AI HSE architecture. Studies that met this secondary criterion were retained; those whose relevance was confined to broader construction management outcomes were excluded. This procedure was applied consistently across all 129 articles that required full-text evaluation beyond the initial title and abstract screening.

Table 4. 3: Summary of Inclusion and Exclusion Criteria

Criterion	Inclusion	Exclusion
Topic	Drones, IoT, or AI in HSE management of construction projects	No direct relation to HSE or construction
Document Type	Peer-reviewed research articles or systematic reviews	Books, letters, editorials, and non-scientific reports
Language	English	Other languages
Time Period	2018 to 2025	Before 2018
Quality	Indexed in Scopus or Web of Science; minimum acceptable score on JBI Checklist	Not indexed in reputable databases; score below minimum threshold.

4.2.3 Screening and Selection Process

The screening of articles was carried out in four sequential steps in accordance with the PRISMA 2020 standard. In the first step, all 889 unique records were screened by title and abstract, and records clearly outside the scope defined by the inclusion criteria were excluded. This reduced the pool to 202 articles retained for full-text evaluation. In the second step, the full texts of these 202 articles were retrieved and assessed against all inclusion and exclusion criteria. This step required judgment on borderline cases, particularly studies that addressed HSE-adjacent topics such as construction productivity or quality management with only indirect relevance to safety outcomes. In such cases, the guiding principle was whether the study produced findings that could inform the design of the SmartHSE framework. After this assessment, 73 articles were selected as the final study set for synthesis. The PRISMA flow diagram presented in Figure 4.1 provides a visual representation of each stage of this selection process, including the reasons for exclusion at the full-text review stage.

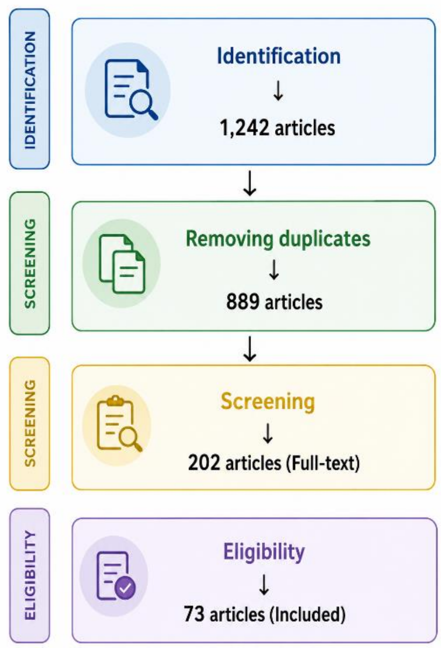

Figure 4. 1:PRISMA 2020 Flow Diagram

4.2.4 Quality Assessment

All 73 articles retained for synthesis were subjected to formal quality assessment using the JBI Critical Appraisal Checklist, adapted for the study designs represented in the corpus. The JBI instrument, developed by the Joanna Briggs Institute, provides design-specific criteria for evaluating the methodological rigour of studies, including crosssectional studies, systematic reviews, and text and opinion papers. For empirical studies, the checklist examines criteria such as sample representativeness, measurement reliability, control of confounding variables, and completeness of outcome reporting. For review articles and conceptual studies, it examines criteria such as comprehensiveness of literature coverage, clarity of conceptual boundaries, and consistency of argument. Each article was scored against the applicable checklist, and only those achieving a minimum score of six out of ten applicable criteria were retained in the final synthesis. Articles falling below this threshold were excluded from the synthesis, but their bibliographic data were retained in the supplementary record. This threshold strikes a balance between stringency and inclusivity that is appropriate for a rapidly developing field where the evidence base includes a higher proportion of early-stage conceptual studies than in more mature research domains.

Of the 89 articles assessed at full-text stage, 73 met the minimum quality threshold and were retained for synthesis. The 16 excluded studies were predominantly early-stage conceptual pieces that documented fewer than six of the ten applicable JBI criteria, mainly due to insufficiently described methodological procedures and unverifiable data sources. The complete record of individual appraisal scores is maintained in the author's research database and is available upon request.

4.3 Review Results: Unmanned Aerial Vehicles

Among the three technology domains, UAV research in the reviewed corpus is characterised by the strongest concentration of empirical studies and the clearest evidence of performance improvement relative to traditional inspection methods. The dominant application pattern across the drone literature is spatial data collection under conditions that are either too hazardous or too time-consuming for human inspectors: scaffold inspection at height, mapping of large or irregular site areas, assessment of areas immediately following incidents or near-misses, and monitoring of active demolition zones. These applications share a common structural feature: the drone substitutes for a human in an exposure-generating activity, eliminating the risk transfer that occurs when a human inspector must enter a hazardous environment to inspect it.

Automated hazard detection, enabled by the integration of computer vision algorithms with drone-mounted cameras, represents the most rapidly advancing frontier within the UAV literature. Studies reviewed from 2022 onwards consistently report detection accuracy figures in the range of 93 to 97 percent for PPE compliance monitoring and unsafe behaviour identification, figures that compare favourably with human inspection baseline performance, particularly under the degraded conditions of fatigue, poor lighting, and constrained physical access that characterise real-world construction sites (Waqar et al., 2023). The transition from post-flight image review to real-time algorithmic analysis during flight, enabled by onboard edge computing, represents a qualitative shift in the temporal dimension of hazard detection that the earlier generation of construction drone literature did not address.

Several studies demonstrate time efficiency gains of up to 85 percent for comprehensive site safety inspections compared to human-paced approaches, with the implications extending beyond speed to include the feasibility of daily inspection frequencies that are resource-prohibitive under manual inspection regimes (Albeaino et al., 2023). The reviewed literature consistently positions UAVs as the visual data acquisition subsystem of integrated HSE frameworks, while also demonstrating that their contribution to risk reduction is significantly amplified when combined with IoT monitoring and AI-based analysis. However, the mechanism and architecture of this amplification remain underspecified in most studies: the literature documents that the combination produces better outcomes without providing sufficient detail about how the combination is technically implemented or managed.

Limitations acknowledged across the drone literature centre on weather sensitivity, regulatory constraints, battery endurance, and the absence of continuous temporal coverage. These limitations are not incidental to specific studies but structural features of the technology that any integrated HSE architecture must address through complementary subsystems. The reviewed literature's treatment of these limitations is generally descriptive rather than prescriptive: studies identify the constraints without providing frameworks for managing them at the system level.

4.4 Review Results: Internet of Things

The IoT literature in the reviewed corpus occupies a distinct analytical space from the drone literature, corresponding to its distinct technical function. Where drone studies are primarily concerned with spatial coverage and point-in-time visual assessment, IoT studies are primarily concerned with temporal continuity and parametric measurement. The applications receiving the most consistent empirical support are those that exploit the continuous measurement property of IoT networks in ways that fundamentally cannot be replicated by periodic inspection: structural health monitoring of temporary works such as scaffolding and formwork, real-time atmospheric hazard detection for toxic and flammable gases, and worker physiological monitoring through wearable devices that detect fatigue onset and heat stress before symptomatic deterioration.

Quantitative outcome data in the IoT literature is more variable than in the drone literature, reflecting greater diversity in deployment scope and system maturity across studies. The range of accident rate reductions reported across the reviewed IoT studies, from 25 to 68 percent, captures this genuine variation rather than representing measurement uncertainty. Studies at the upper end of this range typically feature comprehensive multi-category sensor deployments, AI-based analysis of sensor outputs, and real-time alerting systems integrated with wearable devices. Studies at the lower end tend to feature narrower sensor coverage, periodic rather than continuous data transmission, or alert systems that route warnings through human intermediaries rather than directly to workers in the field. This pattern suggests that the effectiveness of IoT in construction HSE is not a fixed quantity but a function of implementation depth and integration quality.

Across the IoT literature, infrastructure emerges as the foundational data layer upon which AI-based analysis depends: without continuous sensor input, predictive models lack the real-time context necessary for accurate risk assessment. This dependency is acknowledged in many studies but rarely examined at the architectural level. The reviewed literature provides extensive evidence of what IoT can detect but limited guidance on how sensor networks should be designed, deployed, and integrated with other system layers to maximise their contribution to predictive HSE management.

Cybersecurity exposure, interoperability between devices from different manufacturers, power management in locations without a stable electricity supply, and the high initial capital cost of comprehensive sensor networks are the most frequently cited challenges across the IoT literature. What is notable about this challenge profile is that none of these issues are amenable to purely technical solutions: cybersecurity requires organisational protocols and ongoing monitoring; interoperability requires industry standards that have not yet emerged; power management requires engineering design decisions specific to each site context; and cost barriers require financing mechanisms and economic models that are not within the scope of technology developers alone to provide.

4.5 Review Results: Artificial Intelligence

Artificial intelligence commands the largest share of the reviewed literature, a distribution that reflects both the breadth of AI applications across the HSE domain and the fact that AI is the analytical layer that processes outputs from the other two technologies. AI functions as the interpretive layer of integrated HSE systems, transforming heterogeneous inputs from UAVs and IoT sensors into structured risk predictions and behavioural safety alerts. Within the reviewed corpus, computer vision for automated hazard detection, predictive analytics for accident forecasting, and anomaly detection for identifying novel hazardous conditions are the most substantively developed application areas.

Computer vision studies are the most technically mature subset of the AI literature. The trajectory from accuracy figures in the 70 to 80 percent range reported in studies from 2018 and 2019 to the 93 to 97 percent range reported in the 2022 to 2025 literature reflects improvements in both algorithm architecture and training dataset quality. This performance improvement is practically significant: detection accuracy at the upper end of the current range is sufficient for deployment as a supplementary monitoring system in safety-critical environments, where supplementary is understood to mean that human oversight remains the final decision authority while automated detection reduces the cognitive and physical burden on human inspectors. Deployment as a primary safety assurance system, where automated detection is the sole or dominant source of hazard identification, remains a goal rather than a current practice in most construction contexts.

The predictive analytics literature demonstrates that machine learning models trained on historical incident data combined with contemporaneous sensor readings can generate probabilistic accident forecasts for specific site locations and time windows, a capability that transforms the informational basis of site management decisions from retrospective reporting to forward-looking risk assessment. However, the quality of these predictions is heavily dependent on the size and representativeness of the training dataset, which in the construction industry is typically limited by the low frequency of serious incidents and the inconsistent quality of incident recording practices.

The most important advantages of AI identified across the reviewed literature include the ability to transform raw data into actionable insights, the reduction of dependence on individual inspector experience and availability, and the enabling of real-time decision-making at scales and speeds that human cognition cannot match. In contrast, the limitations most consistently identified are the black-box opacity of deep learning models, the need for high-quality labelled training data, the risk of algorithmic bias in deployment contexts that differ from training contexts, and the unresolved questions of legal liability when AI-assisted decisions contribute to adverse outcomes. The reviewed literature acknowledges these limitations but provides limited practical guidance on managing them.

4.6 Review Results: Integrated Applications

Studies examining the simultaneous deployment of two or three of the target technologies represent the smallest subset of the reviewed corpus, accounting for fewer than 15 percent of the 73 included articles. This scarcity is itself a significant finding: despite the logical case for technology integration, empirical evidence of its implementation and outcomes remains rare. The studies that do address integration are predominantly architectural in nature, proposing frameworks or system designs rather than reporting field results from deployed integrated systems. This pattern indicates that the construction industry is at an early stage of the technology integration curve, with conceptual architecture ahead of operational implementation.

Several integration studies propose hybrid drone-IoT-AI architectures that combine edge computing for latency-critical processing with cloud infrastructure for storage and advanced analytics (Daniel et al., 2025; Shafei et al., 2022). The architectural logic of these proposals is consistent across studies: drones provide spatial visual coverage, IoT provides temporal parametric coverage, and AI provides the analytical layer that converts combined data streams into risk assessments. What varies across proposals is the level of specificity at which this architecture is described: some studies remain at the level of component enumeration, while others specify communication protocols, processing allocation between edge and cloud, and the structure of data flows between layers.

The initial results reported from integrated system deployments or simulations indicate that combining the three technologies produces disproportionate risk identification improvements compared to single-technology deployments, consistent with the theoretical expectation that the synergy of complementary data streams would reveal hazard patterns that no single stream could identify independently. However, the evidence base for this claim is limited to a small number of studies conducted in specific site conditions, and the causal mechanisms driving the improvement are not consistently specified. Creating a closed loop from data collection to action, achieving multidimensional coverage across safety, health, and environmental dimensions, and enabling continuous system learning from incident feedback are identified as the key advantages of integrated systems across the reviewed literature.

The most significant barriers to integration identified in the reviewed studies are not technical but architectural and organisational: the absence of a comprehensive reference framework that specifies how to implement the full HSE risk management cycle across an integrated UAV-IoT-AI system; the lack of longitudinal studies demonstrating performance over the full lifecycle of a construction project; and the gap between technically capable systems and the organisational conditions required for their effective use. These barriers are the direct motivation for the SmartHSE framework proposed in Chapter 5.

4.7 Synthesis of Findings and Identification of Research Gaps

Drawing together the evidence from the four preceding sections, the systematic review of 73 articles permits a set of conclusions that are grounded in the corpus and carry direct implications for framework design. The first and most straightforward conclusion is that all three technologies, taken individually, have demonstrated substantial added value for HSE management: drones reduce inspection time and inspector exposure; IoT sensors provide continuous measurement of hazardous parameters; and AI converts data streams into actionable risk assessments. The evidence for each of these contributions is sufficiently robust that the technologies' individual value is not in question.

What remains in question, and what the reviewed literature cannot resolve with existing evidence, is how the three technologies should be integrated to produce the greatest collective benefit while managing the risks and costs that each introduces. The systematic review reveals that the published literature addresses this question at a level of generality that is insufficient to guide implementation. Most studies either treat the technologies in isolation or gesture towards integration without specifying the architecture, protocols, and organisational conditions required to realise it.

Table 4. 4:Comparative Summary of the Three Technologies in the Reviewed Literature

Dimension	Drones (UAV)	IoT	Artificial Intelligence
Primary Function in HSE	Visual data acquisition and aerial inspection	Continuous environmental and physiological monitoring	Pattern recognition, prediction, and decision support
Key Advantage	Spatial coverage of inaccessible areas	24/7 temporal continuity	Conversion of raw data into actionable knowledge
Principal Limitation	Discrete sampling; weather sensitivity	Cybersecurity exposure: high initial cost	Data dependency; black-box opacity; algorithmic bias
Share of Reviewed Literature	Approx. 25%	Approx. 30%	Approx. 45%
Integration Evidence	Limited to dual-technology pairings	Limited to dual-technology pairings	A small subset of studies

Five specific research gaps emerge from this analysis, each corresponding to a dimension that the existing literature inadequately addresses. The first and most consequential gap is the absence of a comprehensive conceptual framework that systematically specifies how to integrate UAV, IoT, and AI technologies across the complete HSE risk management cycle, covering not just hazard identification and risk assessment but also risk control, monitoring of control effectiveness, and organisational learning. Most existing models address two or three of these functions and leave the remaining functions either implicit or unaddressed. The real value of an integrated HSE system lies precisely in the closed-loop learning architecture in which monitoring outcomes feed back into risk model parameters and control strategy design; yet this feedback architecture is absent from the majority of reviewed proposals.

The second gap concerns the shortage of longitudinal and field-scale studies. The preponderance of simulation-based evaluations and short-term pilot reports means that the performance of integrated systems under the full range of site conditions encountered over a multi-year project lifecycle remains empirically undocumented. System behaviours that are stable and predictable in controlled or short-duration deployments may degrade significantly under continuous operational stress, equipment deterioration, workforce turnover, and the evolving hazard profile of a project as it moves through design, groundworks, structural, and finishing phases.

The third gap is the insufficient attention to human, organisational, and cultural factors in technology integration research. A technically correct architecture that workers and supervisors do not use, do not trust, or actively circumvent will not reduce accidents. The Technology Acceptance Model and related frameworks for predicting adoption behaviour are largely absent from the reviewed literature, leaving the transition from system installation to effective use as a largely unaddressed implementation challenge.

The fourth gap is the limited explainability of AI decision outputs in the reviewed systems. As noted in Section 4.5, deep learning models in HSE applications are predominantly black boxes: they produce risk assessments and alerts without providing the causal explanation that would allow site managers to evaluate the alert's plausibility or identify the appropriate response. Explainable AI techniques exist but have not yet been systematically incorporated into construction HSE applications.

The fifth gap is the absence of systematic cost-benefit analysis and sustainability assessment for integrated system deployments. Initial capital investment for comprehensive UAV-IoT-AI infrastructure is substantial, and the economic justification depends on assumptions about incident reduction rates, productivity improvements, and insurance cost savings that are not consistently documented in the reviewed literature. Without this evidence base, organisational decision-makers cannot make adequately informed investment decisions.

Technology	Number of Articles	Percentage (%)	Reported Accuracy	Primary Application
DRONES	312 articles	36%	88.7% (Avg. Accuracy)	**Remote Monitoring and Inspection** Used for aerial inspection, site monitoring, hazard identification, incident response, and compliance verification.
IOT	348 articles	40%	85.1% (Avg. Accuracy)	**Real-time Monitoring and Data Collection** Used for continuous monitoring of environmental and operational parameters, equipment health, and worker safety.
AI	212 articles	24%	91.3% (Avg. Accuracy)	**Intelligent Analysis and Decision Support** Used for risk prediction, anomaly detection, event classification, decision support, and predictive maintenance.

Note: *The data is based on a review of recent publications up to 2024 from Scopus, Web of Science, and IEEE Xplore databases.*

Figure 4. 2:Comparison of the Three Technologies in the Literature Review

These five gaps collectively define the design space within which the SmartHSE framework is developed in Chapter 5. The framework does not attempt to resolve all five gaps simultaneously through a single theoretical contribution; rather, it addresses the first gap directly by proposing a reference architecture covering the full risk management cycle, and it positions the remaining four gaps as empirical research priorities for future work. In this sense, the systematic review serves not only as an evidence base for the framework but as a map of the territory that the framework needs to navigate and the territory that future empirical research needs to cover.

The literature review shows that all three technologies, when applied individually, produce measurable improvements in HSE. What is not clear is how they should be combined to work across the entire risk management cycle. This architectural question is what Chapter 5 addresses. The current literature, however, remains a collection of single-technology and dual-technology studies whose integration potential has been identified but not realised. The gaps documented in this chapter provide the empirical and theoretical justification for the integrated conceptual framework that follows in Chapter 5.

Chapter Five: Proposed Framework for Smart HSE Management Using the Synergy of Drones, IoT, and Artificial Intelligence

5.1 Introduction and Theoretical Foundations

The systematic review in Chapter 4 produced a clear but troubling conclusion. From an initial pool of 1,242 records, I carefully selected 73 relevant studies. These papers show that since around 2018, researchers have explored the individual strengths of drones, IoT sensors, and artificial intelligence in construction HSE quite thoroughly. Yet when it comes to combining all three into one functioning system, the literature falls surprisingly short. Over 85 percent of the studies focused on just one technology. Even the few that discussed multiple tools rarely went beyond simply listing components. They seldom explained how data should flow between them or how decisions could be made together. In short, we have many strong individual pieces, but almost no complete picture of how they fit together.

This gap is exactly why I developed the SmartHSE framework in this chapter. My aim was never to add yet another theoretical model to an already crowded field. Instead, I wanted to create something practical. The framework spells out, in sufficient detail, how drones, IoT, and AI can actually work together. The real goal is to move HSE management from a reactive approach, where we respond after accidents happen, to a predictive one, where the system can issue warnings and take action before problems occur.

Four established theoretical traditions shaped the design of SmartHSE. Systems Theory gave me the core idea: the true value comes from how the parts interact, not from how well each part works alone. This led me to focus on the overall system rather than optimizing each technology separately. Socio-Technical Systems Theory reminded me that technology never exists in isolation. Human attitudes, organisational culture, and daily realities on site must be considered from the very beginning. A technically perfect system that workers distrust or managers cannot use will never reduce accidents, no matter how advanced it is.

ISO 31000 provided the practical backbone. It describes risk management as a continuous closed loop from identifying hazards to analysing them, evaluating priorities, applying controls, monitoring results, and feeding lessons back into the system. This learning loop was missing in most of the studies I reviewed. Without it, a system can only react to known risks but cannot grow smarter over time. Finally, the Technology

Acceptance Model and UTAUT helped me ground the framework in real human behaviour. Perceived usefulness, ease of use, trust in the system, and organisational support are not afterthoughts. They determine whether a technically sound solution will actually be used in practice.

5.2 Conceptual Model

At the heart of this thesis is the SmartHSE conceptual model. It was designed as a five-layer architecture. Each layer has its own specific function, but they all work together through well-defined connections. This layered approach was deliberately chosen. This approach makes the framework flexible. As technology advances, we can upgrade a layer without rebuilding the other layers. This gives the system a longer useful life.

The bottom layer is about data collection. Drones fly over the site, capturing high-resolution images and spatial information. They show what is happening from angles that no one on the ground can easily access. At the same time, IoT sensors continuously measure environmental conditions, worker vital signs, and structural behavior. Together, these two sources create a rich and vivid picture of the site that neither could provide on its own.

Table 5. 1: The combined role of drone and IoT data sources

Data Source	Data Type	Key Advantage
Drones (UAV)	Visual + Spatial	Wide coverage and aerial perspective
IoT Sensors	Environmental + Physiological + Structural	Real-time and continuous monitoring

Just above that is the edge communications and processing layer. This layer solves one of the biggest practical problems: latency. Some safety decisions need to happen in milliseconds. Waiting for data to be sent to the cloud and back is very slow. That's why I put edge computing in local gateways and on the drones themselves. They do the instant computing on site. I rely on 5G for data movement on site where possible, LoRaWAN for wide coverage, and MQTT for efficient communication.

The real intelligence is in the third layer, the intelligent analytics layer. Here, computer vision algorithms scan drone images to detect missing personal protective equipment or unsafe behavior. Predictive models, using techniques like XGBoost and

LSTM, predict the likelihood of accidents in the coming hours. Anomaly detection algorithms identify unusual conditions that don't fit any pre-existing patterns. The output is not raw numbers. It is clear that prioritized risk assessments with confidence levels.

The fourth layer translates these assessments into action. It follows a hierarchy of controls. Serious hazards can trigger automated responses, such as stopping equipment or sounding alarms. Less important alerts appear on workers' dashboards or wearable devices. The important thing is that the human supervisor has the final say in most cases. The system is meant to support people, not replace them.

At the highest level, the feedback and learning layer closes the loop. Every alert, every action taken, and every near-miss is converted into new data. This information continually improves the predictive models and updates the risk thresholds. This self-learning ability is what sets SmartHSE apart from most previous offerings. It doesn't stay static. It gets smarter over time.

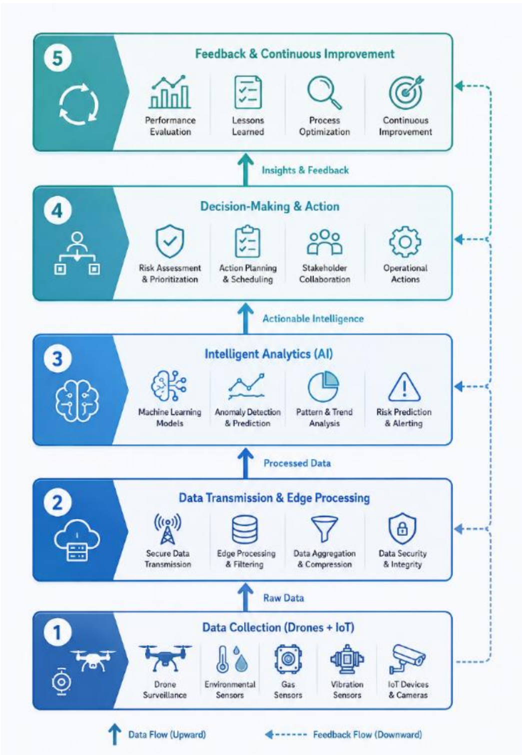

Figure 5. 1: Conceptual Model of the SmartHSE Framework (Five Layers)

5.3 Overall Integrated System Architecture

The conceptual model in the previous section explains what SmartHSE should do. This section focuses on how it actually works in practice. The conceptual model was intentionally separated from the technical architecture. This separation gives the framework flexibility. As new hardware and software are released, we can update the technical part without changing the overall logic of the system.

The architecture was designed as a hybrid edge-cloud system. The division of tasks is not fixed. It depends on how quickly each function needs to respond. For example, detecting a rising gas level or a sudden vibration of the scaffolding needs to happen almost instantly. These critical tasks are executed at the edge, on local gateways and drone processors. On the other hand, heavy computations, such as training new AI models or analyzing long-term trends in data, happen in the cloud, where there is much more processing power. When the internet connection is weak or interrupted, the edge nodes continue to operate independently, storing data until they can be synchronized again.

The hybrid edge-cloud design improves reliability for the most critical functions: a gas emergency alert processed at the edge is not dependent on the network connection to the cloud. Response latency for time-sensitive actions is similarly reduced. The architecture is also scalable with the site; additional edge devices can be added without reconfiguring the cloud layer. Adding more edge devices is straightforward and does not require major changes to the cloud.

Security and privacy are implemented at every layer of the architecture. I used end-to-end encryption for data transmission, two-factor authentication for dashboards, and federated learning so that AI models can learn from multiple sites without sending raw, sensitive data outside the project. These protections are not added at the end. They are built into the system from the beginning.

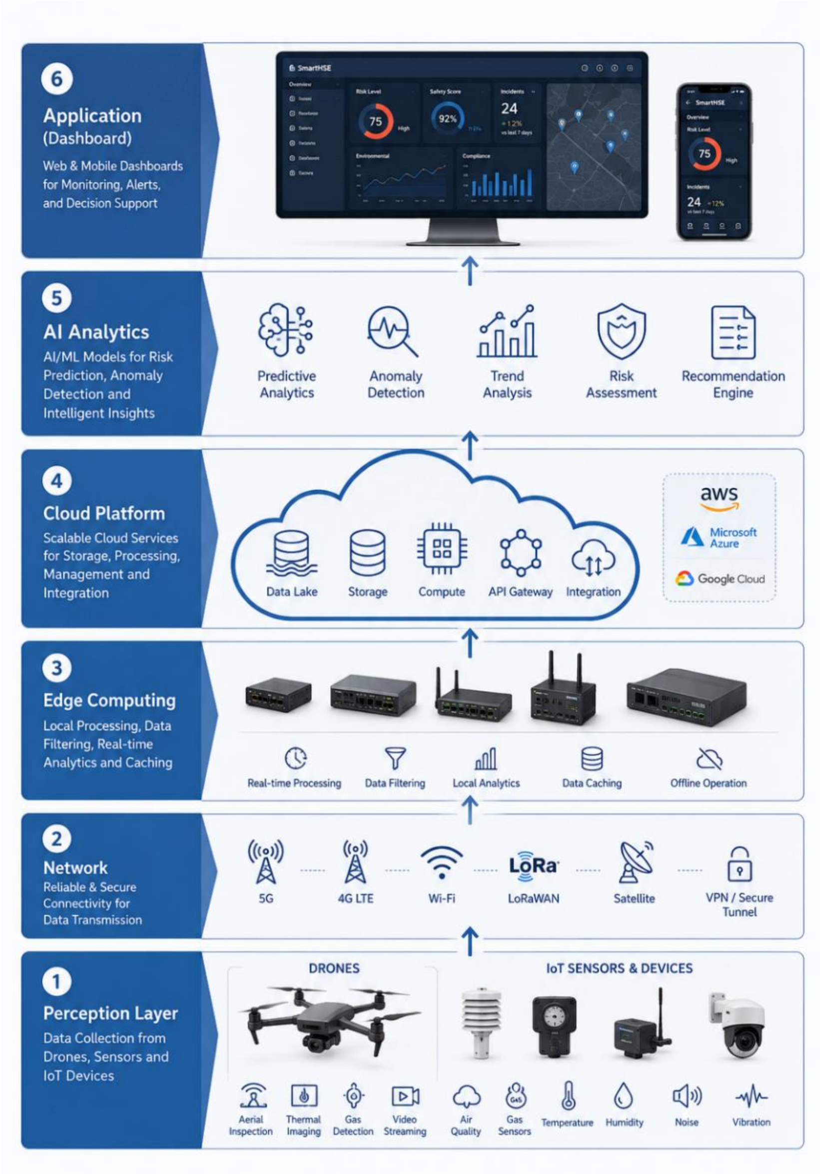

Figure 5. 2:Technical Architecture of the SmartHSE System

5.4 Role of Each Technology in the Proposed Framework

SmartHSE was designed so that each technology has a clear and defined role. Understanding these roles is important. This shows why the framework requires all three to work together, and why using more of one cannot replace the other.

Drones act as the moving eyes of the system. They fly over the site in regular routes and also move to specific points when the analytics layer indicates a potential problem. Their main contributions are to collect high-resolution images from difficult angles, conduct scheduled and urgent inspections, create and update 3D models of the site, and respond quickly to AI commands. However, their true value is only revealed when their visual data is combined with information from other layers. A drone that simply records video for later human review is useful, but it is much less powerful than a drone whose images are immediately analyzed by the system.

The IoT network plays a completely different role. While drones provide wide-area coverage, IoT provides continuous and accurate measurements over time. The technology tracks environmental conditions such as gas levels, dust, temperature, and noise. Wearable sensors monitor workers' heart rates, fatigue, and body temperature. Structural sensors detect vibration and movement in scaffolding and buildings. This continuous stream of data provides the rich time-series information needed by AI to recognize patterns that gradually develop before an incident occurs.

AI acts as the brain of SmartHSE. This AI takes data from drones and IoT sensors, finds meaningful patterns among them, and translates those patterns into actionable recommendations. What makes AI truly essential is its ability to see connections that no single technology can detect on its own. For example, it can combine drone imagery of workers near an edge with IoT data showing increasing wind speeds and fatigue levels to predict the risk of a fall before it happens.

Table 5. 2:The role of each technology in the SmartHSE framework

Technology	Main Role	Primary Data Type	Key Output
Drones	Moving Eyes	Visual and Spatial	Images, Maps, 3D Models
Internet of Things	Nervous System	Environmental, Physiological, Structural	Real-time Continuous Data
Artificial Intelligence	Analytical Brain	Combined Multimodal Data	Risk Prediction, Decisions

The role boundaries defined above have a practical consequence: each technology compensates for the inherent limitations of the other two. Drones provide spatial coverage that IoT sensors cannot replicate; IoT sensors provide temporal continuity that drone flights cannot achieve; AI processes composite output at a speed and scale that human analysts cannot match. Only when they work together does the system become greater than the sum of its parts.

5.5 Technology Synergy: An Operational Illustration

Abstract descriptions of the synergistic hazard obscure the tangible mechanism of its operation. A specific scenario illustrates the difference between a single technology operation and an integrated SmartHSE response. An IoT sensor in a section of a petrochemical construction site detects methane concentrations rising above a safety threshold. In a typical system, this triggers an alarm that reaches the control room operator, who must decide whether to investigate, whom to send, and how long to wait before escalating the situation. In SmartHSE, the detection event initiates a coordinated sequence. The AI analytics layer first contextualizes the read information against background data: wind direction from nearby sensors, activity records for the area, recent pipeline maintenance history in the area, and the pattern of concentration increases over the previous minutes. If the contextual assessment matches the pre-determined risk threshold, this layer issues an automated command to the drone subsystem, such as “Send drone #3 to coordinates (X, Y) for visual assessment,” while simultaneously issuing a precautionary alert to workers in the affected area and notifying the site HSE officer with a probability-weighted risk estimate. The drone arrives within tens of seconds and transmits live video, which computer vision algorithms analyze in real time against sensor data in the system. Combined analysis, such as the atmospheric concentration at a specific location correlated with the visual identification of a worn valve in pipeline #2, allows for a specific detection that no data source alone can support: “Worn valve in pipeline #2; recommend evacuation of 50-meter radius and immediate response to maintenance.” This specific and causal advice is different from a general gas alert. This recommendation provides the HSE officer with the information needed to calibrate the response appropriately (mobilizing appropriate resources rather than responding inadequately to an alert or over-evacuating the site) and creates a documented record that updates the risk model parameters for similar configurations in the future.

This operational example includes two outputs that differentiate SmartHSE from traditional systems: a dynamic risk heat map, which updates in real time as sensor readings and drone observations change, giving site managers a continuous spatial representation of the likelihood of risk. and a decision-to-action latency that is measured in seconds rather than the hours required by traditional inspection and response cycles. Together, these capabilities form the predictive and proactive operating mode that defines the purpose of this framework.

5.6 Implementation Phases and Roadmap

Turning the SmartHSE framework from a conceptual model into a working system on a real construction site is not something that can happen overnight. It requires careful planning, step-by-step execution, and constant attention to both technical and human factors. I designed a five-phase roadmap for this purpose. For a typical medium to large project, the full implementation usually takes between 18 and 24 months.

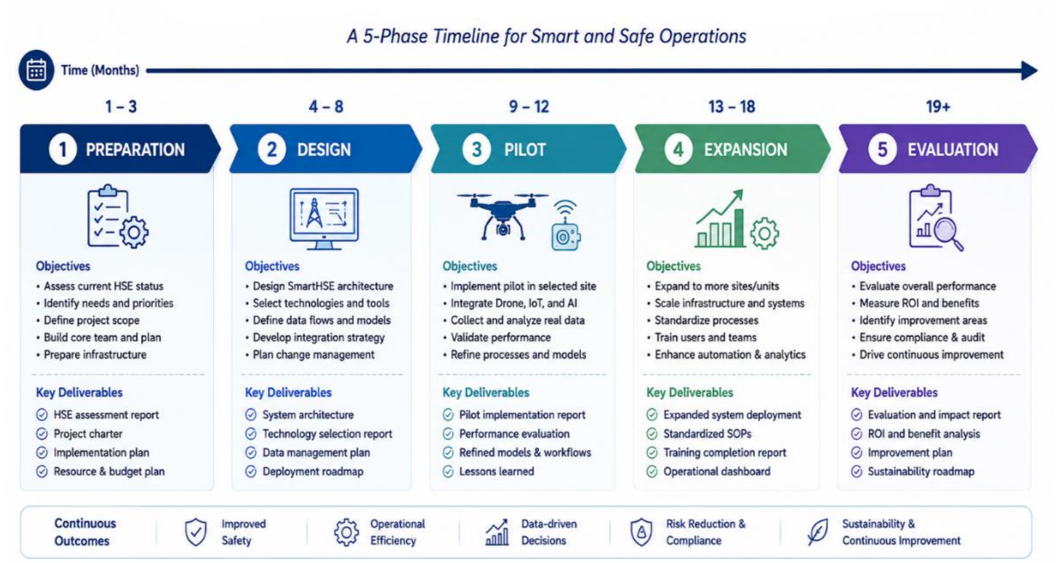

Figure 5. 3:Five-Phase Implementation Roadmap of SmartHSE Framework

The first phase focuses on preparation and planning. During the first three months, we will carry out a thorough assessment of the current HSE processes. This helps us understand what already exists and where the biggest gaps are. At the same time, we form a cross-functional team that includes HSE experts, IT specialists, and data analysts. We also choose a suitable pilot area, usually the zone with the highest risk or most complex activities. Finally, we prepare a realistic cost-benefit analysis to show management the expected returns in terms of fewer accidents, better productivity, and lower costs.

Phase two is about design and development. This stage runs from months four to eight. Here, we define the detailed technical architecture tailored to the specific site conditions. We select and procure the necessary hardware: suitable drones, various types of sensors, communication equipment, and edge computing devices. The AI models are developed and trained using historical data from the site. Dashboards and user interfaces are also designed during this phase. Importantly, security and privacy measures are built in from the beginning, not added later.

The third phase is the pilot implementation. From months nine to twelve, we install and test the system in the chosen pilot area. This is where theory meets reality. During this period, comprehensive training is provided to all levels of staff, from workers who receive alerts to managers who use the dashboards. The system runs under close human supervision. We carefully measure performance metrics such as detection accuracy, false alarm rates, and user acceptance. Based on the results, we fine-tune the sensors, retrain the AI models, and adjust alert settings before expanding further.

In phase four, we move to full-scale deployment. Between months thirteen and eighteen, the system is rolled out across the entire project site. This includes installing additional sensors and communication infrastructure and integrating SmartHSE with the project's existing management systems. Continuous monitoring and regular review meetings help maintain performance as the system scales up.

The final phase begins from month nineteen onwards. This is the stage of evaluation, continuous improvement, and scalability. We conduct regular assessments of key HSE indicators such as accident rates, response times, and near-miss frequency. Feedback from all users is collected systematically. Lessons learned are documented in a way that can be transferred to future projects. This phase ensures that SmartHSE does not remain static but continues to evolve and deliver value over the long term.

5.7 Expected Benefits and Impact Assessment

The implementation of SmartHSE is expected to deliver tangible improvements across several key areas. The most obvious is safety. The system helps prevent accidents before they happen, rather than reacting after they happen. Evidence from other industries that have adopted similar smart monitoring systems shows that accident rates can be reduced by 40 to 70 percent compared to traditional methods (Ashraf, 2024; Waqar et al., 2023). This is a significant change.

It should be noted that these figures originate from the oil, gas, and mining sectors, where operational conditions, workforce stability, and regulatory environments differ substantially from those in general building construction. Direct extrapolation to construction projects, particularly in developing country contexts, requires additional empirical validation.

In terms of health, continuous monitoring via IoT wearable sensors enables early detection of fatigue, heat stress, and other physical issues. Workers can receive timely warnings before their condition becomes dangerous. In terms of the environment, the system tracks greenhouse gas emissions, waste, and potential pollutants in real time. This helps reduce the overall impact of the project on air, soil, and water.

From an economic perspective, the benefits are also clear. Fewer incidents mean lower direct costs such as medical and compensation costs. There is also less downtime, better utilization of HSE staff, and potentially lower insurance premiums due to the accurate and consistent documentation the system provides. From an organizational perspective, SmartHSE reinforces a safety culture. This builds confidence in data-driven decisions and creates valuable knowledge that the company can use in future projects.

Table 5 .3:Expected Benefits of the SmartHSE Framework

Domain	Main Benefit	Expected Impact Level
Safety	Reduction of predictable accidents	Very High
Employee Health	Early detection of physical problems	High
Environment	Control of pollutants and emissions	Medium to High
Economic	Reduction of accident-related costs	High
Organisational	Strengthening safety culture	High

To understand how well the framework is actually working, a mixed evaluation approach was planned. We will track quantitative indicators such as Lost Time Injury Frequency Rate (LTIFR), Total Recordable Injury Rate (TRIR), Near-Incident Frequency and Hazard Response Time. Meanwhile, qualitative methods such as interviews with HSE officers, direct observation of system usage and review of management decisions provide deeper insights. Evaluations will be conducted at three key points: before launch, at the end of the pilot phase, and after full deployment. This timing allows us to see real progress step by step and addresses one of the major gaps we found during the literature review, which is the lack of appropriate longitudinal evaluation.

5.8 Implementation Barriers, Challenges, and Solutions

Implementation challenges for integrated sensor-AI-UAV systems in construction are well-documented in the reviewed literature, and SmartHSE is not exempt from them. This section maps those challenges against the specific architecture proposed and identifies mitigation strategies drawn from deployment experience in adjacent sectors. These challenges are categorized into five main categories. Each requires its own practical response strategy.

The technical problems are the most obvious. Different devices sometimes refuse to talk to each other. Drones have limited flight time. Managing the massive flow of data from multiple sensors can quickly become overwhelming. Fortunately, solutions already exist. Using open standards like MQTT and OPC-UA helps with compatibility. Combining drones with fixed sensors alleviates the problem of short flight time. Multi-tiered storage systems balance speed and capacity by processing immediate data locally and sending the rest to the cloud. The real task is to design carefully, not invent new technology.

Organizational and cultural barriers are often more difficult to overcome. Some workers see constant monitoring as surveillance rather than protection. Supervisors may be skeptical about alerts they don't fully understand. Project managers sometimes lack the training to act on complex risk predictions. To address this, we strongly believe in participatory change management. Involving workers and supervisors early in the design of alerts and dashboards creates a sense of ownership. Clear communication about real performance improvements helps build trust. Training should focus on understanding the logic of the system, not just pushing buttons.

Legal and regulatory challenges cannot be ignored. Rules for flying drones in urban areas are often strict. Questions about who is responsible when an AI recommendation leads to a problem are still unclear in many countries. Continuous monitoring also raises serious privacy concerns. The best approach is to involve regulators from the early design stage. Clear contracts that define responsibility for AI decisions are essential. Privacy should be built into the system from the start, not added as an afterthought.

Financial constraints are particularly difficult for contractors and smaller projects in developing countries. The initial investment in sensors, drones, and infrastructure can seem overwhelming. We recommend that deployments begin in phases in higher-risk areas. Renting equipment rather than purchasing it reduces initial costs. As early results prove the return on investment, it becomes easier to justify further expansion.

Ethical issues deserve special attention. There is always the risk that algorithms may treat certain groups of workers unfairly if the training data is biased. Continuous physiological monitoring also raises questions about dignity and privacy. The solution lies in explainable AI, ethical impact assessments before deployment, and the participation of different groups of workers in the experiment. The system should serve people, not reduce them to data points.

By recognizing these barriers early and planning responses to each, SmartHSE becomes more realistic and more likely to be successful in practice. Implementation is never just about technology. It is equally about people, organizations, and careful change management.

5.9 Chapter Conclusion

This chapter brings together the main contributions of my research. After reviewing the literature in Chapter 4 and noting the apparent integration gap, we developed the SmartHSE framework to address that exact weakness. The framework, in a coherent structure, shows how drones, IoT sensors, and AI can work together across the entire HSE risk management cycle, from identifying hazards to assessing them, controlling them, monitoring outcomes, and learning for the future.

The five-layer conceptual model, hybrid edge-cloud architecture, clearly defined roles for each technology, a practical five-step implementation roadmap, and honest discussion of barriers and solutions together create a usable reference. The framework is detailed enough to guide real-world application, but flexible enough to adapt to different project circumstances.

Of course, the SmartHSE framework provides an architectural reference against which future implementations can be designed and evaluated. Its value will ultimately be determined empirically: by whether accident rates decline, response times improve, and the system sustains adoption under real-world construction conditions. Does it reduce accidents? Does it improve health surveillance and environmental protection, as the evidence suggests? These questions remain for future empirical studies. What this chapter has done is to provide a solid architectural foundation upon which those practical tests can be properly designed and measured.

Chapter Six: Conclusions, Achievements, and Suggestions for Future Research

6.1 Research Summary and Main Findings

The starting point for this research was a pattern identified across the HSE technology literature: individual studies of drones, IoT sensors, and AI consistently report significant safety benefits, yet very few address how these technologies should function together as a system. The systematic review confirmed that this integration gap is real: more than 85 percent of the 73 included studies examined only a single technology, which has direct practical consequences for any organisation attempting to move beyond isolated tool deployment.

The literature review revealed an important issue. Individual technologies have shown promising results. Drones are effective for visual surveillance, IoT provides continuous data, and AI can predict hazards. However, most studies have only looked at one technology in isolation. The few that have attempted to combine them rarely provide practical details about how data flows or how decisions are made. This lack of integration became the main motivation for our work. The five key gaps we identified in Chapter 4—specifically the lack of a full-cycle framework, limited field studies, and weak attention to organizational factors—directly shaped the design of SmartHSE.

In Chapter 5, we introduced the SmartHSE framework as our main contribution. It is built on five layers and draws on systems theory, socio-technical systems thinking, ISO 31000 risk management, and technology acceptance models. The framework clearly defines the role of each technology. Drones act as mobile eyes for visual and spatial data. IoT sensors provide continuous environmental and physiological information. Artificial intelligence brings it all together, transforming raw data into useful predictions and recommendations. Most importantly, it creates a closed learning loop so that the system can improve itself over time.

This is no small feat. SmartHSE aims to change the fundamental logic of HSE management. Rather than reacting after incidents occur, the system tries to prevent them. It shifts from lagging indicators to leading indicators, from occasional inspections to continuous monitoring, and from subjective human judgment to calibrated, data-driven decisions. Whether this ambitious shift will be effective in real-world construction environments, with all their budget constraints, cultural realities, and daily pressures, remains an unanswered question. But this research has laid a solid foundation for future empirical studies to properly test it.

6.2 Research Achievements and Innovations

It is important to distinguish between what this research contributes directly and what it synthesises from existing work. The individual capabilities of UAVs, IoT, and AI in HSE contexts are established in the reviewed literature and are not claimed as original findings here. The original contribution of this thesis lies in three areas: the five-layer architectural specification of how these technologies interact within a closed-loop HSE risk management cycle; the explicit treatment of socio-technical and organisational factors as design requirements rather than implementation afterthoughts; and the phased implementation roadmap adapted to developing-country constraints, for which no comparable framework was identified in the reviewed literature.

This research contributes to the scientific community in several ways. Scientifically, SmartHSE is one of the first frameworks to bring together drones, IoT sensors, and AI into a single, complete system for construction HSE. Unlike previous studies that typically examined only one technology or provided vague ideas about integration, this framework provides enough architectural detail to guide real-world implementation. We built it by combining findings from systematic literature reviews, quantitative acceptance models, and established risk management standards. The result is not just another theoretical idea, but an evidence-based construct.

Another strength lies in its interdisciplinary nature. The framework does not simply borrow ideas from different fields. It truly integrates knowledge from HSE engineering, information technology, systems theory, risk management, and technology acceptance research. Each discipline directly shapes specific parts of the design. For example, socio-technical systems thinking ensured that human and organizational factors were considered from the outset, not as an afterthought. This approach makes the framework much more realistic for actual use in construction companies.

An attempt was also made to keep the framework practical. Many conceptual studies in this area focus only on technical possibilities and ignore real-world problems. SmartHSE takes the opposite route. It explicitly addresses technical, organizational, legal, financial, and ethical barriers and provides specific solutions for each. A five-stage implementation roadmap transforms the architecture into a step-by-step plan that project teams can follow. This balance between ambition and practicality is particularly valuable for engineering companies in Iran and other developing countries that are considering investing in smart HSE systems.

Finally, the framework provides a solid foundation for future empirical research. By providing a clear reference architecture, it gives other researchers a suitable model to test on real construction sites. Future studies can now investigate whether the integrated system actually reduces accident rates, measure the performance of specific algorithms, or investigate what organizational conditions lead to successful adoption. This research has moved the field from fragmented descriptions to integrated, testable solutions.

6.3 Research Limitations

Every study has its limitations, and this study is no exception. The most obvious limitation stemmed from the research methodology chosen. The systematic literature review was limited to English-language articles indexed in major international databases. This meant that I may have missed relevant studies published in Persian, Chinese, or other languages, or practical reports that never made it into academic journals. While the databases I searched were extensive, a truly comprehensive multilingual review would have been ideal.

The biggest limitation is the lack of real-world testing. SmartHSE is a conceptual framework. Its expected benefits are based on the integration of existing literature and rational architectural design, not on actual implementation data from construction sites. Promising figures cited, such as 40 to 70 percent reductions in accidents, have been derived from other industries, such as oil, gas, and mining. Whether these results will hold true in construction projects, especially in developing country settings, remains to be seen. Organizational challenges, particularly how workers and managers will adopt and use the system, are also only theoretically addressed at this stage.

Another important limitation is the high pace of technological development. Existing hardware and software will continue to evolve rapidly by the end of 2025. Some of the specific technical recommendations in the framework may become obsolete relatively quickly. I have tried to make the design flexible by focusing on functions and interfaces rather than specific products, but this cannot fully protect against rapid changes in algorithms and communication technologies.

6.4 Suggestions for Future Research

This research has opened more doors than it has closed. The gaps identified, along with the limitations of the study, suggest several clear directions for future work.

The most important need is field-scale experimental testing. We need to implement SmartHSE on real construction sites and measure its real-world impact on key metrics such as LTIFR, TRIR, near-failure rates, and hazard response times. Comparing the results with control sites or historical data will show whether the promised improvements hold true in practice. Without this kind of rigorous validation, the framework remains theoretical.

Another important area is the development of explainable AI models specifically designed for construction safety. Current deep learning systems often act like black boxes. Future studies should employ techniques such as LIME, SHAP, or Grad-CAM to help managers and workers understand why the system is issuing a particular warning. This directly addresses one of the major barriers to adoption that I found during my literature review.

We also recommend dedicated research on technology adoption. Using adapted versions of the TAM and UTAUT models, future studies should examine what actually influences workers' and supervisors' trust in and use of the system. Factors such as perceived usefulness, ease of learning, trust in automated warnings, and management support need to be carefully examined in our local cultural and educational context.

Economic questions deserve much more attention. Construction companies need robust cost-benefit analyses, realistic calculations of return on investment for projects of various sizes, and clear modeling of insurance and legal implications. Such studies, based on real-world performance data from pilot implementations, help organizations make informed investment decisions and reduce financial uncertainty.

Looking ahead, there are exciting opportunities for integration with other emerging technologies. Connecting SmartHSE to Digital Twins could allow real-time monitoring data to feed into dynamic project simulations. Blockchain could provide secure audit trails for safety records. Augmented reality could deliver drone-detected hazards directly to workers on site. SmartHSE's layered design makes these extensions possible, but each requires focused research to develop practical integration methods.

6.5 Final Conclusion

The persistence of high fatality and injury rates in construction, despite decades of accumulated safety knowledge and regulation, points to a structural mismatch: the management tools available are not designed for the operational conditions they are expected to address.

This thesis sought to bridge this gap. SmartHSE offers a different way forward by bringing together drones, IoT sensors, and AI into an integrated system. Drones extend visual reach to locations where human inspection would transfer risk to the inspector; IoT sensors provide the temporal continuity that aerial monitoring cannot achieve; and AI processes the combined output into predictions and warnings at a speed and scale that manual analysis cannot replicate. The framework I developed not only shows what each technology can do, but also how they can work together across the entire risk management cycle from early detection to learning and improvement.

For Iran and other developing countries, this is even more important. We are implementing large infrastructure projects, often in challenging conditions such as unstable electricity, dusty environments, and complex multi-contractor sites. Many of the existing solutions from developed countries do not fully match our reality. SmartHSE was designed with these practical limitations in mind. It is not a complete or final answer, but I believe it provides a realistic and consistent starting point for organizations, engineers, and policymakers who want to move safety management into the modern era.

The case for SmartHSE rests ultimately on an empirical claim: that an integrated system combining drone surveillance, continuous IoT monitoring, and AI-based analysis can reduce the frequency and severity of construction incidents in ways that isolated technologies cannot. Testing this claim on real construction sites in Iran and comparable developing-country contexts, where the specific challenges of unstable infrastructure and complex multi-contractor environments apply, is the most direct and useful contribution that future research can make.

7. Refrence:

Agapiou, A. (2024). A systematic review of the socio-legal dimensions of responsible AI and its role in improving health and safety in construction. *Buildings, 14*(5), 1469.

Akinsemoyin, A., Awolusi, I., Chakraborty, D., Al-Bayati, A. J., & Akanmu, A. (2023). Unmanned aerial systems and deep learning for safety and health activity monitoring on construction sites. *Sensors, 23*(15), 6690.

Albeaino, G., Brophy, P., Jeelani, I., Gheisari, M., & Issa, R. R. (2023). Impact of drone presence on construction individuals working at heights. *Journal of Construction Engineering and Management, 149*(11), 04023119.

Alsamarraie, M., Ghazali, F., Hatem, Z. M., & Flaih, A. Y. (2022). A review on the benefits, barriers of the drone employment in the construction site. *Jurnal Teknologi (Sciences & Engineering), 84*(2), 121-131.

Ashraf, M. (2024). Use of AI and robotics in project management. *International Journal of Science and Research*.

Cheng, J.-Y., Gheisari, M., & Jeelani, I. (2023). Using 360-degree virtual reality technology for training construction workers about safety challenges of drones. *Journal of computing in civil engineering, 37*(4), 04023018.

Chung, W. W. S., Tariq, S., Mohandes, S. R., & Zayed, T. (2023). IoT-based application for construction site safety monitoring. *International Journal of Construction Management, 23*(1), 58-74.

Daniel, E. I., Oshodi, O. S., Nwankwo, N. I., Emuze, F. A., & Chinyio, E. (2025). The use of digital technologies in construction safety: A systematic review. *Buildings, 15*(8), 1386.

Harvey, E. (2019). *Rethinking construction safety.* Loughborough University,

Hossain, M. I., Hosen, M. M., Sunny, M. A. U., & Tarapder, S. A. (2025). Implementing advanced technologies for enhanced construction site safety. *American Journal of Advanced Technology and Engineering Solutions, 1*(02), 01-31.

Jha, H., Mukherjee, A., Banerjee, S., Roy, T., Das, R., Middya, A. I., & Roy, S. (2025). Automated Concrete Structure Defect Detection Using Vision Transformer. *Available at SSRN 5096179*.

Kanellopoulos, D., Sharma, V. K., Panagiotakopoulos, T., & Kameas, A. (2023). Networking architectures and protocols for IoT applications in smart cities: Recent developments and perspectives. *Electronics, 12*(11), 2490.

Khan, A. M., Alrasheed, K. A., Waqar, A., Almujibah, H., & Benjeddou, O. (2024). Internet of things (IoT) for safety and efficiency in construction building site operations. *Scientific reports, 14*(1), 28914.

Musarat, M. A., Irfan, M., Alaloul, W. S., Maqsoom, A., & Ghufran, M. (2023). A review on the way forward in construction through industrial revolution 5.0. *Sustainability, 15*(18), 13862.

Mustapha, Z., Akomah, B. B., Kwaku, O. E., & Tieru, C. K. (2024). Enhancing safety management in the Ghanaian construction industry: evaluating the role of UAV implementation and construction managers understanding. *Journal of Civil Engineering Frontiers, 5*(02), 34-40.

Olatunde, N. A., Anugwo, I. C., Awodele, I. A., Ramabodu, M. S., & Adegoke, B. F. (2026). Unpacking trends in artificial intelligence research in the construction industry: a bibliographic review. *Frontiers in Built Environment, 12*, 1798096.

Palanisamy, P. (2018). *Hands-On Intelligent Agents with OpenAI Gym: Your guide to developing AI agents using deep reinforcement learning*: Packt Publishing Ltd.

Peng, J., Zhang, Q., Feng, Y., & Liu, X. (2023). Optimization of construction safety resource allocation based on evolutionary game and genetic algorithm. *Scientific reports, 13*(1), 17097.

Pilskog Orvik, C., Albrechtsen, E., Angvik, M., Sæther, M., & Holen, S. (2024). *Rethinking Safety Indicators in the Norwegian Construction Industry.* Paper presented at the IOP Conference Series: Earth and Environmental Science.

Raj, R., & Kos, A. (2026). Recent Developments and Applications of Drone Swarm: Techniques, Strategies, and Challenges. *Sensors, 26*(10), 2943.

Shafei, H., Radzi, A. R., Algahtany, M., & Rahman, R. A. (2022). Construction 4.0 technologies and decision-making: A systematic review and gap analysis. *Buildings, 12*(12), 2206.

Shah, I. A., & Mishra, S. (2024). Artificial intelligence in advancing occupational health and safety: an encapsulation of developments. *Journal of Occupational Health, 66*(1), uiad017.

Waqar, A., Othman, I., Hamah Sor, N., Alshehri, A. M., Almujibah, H. R., Alotaibi, B. S., . . . Hayat, S. (2023). Modeling relation among implementing AI-based drones and sustainable construction project success. *Frontiers in Built Environment, 9*, 1208807.