

Generative Artificial Intelligence for Executive Visual Communication in Saudi Giga-Projects: A Framework for Decision Support under Vision 2030.

Abstract

Generative artificial intelligence is now beginning to change the way large infrastructure organisations turn complex project information into actions for executives. In the case of Saudi giga-projects, the issue is not just about producing more analysis; it is about creating visual narratives that are reliable, traceable and ready to support decisions in projects that are of a huge scale, highly interdependent, subject to rapid change and involve a wide range of stakeholder interests. This review brings together research that has been published between 2020 and 2025 on generative AI, human-AI decision making, explainable AI, the digitalisation of construction, digital twins and project readiness in Saudi Arabia. Its aim is to find out how generative systems can assist with executive visual communication without replacing accountable human judgement. The evidence shows four repeated opportunities: the rapid synthesis of heterogeneous records, multimodal reporting, scenario-based explanation and adaptive visualisation. But these advantages depend on the quality of the data, the provenance of the data, the organisation's readiness, model governance and on having explicit control over uncertainty. The paper puts forward an Executive Visual Decision Support framework which includes five layers: project data, assurance, generative reasoning, visual communication and executive governance. The framework views generative AI as an interface between project intelligence and leadership deliberation rather than as an independent decision-maker. For environments associated with the delivery of Vision 2030, the main point is that value is obtained when the visual outputs can be audited against the source data, show confidence levels and alternatives, and keep the escalation paths available for decisions that are in dispute. The review ends with a proposed implementation programme and some research suggestions for testing in the context of Saudi giga-projects.

1. Introduction

The large-scale project environment in Saudi Arabia presents a particularly challenging situation for executive communication. Since such major programmes involve a wide range of information—including design data, schedules, cost forecasts, procurement records, risk registers, site observations, commitments from stakeholders, environmental indicators, and strategic objectives—this information is constantly changing and is often stored in separate systems. As a result, top-level decision-makers have to solve the problem of synthesising the information before making a decision: the relevant evidence must first be summarised, brought together and communicated in a way that enables timely judgement. Research into construction AI indicates that AI can be of assistance in the areas of prediction, optimisation, information fusion, natural-language processing and risk management, while digital twins can enhance the continuity between physical assets and their data representations (Pan and Zhang, 2021; Abioye et al., 2021; Boje et al., 2020; Opoku et al., 2021). Generative AI builds on this by creating text, images, structured summaries and multimodal outputs from contextual information.

The situation with executive use cases is different from that of routine automation. An executive summary is only useful if it addresses the correct question, separates evidence from inference, makes uncertainty clear, and enables a decision-maker to question the reasoning that underlies the presentation. Research into generative AI in professional settings has found significant improvements in productivity, particularly in tasks that involve a lot of language, although it also shows that augmentation alters the way expertise is distributed and the way users depend on advice from machines (Noy and Zhang, 2023; Brynjolfsson et al., 2025). Management studies cautions that automation and augmentation are interdependent rather than simple substitutes, and that failing to critically evaluate machine advice can limit independent human judgement (Raisch and Krakowski, 2021; Fugener et al., 2021). In the context of giga-project governance—where a visual summary can affect decisions regarding the allocation of capital, schedule recovery, stakeholder escalation or risk acceptance—communication quality therefore becomes an aspect of decision quality.

The situation in Saudi Arabia makes this requirement even more important. Felemban's research into AI readiness within the public construction sector in Saudi Arabia finds that technology

readiness, the state of the data, support from senior management, organisational maturity, absorptive capacity and government support are all essential for the adoption of AI. The study also highlights the problems relating to the availability, accuracy, access, sharing and quality of data, all of which place constraints on AI-assisted planning (Felemban, 2025). Further evidence from the NEOM case shows that there is a connection between effective communication, readiness for change and involvement of stakeholders and project performance, meaning that technical information must pass through social and organisational processes before it can be put into action (Turi et al., 2025). Executive visual communication should thus be regarded as a socio-technical capability, not merely as a task involving the design of a dashboard.

Generative AI is relevant since it is able to convert project evidence from one form to another. A system could, for example, retrieve a schedule variance, link it to correspondence and site imagery, summarise the causal explanations, generate recovery scenarios and produce a concise executive visual. Research in the field of construction has shown the integration of generative models with computer vision for automated daily reporting, and various reviews mention new applications in the areas of design, planning, documentation and project management (Xiao et al., 2024; Alwashah et al., 2025). However, fluent generation also introduces new risks. Large language models are capable of making unsupported statements, omitting qualifiers or stating uncertain conclusions with confident assurance (Ji et al., 2023; Dwivedi et al., 2023). The more refined the output, the greater the chance that people might confuse the quality of the presentation with the strength of the evidence.

The review examines the way in which generative AI can enhance executive visual communication in Saudi giga-projects without compromising traceability, challenge, and accountable judgement. It views generative AI as a communication layer that is governed and which operates based on verified project information, and incorporates evidence from collaborations between humans and AI, from explanations provided by AI, from the organisation's overall AI capabilities, from construction digital twins, and from studies of Saudi projects into a decision-support framework for the achievement of Vision 2030.

88

89 **2. Aim and objectives**

90

91 This review aims to combine the evidence from 2020 to 2025 regarding the potential of
92 generative AI to transform the information from heterogeneous giga-projects into executable,
93 visually organised support for executive decision-making within the context of Saudi Vision
94 2030, and to establish a governance-aware framework for its responsible use.

95

96 The analysis is based on five objectives. The first is to draw a map of generative and related AI
97 functions throughout the executive decision cycle, covering areas such as sensemaking,
98 exception detection, scenario framing, option comparison, escalation and decision recording. The
99 second objective is to identify the design requirements for executive visuals so that they retain
100 provenance, uncertainty and contextual nuance rather than just simplifying the information.
101 Third, the review looks at the socio-technical conditions that affect usefulness, such as the level
102 of data maturity, the integration with digital twins, organisational capability, trust and the
103 allocation of roles between humans and AI. Fourth, it creates an Executive Visual Decision
104 Support framework which links data, assurance, generative reasoning, visual explanation and
105 human governance. Fifth, it sets out an implementation and research programme for testing the
106 framework in actual Saudi giga-project environments using criteria relating to speed,
107 comprehension, traceability, decision confidence and accountability.

108

109 **3. Review methodology**

110

111 An integrative review based on a structured approach was chosen since the research problem
112 spans a number of established and developing fields of study which do not have a common unit
113 of analysis. The relevant evidence can be found in the areas of construction engineering,
114 information systems, management, human-computer interaction, public-sector AI governance
115 and generative AI research. A review that is too technically focused would fail to address the

organisational and cognitive conditions, whereas one that is purely managerial would fail to take into account the data architectures via which visual decision support becomes possible. The design of the integrative review therefore centres on conceptual convergence and on mechanisms that are transferable rather than on estimating a pooled effect.

The time frame for the review was set to span from 2020 to 2025 so as to include current AI capabilities while still ensuring a clear and consistent technological period. The search strategy included four categories: artificial intelligence or generative AI; construction, architecture-engineering-construction or digital twin; explainability, visualisation, decision support or human-AI collaboration; and Saudi Arabia, Vision 2030, giga-project or large-project governance. The potential studies were located using the publisher's online platforms, scholarly discovery services, and by carrying out backward and forward citation checks. In order to be included, the research had to be in English, fall within the date range, have a verifiable DOI, and be clearly relevant to at least two of the following areas: AI-enabled knowledge generation; executive or human-AI decision processes; explainability and visual communication; construction data and digital twins; organisational readiness or AI governance; and the Saudi project environment. The doctoral study by Felemban (2025) was kept in because it offers direct, context-specific evidence regarding public construction AI readiness and has a persistent DOI.

The set of evidence ending with thirty sources was selected with a preference for peer-reviewed empirical studies, high-quality reviews and well-established conceptual contributions. Purposefully, the collection includes both studies on generative AI and earlier work on explainability and human-AI collaboration since executive visual support relies on more than just the capabilities of the model. For instance, research into trust, delegation and interpretability offers explanations for why an accurate model can still lead to bad organisational outcomes if users place too much reliance on it, fail to understand its limitations or are unable to examine its reasoning (Glikson and Woolley, 2020; Liao et al., 2020; Mohseni et al., 2021; Fugener et al., 2022).

The sources were coded in a concept-focused manner rather than by summarising them article by article. Five synthesis areas were employed: information compression and generation; visual explanation and provenance; human-AI judgement and trust; construction data integration; and governance and organisational readiness. In each of these areas, the analysis noted the claimed mechanism, the type of evidence, the boundary conditions and the implications for an executive readership. Contradictory findings were kept. For example, studies on productivity support the efficiency potential of generative assistance, while those on collaboration indicate that machine advice can reduce the diversity of human judgement or lead to delegation errors (Noy and Zhang, 2023; Fugener et al., 2021; Fugener et al., 2022). This kind of tension is regarded as a design condition and is not settled by adopting either technological optimism or algorithmic aversion.

The process took place in three stages. In the first stage, common functional capabilities were found throughout the corpus, for example summarisation, multimodal fusion, prediction, explanation and scenario generation. In the second stage, the enabling conditions and failure modes were linked to these functions, involving data quality, hallucination, opacity, cognitive overload, trust calibration and governance. In the third stage, the relationships were converted into a layered framework designed to assist with visual decision-making at an executive level. The framework is therefore an analytical proposition based on converging evidence and not an empirically validated causal model; its worth consists in defining testable relationships and implementation controls which can be looked at in future field studies.

Table 1. Evidence synthesis across the reviewed domains

Evidence domain	Representative sources	Synthesis	Implication for executive use
Generative work	Noy and Zhang (2023); Brynjolfsson et al. (2025); Alwashah et al. (2025)	Faster drafting, synthesis and knowledge transfer are plausible, but gains depend on task design and source grounding.	Measure reduced decision latency together with factual integrity.

Evidence domain	Representative sources	Synthesis	Implication for executive use
Explainability and trust	Arrieta et al. (2020); Liao et al. (2020); Glikson and Woolley (2020); Mohseni et al. (2021)	Explanations are audience- and task-dependent; trust can be useful or miscalibrated.	Show provenance, uncertainty and drill-down rather than only a recommendation.
Human-AI collaboration	Fugener et al. (2021, 2022); Seeber et al. (2020); Raisch and Krakowski (2021)	AI can augment judgement but may alter delegation, conformity and collective diversity.	Preserve independent challenge, human ownership and override.
Construction data systems	Boje et al. (2020); Opoku et al. (2021); Long et al. (2024); Xiao et al. (2024)	Digital twins and multimodal reporting can connect project evidence across lifecycle activities.	Ground generation in authoritative project systems and versioned data.
Saudi readiness and governance	Felemban (2025); Turi et al. (2025); Birkstedt et al. (2023); Wirtz et al. (2022)	Data readiness, management support, communication, stakeholder processes and governance shape adoption.	Treat visual decision support as a socio-technical capability with explicit controls.

4. Thematic synthesis

4.1 From project information to executive sensemaking

170

171 AI research in the field of construction always presents value in the form of transforming
172 scattered project data into predictions, pattern recognition, and optimisation. Pan and Zhang
173 (2021) see information fusion, knowledge reasoning and natural-language processing as key
174 abilities for construction management, whereas Abioye et al. (2021) emphasise risk management,
175 monitoring and resource optimisation. The literature on digital twins provides a systems
176 approach in that its aim goes beyond simply gathering more data to instead keep a useful
177 connection between the physical conditions, the digital representations and the decision-making
178 processes (Boje et al., 2020; Opoku et al., 2021; Long et al., 2024). Accordingly, executives
179 ought to have generative AI positioned above the governed project-data layer rather than have it
180 work from isolated documents.

181

182 Generative models include a translation feature in this architecture. They are able to restate
183 technical records in decision-focused language, produce comparisons, link text with images and
184 tailor the output to the relevant stakeholders. For instance, Xiao et al. (2024) combine computer
185 vision with a generative model in order to create daily construction reports based on site video,
186 demonstrating how a machine-generated narrative can be based on project observations. Reviews
187 indicate that the use of generative AI in the construction sector is growing in the areas of design,
188 planning, documentation and management, and built-environment research points to user-centred
189 and data-driven generative approaches (Alwashah et al., 2025; Zhang and Zhang, 2025). The
190 executive opportunity is not a general chatbot but an orchestration layer which retrieves the
191 approved evidence and reassembles it into a decision object.

192

193 The literature on productivity helps to account for this appeal. Noy and Zhang (2023) show that
194 generative assistance can cut down on the time taken and at the same time improve the quality
195 when it comes to professional writing, while Brynjolfsson et al. (2025) note productivity gains in
196 a large-scale workplace implementation. Even though these situations are different from those in
197 the construction industry, both suggest that language models can reduce the cost of turning
198 knowledge into usable communication. In the case of a giga-project, this can reduce the gap

between analysis in the control room and leadership discussions. The advantage here is that decision latency can be reduced without compromising evidential integrity.

4.2 Visual explanation, provenance and uncertainty

Executive visuals are intentionally selective. For example, a dashboard, a risk map or a scenario card simplifies complex information so that attention can be focused on the most important aspects. With the help of generative AI this selectivity becomes dynamic since the system can decide what to summarise, how to present it and which type of representation to generate. Research into explainability demonstrates the need for audience-specific design. Arrieta et al. (2020) claim that explanations depend on the audience and the purpose; Liao et al. (2020) show that users pose different questions when they are trying to understand AI outputs; and Mohseni et al. (2021) link explainability objectives with different user groups and evaluation methods. An executive visual does not have to include all the technical details, but it does have to reveal only what is necessary for decision-making.

The following three requirements apply. The first is that provenance must be visible—each claim which has an impact on a decision should be able to be traced back to the original records, models or assumptions. The second is that uncertainty should be shown using confidence ranges, scenario bands, data-quality indicators or by indicating explicitly any unresolved issues. The third is that alternatives should stay inspectable; since a single recommended course of action can have an anchoring effect, executives should be able to compare plausible options and understand the factors that distinguish them.

Research into social transparency supports this point. Ehsan and others (2021) found that explanations are situated in a social context and can be enhanced by making available information regarding the people, processes and organisational background associated with an AI-assisted decision. Kaur and colleagues (2020) warn that interpretability tools do not by

themselves lead to a correct understanding; users may misuse them or form incorrect mental models. It is for this reason that in the context of giga-project communication provenance should not only include the dataset but also the assumption, the responsible function, the time of update and the decision rule. A technically visually attractive output may be transparent in technical terms yet still remain opaque in an organisational sense.

The need for provenance to be established becomes absolute. Ji et al. (2023) show that natural-language generation is capable of producing content that is not supported by the source material. Dwivedi et al. (2023) extend the issue to include credibility, transparency, bias, and governance when the technology is used in an organisation. In the case of executive reporting, generation must be limited to approved project repositories, numerical values must be obtained from authoritative systems, and any content without support should be either blocked or marked as hypothetical. Visual communication should serve as an assurance surface that enables users to see the limitations of the model rather than concealing them.

4.3 Human-AI judgement and calibrated reliance

The best argument for the use of generative AI in executive communication is that of augmentation, but this does not mean that augmentation is safe. Raisch and Krakowski (2021) point out that automation and augmentation constitute an organisational paradox since tasks can shift between humans and machines over time. Shneiderman (2020) advocates for systems that combine a high degree of automation with meaningful human control. In the case of giga-projects, generative AI should be used to automate the synthesis whereas material judgement, accountability and the acceptance of exceptions should remain with authorised people.

Trust should be adjusted rather than pushed to the highest level. Glikson and Woolley (2020) indicate that trust in AI is influenced by the features of the technology and the way it is presented, and Fugener et al. (2022) show that inadequate metaknowledge can hinder human

delegation even if the users are willing to work with AI. Furthermore, Fugener et al. (2021) discover that advice from AI can increase individual accuracy but at the same time decrease the diversity of human judgement, which may in turn weaken collective problem solving. When an executive committee sees the same machine-generated framing prior to forming their own independent views, the debate might be narrowed.

A more effective workflow should keep the process of machine synthesis separate from human challenge. The system can prepare the evidence, detect any conflicts and suggest various scenarios, and at the same time the interface should encourage counter-interpretations. Executives should be able to record their own independent assessments when making decisions of high consequence, check out the synthesis produced by the AI and then note the points at which their judgement differs. Seeber et al. (2020) view AI as a possible teammate whose impact will vary according to the machine, the collaborative process and the institutional context. Makarius et al. (2020) do similarly stress the cognitive, relational and structural conditions involved in organisational integration. A visual summary should indicate who is responsible for the decision, who has validated the evidence and where the generated narrative has been changed as a result of human challenge.

Kellogg and others in 2020 introduced a governance aspect by demonstrating the way in which algorithms alter the mechanisms of control within organisations. As part of a large-scale project, generative reporting can affect the issues that come to the attention of senior management and the interpretations that are given greater prominence. Prompt design, data access and the rules for escalating issues are thus examples of organisational controls. The support system must keep a record that can be audited of all the information retrieved, transformed, user edits made and the final decisions reached, so that the chain of influence can be examined.

4.4 Organisational readiness and the Saudi giga-project context

AI capability is in part a matter of how organisational resources are arranged. Mikalef and Gupta (2021) view AI capability as made up of tangible, human and intangible resources, which means that having access to a model by itself does not lead to improved performance. Felemban (2025) reaches a similar conclusion in the context of public construction in Saudi Arabia, noting that readiness is linked to technology infrastructure, the state of the data, support from senior management, organisational maturity, absorptive capacity, training, collaboration and government support; nevertheless, poor upstream readiness cannot be compensated for by better graphics.

Data maturity forms the initial boundary condition. When a large-scale project deals with information, it involves design platforms, document-control systems, schedule tools, cost systems, geographic information, sensors, correspondence and reports from contractors. Even if the definitions differ, a generative interface can make these various sources appear to be unified. Before any generation takes place, the assurance layer has to resolve the identifiers, timestamps, versions and ownership. Research into digital twins supports this focus on interoperability and the integration of data throughout the lifecycle (Boje et al., 2020; Long et al., 2024). The model must not invent consistency in areas where the project has not already established it.

The second boundary condition is governance. Birkstedt and others (2023) point out that technology, the various stakeholders and the context, regulation and processes are recurring aspects of organisational AI governance, and they stress the importance of senior management commitment and operationalisation. Wirtz and co-authors (2022) also suggest a governance approach that is oriented towards risk, covering technological, data, informational, social, ethical and legal issues. When these findings are applied to the Saudi giga-projects, this leads to the implementation of role-based access, approved model configurations, confidentiality controls, documented escalation thresholds, model review and human approval for any recommendations that have significant consequences.

The third criterion is communication readiness. As Turi and others (2025) have found in their study of projects within the NEOM context, there is a connection between effective communication, change readiness and stakeholder engagement and the different performance pathways of projects. Executive visualisation thus plays a role in the process by which requirements, risks and changes are interpreted. A generative system should allow for multiple views of the same verified evidence—such as a high-level narrative on exceptions for the board, a view for the programme director showing the actions needed, and a technical level of detail. The evidence must remain consistent even when the level of detail changes.

These themes together show a key principle, namely that generative AI creates value by reducing representational friction without at the same time reducing epistemic discipline; it should make project intelligence easier to understand while still keeping the source, the uncertainty, the assumptions, the alternatives and responsibility.




Figure 1. Visual intelligence pipeline from verified project evidence to accountable executive action.

5. Executive Visual Decision Support framework

The Executive Visual Decision Support (EVDS) framework consists of five interconnected layers together with a feedback loop; its aim is to transform project evidence into visual presentations that are ready for use by executives while at the same time keeping assurance and ensuring that there is accountable human control.

The project evidence layer is known as Layer 1. It consists of the approved schedule, cost, risk, design, procurement, site, environmental, stakeholder and performance data, along with the relevant narrative records. Although digital twins and building-information models offer spatial and asset context, they are still considered sources and not replacements for governance. Each source must have an owner, a time of update, a version number and an access classification. This arrangement stops the generative model from becoming an unauthorised system of record (Pan and Zhang, 2021; Long et al., 2024).

The assurance layer is made up of Layer 2. Before information gets to a generative model, both automated and manual controls are used to verify its completeness, consistency, provenance, and relevance for decision-making. Numerical fields are checked against authoritative systems; conflicting values are identified and brought to light rather than being quietly averaged; missing data are noted; and records that are restricted are filtered according to the user's role. For outputs that have high consequences, higher evidence thresholds are applied and specialist review is required. This layer is particularly important in Saudi public-project environments according to Felemban's (2025) findings on data, technology readiness and organisational maturity.

Layer 3 is the layer responsible for generative reasoning. The models carry out retrieval, summarisation, comparison, anomaly narration, the generation of causal hypotheses and the drafting of scenarios. The system arranges the evidence but is not regarded as having independent managerial judgement. Retrieval is restricted to approved sources and the prompts

have to include citation of evidence, explicit assumptions and a distinction between observed facts and hypotheses. The research into hallucinations supports this kind of control since fluency cannot be used to establish factual support (Ji et al., 2023). Scenarios must be presented as conditional possibilities unless predictive models have been validated and provide probabilities.

Layer 4 serves as the visual communication layer and transforms the analysis results into decision-making objects such as exception cards for matters that need to be acted upon, scenario panels which compare different options and their consequences, temporal deltas that show what has changed, spatial views that link risks or progress to specific locations, and decision trails that connect recommendations to the evidence, assumptions, and previous commitments. The visual grammar should be stable so that executives can learn where to find information on provenance, uncertainty, responsibility, and the next action. Research into explainability indicates that the design should be aligned with the questions of the users rather than aiming to maximise technical detail (Arrieta et al., 2020; Liao et al., 2020).

Layer 5 functions as the executive governance layer; humans have the option of accepting, rejecting, modifying or escalating the synthesis produced by the machine. For issues relating to capital, safety, contractual agreements, regulation and strategy, decision-making powers remain clearly defined. The system keeps a record of the final decision, the authority responsible for it, the reasons given and any deviation from the recommended action. Research into human-centred AI suggests that combining computational support with clear human control is effective, and studies in organisational behaviour indicate that the integration of AI depends on relational and institutional design (Shneiderman, 2020; Makarius et al., 2020).

The decision outcomes are sent back to the evidence and assurance sections. In the case where an executive rejects an interpretation on the grounds that a key dependency was missing, the system should call for a correction to the data or rules rather than making only a change to the prompt. Whenever there are repeated requests for a drill-down, this can be used to improve the visual design. Relevant indicators are briefing preparation time, the proportion of claims that can be

traced back to authoritative sources, the number of data conflicts detected before review, executive comprehension, the number of unsupported outputs, the length of the decision cycle, the number of human overrides and the number of actions that are closed.

The implementation must also proceed in stages, with each stage being approved by the governing body. In the first phase, a pilot project should set up authoritative data sources, identify the permitted use cases, assign responsible owners and establish procedures for dealing with failures. In the second stage, the generated briefs should be checked against the existing executive packs, noting any factual discrepancies, missing context, user corrections and the time saved. Only when the system has performed satisfactorily should it be used for scenario comparison or for drafting recommendations. In the third stage, feedback from executives, project controls, legal, cybersecurity and domain specialists should be incorporated into a controlled operating model. Periodic audits should then look at the source coverage, any changes to the model, access patterns, unsupported claims and the patterns of human override. This approach transforms the adoption of the system from a simple technology rollout into a continuously monitored process of organisational learning, one that is in line with the governance and readiness evidence found in the reviewed literature.

The approach does not rely on a single overall AI score since a quick system might be unsafe, a system that is highly trusted might be depended on too much, and an appealing system might hide missing evidence. Instead, effectiveness should be evaluated in terms of speed, understanding, traceability, the level of challenge and accountability.




Figure 2. Governed human-AI decision loop with provenance, uncertainty, access and audit guardrails.

Table 2. EVDS design requirements across the executive decision cycle

Decision stage	Generative function	Preferred visual form	Required assurance	Human responsibility
Sensemakin g	Retrieve and summarise current evidence	Exception card + source links	Freshness, completeness, source authority	Confirm relevance and materiality
Diagnosis	Compare records and describe plausible drivers	Causal map + evidence panel	Separate observation from hypothesis	Challenge causal interpretation
Option	Draft	Scenario	Explicit	Define feasible

Decision stage	Generative function	Preferred visual form	Required assurance	Human responsibility
framing	alternatives and consequences	comparison panel	assumptions; no invented probabilities	options and constraints
Decision	Condense trade-offs and unresolved issues	Decision card + uncertainty markers	Trace every material claim and number	Accept, reject, modify or escalate
Action	Translate decision into owned next steps	Action register + milestone delta	Owner, deadline and system-of-record link	Assign accountability and monitor closure
Learning	Compare outcomes with assumptions	Decision trail + variance view	Preserve history and overrides	Review lessons and update controls

6. Discussion and implications

The review treats executive visual communication as an information-governance issue using a visual interface, rather than as a graphics problem involving an AI engine. Although generative systems are very good at producing coherent representations, coherence can lessen cognitive effort but may also mask disagreements between sources. That is why the EVDS framework puts assurance before generation and governance afterwards: executives only gain the advantage of compression once the organisation has decided what can legitimately be compressed.

What this means for programme leaders is that they should start by making limited decisions rather than implementing them across the entire organisation. Good examples of early

421 applications would be regular executive reviews in which the evidence is organised, for instance
422 by means of schedule exceptions, change-control summaries or risk escalations. In such
423 situations the outputs produced automatically can be compared with the existing reports after
424 being examined by a person. Research into automated construction reporting shows that
425 generative models are able to reduce the amount of documentation needed when they are linked
426 to observable project evidence (Xiao et al., 2024). Expansion should come only after traceability
427 and comprehension have been demonstrated, not because of the visual appeal.

428 A further consequence is that systems should be designed with a challenge in mind.
429 Explainability must allow an executive to find out what has changed, what evidence backs up a
430 particular statement, what aspects are uncertain, what alternative options were examined and
431 who is responsible for the next step. Ehsan et al. (2021) demonstrate that organisational context
432 is an important element of an explanation, whereas Kaur et al. (2020) show that interpretability
433 tools can be misinterpreted. A helpful visual integrates a brief narrative with the ability to drill
434 down, shows provenance and includes the decision history. Complexity should be kept available
435 when needed rather than being discarded.

436 The third point relates to the development of capability. Organizations require individuals who
437 are able to link together project controls, data engineering, model governance, information
438 design and the processes involved in executive decision-making. This is in line with the finding
439 that AI capability relies on a range of complementary resources rather than on technology alone
440 (Mikalef and Gupta, 2021). Felemban (2025) reaches a similar conclusion, identifying absorptive
441 capacity, training, collaboration and management support as readiness factors. A centre of
442 enablement could establish approved patterns, test models, keep the retrieval controls, provide
443 training to users and examine incidents, while the business owners would remain accountable.

444 The fourth point is about institutional memory. By means of a governed generative layer, current
445 exceptions can be connected with previous decisions, assumptions and outcomes, thus enhancing
446 continuity during changes in leadership and when working with contractors. It is necessary to
447 retain disagreements and outdated reasoning rather than rewrite history into a single story. As
448 algorithmic-control research has shown, information systems have an effect on what is visible
449 and on organisational power (Kellogg et al., 2020). Governance should therefore look into which

data are included, which issues attract the attention of executives and the way in which the outputs influence escalation.

In terms of research, the unit of analysis should change from model accuracy to decision-system performance. Future studies could compare traditional reporting with EVDS-supported reporting by looking at comprehension, decision time, the recall of uncertainty, the quality of challenges, reliance behaviour, action closure and traceability. Experiments might be carried out to find out whether making an independent judgement before being exposed to AI reduces the conformity effects suggested by research into human-AI collaboration (Fugener et al., 2021). Field studies could investigate how Saudi construction readiness factors affect adoption and whether greater digital-twin integration leads to better factual grounding.

7. Limitations and future research

There are three limitations to this review. The first is that there has been only a limited amount of direct empirical research into the use of generative AI for executive visual communication in Saudi giga-projects; the framework has therefore incorporated evidence from related areas such as construction AI, explainability, organisational studies and professional knowledge work. It is necessary to test, not merely assume, the transferability of this evidence. The second limitation is that the rapid development of multimodal models means that specific technical capabilities could change more quickly than organisational evidence does. As a result, the framework places a focus on durable controls—namely provenance, assurance, human challenge and auditability—rather than on particular model brands or architectures. The third limitation is that the review makes use of a curated integrative corpus instead of carrying out a bibliometric census, so it has been designed for conceptual depth rather than for complete coverage.

Future research should verify the EVDS framework by means of longitudinal pilot studies. One such study could look at the briefing cycles both before and after the implementation has taken place, recording the time taken for preparation, the ability to trace sources, levels of comprehension, and the time lag before a decision is made. Another study could vary the visual displays of uncertainty in order to find out their impact on executive confidence and on the tendency to escalate. A third study could test out different sequences involving humans and AI, for example by having a human make an independent judgement before the machine carries out synthesis, in order to examine the effects on conformity and diversity. The research should also

investigate Arabic-English bilingual generation, cross-organisational data rights, confidentiality, model drift, and the way in which generative interfaces interact with digital twins. The most compelling evidence will be obtained from actual project decisions where the quality of communication can be connected with the subsequent actions and outcomes.

8. Conclusion

Generative AI has the potential to serve as a useful communication layer for the leadership team involved in Saudi's giga-project, but the extent of its contribution will depend on it being carefully integrated with project evidence and governance procedures. As the literature indicates, AI can enhance the processes of synthesis, reporting and professional productivity, since the digitalisation of construction is creating richer data environments for decision support. On the other hand, hallucinations, opacity, misaligned trust, a tendency towards conformity and insufficient organisational readiness can cause efficient content generation to turn into unreliable communication for management purposes.

To deal with this tension the EVDS framework connects five layers: project evidence, assurance, generative reasoning, visual communication and executive governance. Its main idea is that generative AI should simplify complexity without at the same time reducing accountability. Executive visuals must therefore show provenance, uncertainty, alternatives, ownership and decision history even if the main presentation stays concise. Final judgement is still the responsibility of humans, alongside the system providing support for synthesis, challenge and learning.

With Vision 2030, the value of this approach is not so much due to any technological innovation as rather to the extent and level of coordination required when delivering giga-projects. A generative interface that is properly governed can enable leaders to transition from scattered information to quicker, more coherent and more inspectable deliberations. At present the focus of the research is on empirical evidence to determine whether such systems reduce decision-making delays and enhance understanding while at the same time preserving independent judgement, traceability and the ability for responsible escalation. Only if these conditions are satisfied can generative visual communication strengthen support for executive decision-making without turning into an unaccountable replacement for it.

It also ensures that there is a consistent level of executive action across various programmes by taking as the deciding factor for confidence, challenge, escalation and any final decision the quality of the evidence rather than the fluency of the language generated.

References

1. Abioye, S.O., Oyedele, L.O., Akanbi, L., Ajayi, A., Davila Delgado, J.M., Bilal, M., Akinade, O.O. and Ahmed, A. (2021) 'Artificial intelligence in the construction industry: A review of present status, opportunities and future challenges', *Journal of Building Engineering*, 44, 103299. <https://doi.org/10.1016/j.jobbe.2021.103299>
2. Alwashah, Z., Xiao, B., Liu, H., Mueller, S.T. and Shao, X. (2025) 'Generative artificial intelligence for construction: Use cases, trends, challenges, and opportunities', *Journal of Building Engineering*, 112, 113802. <https://doi.org/10.1016/j.jobbe.2025.113802>
3. Arrieta, A.B., Diaz-Rodriguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R. and Herrera, F. (2020) 'Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI', *Information Fusion*, 58, pp. 82-115. <https://doi.org/10.1016/j.inffus.2019.12.012>
4. Birkstedt, T., Minkkinen, M., Tandon, A. and Mantymaki, M. (2023) 'AI governance: themes, knowledge gaps and future agendas', *Internet Research*, 33(7), pp. 133-167. <https://doi.org/10.1108/INTR-01-2022-0042>
5. Boje, C., Guerriero, A., Kubicki, S. and Rezgui, Y. (2020) 'Towards a semantic Construction Digital Twin: Directions for future research', *Automation in Construction*, 114, 103179. <https://doi.org/10.1016/j.autcon.2020.103179>
6. Brynjolfsson, E., Li, D. and Raymond, L.R. (2025) 'Generative AI at Work', *The Quarterly Journal of Economics*, 140(2), pp. 889-942. <https://doi.org/10.1093/qje/qjae044>
7. Dwivedi, Y.K. et al. (2023) 'Opinion Paper: So what if ChatGPT wrote it? Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy', *International Journal of Information Management*, 71, 102642. <https://doi.org/10.1016/j.ijinfomgt.2023.102642>

8. Ehsan, U., Liao, Q.V., Muller, M., Riedl, M.O. and Weisz, J.D. (2021) 'Expanding Explainability: Towards Social Transparency in AI Systems', Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, pp. 1-19. <https://doi.org/10.1145/3411764.3445188>
9. Felemban, H.M.H. (2025) Organisational Readiness for AI in the Front-end Planning of Public Construction Projects. Doctoral thesis, Loughborough University. <https://doi.org/10.26174/thesis.lboro.30415369.v1>
10. Fugener, A., Grahl, J., Gupta, A. and Ketter, W. (2021) 'Will Humans-in-the-Loop Become Borgs? Merits and Pitfalls of Working with AI', MIS Quarterly, 45(3), pp. 1527-1556. <https://doi.org/10.25300/MISQ/2021/16553>
11. Fugener, A., Grahl, J., Gupta, A. and Ketter, W. (2022) 'Cognitive Challenges in Human-Artificial Intelligence Collaboration: Investigating the Path Toward Productive Delegation', Information Systems Research, 33(2), pp. 678-696. <https://doi.org/10.1287/isre.2021.1079>
12. Glikson, E. and Woolley, A.W. (2020) 'Human Trust in Artificial Intelligence: Review of Empirical Research', Academy of Management Annals, 14(2), pp. 627-660. <https://doi.org/10.5465/annals.2018.0057>
13. Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y.J., Madotto, A. and Fung, P. (2023) 'Survey of Hallucination in Natural Language Generation', ACM Computing Surveys, 55(12), Article 248, pp. 1-38. <https://doi.org/10.1145/3571730>
14. Kaur, H., Nori, H., Jenkins, S., Caruana, R., Wallach, H. and Vaughan, J.W. (2020) 'Interpreting Interpretability: Understanding Data Scientists' Use of Interpretability Tools for Machine Learning', Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, pp. 1-14. <https://doi.org/10.1145/3313831.3376219>
15. Kellogg, K.C., Valentine, M.A. and Christin, A. (2020) 'Algorithms at Work: The New Contested Terrain of Control', Academy of Management Annals, 14(1), pp. 366-410. <https://doi.org/10.5465/annals.2018.0174>
16. Liao, Q.V., Gruen, D. and Miller, S. (2020) 'Questioning the AI: Informing Design Practices for Explainable AI User Experiences', Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, pp. 1-15. <https://doi.org/10.1145/3313831.3376590>

17. Long, W., Bao, Z., Chen, K., Ng, S.T. and Wuni, I.Y. (2024) 'Developing an integrative framework for digital twin applications in the building construction industry: A systematic literature review', *Advanced Engineering Informatics*, 59, 102346. <https://doi.org/10.1016/j.aei.2023.102346>
18. Makarius, E.E., Mukherjee, D., Fox, J.D. and Fox, A.K. (2020) 'Rising with the machines: A sociotechnical framework for bringing artificial intelligence into the organization', *Journal of Business Research*, 120, pp. 262-273. <https://doi.org/10.1016/j.jbusres.2020.07.045>
19. Mikalef, P. and Gupta, M. (2021) 'Artificial intelligence capability: Conceptualization, measurement calibration, and empirical study on its impact on organizational creativity and firm performance', *Information & Management*, 58(3), 103434. <https://doi.org/10.1016/j.im.2021.103434>
20. Mohseni, S., Zarei, N. and Ragan, E.D. (2021) 'A Multidisciplinary Survey and Framework for Design and Evaluation of Explainable AI Systems', *ACM Transactions on Interactive Intelligent Systems*, 11(3-4), pp. 1-45. <https://doi.org/10.1145/3387166>
21. Noy, S. and Zhang, W. (2023) 'Experimental evidence on the productivity effects of generative artificial intelligence', *Science*, 381(6654), pp. 187-192. <https://doi.org/10.1126/science.adh2586>
22. Opoku, D.-G.J., Perera, S., Osei-Kyei, R. and Rashidi, M. (2021) 'Digital twin application in the construction industry: A literature review', *Journal of Building Engineering*, 40, 102726. <https://doi.org/10.1016/j.jobbe.2021.102726>
23. Pan, Y. and Zhang, L. (2021) 'Roles of artificial intelligence in construction engineering and management: A critical review and future trends', *Automation in Construction*, 122, 103517. <https://doi.org/10.1016/j.autcon.2020.103517>
24. Raisch, S. and Krakowski, S. (2021) 'Artificial Intelligence and Management: The Automation-Augmentation Paradox', *Academy of Management Review*, 46(1), pp. 192-210. <https://doi.org/10.5465/amr.2018.0072>
25. Seeber, I., Bittner, E., Briggs, R.O., de Vreede, T., de Vreede, G.-J., Elkins, A., Maier, R., Merz, A.B., Oeste-Reiss, S., Randrup, N., Schwabe, G. and Sollner, M. (2020) 'Machines as teammates: A research agenda on AI in team collaboration', *Information & Management*, 57(2), 103174. <https://doi.org/10.1016/j.im.2019.103174>

- 600 26. Shneiderman, B. (2020) 'Human-Centered Artificial Intelligence: Reliable, Safe &
601 Trustworthy', *International Journal of Human-Computer Interaction*, 36(6), pp. 495-504.
602 <https://doi.org/10.1080/10447318.2020.1741118>
- 603 27. Turi, J.A., Al Kharusi, S., Khwaja, M.G., Razzak, M.R. and Bashir, S. (2025)
604 'Requirement engineering and project performance: mediating effect of communication
605 effectiveness, change readiness, stakeholder engagement and moderating effect of team
606 cohesion', *Engineering, Construction and Architectural Management*, 32(8), pp. 5484-
607 5507. <https://doi.org/10.1108/ECAM-12-2023-1273>
- 608 28. Wirtz, B.W., Weyerer, J.C. and Kehl, I. (2022) 'Governance of artificial intelligence: A
609 risk and guideline-based integrative framework', *Government Information Quarterly*,
610 39(4), 101685. <https://doi.org/10.1016/j.giq.2022.101685>
- 611 29. Xiao, B., Wang, Y., Zhang, Y., Chen, C. and Darko, A. (2024) 'Automated daily report
612 generation from construction videos using ChatGPT and computer vision', *Automation in
613 Construction*, 168, 105874. <https://doi.org/10.1016/j.autcon.2024.105874>
- 614 30. Zhang, H. and Zhang, R. (2025) 'Generative artificial intelligence (AI) in built
615 environment design and planning - A state-of-the-art review', *Progress in Engineering
616 Science*, 2(1), 100040. <https://doi.org/10.1016/j.pes.2024.100040>