

AI-Driven Predictive HSE and Environmental Risk Management Framework for Sustainable Oil & Gas and Construction Projects in Saudi Arabia.

Abstract

Saudi Arabia's oil and gas and construction sectors are expanding and digitalising, while facing increased expectations for occupational safety, process safety, environmental protection, and sustainable project delivery. This review introduces an integrated, AI-driven predictive health, safety, and environment (HSE) and environmental risk management framework designed for these high-risk environments. Using a structured integrative approach, the review synthesises recent peer-reviewed evidence on machine learning, deep learning, computer vision, IoT sensing, digital twins, predictive maintenance, natural-language processing, and decision-support systems. The evidence shows that AI can shift HSE management from periodic inspections and lagging indicators to continuous risk sensing, probabilistic forecasting, dynamic prioritisation, and earlier intervention. In construction, research is most advanced in visual hazard recognition, PPE monitoring, injury prediction, and digital safety workflows.

In oil and gas, research has advanced asset integrity, pipeline risk assessment, methane monitoring, corrosion prediction, and incident analysis. However, technical accuracy alone does not guarantee safer or more sustainable outcomes. Ongoing challenges include data quality, class imbalance, interoperability, model drift, explainability, privacy, cybersecurity, governance, and human oversight. This review proposes a five-layer framework that integrates multisource data acquisition, data governance, predictive analytics, risk intelligence, and controlled intervention, with continuous feedback to improve models and organisational learning. The framework is tailored to Saudi conditions, addressing harsh climates, large-scale projects, complex contractor networks, dispersed assets, and the need to align HSE performance with environmental sustainability. The review concludes that AI should support, not replace, expert HSE judgement and must be embedded in auditable management systems, risk ownership structures, and regulatory assurance processes.

Keywords: artificial intelligence; predictive HSE; environmental risk; oil and gas; construction safety; Saudi Arabia; digital twins; machine learning; sustainability

1.Introduction

Health, safety, and environment (HSE) management in oil and gas and construction projects has traditionally focused on hazard identification, risk assessment, permit-to-work controls, inspections, audits, incident investigation, and corrective action. While these mechanisms remain essential, they are often episodic and rely heavily on human observation. In large-scale projects, risk conditions can evolve more rapidly than conventional reporting cycles. Dynamic interactions among workforce movement, simultaneous operations, temporary works, lifting, excavation, confined-space entry, heavy equipment, hydrocarbons, pressure systems, weather, dust, noise, and emissions further complicate risk management. The primary challenge is not only to identify known hazards but also to anticipate when multiple weak signals may converge, resulting in unsafe or environmentally damaging conditions.

Artificial intelligence (AI) offers a different logic of risk management. Rather than waiting for a threshold breach or incident, machine-learning models can combine historical and streaming data to estimate the probability, location, severity, or timing of adverse events. Construction research increasingly uses computer vision, deep learning, natural-language processing, smart sensing, and integrated digital delivery to recognise unsafe acts and conditions, quantify risk, and automate reporting [1–9]. Studies have also demonstrated injury-type prediction, design-stage fall-hazard identification, and deeper integration between BIM and AI-enabled construction management [10–13]. These developments suggest that predictive safety can become a continuous function embedded across the project lifecycle rather than a separate compliance activity.

The oil and gas context extends the problem from occupational safety to process safety, asset integrity, and environmental consequence. AI has been examined across upstream operations, where uncertainty, asset intensity, and data heterogeneity persist [17]. Data-driven asset integrity frameworks show how incident records and reliability variables can support integrity risk prioritisation [18]. Digital twins and machine-learning models have been applied to pipeline risk estimation [19], while recent work has analysed oil and gas accidents using AI and machine learning [20]. Environmental monitoring is also becoming more predictive: machine-learning calibration improves continuous methane sensing [21], and pipeline studies link environmental consequences to accident probability and remediation exposure [27]. These streams of

63 evidence indicate that HSE and environmental management should not be treated as separate digital
64 domains.

UNDER PEER REVIEW IN IJAR

Saudi Arabia provides a particularly relevant setting for integration. Major oil, gas, infrastructure, industrial, and urban-development programmes combine large contractor networks, complex logistics, high ambient temperatures, dust exposure, geographically dispersed facilities, and demanding schedules. These conditions amplify the value of real-time sensing but also create challenges in communications, interoperability, workforce acceptance, and model transferability. The national emphasis on digital transformation and sustainability also increases the strategic importance of systems that can jointly reduce injuries, unplanned loss of containment, environmental releases, waste, rework, downtime, and resource inefficiency.

The existing literature is fragmented across sectors, technologies, and risk types. Construction reviews frequently categorise AI applications by computer vision, text, audio, BIM, or robotics [1–4], whereas oil and gas studies typically focus on operational optimisation, integrity, pipelines, or emissions [17–21]. A comprehensive cross-sector HSE framework remains underdeveloped. This review addresses this gap by critically synthesising the evidence into a unified predictive-risk architecture appropriate for high-risk Saudi projects. The focus is not on positioning AI as an autonomous safety authority, but rather on delineating where predictive analytics can enhance existing risk controls, where human verification is essential, and how environmental and occupational risk signals can be jointly governed.

2.Aim of the Study

The study aims to develop a critically grounded, AI-driven predictive HSE and environmental risk management framework for sustainable oil and gas and construction projects in Saudi Arabia. It seeks to connect technologies often discussed in isolation—such as computer vision, machine learning, IoT sensing, digital twins, predictive maintenance, and natural-language processing—with established HSE decision processes. The intended outcome is a practical conceptual architecture that explains how data can be converted into auditable risk intelligence, how predictive alerts should be prioritised, and how human decision-makers can intervene before occupational, process-safety, or environmental losses occur.

3.Research Objectives

The review has five objectives. First, it evaluates how contemporary AI techniques are being used for predictive HSE and environmental risk detection in construction and oil and gas contexts. Second, it compares the data requirements, strengths, and limitations of major AI applications, including vision systems, predictive models, digital twins, sensing networks, and text analytics. Third, it identifies organisational, ethical, technical, and governance barriers that can prevent accurate models from producing reliable operational decisions. Fourth, it synthesises the evidence into an integrated framework tailored to high-risk Saudi project conditions. Fifth, it identifies research priorities for validating predictive HSE systems against real operational outcomes rather than model accuracy alone.

4.Research Questions

This review addresses four questions. RQ1: Which AI technologies and data sources provide the strongest current basis for predicting occupational, process-safety, and environmental risks in oil and gas and construction projects? RQ2: How can heterogeneous indicators be integrated into a common risk-intelligence workflow that supports early intervention? RQ3: Which technical, human, governance, and ethical constraints most affect the trustworthiness and transferability of AI-enabled HSE systems? RQ4: What design principles should guide an AI-driven predictive HSE and environmental framework for sustainable Saudi projects?

5.Review Methodology

A structured integrative review design was adopted because the research problem spans multiple industries, risk domains, and methodological traditions. The approach follows the transparent logic of systematic reviewing while allowing conceptual integration across empirical studies, technical reviews, and implementation-oriented research. The uploaded reference article informed only the manuscript architecture, particularly the separation of data sources, search strategy, selection criteria, and risk-based synthesis. We did not treat it as substantive evidence for this review, and we did not reproduce its wording, findings, tables, or figures.

The literature search focused on Scopus, Web of Science, ScienceDirect, SpringerLink, IEEE Xplore, PubMed (where relevant), and publisher databases. Search strings combined terms such as “artificial intelligence”, “machine learning”, “deep learning”, “computer vision”, “digital twin”, “predictive maintenance”, “construction safety”, “HSE”, “process safety”, “oil and gas”, “pipeline integrity”, “methane monitoring”, “environmental risk”, and “Saudi Arabia”. Boolean combinations were adjusted by database. We prioritised work published from 2020 to 2025, while retaining a small number of earlier or 2020 boundary studies where they provide directly relevant methodological foundations.

Inclusion criteria required peer-reviewed journal articles with a clear relationship to predictive or data-driven risk management; relevance to construction, energy infrastructure, oil and gas, industrial safety, or environmental monitoring; sufficient methodological detail to evaluate the technology or framework; and English-language availability. Studies were excluded when they were purely conceptual without a defensible method, focused on unrelated AI applications, duplicated another publication, or offered only generic commentary. We excluded non-peer-reviewed conference abstracts, editorials, and grey literature; full peer-reviewed journal or proceedings articles were eligible when methods and findings were sufficiently reported. Because this paper is an integrative review rather than a PRISMA-style meta-analysis, we report no fabricated record counts, screening totals, or pooled effect estimates.

The quality assessment considered six dimensions: clarity of the research question; transparency of data sources; appropriateness of the model or analytical method; adequacy of validation; relevance of outcome measures to HSE or environmental performance; and discussion of limitations or deployment conditions. Studies were not converted into a numerical quality score, because such scoring can create false precision across heterogeneous designs. Instead, we weighted studies more heavily that combined transparent data provenance, external or realistic validation, explainability, and operationally meaningful outcomes.

The synthesis used a risk-function matrix. Evidence was coded conceptually according to data source, AI method, predicted risk, decision horizon, intervention type, and implementation constraint. The analysis then compared construction and oil and gas applications to identify transferable mechanisms. For example, construction computer vision and oil-pipeline integrity analytics differ in sensors and timescales, but both require reliable data pipelines, calibrated risk thresholds, explainable alerts, and human escalation. This cross-sector logic informed the proposed framework.

Methodological limitations should be recognised. The review does not claim exhaustive capture of every AI-HSE paper, and rapidly evolving 2025 research may not yet have accumulated independent validation. Publication bias may favour successful models. Reported performance metrics are difficult to compare because datasets, class definitions, and operational environments differ. Many studies evaluate algorithmic accuracy rather than prospective reductions in injuries, releases, or losses. Consequently, the synthesis emphasises transferable design principles and evidential convergence rather than ranking individual algorithms.


165 **Table 1. AI technologies, predictive HSE functions and decision horizons**

Technology	Typical data	Primary predictive function	Decision horizon	Key limitation
Computer vision	Video/images	Unsafe acts, PPE, proximity, site hazards	Seconds–minutes	Occlusion, lighting, privacy
Supervised ML/DL	Incidents, exposure, operations	Event probability, severity, injury type	Hours–weeks	Label quality, imbalance, drift
IoT / environmental sensing	Gas, dust, heat, noise, weather	Exposure, leak and threshold prediction	Seconds–hours	Calibration, connectivity
Digital twins / prognostics	Condition, process, inspection	Failure and degradation forecasting	Days–months	Model fidelity, synchronisation
NLP / knowledge graphs	Reports, permits, procedures	Hazard extraction and compliance intelligence	Minutes–days	Language ambiguity, explainability

166

167

168 **Figure 1. Five-layer AI-driven predictive HSE and environmental risk management**
169 **framework**



170

171 Figure 1 integrates sensing, data governance, predictive analytics, risk intelligence and controlled
172 intervention. The feedback path ensures that verified field outcomes improve calibration and
173 organisational learning rather than leaving the model static.

6. Thematic Literature Review and Critical Analysis The first major theme is the transition from reactive reporting to predictive occupational safety. Deep-learning and machine-learning studies show that incident narratives, site imagery, and contextual variables can be converted into leading risk indicators. Liu et al. synthesise deep-learning safety analytics and highlight performance dependence on data representation and quality [2]. Ajayi et al. demonstrate that health-and-safety records can support classification and regression models for infrastructure risk prediction [25]. Alkaissy et al. similarly apply machine learning to construction injury classification, reinforcing the value of historical incident data when labels are sufficiently consistent [10]. The limitation is that historical datasets encode past reporting behaviour: under-reporting, inconsistent taxonomies, and changing work methods can distort future predictions. Saudi construction evidence also shows the enduring value of structured risk assessment for prioritising accident causes and preventive controls [16].

Computer vision is currently one of the most mature AI safety applications in construction. Rabbi and Jeelani show that vision, text, and audio applications collectively expand AI's role in safety management [1]. Hou et al. move beyond object detection by computing and visualising risk from spatial relationships between hazards and exposed objects [6]. Park et al. demonstrate multi-label recognition of material and human risk factors, illustrating the shift from single-object detection to simultaneous risk interpretation [15]. Wen et al. combine real-time detection with visual question answering and automated safety reporting [7]. These systems can reduce inspection latency, yet performance can degrade under occlusion, poor lighting, dust, camera movement, atypical PPE, and congested work fronts. High validation accuracy therefore should not be interpreted as proof of universal field reliability.

Synthetic data and multimodal methods offer one response to data scarcity. Kim and Yi show that text-to-image generation can create hazardous construction scenarios for training deep neural networks [8]. This is useful where serious hazards are rare and unsafe situations cannot ethically be staged. However, synthetic datasets may reproduce model assumptions rather than real-world variability. They should supplement, not replace, site-specific validation. Knowledge-based and text-oriented approaches provide a complementary route: Wang and El-Gohary use deep learning and knowledge graphs to structure safety requirements [9]. Such approaches can connect legal, procedural, and operational knowledge, which is particularly relevant for multilingual contractor environments and complex permit systems.

A second theme concerns digital integration. Safety applications are more useful when linked to BIM, IoT, edge computing, and project-delivery workflows rather than deployed as isolated tools. Maali et al. show that safety technologies span prevention through design, safety management, sensing, robotics, and computing [4]. Larbi et al. propose integrated digital delivery for construction HSE management [5]. Pan and Zhang describe the broader value of BIM-AI integration across smart construction management [13]. At the same time, Kochovski and Stankovski demonstrate an edge-to-cloud architecture that can support smart and safe construction applications [26]. These studies imply that predictive HSE requires an information architecture, not merely an algorithm.

The third theme is lifecycle prevention. Tözer et al. show that BIM can identify fall hazards earlier in design, moving intervention upstream [11]. Shafei et al. identify IoT, autonomous construction, big data and predictive analytics, AI, BIM, and immersive technologies as important Construction 4.0 capabilities for safety and health [29]. Dobrucali et al. link emerging technologies with multiple safety-performance dimensions, including planning, training, inspection, audits, and leadership [14]. Sarvari et al. contrast traditional Safety-I thinking, which focuses on failure prevention, with broader resilience-oriented approaches [12]. For Saudi megaprojects, this implies that AI should influence design reviews, constructability, temporary works, logistics, and method statements before risk reaches the workforce.

A fourth theme is environmental and process risk in oil and gas. Koroteev and Tekic describe AI's growing role in upstream oil and gas while emphasising data, people, and collaboration challenges [17]. Sattari et al. show how data-driven AI can strengthen asset-integrity decision-making by linking incident experience with integrity and reliability factors [18]. Priyanka et al. integrate digital twins, prognostics, and machine learning for oil-pipeline risk estimation [19]. Wasim and Djukic review pipeline corrosion mechanisms and predictive prevention, underscoring the importance of degradation mechanisms that evolve over long timescales [28]. These applications differ from construction vision analytics: they often require time-series data on inspection, corrosion, pressure, flow, and maintenance histories, with predictions supporting maintenance windows rather than immediate behavioural intervention.

Environmental risk extends predictive HSE beyond worker protection. Lu et al. quantify environmental risk from pipeline accidents by combining event probability with remediation consequence [27]. Andrews et al. demonstrate machine-learning calibration of methane sensors for continuous emissions monitoring [21]. These approaches are especially relevant to sustainability because they connect operational risk with greenhouse-gas emissions, soil and water contamination, and resource loss. Predictive environmental management should therefore combine equipment condition with meteorology, dispersion conditions, leak detection, and sensitive-receptor data rather than treating environmental compliance as periodic sampling alone.

The fifth theme is digital twins and predictive maintenance. Van Dinter et al. show that digital twins can support predictive maintenance but also face challenges in model complexity, data variety, computational burden, and standardisation [23]. Priyanka et al.'s oil-pipeline digital-twin study illustrates how virtual representations can combine operational data with prognostic models [19]. In construction and industrial projects, the same principle applies to cranes, temporary power systems, dewatering, ventilation, lifting equipment, fire systems, process packages, and environmental-control assets. However, a digital twin is only as trustworthy as the synchronisation between physical and virtual states. Stale sensor feeds or unrecorded field modifications can create a false sense of control.

The sixth theme is context sensitivity and governance. Datta et al. demonstrate that AI and machine learning are distributed unevenly across construction lifecycle stages [24]. At the same time, Baduge et al. highlight AI's dependence on sensing, data cleaning, storage, and model development [3]. Kim et al. show that natural-hazard indicators can strengthen construction safety prediction, which is significant where extreme weather alters exposure [30]. In Saudi Arabia, heat, wind, dust, remote operations, and large-scale logistics should therefore be explicit model variables rather than background conditions. Finally, Gazos et al. argue that AI safety must be treated socio-technically, because risk can emerge from interactions between algorithms, organisations, and technical systems [22]. This point is crucial: an alert that is accurate but poorly routed, poorly explained, or routinely overridden may have little preventive value.

7. Discussion

The literature supports a clear conclusion: predictive HSE is not a single model but a closed-loop decision system. Construction research demonstrates rapid hazard recognition and behavioural monitoring; oil and gas research demonstrates slower-horizon integrity, maintenance, and environmental prediction. Combining these perspectives enables a more complete architecture that monitors immediate hazards, degradation processes, and environmental consequences simultaneously. Figure 1 conceptualises this architecture as five connected layers: sensing, governed data, predictive analytics, risk intelligence, and controlled intervention.

At the sensing layer, the framework combines cameras, wearables, equipment telematics, gas and dust sensors, weather stations, inspection records, permits, incident narratives, BIM, geospatial data, and process-control streams. At the governance layer, raw signals are time-aligned, quality-checked, access-controlled, and mapped to common asset, location, task, and hazard taxonomies. This is essential because predictive accuracy is easily overstated when timestamps, asset identifiers, contractor labels, or incident categories are inconsistent. The analytics layer then applies fit-for-purpose methods: computer vision for visible hazards, supervised models for incident probability, anomaly detection for rare deviations, natural-language processing for reports and procedures, and digital twins for state estimation and degradation forecasting.

The fourth layer converts model output into risk intelligence, not raw probability. A safety manager does not need hundreds of unranked alerts; the system should combine likelihood, severity, exposure, control effectiveness, uncertainty, and environmental consequence. This transforms AI output into a dynamic risk register. The intervention layer then links alerts to defined actions: field verification, stop-work authority, maintenance work orders, permit review, exclusion-zone changes, environmental containment, inspection escalation, or management-of-change. Human approval should remain mandatory where actions can materially affect operations, workforce rights, or safety-critical plant.

Table 1 shows that different AI technologies fit different risk horizons. Computer vision is strongest for seconds-to-minutes detection, while digital twins and predictive maintenance can operate over days or months. NLP and knowledge graphs are particularly useful for connecting procedures, incidents, and regulatory requirements, whereas environmental sensors provide continuous evidence on emissions and exposure. The framework therefore rejects a “one-model” approach. A portfolio of models is more defensible because each technology addresses a different uncertainty structure.

Sustainability strengthens rather than competes with HSE. Preventing leaks avoids both safety events and environmental loss. Predictive maintenance can reduce catastrophic failure while also extending asset life and avoiding unnecessary replacement. Better logistics and equipment analytics can reduce idling, fuel use, and exposure to vehicle interactions. Vision-supported housekeeping can identify waste accumulation and blocked access at the same time. This shared-risk perspective is particularly valuable on Saudi projects where construction, commissioning, operation, and environmental management often overlap.

The proposed Saudi implementation model should be phased. Initial deployment should focus on high-consequence, data-rich use cases such as mobile-equipment interaction, work at height, PPE compliance, heat exposure, gas detection, methane leakage, critical rotating equipment, and pipeline integrity. Each use case should have a documented baseline, a clearly defined decision owner, an alert threshold, a verification process, and an outcome metric. Measure success through earlier intervention, reduced exposure duration, avoided releases, improved leading indicators, and verified control effectiveness—not only model precision.

Trust is the decisive implementation condition. Workers and supervisors must understand what is monitored, why it is monitored, and how predictions are used. Excessive surveillance can damage reporting culture if AI is perceived primarily as a disciplinary tool. Conversely, opaque automation can create overreliance. Explainable outputs, visible uncertainty, escalation rules, and periodic human challenge are therefore necessary. Cybersecurity also belongs inside HSE governance because compromised cameras, sensors, or digital twins can generate false risk states or suppress genuine alarms. In safety-critical environments, data integrity is a safety barrier in itself.

Figure 2 presents the operational feedback loop. Predictions trigger verification and intervention; the resulting field outcome is then captured as new evidence. Confirmed hazards, false alarms, near misses, maintenance findings, and environmental measurements return to the learning layer. This feedback enables recalibration and drift detection. Without it, models gradually become disconnected from changing equipment, contractors, work methods, seasons, and project phases. The most sustainable architecture is therefore not a static AI product but a governed learning system embedded in management review.

A further design rule is to distinguish prediction confidence from risk severity. A low-confidence signal may still justify verification when the potential consequence is catastrophic, while a high-confidence signal about a minor deviation may require only routine correction. Escalation matrices should make this asymmetry explicit so model uncertainty does not inadvertently suppress precautionary action.

Model governance should also define a safe fallback mode. If connectivity fails, sensors drift, or an analytics service becomes unavailable, established HSE barriers, alarms, inspections, and permit controls must continue to function independently. AI should add resilience, not become a single point of failure.

Cross-project learning can be improved through common event taxonomies and shared definitions of exposure, control status, near miss, high-potential event, release, and environmental consequence. Standardisation would make model performance more comparable and support learning across contractors without forcing identical technology stacks.

For environmental prediction, consequence modelling should include receptor sensitivity in addition to release magnitude. The same leak can have different significance depending on proximity to people, groundwater, marine systems, protected habitats, or ignition sources. Contextual consequence data therefore belong inside the predictive risk model rather than in a separate post-event assessment.

Finally, leading indicators should support action. A dashboard that increases alert volume without clarifying ownership, urgency, and expected response can increase cognitive load. The practical objective is fewer, better-prioritised signals that direct attention to deteriorating controls before loss occurs.

354 **Table 2. Priority Saudi use cases and governance requirements**

Use case	Predictive inputs	Operational action	Governance requirement
Work at height and mobile equipment	Vision, geofencing, task data	Verify hazard; stop or reconfigure work	Privacy, false-alarm control, clear authority
Heat and exposure risk	Weather, wearables, workload, shift data	Rest, hydration, rescheduling	Worker consent, clinical boundaries
Gas / methane release	Gas sensors, process data, weather	Isolation, leak survey, containment	Calibration, alarm integrity, cyber resilience
Rotating equipment / integrity	Vibration, temperature, maintenance, inspection	Inspection or maintenance work order	Model drift, asset configuration control
Pipeline environmental risk	Integrity, pressure, corrosion, receptor data	Inspection, pressure management, emergency readiness	Traceable consequence model, geospatial accuracy


355

356

357

358

359 **Figure 2. Operational feedback and model-governance loop**



Continuous learning loop: model drift and changing project conditions are managed through verified operational feedback

360

361 Figure 2 emphasises that prediction must be followed by human verification, intervention and
362 outcome capture. Recalibration closes the loop and enables detection of model drift as project
363 phases, assets and contractors change.

364 **8.Research Gaps and Future Research**

365 The first research gap is outcome validation. Many studies report accuracy, precision, recall, F1
366 score, or error reduction, but few prospectively test whether deployment reduces serious injuries,
367 process-safety events, uncontrolled releases, or high-potential near misses. Future Saudi studies
368 should use controlled pilots, interrupted time-series designs, or matched-project comparisons
369 where feasible, while protecting workers from experimental exposure.

The second gap is integrated HSE-environment modelling. Occupational safety, process integrity, and environmental monitoring are usually studied separately even though they share causal pathways. Future research should develop multi-hazard models that connect equipment degradation, operating conditions, worker exposure, meteorology, and environmental receptors. Physics-informed machine learning is promising where purely data-driven models lack enough failure examples.

Third, transferability remains weak. Models trained on one contractor, camera angle, asset class, or climate can fail elsewhere. Research is needed on domain adaptation, federated learning, uncertainty quantification, and Saudi-specific benchmark datasets that preserve privacy while representing heat, dust, multilingual workforces, remote assets, and local construction practices. Synthetic data may help, but must be validated against field distributions [8].

Fourth, explainability and governance require stronger empirical evidence. Future studies should test which explanation formats actually improve supervisor decisions and reduce inappropriate overrides. Research should also examine algorithmic accountability, worker consent, data retention, cyber resilience, and how AI alerts interact with formal HSE responsibilities. Finally, cost-benefit research should quantify whether predictive systems deliver net value after accounting for sensors, connectivity, integration, specialist staff, false alarms, retraining, and lifecycle maintenance.

9. Practical, Managerial and Policy Implications

For practitioners, the framework suggests beginning with risk ownership rather than technology procurement. Each predictive use case should be attached to a hazard register, critical control, responsible role, and response procedure. Data engineers and HSE professionals should jointly define labels and thresholds so that models reflect operational meaning. Field verification must be easy to record, because confirmation data are essential for recalibration.

For managers, AI governance should sit within existing assurance structures. Managers should review model performance, false-alarm rates, missed hazards, data quality, cyber incidents, and override patterns alongside conventional HSE indicators. Procurement specifications should require interoperability, audit logs, role-based access, decision-appropriate explainability, and exportable data. Contractors should not be incentivised to suppress adverse data; instead, governance should reward early detection and transparent reporting.

For policymakers and major clients, the priority is to create assurance expectations without prescribing one algorithm. Guidance can require traceable data provenance, documented validation, human oversight, cybersecurity controls, privacy protection, and periodic revalidation when site conditions change. Shared taxonomies and interoperable reporting standards would improve cross-project learning. Public or consortium-led benchmark datasets, appropriately anonymised, could accelerate model validation while reducing duplicated effort. Policy should also recognise that AI-generated alerts do not transfer legal or managerial accountability away from competent duty holders.

10.Limitations

This review has several limitations. It synthesises heterogeneous studies whose datasets, sectors, metrics, and validation procedures are not directly comparable. The evidence base is stronger for construction safety detection than for integrated environmental prediction, and some oil and gas AI literature focuses on operational optimisation rather than HSE outcomes. The review is conceptual and does not statistically pool effects. It also cannot establish that reported laboratory or pilot performance will persist across Saudi sites. Finally, the proposed framework requires empirical validation through real projects, including assessment of human factors, false-alarm burden, cybersecurity, cost, and regulatory acceptability.

418 11.Conclusion

419 AI can materially strengthen HSE and environmental risk management when it is used to
420 anticipate changing risk rather than automate compliance paperwork. The reviewed evidence
421 shows complementary capabilities: computer vision and multimodal AI can detect unsafe
422 conditions in near real time; machine learning can extract predictive value from incidents and
423 operational histories; digital twins can forecast degradation and maintenance needs; and
424 environmental sensors can identify emissions or releases earlier. Yet these capabilities remain
425 fragmented unless they are connected through common data governance, risk prioritisation,
426 human verification, and feedback.

427 The framework developed in this review integrates those functions into five layers: multisource
428 sensing, governed data, predictive analytics, risk intelligence, and controlled intervention. Its
429 central principle is that AI output must be translated into an auditable HSE decision with a
430 defined owner and feedback pathway. For Saudi oil and gas and construction projects, this
431 architecture is particularly relevant because high-risk work, harsh environmental conditions,
432 large contractor ecosystems, and ambitious sustainability targets coexist. The framework
433 therefore links occupational safety, process integrity, and environmental protection instead of
434 optimising them separately.

435 Future value will depend less on marginal gains in benchmark accuracy and more on trustworthy
436 deployment. Saudi organisations should validate models under local conditions, monitor drift,
437 protect privacy, secure data pathways, and maintain competent human oversight. Research
438 should focus on prospective outcome evidence, integrated multi-hazard models, explainable
439 decision support, and interoperable datasets. Under these conditions, AI can support a transition
440 from reactive HSE management toward predictive, resilient, and sustainability-oriented risk
441 governance without displacing the professional accountability on which safe projects ultimately
442 depend.

Assess implementation maturity across reliability, organisational readiness, and risk-control effectiveness. Reliability covers sensor health, calibration, cybersecurity, and model drift. Organisational readiness covers competence, ownership, escalation discipline, contractor integration, and worker acceptance. Risk-control effectiveness asks whether predictions shorten exposure, strengthen controls, prevent releases, and improve maintenance decisions. Saudi assurance should evaluate the sociotechnical chain from data capture to control recovery rather than treating validation accuracy as sufficient evidence of operational value.

References

- [1] Rabbi, A.B.K., & Jeelani, I. (2024). AI integration in construction safety: Current state, challenges, and future opportunities in text, vision, and audio based applications. *Automation in Construction*, 164, 105443. <https://doi.org/10.1016/j.autcon.2024.105443>
- [2] Liu, J., Luo, H., & Liu, H. (2022). Deep learning-based data analytics for safety in construction. *Automation in Construction*, 140, 104302. <https://doi.org/10.1016/j.autcon.2022.104302>
- [3] Baduge, S.K., Thilakarathna, S., Perera, J.S., Arashpour, M., Sharafi, P., Teodosio, B., Shringi, A., & Mendis, P. (2022). Artificial intelligence and smart vision for building and construction 4.0: Machine and deep learning methods and applications. *Automation in Construction*, 141, 104440. <https://doi.org/10.1016/j.autcon.2022.104440>
- [4] Maali, O., Ko, C.-H., & Nguyen, P.H.D. (2024). Applications of existing and emerging construction safety technologies. *Automation in Construction*, 158, 105231. <https://doi.org/10.1016/j.autcon.2023.105231>
- [5] Larbi, J.A., Tang, L.C., Larbi, R.A., Abankwa, D.A., & Danquah, R.D. (2024). Developing an integrated digital delivery framework and workflow guideline for construction safety management in a project delivery system. *Safety Science*, 175, 106486. <https://doi.org/10.1016/j.ssci.2024.106486>
- [6] Hou, X., Li, C., & Fang, Q. (2023). Computer vision-based safety risk computing and visualization on construction sites. *Automation in Construction*, 156, 105129. <https://doi.org/10.1016/j.autcon.2023.105129>

- [7] Wen, S., Park, M., Tran, D.Q., Lee, S., & Park, S. (2024). Automated construction safety reporting system integrating deep learning-based real-time advanced detection and visual question answering. *Advances in Engineering Software*, 198, 103779. <https://doi.org/10.1016/j.advengsoft.2024.103779>
- [8] Kim, H., & Yi, J.-S. (2024). Image generation of hazardous situations in construction sites using text-to-image generative model for training deep neural networks. *Automation in Construction*, 166, 105615. <https://doi.org/10.1016/j.autcon.2024.105615>
- [9] Wang, X., & El-Gohary, N. (2023). Deep learning-based relation extraction and knowledge graph-based representation of construction safety requirements. *Automation in Construction*, 147, 104696. <https://doi.org/10.1016/j.autcon.2022.104696>
- [10] Alkaissy, M., Arashpour, M., Golafshani, E.M., Hosseini, M.R., Khanmohammadi, S., Bai, Y., & Feng, H. (2023). Enhancing construction safety: Machine learning-based classification of injury types. *Safety Science*, 162, 106102. <https://doi.org/10.1016/j.ssci.2023.106102>
- [11] Tözer, K.D., Gürcanlı, G.E., Çelik, T., & Akboğa Kale, Ö. (2024). Safer designs with BIM based fall hazards identification and accident prevention. *Safety Science*, 169, 106346. <https://doi.org/10.1016/j.ssci.2023.106346>
- [12] Sarvari, H., Edwards, D., Rillie, I., & Posillico, J.J. (2024). Building a safer future: Analysis of studies on Safety-I and Safety-II in the construction industry. *Safety Science*, 178, 106621. <https://doi.org/10.1016/j.ssci.2024.106621>
- [13] Pan, Y., & Zhang, L. (2023). Integrating BIM and AI for smart construction management: Current status and future directions. *Archives of Computational Methods in Engineering*, 30, 1081–1110. <https://doi.org/10.1007/s11831-022-09830-8>
- [14] Dobrucali, E., Demirkesen, S., Sadikoglu, E., Zhang, C., & Damci, A. (2024). Investigating the impact of emerging technologies on construction safety performance. *Engineering, Construction and Architectural Management*, 31(3), 1322–1347. <https://doi.org/10.1108/ECAM-07-2022-0668>
- [15] Park, J., Seong, S., Park, S., Kim, M., & Kim, H.Y. (2024). Multi-label material and human risk factors recognition model for construction site safety management. *Journal of Safety Research*, 91, 354–365. <https://doi.org/10.1016/j.jsr.2024.10.002>

- [16] Sanni-Anibire, M.O., Mahmoud, A.S., Hassanain, M.A., & Salami, B.A. (2020). A risk assessment approach for enhancing construction safety performance. *Safety Science*, 121, 15–29. <https://doi.org/10.1016/j.ssci.2019.08.044>
- [17] Koroteev, D., & Tekic, Z. (2021). Artificial intelligence in oil and gas upstream: Trends, challenges, and scenarios for the future. *Energy and AI*, 3, 100041. <https://doi.org/10.1016/j.egyai.2020.100041>
- [18] Sattari, F., Lefsrud, L., Kurian, D., & Macciotta, R. (2022). A theoretical framework for data-driven artificial intelligence decision making for enhancing the asset integrity management system in the oil & gas sector. *Journal of Loss Prevention in the Process Industries*, 74, 104648. <https://doi.org/10.1016/j.jlp.2021.104648>
- [19] Priyanka, E.B., Thangavel, S., Gao, X.-Z., & Sivakumar, N.S. (2022). Digital twin for oil pipeline risk estimation using prognostic and machine learning techniques. *Journal of Industrial Information Integration*, 26, 100272. <https://doi.org/10.1016/j.jii.2021.100272>
- [20] Khan, A., Al Busafi, M., Al Thuhli, A., & Siddiqui, S. (2024). Accident and incident analysis in the oil and gas sector using artificial intelligence and machine learning. *Procedia Computer Science*, 251, 295–302. <https://doi.org/10.1016/j.procs.2024.11.113>
- [21] Andrews, B., Chakrabarti, A., Dauphin, M., & Speck, A. (2023). Application of machine learning for calibrating gas sensors for methane emissions monitoring. *Sensors*, 23(24), 9898. <https://doi.org/10.3390/s23249898>
- [22] Gazos, A., Kahn, J., Kusche, I., Büscher, C., & Götz, M. (2025). Organising AI for safety: Identifying structural vulnerabilities to guide the design of AI-enhanced socio-technical systems. *Safety Science*, 184, 106731. <https://doi.org/10.1016/j.ssci.2024.106731>
- [23] van Dinter, R., Tekinerdogan, B., & Catal, C. (2022). Predictive maintenance using digital twins: A systematic literature review. *Information and Software Technology*, 151, 107008. <https://doi.org/10.1016/j.infsof.2022.107008>
- [24] Datta, S.D., Islam, M., Sobuz, M.H.R., Ahmed, S., & Kar, M. (2024). Artificial intelligence and machine learning applications in the project lifecycle of the construction industry: A comprehensive review. *Heliyon*, 10(5), e26888. <https://doi.org/10.1016/j.heliyon.2024.e26888>

- [25] Ajayi, A., Oyedele, L., Owolabi, H., Akinade, O., Bilal, M., Davila Delgado, J.M., & Akanbi, L. (2020). Deep learning models for health and safety risk prediction in power infrastructure projects. *Risk Analysis*, 40(10), 2019–2039. <https://doi.org/10.1111/risa.13425>
- [26] Kochovski, P., & Stankovski, V. (2021). Building applications for smart and safe construction with the DECENTER Fog Computing and Brokerage Platform. *Automation in Construction*, 124, 103562. <https://doi.org/10.1016/j.autcon.2021.103562>
- [27] Lu, H., Xi, D., & Qin, G. (2023). Environmental risk of oil pipeline accidents. *Science of the Total Environment*, 874, 162386. <https://doi.org/10.1016/j.scitotenv.2023.162386>
- [28] Wasim, M., & Djukic, M.B. (2022). External corrosion of oil and gas pipelines: A review of failure mechanisms and predictive preventions. *Journal of Natural Gas Science and Engineering*, 100, 104467. <https://doi.org/10.1016/j.jngse.2022.104467>
- [29] Shafei, H., Rahman, R.A., & Lee, Y.S. (2025). Evaluating Construction 4.0 technologies in enhancing safety and health: Case study of a national strategic plan. *Journal of Engineering, Design and Technology*, 23(4), 1211–1242. <https://doi.org/10.1108/JEDT-08-2023-0330>
- [30] Kim, J.-M., Yum, S.-G., Adhikari, M.D., & Bae, J. (2024). A deep-learning approach to leveraging natural hazard indicators for improved safety on construction sites. *Safety Science*, 177, 106596. <https://doi.org/10.1016/j.ssci.2024>