© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723505

Predictive Risk Intelligence and Early-Warning Systems for Saudi Giga-Projects: Strengthening Business Continuity and Adaptive Crisis Response

QAMAR SHAHZAD

Abstract- ***Saudi Arabia's giga-project programme concentrates exceptional capital, technological novelty, inter-organisational dependence and delivery pressure within projects that must remain operationally credible while their scope and environment continue to evolve. Conventional risk registers are necessary but insufficient in such settings because they largely describe known exposures at discrete review points. This review develops a predictive risk intelligence perspective in which heterogeneous project data are continuously converted into weak signals, forecast risk states and decision triggers that support business continuity before disruption becomes crisis. Thirty peer-reviewed studies published between 2020 and 2025 were synthesised across construction risk analytics, digital twins, real-time monitoring, megaproject crisis precursors, supply-chain resilience and organisational resilience. The synthesis shows that effective early warning depends less on a single algorithm than on a socio-technical architecture connecting reliable data, explainable models, threshold governance, escalation authority and pre-agreed continuity responses. Machine learning can improve prediction of delay, systemic risk and cyber exposure, while digital twins and sensor systems can provide near-real-time situational awareness. However, fragmented data ownership, model drift, weak interpretability and organisational reluctance to act on uncomfortable signals can neutralise technical capability. A five-layer Predictive Risk Intelligence and Adaptive Continuity framework is therefore proposed for Saudi giga-projects. It links sensing, risk inference, warning governance, continuity activation and organisational learning. The paper argues that predictive intelligence should be treated as a governance capability rather than a dashboard feature, with performance judged by decision lead time, false-warning control, response effectiveness and recovery adaptability.***

Keywords: predictive risk intelligence; early-warning systems; Saudi giga-projects; business continuity; crisis response; digital twins; machine learning; organisational resilience; Vision 2030

I. INTRODUCTION

Saudi Arabia's Vision 2030 development portfolio has created an infrastructure environment in which projects are not simply large construction schemes but interconnected economic, technological and institutional systems. Giga-projects combine transport, utilities, tourism, real estate, digital infrastructure and public services while depending on extensive contractor networks, international supply chains and evolving regulatory arrangements. Research on Saudi construction already documents persistent exposure to schedule and cost overrun risks, while studies of NEOM identify contractual, technological, financial, human-resource and operational uncertainties that interact across project boundaries [20,21]. These characteristics increase the probability that a local deviation will propagate through interfaces and become a programme-level disruption.

Traditional project risk management generally identifies threats, estimates probability and impact, allocates ownership and records mitigation actions. That architecture remains valuable, but its review-cycle orientation is poorly matched to conditions in which risk changes faster than reporting cadence. Schedule deterioration may first appear as small productivity variance; supply instability may be visible in purchase-order ageing before material shortage; cyber vulnerability may emerge through anomalous access patterns; and stakeholder conflict may surface through delayed approvals or unresolved design queries. The managerial problem is therefore not merely identifying risks, but detecting changes in risk trajectory early enough to preserve options.

Predictive risk intelligence is defined here as the organisational capability to integrate historical, current and contextual data in order to estimate

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723505

emerging risk states, explain the drivers of deterioration, and translate forecasts into governed action. It differs from ordinary risk identification because it is forward-looking and continuously refreshed. It differs from early warning because warning is only one output: predictive intelligence also supports scenario interpretation, response selection and learning. It differs from business continuity management, which focuses on maintaining critical functions through disruption, although the two must be connected. Crisis management concerns command and coordination once disruption is imminent or realised, whereas adaptive resilience describes the capacity to absorb shocks, reconfigure resources and improve through experience.

Recent research provides components of this capability but remains fragmented. Digital-twin scholarship emphasises synchronised cyber-physical representation and the possibility of proactive control [1–4]. Machine-learning studies show that delays, risk interactions and cyber exposures can be predicted from project data [5–7]. Early-warning research identifies weak signals and crisis precursors, yet also shows that management may notice warnings without acting adequately [8–11]. Resilience studies emphasise network reconfiguration, organisational adaptation and supply-chain integration [12–15]. The unresolved issue is how these strands can be assembled into an operating model suitable for complex Saudi giga-projects.

Accordingly, this review aims to develop an integrated framework for predictive risk intelligence and early warning that strengthens business continuity and adaptive crisis response in Saudi giga-project environments. The objectives are to clarify the conceptual relationship among predictive analytics, weak signals, warning thresholds, continuity and resilience; synthesise recent evidence on enabling technologies and governance practices; identify implementation barriers relevant to multi-project Saudi programmes; and propose a decision-centred architecture that converts data into timely, accountable intervention. The contribution is not another catalogue of risks. It is a synthesis of how risk information can move from sensing to prediction, from prediction to escalation, and from escalation to continuity action before performance loss becomes systemic crisis.

II. REVIEW METHODOLOGY

The study adopts a structured integrative review suitable for an interdisciplinary topic spanning project management, construction informatics, safety, resilience and crisis studies. The review boundary was restricted to peer-reviewed work published from 2020 through 2025 so that the synthesis reflects contemporary digital risk-management capability. Search concepts combined terms related to construction or infrastructure projects with predictive risk, machine learning, digital twins, real-time monitoring, early warning, weak signals, crisis precursors, business continuity, supply-chain resilience and organisational resilience. Saudi-specific studies were deliberately included to anchor the synthesis in local project conditions, while international evidence was retained where it addressed mechanisms transferable to giga-project delivery.

Studies were included when they contributed evidence or a defensible conceptual mechanism for at least one stage of the risk-intelligence chain: sensing, prediction, interpretation, warning, response or learning. Purely descriptive technology papers with no project-management relevance were excluded, as were studies outside the specified publication period. The final set comprised thirty references. Sources were read comparatively rather than treated as independent endorsements of a single model. Evidence was coded into six themes: risk signal formation; predictive modelling; digital-twin and sensor integration; warning governance; continuity and crisis response; and resilience learning. Particular attention was given to limitations reported by the studies because deployment quality in giga-projects depends on data reliability, human interpretation and institutional authority as much as analytical accuracy.

III. FROM STATIC RISK REGISTERS TO PREDICTIVE RISK INTELLIGENCE

Risk registers compress uncertainty into named items, owners and ratings. Their weakness is not conceptual simplicity but temporal discontinuity. A register can show that schedule delay is “high”, yet reveal little about whether the exposure is stabilising, accelerating or migrating to another interface. Predictive

intelligence introduces trajectory. It asks which observable variables are changing, how strongly they precede undesirable outcomes, and how much decision time remains before available responses narrow.

This distinction is visible in recent construction analytics. Egwim et al. demonstrate that ensemble machine-learning models can estimate delay risk from influential project factors, moving risk assessment from retrospective explanation toward forecast classification [6]. Alsulamy's Saudi study similarly compares gradient-boosting approaches for construction delay prediction and reports strong predictive performance across time-overrun categories [22]. Moussa et al. extend the logic by modelling interactions and systemic risks rather than isolated risk items, an important advance because infrastructure losses are often produced by dependencies between schedule, resources, interfaces and organisational decisions [5]. These studies imply that risk intelligence should model relationships and direction of change, not simply update probability scores.

A predictive system, however, is only as reliable as its information architecture. Construction data are heterogeneous: programme schedules, earned-value measures, procurement status, design changes, requests for information, quality defects, safety observations, sensor streams, commercial correspondence and external indicators may use different identifiers and update frequencies. Digital-twin research addresses this fragmentation by linking physical assets, process states and digital information. Boje et al. argue for a semantic construction digital twin that progresses beyond static BIM toward sensor integration and AI-enabled reasoning [1]. Sacks et al. similarly frame digital-twin construction as a data-centric management mode in which site information supports proactive analysis and optimisation [2]. Opoku et al. find that construction digital-twin applications remain uneven across lifecycle stages, highlighting the gap between conceptual potential and operational integration [3].

For Saudi giga-projects, the implication is that predictive intelligence should sit above, not replace, established control systems. A common data layer should reconcile project coding structures and preserve traceability from raw observation to risk inference. The analytics layer can then detect anomalies, classify emerging states and estimate downstream consequences. The managerial layer must translate these outputs into decisions. Without that chain, organisations risk creating visually sophisticated dashboards that reproduce the same lagging information already available in monthly reports.

IV. WEAK SIGNALS AND EARLY-WARNING LOGIC

Early warning is the bridge between prediction and action. The relevant signal is rarely a dramatic threshold breach; it is often a weak pattern whose significance becomes visible only when several indicators are interpreted together. Research on megaproject crises identifies information uncertainty, participant capability, regulatory rigidity and benefit-cost imbalance as precursors that can be observed before full crisis manifestation [9]. In building projects, empirical work on early-warning signs similarly shows that project failure is preceded by recognisable managerial and performance symptoms [10]. Fang et al. demonstrate that time-sensitive social risk associated with large construction projects can be modelled using long short-term memory networks, illustrating how weak trends can be converted into forward signals [11].

Weak-signal research also exposes a governance problem: detection does not guarantee response. Pasman shows that serious incidents may be preceded by signals that are recognised but insufficiently acted upon because leadership, culture or organisational structure blocks escalation [8]. This finding is crucial for giga-projects. When commercial, political or schedule pressure is high, teams may normalise adverse variance, reinterpret warnings as temporary noise, or delay escalation to avoid reputational consequences. An early-warning system therefore requires explicit rules defining who can issue a warning, what evidence is required, who must respond, and when a risk must be elevated across organisational boundaries.

Warning design should use multiple levels rather than a binary alarm. An observation level can flag unusual

movement without requiring intervention. A caution level can trigger validation and scenario analysis. A warning level can activate agreed mitigations and increase reporting frequency. A critical level can invoke continuity arrangements, executive command and resource reallocation. Thresholds should combine absolute limits with rate-of-change measures because a modest deviation that is accelerating may be more dangerous than a larger but stable variance. For interdependent risks, composite indicators should also capture correlation and propagation potential.

False alarms and missed alarms must be treated as governance metrics. Excessive warnings create alert fatigue and erode trust, whereas conservative thresholds preserve dashboard credibility at the cost of decision time. Model performance should therefore be judged by operational consequences, not accuracy alone. Useful measures include warning lead time, precision of critical alerts, percentage of warnings receiving action within target time, avoided disruption, and the proportion of post-event root causes that had detectable precursors. This shifts evaluation from "Was the model correct?" to "Did the system create a better decision while options remained available?"

Table 1. Early-warning dimensions for giga-project risk intelligence

Dimension	Leading evidence	Analytical treatment	Governed response
Schedule	Constraint ageing, productivity trend, float erosion	Trend, classification, scenario simulation	Replan sequence; protect critical interfaces
Supply chain	PO ageing, factory progress, logistics disruption	Supplier risk score; anomaly detection	Alternate sourcing; expedite; reserve logistics
Design/interface	RFI backlog, change density, unresolved clashes	Network dependency and text analytics	Interface workshop; freeze/release priorities
Safety/environment	Near misses, sensor anomalies, work density	Computer vision; threshold and sequence models	Stop/modify work; increase controls and supervision
Cyber-physical	Access anomalies, device faults, unusual traffic	ML-based cyber risk assessment	Isolate systems; invoke cyber continuity plan

V. DIGITAL TWINS, SENSORS AND REAL-TIME SITUATIONAL AWARENESS

Predictive intelligence depends on sufficiently current evidence. Real-time monitoring research shows how cameras, location systems, wearables, laser scanning, drones and environmental sensors can create continuous representations of construction conditions [16]. Rao et al. categorise sensor methods for mapping, positioning, tracking and hazard identification, but also emphasise implementation challenges in real-world sites [16]. Shanti et al. demonstrate drone and deep-learning methods for detecting work-at-height hazards in near real time [17]. Tian and Meng integrate neural-network prediction with BIM visualisation to provide intelligent warnings for excavation effects on metro tunnels [18]. Together, these studies show that sensing becomes valuable when it is linked to a defined risk mechanism and decision pathway.

Digital twins provide a wider integration layer. Kamari and Ham use deep-learning-based digital twinning for disaster preparedness and risk-informed construction decisions [19]. Liu et al. apply digital twins to control risks during prefabricated hoisting, demonstrating the value of synchronising operational status with a risk model [4]. Reviews of civil-infrastructure and emergency-management applications show that twins can support simulation, monitoring and response under extreme events [23,24]. More recent work frames digital twins as a basis for urban resilience and synchronised construction management, where real-time data can support dynamic replanning rather than static visualisation [25,26].

For giga-projects, digital-twin value lies in selective operational fidelity. Attempting to mirror every asset and process at maximum detail can create expensive data accumulation without proportional risk insight. A risk-oriented twin should prioritise critical interfaces: energisation sequences, logistics bottlenecks, temporary works, major lifts, utility tie-ins, commissioning dependencies, workforce density, environmental constraints and cyber-physical systems. These areas can be connected to risk indicators and scenario models so that the twin functions as a decision environment rather than a three-dimensional archive.

Data quality remains a limiting factor. Sensors drift, project codes change, subcontractors use incompatible systems, and manual status updates may be delayed or strategically optimistic. Risk models should therefore carry data-quality metadata and confidence scores. A warning generated from sparse or stale evidence should be distinguishable from one supported by convergent real-time indicators. This is particularly important where automated models could influence high-consequence actions such as stopping work, resequencing interfaces or activating emergency resources.

VI. PREDICTIVE ANALYTICS FOR INTERDEPENDENT PROJECT RISKS

Machine learning expands the ability to recognise nonlinear risk patterns, but its strongest contribution is not replacing expert judgement. It is increasing the number of relationships that can be examined consistently and quickly. Construction risk is inherently multivariate: delayed design approval may increase procurement compression, which reduces supplier choice, accelerates logistics, raises quality exposure and creates rework that further damages schedule. Moussa et al. show the importance of modelling risk interactions and systemic effects in infrastructure projects [5]. Darko et al. similarly argue that digital technologies should support integrated risk management rather than isolated lifecycle applications [27].

Different models are suitable for different warning problems. Gradient-boosted trees are effective for structured project-control data and can provide feature-importance explanations. Recurrent neural networks can model sequences where the order of events matters. Computer vision can infer unsafe states from imagery. Natural-language processing can screen correspondence, meeting minutes and non-conformance narratives for emerging themes. Probabilistic models remain useful where uncertainty itself must be represented explicitly. The key architectural decision is therefore model-task fit rather than adoption of the most complex algorithm.

Explainability is essential in projects where warnings can trigger costly intervention. A planner is more likely to trust a delay forecast when the system identifies the contributing pattern—such as declining constraint removal, design ageing and subcontractor productivity—than when it returns a probability without rationale. Research reviews of AI in construction risk management identify interpretability, fragmented implementation, data limitations and organisational resistance as continuing barriers [7,28]. Predictive systems should therefore expose the variables driving each alert, show comparison with prior states, and allow human review before major escalation.

Model drift is another concern. Giga-project phases change substantially from enabling works to structural delivery, systems installation, testing and operation. A model calibrated during civil works may become unreliable during commissioning because the relevant signals and dependencies change. Models should be monitored for performance degradation and recalibrated at phase gates. Human feedback should also be captured: when managers override a warning, the rationale and subsequent outcome become valuable learning data. This creates a controlled human-in-the-loop system rather than an opaque automated authority.

VII. BUSINESS CONTINUITY AND ADAPTIVE CRISIS RESPONSE

Prediction has strategic value only if it changes preparedness. Business continuity provides the mechanism for converting risk intelligence into protection of critical functions. In a giga-project, continuity does not necessarily mean maintaining every activity. It means preserving essential delivery

capabilities, decision rights, safety functions, digital access, cash-flow control, critical supply, temporary utilities and stakeholder communication while degraded conditions are stabilised. Organisational resilience research indicates that crisis performance depends on the capacity to adapt structures and practices rather than merely execute a fixed plan [13]. Project portfolio research likewise links resilience to robust prioritisation and resource allocation under uncertainty [14].

Supply-chain resilience is particularly important because giga-projects aggregate specialised materials, long-lead equipment and international vendors. Malik et al. show that resilient capabilities help manage integration complexity in construction supply chains [15], while Liao et al. identify resilience as a developing direction in construction supply-chain management [29]. Predictive indicators should therefore extend upstream: supplier financial health, manufacturing progress, logistics congestion, customs risk, commodity volatility, design-release maturity and alternate-source readiness can be monitored before shortage appears at site.

Adaptive crisis response differs from ordinary mitigation because it assumes that plans may need to be reconfigured as conditions evolve. A disruption may invalidate the original recovery sequence, creating the need to shift labour, protect critical path interfaces, defer nonessential scopes or establish temporary operating modes. Digital twins and scenario models can support this adaptation by testing alternative sequences and resource allocations before action [23–26]. The crisis structure should receive the same predictive evidence used by project controls, preventing a break between pre-crisis intelligence and response command.

Continuity planning should therefore be trigger-based. For each critical risk class, the organisation should define activation criteria, minimum service levels, alternate resources, decision owners and recovery priorities. For example, a procurement warning concerning a critical switchgear package should not merely increase a red score. It should trigger confirmation of manufacturing status, evaluation of alternate suppliers, review of temporary energisation options, commercial authority for expediting and schedule simulation of resequencing. This converts early warning from notification into an executable decision package.

VIII. SAUDI GIGA-PROJECT CONTEXT AND GOVERNANCE IMPLICATIONS

Saudi evidence reinforces the need for a multi-dimensional model. Algumzi identifies technology, human resources, contractual and financial factors among material risks associated with NEOM [20]. Alshihri et al. document risk factors contributing to cost and time overruns in Saudi building projects [21]. Saudi megaproject research also identifies a broad portfolio of risks requiring systematic prioritisation [30]. More recent predictive work on Saudi construction demonstrates that machine-learning models can classify delay exposure from project attributes [22]. Collectively, these studies indicate that risk intelligence in the Kingdom should span technical delivery and institutional coordination rather than focus narrowly on site productivity.

Governance must account for the scale of programme interfaces. Individual contractors may optimise local performance while shifting risk to shared infrastructure, access routes, utility corridors or commissioning systems. A programme-level risk intelligence function should therefore maintain cross-project indicators and identify correlated exposures. It should also control definitions so that terms such as “late design”, “procurement at risk” or “critical supplier” have consistent thresholds across packages. Without common semantics, aggregation can create false confidence because similar colours represent different underlying criteria.

Decision rights are equally important. Predictive analytics can challenge established reporting narratives by exposing deterioration earlier than formal forecasts. Leadership should protect the independence of risk intelligence teams while maintaining accountability for model quality. Warning owners require authority to request evidence, convene cross-functional review and escalate unresolved risks. At the same time, automated models should not bypass engineering judgement or contractual governance. The appropriate model is augmented decision-making:

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723505

algorithms surface patterns; domain experts validate mechanism and consequence; authorised leaders commit resources and actions.

IX. PROPOSED PREDICTIVE RISK INTELLIGENCE AND ADAPTIVE CONTINUITY FRAMEWORK

The synthesis supports a five-layer framework. Layer one is the sensing and evidence layer. It integrates project controls, BIM, procurement, commercial data, quality and safety systems, correspondence, sensors, external market indicators and stakeholder information. Each source should carry ownership, refresh frequency and quality metadata. Layer two is the risk inference layer. Here, statistical rules, machine-learning models, scenario logic and expert judgement transform observations into risk trajectories. Outputs should include probability, severity, propagation potential, confidence and principal drivers rather than a single colour rating.

Layer three is warning governance. Risk trajectories are compared with tiered thresholds that consider level, acceleration and interdependency. Alerts are routed according to pre-defined accountability, with time limits for validation and response. This layer records whether a warning is accepted, challenged, deferred or overridden. The purpose is to prevent analytical output from disappearing into informal discussion. Layer four is continuity and crisis activation. Validated warnings are linked to response playbooks containing minimum service objectives, alternate resources, sequencing options, communications and executive authorities. High-impact signals therefore produce executable choices rather than generic recommendations.

Layer five is learning and recalibration. Actual outcomes are compared with forecasts, response effectiveness is evaluated, thresholds are adjusted and models are retrained when necessary. Lessons are transferred across packages and project phases. The feedback loop also monitors whether organisational behaviour undermines the system: repeated late acknowledgement, systematic override of specific warning types or persistent false alarms become governance risks in their own right.

Figure 1 represents the architecture as a closed decision loop, while Figure 2 shows the escalation path from weak signal to adaptive response. The framework differs from a conventional dashboard because it explicitly connects analytical confidence with decision rights and continuity actions. It also recognises that prediction can be useful even when exact outcomes are uncertain. A credible indication that risk is accelerating may justify preserving options, increasing monitoring or pre-positioning resources before a definitive forecast is possible.

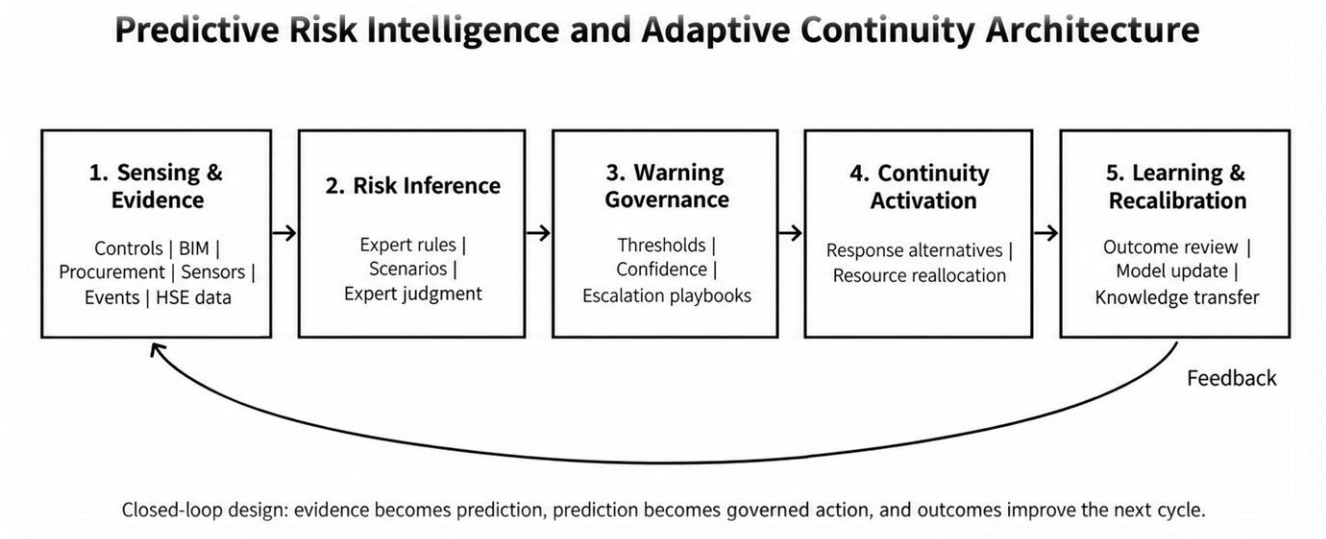

Figure 1. Five-layer predictive risk intelligence and adaptive continuity architecture.

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723505

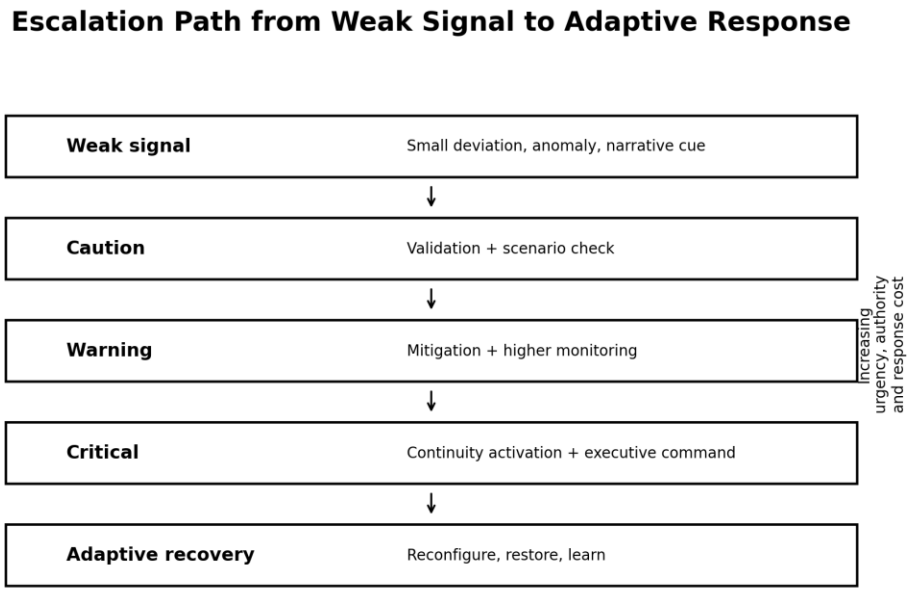

Figure 2. Escalation logic from weak signal detection to adaptive recovery.

X. IMPLEMENTATION PRIORITIES AND CONTROL MEASURES

Implementation should begin with a limited set of high-value risk use cases rather than enterprise-wide technology deployment. Candidate use cases are schedule slippage, long-lead procurement, design-interface congestion, critical safety exposure, major temporary works, commissioning dependency and cyber-physical disruption. Each use case should have a measurable outcome, such as increased warning lead time or reduced unplanned stoppage. This prevents the programme from equating digitalisation with risk improvement.

The second priority is data governance. A canonical project coding structure, interface identifiers, time stamps and controlled master data are prerequisites for reliable analytics. Data ownership must be contractually clear, particularly where joint ventures and specialist suppliers operate separate platforms. Access controls should support cybersecurity without preventing authorised cross-functional analysis. Because construction data may contain commercially sensitive information, risk models should also maintain audit trails showing who accessed evidence and how predictions were produced.

Third, organisations should build threshold governance before advanced modelling. Many programmes already possess enough data to detect deterioration, but lack agreed escalation rules. Historical projects can be used to back-test candidate indicators and define practical thresholds. Once the warning process is working, more sophisticated models can improve sensitivity and lead time. This sequence avoids a common failure mode in which complex algorithms are introduced into an organisation that has no disciplined response mechanism.

Fourth, predictive intelligence should be integrated with routine governance. Risk review, planning, procurement, interface management and continuity teams should view the same core signals. Executive reporting should show trend and decision status, not only current severity. Critical alerts should contain a concise explanation of evidence, forecast consequence, confidence, affected interfaces, required decision and deadline. This structure allows leaders to act without reading the entire analytical model.

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723505

Table 2. Implementation roadmap and assurance measures

Stage	Primary objective	Key control	Performance evidence
1. Use-case selection	Focus on high-consequence decisions	Define risk, decision owner and response horizon	Clear lead-time or disruption-reduction metric
2. Data foundation	Create trusted evidence	Common coding, ownership, quality metadata	Completeness, freshness and lineage
3. Warning governance	Make escalation executable	Tiered thresholds and time-bound responses	Action rate, false alarms, acknowledgement time
4. Predictive integration	Improve sensitivity and foresight	Explainable models with human validation	Precision, recall, lead time, model stability
5. Adaptive continuity	Protect critical delivery functions	Trigger-based playbooks and scenario testing	Service continuity and recovery performance
6. Learning	Institutionalise improvement	Post-event review and recalibration	Repeated precursor recognition; reduced recurrence

XI. DISCUSSION

The literature converges on a central conclusion: technical prediction and organisational resilience are complementary, not interchangeable. Digital twins, sensors and machine learning can improve visibility and forecasting, but they cannot decide which risk is tolerable, how much disruption a programme can absorb, or when commercial pressure should yield to continuity protection. Conversely, experienced managers cannot reliably process the volume and velocity of data generated by contemporary giga-projects without analytical support. Effective predictive risk intelligence is therefore a socio-technical capability.

This perspective also changes the meaning of "early". A warning is not early simply because it appears before an incident. It is early when sufficient response options remain. A procurement alert issued two weeks before a critical delivery may technically precede shortage but still be operationally late if alternate manufacturing requires four months.

A second insight concerns uncertainty. Decision-makers often demand highly accurate forecasts before acting, yet the highest-value interventions may occur when evidence is incomplete. The framework addresses this by separating confidence from consequence. A low-confidence signal with catastrophic downside can justify inexpensive option-preserving action, such as validating a supplier, reserving alternate logistics capacity or increasing monitoring. More disruptive interventions can wait for stronger evidence. This graded logic supports proportionality without ignoring weak signals.

A third insight is that crisis prevention requires institutional memory. Recording signal, interpretation, decision and outcome creates a reusable dataset for both machine learning and management development. Over time, this can reveal recurring precursors, biased thresholds and response patterns. For Saudi giga-projects, shared learning across programmes could become a strategic national asset, improving the maturity of future complex projects rather than treating each development as a unique experiment.

XII. LIMITATIONS AND FUTURE RESEARCH

This review is limited by the heterogeneity of the evidence base. Studies use different definitions of risk, resilience and digital twins, and many predictive applications are validated on bounded datasets rather than live mega-project programmes. Transferability to Saudi giga-projects should therefore be tested through longitudinal implementation studies. The review also focuses on publications from 2020 to 2025,

intentionally excluding older foundational theories except where they are reflected through contemporary sources.

Future research should evaluate multi-source warning models across complete project phases and measure decision lead time rather than predictive accuracy alone. Studies should examine how model explainability influences escalation behaviour, how contractual arrangements affect willingness to share risk data, and how false alarms alter managerial trust. Research is also needed on portfolio-level contagion: a disruption in one giga-project may affect shared labour, suppliers, logistics corridors or regulatory capacity elsewhere. Digital twins could be extended from asset models to programme-network simulations that test such propagation.

Another priority is continuity optimisation. Current research often predicts a problem without evaluating which intervention best protects critical functions. Combining predictive models with optimisation and scenario simulation could recommend resource reallocations, resequencing or alternative sourcing under explicit cost and resilience constraints. Finally, governance research should examine the independence and accountability of risk-intelligence functions in politically visible programmes, where incentives to maintain positive reporting may conflict with the purpose of early warning.

XIII. CONCLUSION

Saudi giga-projects require a risk-management architecture capable of operating at the speed and interconnectedness of their delivery environment. The reviewed evidence shows that predictive analytics, digital twins, real-time sensing and weak-signal analysis can extend visibility beyond conventional risk registers, but technology alone does not create resilience. The decisive capability is the conversion of emerging evidence into governed decisions while response options remain available.

The proposed Predictive Risk Intelligence and Adaptive Continuity framework links five functions: sensing, risk inference, warning governance, continuity activation and learning. Its practical value lies in joining project controls with business continuity and crisis response rather than maintaining them as separate disciplines. Risk trajectories are detected from multiple sources, interpreted through analytical and expert reasoning, escalated through tiered thresholds, connected to executable response playbooks and evaluated after outcomes become known.

For Vision 2030 giga-projects, this approach can strengthen schedule reliability, supply continuity, safety, cyber resilience and interface management while building institutional learning across projects. The most important design principle is that an early-warning system should be judged by the decisions it enables, not the sophistication of its dashboard. A mature system creates credible lead time, explains why risk is changing, identifies who must act, preserves continuity of critical functions and improves after each intervention. Predictive risk intelligence thus becomes an organisational resilience capability: a disciplined means of seeing earlier, deciding faster and adapting before disruption becomes crisis.

REFERENCES

[1] C. Boje, A. Guerriero, S. Kubicki, and Y. Rezgui, “Towards a semantic Construction Digital Twin: Directions for future research,” *Automation in Construction*, vol. 114, Art. no. 103179, 2020. ScienceDirect

[2] R. Sacks, I. Brilakis, E. Pikas, H. S. Xie, and M. Girolami, “Construction with digital twin information systems,” *Data-Centric Engineering*, vol. 1, Art. no. e14, 2020. Cambridge

[3] D.-G. J. Opoku, S. Perera, R. Osei-Kyei, and M. Rashidi, “Digital twin application in the construction industry: A literature review,” *Journal of Building Engineering*, vol. 40, Art. no. 102726, 2021. ScienceDirect

[4] Z. Liu, A. Li, Z. Sun, G. Shi, and X. Meng, “Digital Twin-Based Risk Control during Prefabricated Building Hoisting Operations,” *Sensors*, vol. 22, Art. no. 2522, 2022. MDPI

[5] A. Moussa, M. Ezzeldin, and W. El-Dakhakhni, “Predicting and managing risk interactions and systemic risks in infrastructure projects using

machine learning," *Automation in Construction*, vol. 168, Art. no. 105836, 2024. ScienceDirect

[6] A. N. Egwim, H. Alaka, L. O. Toriola-Coker, H. Balogun, and F. Sunmola, "Applied artificial intelligence for predicting construction projects delay," *Machine Learning with Applications*, vol. 6, Art. no. 100166, 2021. ScienceDirect

[7] K. Tian, Z. Zhu, J. Mbachu, M. Moorhead, and A. Ghanbaripour, "Artificial intelligence in construction risk management: a decade of developments, challenges, and integration pathways," *Journal of Risk Research*, vol. 28, pp. 1078–1110, 2025. Taylor & Francis

[8] H. J. Pasman, "Early warning signals noticed, but management doesn't act adequately or not at all: a brief analysis and direction of possible improvement," *Journal of Loss Prevention in the Process Industries*, vol. 70, Art. no. 104272, 2020. ScienceDirect

[9] A. Wang *et al.*, "An investigation of key precursors of megaproject crises," *International Journal of Managing Projects in Business*, vol. 15, pp. 938–959, 2022. Emerald

[10] D. Aghimien *et al.*, "Assessment of early warning signs predisposing building projects to failure in Nigeria," *Journal of Engineering, Design and Technology*, vol. 18, pp. 1403–1423, 2020. Emerald

[11] J. Fang, D. Huang, and J. Xu, "Social Risk Early Warning of Environmental Damage of Large-Scale Construction Projects in China Based on Network Governance and LSTM Model," *Complexity*, vol. 2020, Art. no. 8863997, 2020. Wiley

[12] R. Sun, Y. Liu, and J. Zhao, "Innovation Network Reconfiguration Makes Infrastructure Megaprojects More Resilient," *Computational Intelligence and Neuroscience*, vol. 2022, Art. no. 1727030, 2022. Wiley

[13] B. Mokline, "Organizational resilience as response to a crisis: case of COVID-19 crisis," *Continuity & Resilience Review*, vol. 3, pp. 232–247, 2021. Emerald

[14] A. Mahmoudi, M. Abbasi, and X. Deng, "A novel project portfolio selection framework towards organizational resilience: Robust Ordinal Priority Approach," *Expert Systems with Applications*, vol. 188, Art. no. 116067, 2022. ScienceDirect

[15] Malik, K. I. A. Khan, S. Qayyum, F. Ullah, and A. Maqsoom, "Resilient Capabilities to Tackle Supply Chain Risks: Managing Integration Complexities in Construction Projects," *Buildings*, vol. 12, Art. no. 1322, 2022. MDPI

[16] S. Rao, M. Radanovic, Y. Liu, S. Hu, Y. Fang, K. Khoshelham, M. Palaniswami, and T. Ngo, "Real-time monitoring of construction sites: Sensors, methods, and applications," *Automation in Construction*, vol. 136, Art. no. 104099, 2022. ScienceDirect

[17] M. Z. Shanti, C.-S. Cho, B. Garcia de Soto, Y.-J. Byon, C. Y. Yeun, and T. Y. Kim, "Real-time monitoring of work-at-height safety hazards in construction sites using drones and deep learning," *Journal of Safety Research*, vol. 83, pp. 364–370, 2022. ScienceDirect

[18] W. Tian, J. Meng, X.-J. Zhong, and X. Tan, "Intelligent Early Warning System for Construction Safety of Excavations Adjacent to Existing Metro Tunnels," *Advances in Civil Engineering*, vol. 2021, Art. no. 8833473, 2021. Wiley

[19] M. Kamari and Y. Ham, "AI-based risk assessment for construction site disaster preparedness through deep learning-based digital twinning," *Automation in Construction*, vol. 134, Art. no. 104091, 2022. ScienceDirect

[20] A. Algumzi, "Risks and Challenges Associated with NEOM Project in Saudi Arabia: A Marketing Perspective," *Journal of Risk and Financial Management*, vol. 15, Art. no. 381, 2022. MDPI

[21] S. Alshihri, K. Al-Gahtani, and A. Almohsen, "Risk Factors That Lead to Time and Cost Overruns of Building Projects in Saudi Arabia," *Buildings*, vol. 12, Art. no. 902, 2022. MDPI

[22] S. Alsulamy, "Predicting construction delay risks in Saudi Arabian projects: A comparative analysis of CatBoost, XGBoost, and LGBM," *Expert Systems with Applications*, vol. 268, Art. no. 126268, 2025. ScienceDirect

[23] Y. Han *et al.*, "Literature review of digital twin technologies for civil infrastructure," *Journal of Infrastructure Intelligence and Resilience*, vol. 2, Art. no. 100050, 2023. ScienceDirect

[24] F. Jiang *et al.*, "A Review of Digital Twin Applications in Civil and Infrastructure

Emergency Management," *Buildings*, vol. 13, Art. no. 1143, 2023. MDPI

[25] Y. Wang, Q. Yue, X. Lu, D. Gu, Z. Xu, Y. Tian, and S. Zhang, "Digital twin approach for enhancing urban resilience: A cycle between virtual space and the real world," *Resilient Cities and Structures*, vol. 3, pp. 34–45, 2024. ScienceDirect

[26] X. Li *et al*., "Digital twin-enabled synchronized construction management: A roadmap from construction 4.0 towards future prospect," *Developments in the Built Environment*, vol. 19, Art. no. 100512, 2024. ScienceDirect

[27] A. Darko, A. P. C. Chan, Y. Yang, and M. O. Tetteh, "Building information modeling (BIM)-based modular integrated construction risk management - Critical survey and future needs," *Computers in Industry*, vol. 123, Art. no. 103327, 2020. ScienceDirect

[28] Y. Junjia, A. H. Alias, N. A. Haron, and N. Abu Bakar, "Machine learning algorithms for safer construction sites: Critical review," *Building Engineering*, vol. 2, Art. no. 544, 2024. Building Engineering

[29] L. Liao, C. Yang, and L. Quan, "Construction supply chain management: A systematic literature review and future development," *Journal of Cleaner Production*, vol. 382, Art. no. 135230, 2023. ScienceDirect

[30] A. M. Abdulmoneim, A. Al Samadony, and I. A. Nosair, "Identification and Ranking the Most Significant Risks of the Mega Construction Projects in Saudi Arabia," *Saudi Journal of Civil Engineering*, vol. 5, pp. 35–49, 2021. Saudi Journal of Civil Engineering