© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723455

AI-Augmented Executive Dashboards: Integrating Generative Artificial Intelligence and Data Visualization for Project Decision-Making in Saudi Arabia

MUHAMMAD KHURRAM TOOR
ORCiD: 0009-0003-3729-5252

Abstract-* Executive project dashboards are no longer expected to do merely than display historical key performance indicators. In the case of complex portfolios, managers need to interpret weak signals, connect both structured and unstructured evidence, evaluate various courses of action, and have to justify their decisions all while under time pressure. This study examines how generative artificial intelligence can be combined with data visualisation in order to create AI-augmented executive dashboards for project decision-making in Saudi Arabia. A structured integrative review was conducted in order to gather thirty peer-reviewed studies published between 2020 and 2025 in the fields of project management, decision support, visual analytics, human-AI collaboration, generative AI, and Saudi digital transformation. The results indicate that the most value is not derived from making autonomous decisions, but rather from having a multi-layered interface in which validated project data, predictive models, visual representations, and a conversational generative element support executive judgement. The quality of the dashboard is assessed according to the currency of the information, its completeness, the suitability of the visual design, the degree of cognitive load, semantic consistency, and the clear communication of uncertainty. Generative AI can improve dashboards by allowing natural-language queries, giving narrative explanations, developing scenarios, carrying out cross-source synthesis, and providing interactive visualisations; nevertheless, these features involve risks such as hallucination, automation bias, opaque reasoning, and the diffusion of accountability. The level of adoption in Saudi Arabia is also affected by organisational readiness, data governance, the amount of attention given by leadership, the digital infrastructure, and the existing project delivery practices. The review proposes a human-governed framework which divides the stages of evidence, analytics, generation, presentation, and decision-making, while at the same time ensuring that traceability is maintained for each generated statement back to the source data. This framework provides a design approach for executive project intelligence in Saudi organisations and establishes a research agenda focused on decision quality, calibrated trust, multilingual interaction, portfolio-scale validation, and governance-by-design.

I. INTRODUCTION

The kingdom is pursuing a number of major and interdependent projects in the areas of infrastructure, technology, public services, tourism, energy, and urban development, which requires executives to adjust their strategic objectives in light of changing operational evidence. The data are dispersed among a variety of sources including schedules, cost systems, risk registers, contract repositories, progress reports, sensors, emails, meeting minutes, and specialised platforms. Even though traditional dashboards attempt to cope with this level of complexity by displaying performance through charts and indicators, they generally fail to keep up with current developments, concentrate primarily on metrics, and assume that the user already knows the questions they should be asking. The present opportunity is to combine visual analytics with generative artificial intelligence in such a way that the dashboard becomes an interactive environment for making decisions rather than remaining just a simple reporting tool.

Research into the application of artificial intelligence in project management has already seen a shift from manual estimation to prediction based on data, together with improvements in classification, optimisation, and assistance with decision-making. The reviews point out the different uses in the fields of planning, uncertainty management, monitoring, risk, cost, and schedule, while also underlining the limitations concerning data quality, adoption, and governance (Bento et al., 2022; Taboada et al., 2023; Mohammad & Chirchir, 2024; Nenni et al., 2025). In the construction industry, AI has been associated with forecasting, automation, pattern recognition, and improved management information, even though fragmented data and organisational capability continue to be major difficulties (Abioye et al., 2021; Pan & Zhang, 2021).

Executive dashboards go beyond being simple displays of figures; the quality of the information and the manner in which it is presented influence the way people think. Hjelle et al. (2024) show that the format, the currency, and the completeness of the information have an indirect effect on the quality of the decisions by influencing people's perception of the complexity of the task and their level of satisfaction with the information. Similarly, Burnay et al. (2024) reach the same conclusion, connecting the use of dashboards to cognitive loads that are informational, representational, and non-informational. Although guidance from visual analytics can assist users in dealing with complicated analytical tasks, this kind of guidance must remain responsive to the context and to the user's degree of control (Han & Schulz, 2023). The fact that these findings have been made indicates that if the visual and semantic structure beneath a dashboard is weak, merely adding a conversational model to it may make the information problems worse rather than improving them.

Large language models can retrieve, summarise, transform, and generate natural-language versions of organisational information, thereby allowing executives to ask further questions, request explanations, compare different scenarios, or create a brief summary on the basis of an overview of the entire portfolio. Research on generative AI shows that productivity is increased in knowledge work and in customer-facing tasks, but it also reveals varied effects and that performance is influenced by both the user's level of expertise and the structure of the task (Noy & Zhang, 2023; Brynjolfsson et al., 2025). Reviews highlight that generative models provide new possibilities for organisational analysis but at the same time lead to problems concerning accuracy, governance, ethics, and verification (Dwivedi et al., 2023; Feuerriegel et al., 2024). In the area of visualisation, generative methods can be of assistance with data enhancement, visual mapping, styling, and interaction, although reliable evaluation and the preservation of faithful connections between the generated outputs and the source data remain unresolved issues (Ye et al., 2024).

Research into the incorporation of AI into public-sector organisations in Saudi Arabia shows that the amount of attention given by leaders, the organisation's institutional priorities and the internal conditions of the organisation have an effect on the manner in which AI moves from experimental phases to operational application (Alshahrani et al., 2022). Research on digital transformation in Saudi institutions has identified top-management support, strategy, technology, organisational culture, resources, data and implementation capability as both key success factors and as challenges (Alojail et al., 2023). Evidence relating to the adoption of analytics in Saudi Arabia also highlights the interaction between technological, organisational and competitive conditions (Alaskar et al., 2021). More recently, practical evidence from AI projects in Saudi Arabia has demonstrated that agile and hybrid project methods, a clear definition of roles and governance suitable to the aim of the initiative are important for AI projects (Alsumari & Alrwais, 2025). It therefore follows that an executive dashboard for Saudi projects should be regarded as a socio-technical decision architecture, not merely as a software feature.

The article examines a gap that arises when the four different fields usually studied separately are considered: AI-assisted project management, executive dashboards and visualisation, generative decision support, and Saudi digital transformation. Its key argument is that the decision value of an AI-enhanced dashboard is the result of a disciplined method involving both the separation and integration of five layers: trusted project evidence, analytical inference, generative interpretation, visual presentation, and human decision authority. It thus suggests a governance-aware framework for the design and evaluation of executive dashboards which use generative AI to enhance rather than replace project leadership.

II. AIM, OBJECTIVES AND REVIEW QUESTIONS

The aim of this review is to make a critical examination of how generative artificial intelligence can be integrated into executive project dashboards in order to include an interpretive and conversational element alongside visual analytics, and to set up a governance-aware framework for project decision-making in Saudi Arabia; the focus is on the whole executive decision interface, that is, on the processes involved in gathering, presenting, explaining, questioning, challenging and turning evidence into accountable action.

The review first distinguishes the role of generative AI from that of traditional business intelligence and predictive analytics. It then examines how visual design and conversational interaction affect cognitive load, trust, and the quality of the decisions made. Following this, it connects the functions of AI-augmented dashboards to various kinds of project decisions involving cost, schedule, risk, resources, scope, stakeholders, and portfolio priorities. Then it looks at the Saudi-specific circumstances relating to organisational readiness and governance which might either facilitate or obstruct implementation. Finally, it proposes a human-governed architecture and a research programme aimed at the rigorous evaluation of these approaches in actual project settings.

Consequently, the following four questions are raised: What complementary roles should generative AI and visual analytics play in an executive project dashboard? What mechanisms exist by which their integration can improve the speed, quality, and explainability of project decisions without causing overreliance? Which data, organisational, and governance conditions are most important for adoption in Saudi project settings? What design principles are necessary in order to keep executive accountability while at the same time enabling conversational and generative features to add decision value?

III. REVIEW METHODOLOGY

An integrative review using a structured method was selected since the research problem involves a number of well-established but methodologically diverse areas. A narrow meta-analysis would not be appropriate because the studies on project AI, dashboards, human-AI collaboration, generative models, and Saudi digital transformation make use of different constructs, outcomes, datasets, and research designs. The review therefore includes a definite search procedure along with thematic synthesis, with the aim of achieving conceptual integration rather than that of obtaining a pooled effect size.

The time frame for the review was defined so as to encompass the period from 2020 to 2025, in order that it might include recent research on dashboards, the rapid development of generative AI, and the current level of organisational adoption. The searches were divided into four groups of keywords: 'artificial intelligence' or 'generative AI' combined with 'project management'; 'dashboard' or 'data visualisation' combined with 'decision support'; 'human-AI collaboration' together with 'trust', 'explainability', or 'decision-making'; and 'Saudi Arabia' combined with 'artificial intelligence', 'analytics', or 'digital transformation'. Preference was given to studies that had been peer reviewed, were written in English, included a DOI, and made a direct contribution to at least one of the analytical areas. Studies were excluded if they were purely promotional, did not have a clear record of scholarly publication, focused on technical performance without any consequences for decision-making, or fell outside the date range. The final analytical corpus consisted of thirty core studies.

For each study the research context, decision task, AI capability, data type, visual or interaction mechanism, human role, reported benefit, implementation condition, and governance risk were all given a code.

The first thing we did was compare the concepts in each of the literature streams in order to identify the areas of agreement and the areas of tension. The second thing involved examining the relationships between the streams, for instance the relationship between the completeness of the information presented in the dashboard and the grounding requirements of a generative model, or the connection between AI explanation and the risk of unwarranted reliance. The third step was to combine these relationships and create a layered conceptual framework for executive project dashboards.

The methodological quality was assessed in terms of relevance, transparency, publication status, and the extent to which the evidence was appropriate for the claim made. The experimental findings were not generalised as though they applied to all possible project situations, and the Saudi studies were not regarded as being interchangeable with evidence from other areas. It is necessary to have this boundary discipline since an AI-augmented dashboard is a complex intervention: its impact depends at the same time on the quality of the data, the behaviour of the model, the visual design, the users' level of competence, the organisation's routines, and the stakes involved in the decisions.

Table 1. Review domains and analytical contribution

Evidence domain	Primary contribution to the review	Representative studies
Project AI and project control	Planning, forecasting, uncertainty, risk, schedule, resources and control	Bento et al. (2022); Taboada et al. (2023); Mohammad & Chirchir (2024); Nenni et al. (2025)
Dashboard and visual analytics	Information quality, cognitive load, guidance and decision representation	Han & Schulz (2023); Hjelle et al. (2024); Burnay et al. (2024)
Human-AI collaboration	Delegation, trust calibration, explainability, overreliance and human control	Glikson & Woolley (2020); Puranam (2021); Bansal et al. (2021); Buçinca et al. (2021); Fügener et al. (2022)
Generative AI and decision support	Natural-language synthesis, option generation, productivity and generative visualisation	Noy & Zhang (2023); Dwivedi et al. (2023); Feuerriegel et al. (2024); Handler et al. (2024); Ye et al. (2024); Brynjolfsson et al. (2025)
Saudi adoption context	AI assimilation, analytics adoption, digital transformation and AI project practices	Alaskar et al. (2021); Alshahrani et al. (2022); Alojail et al. (2023); Alsumari & Alrwais (2025)

IV. CONCEPTUAL FOUNDATIONS: FROM REPORTING DASHBOARDS TO GENERATIVE DECISION INTERFACES

A typical executive dashboard presents the status of a project in just a few indicators, trends, variances, alerts, and detailed views that allow further investigation. The advantage of such a dashboard is what is known as perceptual economy: because of the use of visual representations, users are able to spot outliers, compare values, and keep track of changes more quickly than by reading several reports. Although a red status indicator can be seen at once, the assumptions, dependencies, or narrative evidence which led to it are hidden. Hjelle et al. (2024) show that support for decision-making depends not only on the availability of data but also on how up-to-date, complete, and appropriately formatted the information is. Burnay et al. (2024) also point out that badly designed content can result in cognitive load, which in turn reduces acceptance and usefulness.

Rather than simply going through pre-set filters, an executive could ask why a particular milestone has shifted, what assumptions are leading to the cost exposure, what new risks have been mentioned in the correspondence with the contractor, or how a two-month delay in the procurement process might impact the downstream packages. The model is able to convert natural-language questions into retrieval and analytical operations, combine the results, and provide a narrative together with the appropriate charts. Likewise, Handler et al. (2024) regard large language models as belonging to a new category of decision-support tools since they can take part in all stages of intelligence gathering, option development, and decision preparation.

The key point lies in the difference between generation and evidence. Since language models are built to create plausible sequences rather than to act as authoritative project databases, an executive dashboard has to base its generative outputs on controlled data sources and maintain a clear link from each statement to the relevant evidence. Instead of calculating an earned-value figure from memory when a regulated project-control service is able to do so in a deterministic way, the generative component should explain the metric, link it to related evidence, point out any uncertainty, pose further questions, and assist the user in exploring different scenarios. This approach illustrates the broader automation-augmentation paradox: organisational value does not arise from giving every task to AI, but from deciding which tasks should be automated, which should be augmented, and where human judgement must still play an essential role (Raisch & Krakowski, 2021).

Puranam (2021) presents the issue of collaborative decision-making as one that relates to the design of an organisation, specifically concerning the allocation of tasks and the coordination between humans and AI. Fügener et al. (2022) indicate that productive delegation is cognitively demanding since individuals have to decide whether algorithmic advice is useful or whether their own expertise should take priority. Glikson and Woolley (2020) show that trust in AI has cognitive and emotional roots and depends on how the system is represented and on its capabilities. Bansal et al. (2021) found that explanations can lead to an increased acceptance of AI recommendations even if such acceptance is not justified, while Buçinca et al. (2021) demonstrate that cognitive forcing mechanisms can reduce overreliance by getting users to think before accepting AI advice. Shin (2021) connects explainability and causability with trust and acceptance, but the result for executive dashboards is that explanation should facilitate verification rather than persuasion.

A human-centred design approach therefore calls for reliable data, transparent boundaries, reversible interaction, and meaningful user control (Shneiderman, 2020). Interpretability research also cautions against considering explanation tools as being universally self-explanatory since users have different objectives and levels of expertise when carrying out interpretation tasks (Kaur et al., 2020). This implies that for project executives the dashboard should adjust the level of detail rather than reducing every issue to a single narrative. While a board-level summary might include just one sentence on exposure and the urgency of decisions, a project director may need the underlying schedule path, the evidence extracts, the model assumptions, and the sensitivity analysis. Generative AI is valuable exactly because it can vary the level of explanation, provided that the source layer and access controls stay stable.

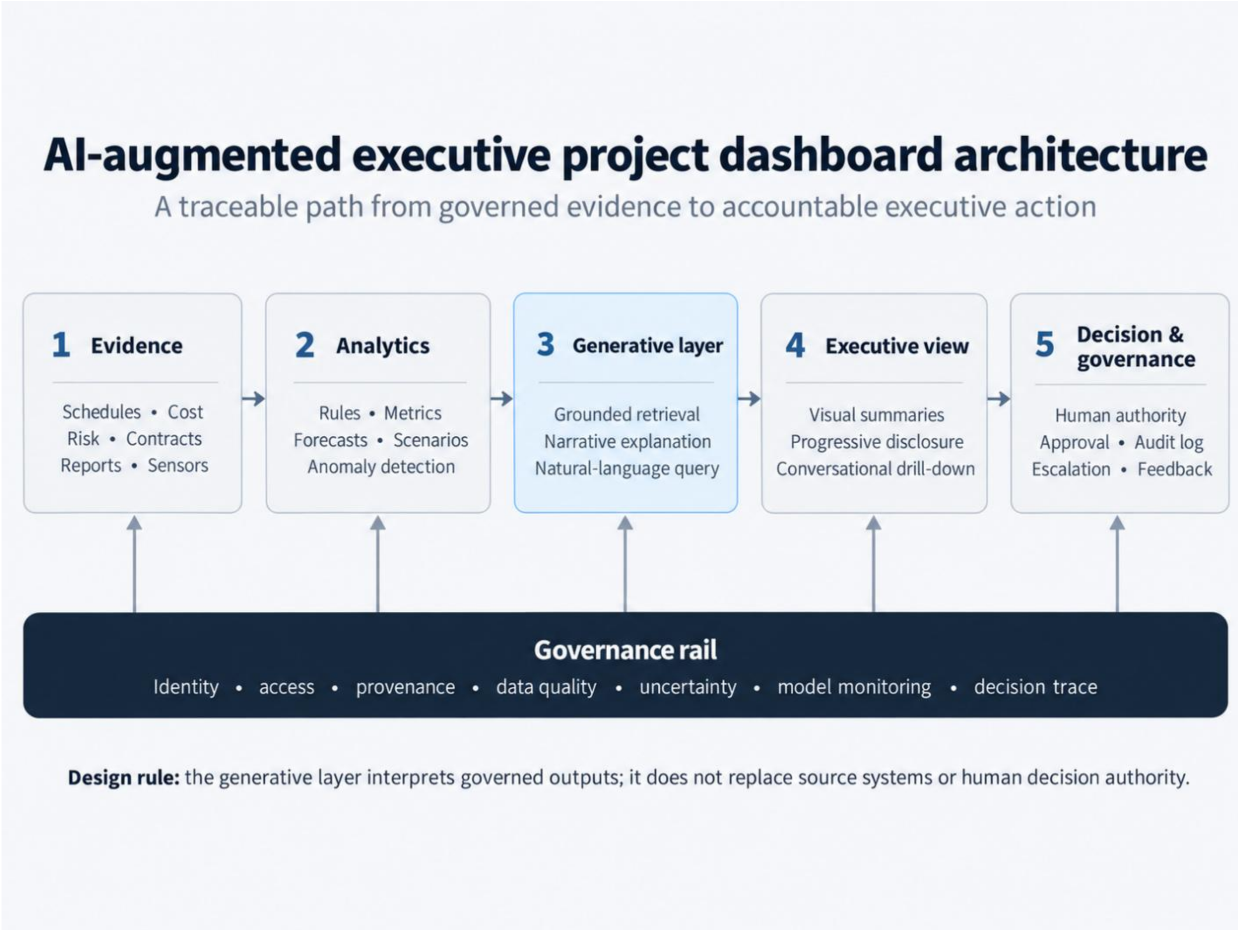

Figure 1. Layered architecture for a governed AI-augmented executive project dashboard

V. THEMATIC SYNTHESIS AND DISCUSSION

5.1 Data and semantic foundations

The literature continually indicates that both AI performance and managerial usefulness come as a result of data quality. Project settings are especially challenging since the information is spread out among structured systems and narrative records. In the areas of project and construction management, AI reviews point out that data availability, standardisation, integration, and organisational maturity are repeated limitations (Abioye et al., 2021; Pan & Zhang, 2021; Mohammad & Chirchir, 2024). Hence, an executive dashboard which combines generative AI with visual analytics must have a governed semantic layer that specifies entities such as project, work package, milestone, contract, risk, change, supplier, cost account, and decision. If there are no common definitions, the language model might generate smooth comparisons between metrics that are in fact not comparable.

The quality of executive decisions is influenced by when a datum was valid, when it was reported, and whether or not it has been superseded. Hjelle et al. (2024) demonstrate that the currency of information affects the quality of decisions, which means that the visibility of timestamps should be considered a design requirement and not just a technical one. Statements that are generated should differentiate between live data and approved baselines, between forecasts and actual figures, and between model inference and human judgement. In cases where data are missing, the system must indicate the gap rather than quietly assign a value.

5.2 Visualisation, cognition and conversational interaction

Charts are useful for carrying out comparisons, identifying trends, showing distributions, and detecting anomalies; language is useful for providing explanations, making qualifications, achieving synthesis, and enabling dialogue. When the two are combined, the amount of effort needed to go from seeing a problem to understanding its possible causes is reduced. As Burnay et al. (2024) point out, both the informational and representational loads are important for dashboard adoption, whereas Han and Schulz (2023) stress the importance of guidance that supports rather than restricts analytical reasoning. The design aim must therefore be progressive disclosure, that is, only the minimum amount of evidence necessary for orientation should be displayed, after which the user should be allowed to request information about the causes, the assumptions, the source records, the alternatives, or more detailed visual information.

A user could for instance enquiry as to which five projects have the greatest impact on portfolio schedule exposure and then ask for a visual comparison, an explanation of the factors involved, or for the source records relating to one project. Ye et al. (2024) refer to this kind of possibility in the areas of data enhancement, visual mapping, styling, and interaction. The system must choose the appropriate chart types on the basis of established design rules, maintain the scales and units, avoid misleading truncation, and clearly label the synthetic scenarios.

5.3 Generative AI as an interpretive layer

The best role for generative AI is to provide interpretation in the context of governed analytics. Although predictive models are able to estimate slippage or cost exposure, the generative component can convert the results into language suitable for executives, compare this with recent narrative evidence and suggest questions that warrant further investigation. This approach safeguards the integrity of the numerical data while making use of the model's strengths in handling language. Feuerriegel et al. (2024) define generative AI as a type of system capable of producing new content across different modalities, whereas Dwivedi et al. (2023) highlight both the organisational opportunities and the policy challenges. In relation to dashboard design, it follows that each generated narrative should include provenance, that is to say the sources consulted, the time at which they were retrieved, the model or rule that was applied, and the level of uncertainty.

Rather than making a recommendation, the system could present a number of alternatives such as delaying the package, reallocating a resource, speeding up the procurement process, or accepting a specific risk. It could then make use of deterministic or predictive services to assess the consequences where this is possible and to summarise those considerations which cannot be quantified. Handler et al. (2024) point out that the development of options is one field in which language models can enhance decision support. This is especially useful in projects since many executive decisions do not involve prediction but instead consist of trade-offs between cost, time, quality, strategic value, contractual exposure, and stakeholder consequences.

Noy and Zhang (2023) observed improvements in both writing speed and the quality of output when it came to professional writing tasks, whereas Brynjolfsson et al. (2025) stated that the productivity gains in customer support depended on the level of experience of the workers. Moreover, they also introduce a risk whereby high-quality output might be taken for expert judgement, particularly in cases where the system does not have adequate context.

5.4 Project decision domains

Research into Project AI suggests that decision augmentation should be based on managerial questions not on technologies. Taboada et al. (2023) find planning, measurement, and uncertainty to be the key areas in which AI is being used in project management, and Bento et al. (2022) and Nenni et al. (2025) record applications in the areas of prediction, risk, scheduling, resource management, and control.

When making cost decisions, the dashboard should incorporate the approved budget, the commitments, the actual figures, the forecast at completion, the changes, the claims, and the contingency consumption. Generative AI is able to identify the main causes of variance and retrieve the contractual or operational narratives related to these causes, even

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723455

though the quantitative calculations stay deterministic. In the case of schedule decisions, it can link together information on critical-path movements, milestone confidence, constraint logs, and procurement status, then summarise the dependencies that make the forecast vulnerable. As for risk, the system can pick up on new themes in the narrative sources and associate them with registered risks, but decisions regarding risk acceptance and ownership must still be made explicitly by humans.

The executive dashboard's value lies not just in displaying each project separately, but in revealing the common bottlenecks, the competing demands for resources, the correlated exposure to suppliers, and the strategic interdependencies. With the aid of generative querying, executives can shift from looking at a portfolio heat map to formulating a specific question as to why a number of projects are deteriorating at the same time. When it comes to scope and change decisions, it is helpful to have traceability between the requested changes, the approved baselines, the effects on cost and schedule, and the commitments made to stakeholders. Although stakeholder decisions can make use of summarisation to combine the issue logs and communications, sensitive sentiment or any inferred intent should be handled with care since language models might overinterpret incomplete evidence.

The level of confidence one can have in a deterministic calculation, in a statistical forecast, in the classification of a document, and in a generative narrative are all different things. The dashboard should present the uncertainty in a form that is suitable to the type of output obtained: by showing data completeness for a metric, by giving a prediction interval for a forecast, by indicating classification probability where this is meaningful, and by showing evidence coverage for a generated summary.

5.5 Saudi organisational and project context

Alshahrani and colleagues in 2022 indicate that the extent to which AI is adopted in the public organisations of Saudi Arabia is determined by the way decision-makers distribute their attention between technological, organisational, and environmental matters. Similarly, Alaskar and co-authors in 2021 find that the adoption of analytics in Saudi supply chains is based on a mix of organisational and technological conditions when there is competitive pressure. Alojail and others in 2023 identify leadership support, strategy, technology, vendors, budget, culture, infrastructure, data, and resistance as key factors affecting digital transformation in a Saudi ministry.

Concerns regarding data readiness include common identifiers, master data, access rights, historical depth, and quality controls. With regard to analytical readiness, validated models, the definitions of metrics, and integration services are important. Organisational readiness involves decision rights, leadership sponsorship, user competence, and workflows that are based on the outputs from dashboards (Jöhnk et al., 2021). Governance readiness covers issues relating to privacy, auditability, model oversight, cybersecurity, and accountability. Even if a highly capable language model is used, it cannot make up for the lack of ownership of schedule baselines, and a powerful data platform will not be able to generate value if executives keep on depending on separate manual reports.

Evidence from the Saudi AI projects also shows how important adaptive delivery methods are. Alsumari and Alrwais (2025) state that agile approaches are widely used among the Saudi AI practitioners who were surveyed, while at the same time mentioning the difficulties associated with leadership, roles, and the use of hybrid methods. A reasonable approach would be to start with one area of decision-making, for example schedule and risk, to define a small number of high-value questions for senior management, to set up governed data links, and then to check whether the interface improves the quality of decisions. Expansion to the areas of cost, contracts, resources, and portfolio optimisation should only take place after groundings, latency, and governance controls have been demonstrated.

Verhoef and others in 2021 claim that digital transformation affects the way value is created, the organisation's structure, and its strategic processes, rather than simply altering technology assets. In the context of the project in Saudi Arabia, the dashboard should therefore be incorporated into portfolio reviews, governance meetings, escalation procedures, and decision records. The most important implementation metric is not the number of prompts that are submitted, but whether decisions become more timely, better supported by evidence, more traceable, and more in line with strategic priorities.

5.6 Trust, explainability and accountable human oversight

Trust should be adjusted rather than pushed to the highest level. Glikson and Woolley (2020) indicate that trust in AI emerges as a result of various cognitive and emotional mechanisms, meaning that a nicely designed dashboard can have an effect on confidence without any connection to real reliability. Bansal et al. (2021) show that explanations can lead to greater acceptance even if they do not result in an improvement of complementary team performance. It follows that the dashboard should make it easy to have disagreements and to carry out verification.

Practical safeguards involve asking the system to refer to the original records when making claims, showing alternative interpretations in cases where the evidence is unclear, distinguishing between facts and the model's inferences, making the assumptions visible, and keeping a clear space for a human decision. According to Buçinca et al. (2021), cognitive forcing can reduce excessive dependence on the system; in practice, a similar mechanism could require the executive to make an initial judgement before seeing the AI's recommendation in relation to important decisions, or could mandate a careful examination of the contrary evidence. Shin (2021) states that both explainability and causability influence trust and acceptance, but the way in which explanations are provided should be adapted to the particular decision; the project director may require causal factors and sensitivity, while the audit function may need provenance and control evidence.

The principle of human-centred AI also includes the idea of bounded autonomy. Shneiderman (2020) stresses the importance of systems that are reliable, safe, and trustworthy, with a strong element of human control. Puranam (2021) and Fügener et al. (2022) also point out that performance involving humans and AI depends on how tasks are allocated and on productive delegation. Hence, an executive dashboard may act autonomously in retrieving data, preparing summaries, highlighting anomalies, or drafting a briefing, but it should not act autonomously in approving changes, accepting risk, rebaselining a project, committing money, or making consequential decisions unless there is a separate workflow which specifically authorises such actions. This boundary does not mean rejecting automation; rather, it is a kind of control design in which the authority of the system is aligned with organisational accountability.

VI. PROPOSED FRAMEWORK FOR AI-AUGMENTED EXECUTIVE DASHBOARDS

The synthesis proposes a five-layer model for Saudi project environments. The first of these layers is the evidence layer, which includes governed project systems, documents, and approved external data; each item must have an assigned owner, include timestamps, have an access classification, and be accompanied by quality metadata. The second layer is the analytical layer, in which deterministic calculations, statistical models, and rules are used to turn the evidence into metrics, forecasts, anomalies, and scenario outputs. The third layer is the generative layer, which interprets the governed outputs, retrieves the relevant evidence, answers questions posed in natural language, and produces draft explanations or scenario narratives. This layer should be based on retrieval and must not present content that is not supported as fact.

The fourth layer is that of the executive experience layer: visualisations enable a quick understanding of the current state of the portfolio whilst conversational interaction allows for flexible questioning. The interface must make use of progressive disclosure, adhere to consistent visual conventions, specify the units being used, show the recency of the data and include clear indications of uncertainty. Each statement that is generated must provide a way of accessing the underlying records. The fifth layer concerns governance and decision-making, and it is in this layer that roles, approvals, audit logs, model monitoring, escalation rules and decision records are kept. Human decision-making authority is located here, outside the generative model, and feedback regarding the decisions and their outcomes then goes back to the evidence and analytical layers, thus

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723455

enabling controlled learning without the need to rewrite the historical records.

The framework outlines four design principles. The first of these is to ground generation by having the system obtain approved evidence prior to giving an explanation. The second is to show rather than tell: important claims should be accompanied by the relevant visual or source view rather than being given solely in the form of prose. The third principle is to separate advice from authority since although the AI can provide options it is the various organisational roles that decide who has the power to make decisions. The fourth is to keep contestability, so that users should be able to examine, challenge, correct and add comments to the AI's outputs. Thus, these principles turn responsible AI from a policy statement into actual interface behaviour.

Evaluation should also be multi-dimensional since although technical accuracy is necessary it is not enough. Field studies need to assess decision time, decision quality, the extent of evidence coverage, user calibration, the number of times outputs are challenged, the amount of rework involved, the quality of escalations, and the outcomes of downstream projects. Adoption figures must be carefully interpreted since frequent use can still go hand in hand with poor judgement. The most convincing evidence would be to compare conventional dashboards with AI-augmented ones based on identical project decisions and at the same time note whether the benefits vary according to the level of executive experience, project complexity, and data maturity.

Table 2. Design and governance requirements by executive decision domain

Decision domain	Governed evidence	Generative function	Visual cue	Human control
Cost	Budget, commitments, actuals, EAC, change and claims	Explain major variance drivers; retrieve supporting narratives	Variance waterfall; trend; scenario bands	No autonomous commitment or reforecast approval
Schedule	Baseline, critical path, milestone confidence, constraints, procurement	Explain slippage drivers; connect dependencies; formulate what-if questions	Milestone confidence; dependency map; exception trend	Rebaseline requires authorised human approval
Risk	Risk register, issues, correspondence, indicators	Surface emerging themes; link evidence to registered risks	Risk trajectory; exposure matrix; evidence coverage	Risk acceptance and ownership remain human decisions
Portfolio & resources	Cross-project demand, suppliers, capacity, strategic value	Synthesize interdependencies and shared bottlenecks	Portfolio heat map; resource contention view	Priority changes require governance forum approval
Scope & change	Change requests, baseline scope, contract and impact assessment	Summarise change rationale and cross-domain effects	Change pipeline; impact comparison	AI may draft, not approve, formal change
Stakeholders & governance	Issue logs, decisions, communications, approvals	Prepare evidence-grounded briefings; identify unresolved actions	Decision trail; escalation view	Sensitive inference must be reviewable and contestable

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723455

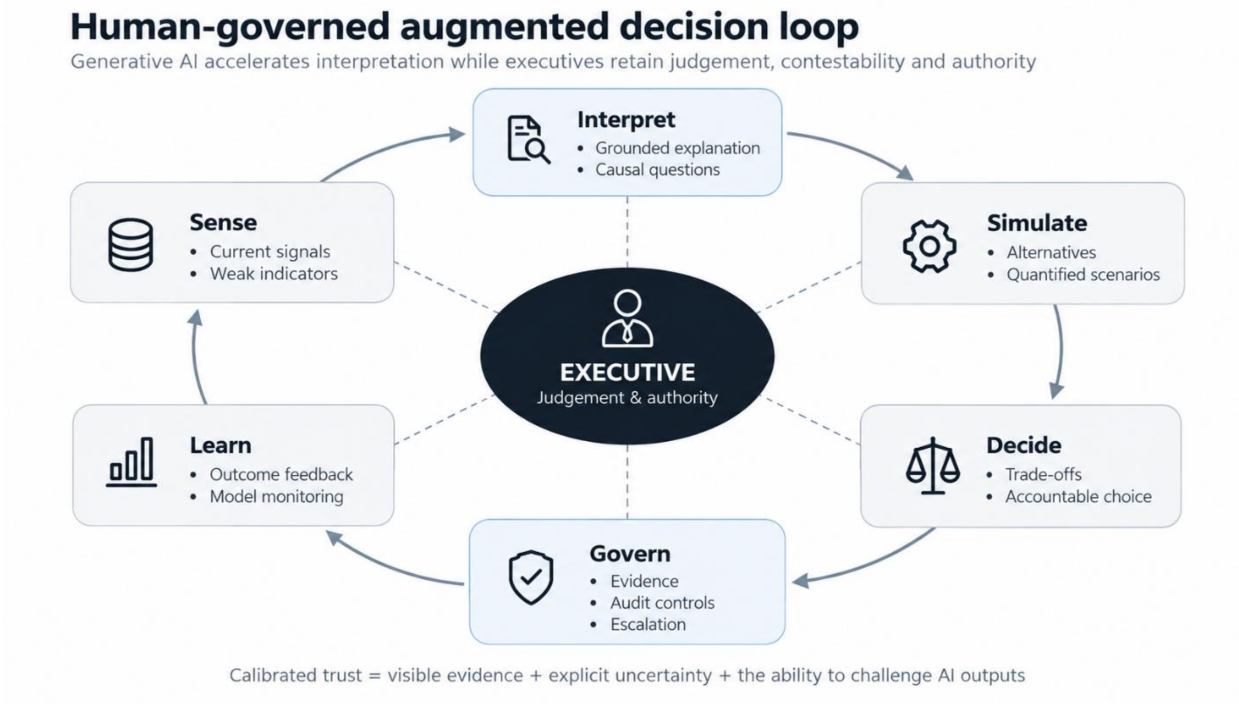

Figure 2. Human-governed augmented decision loop for executive project decisions

VII. RESEARCH AGENDA AND LIMITATIONS

There are several research gaps that need to be addressed through empirical investigation. First, the existing literature does not contain sufficient field evidence to show whether conversational dashboards lead to better executive project decisions than well-designed traditional dashboards. Controlled studies should be carried out to separate out the individual contributions of visual quality, generative explanation, and predictive analytics rather than considering them as part of a single feature set. Second, trust calibration must be assessed behaviourally, taking into account the situations in which users reject wrong AI outputs, and not merely by means of satisfaction surveys. Third, research should be conducted in Saudi Arabia focusing on Arabic-English interaction, local project terminology, and the impact of bilingual evidence on the quality of retrieval and summarisation.

Fourth, large-scale portfolio studies are required since the majority of AI research focuses on individual tasks or models. The executive value that may result from linking weak signals across different projects also leads to greater difficulties regarding privacy, access control, and semantic integration. Fifth, the various governance mechanisms should be subject to experimental comparison, as provenance displays, counter-evidence prompts, confidence communication, and compulsory human confirmation might have different impacts on speed and accuracy. Finally, long-term research should look into whether generative dashboards lead over time to changes in the organisation's attention, its meeting behaviour, its escalation procedures, and the way decision authority is distributed.

Empirical validation should therefore look not just at whether executives make the right decisions, but also at whether the interface helps in the search for evidence, brings disagreements to light earlier, enables defensible escalation, and ensures that responsibility is maintained when the data are incomplete, in dispute, delayed, or interpreted differently at various levels of the organisation.

It creates a deliberately limited corpus covering the period 2020 to 2025 and gives preference to studies that have traceable DOI entries, which means that some relevant evidence from practitioners and more recent work might be left out. The body of evidence is drawn from a number of different disciplines and therefore does not allow for a statistical estimate of the effect. Various mechanisms are deduced by linking the results from studies on dashboards, those involving humans and AI, and those on generative AI rather than by referring to actual trials within Saudi project portfolios. It must therefore be regarded as a design proposal based on theoretical grounds that needs to be validated through real-world application, not as evidence that AI-augmented dashboards will improve every project decision.

VIII. CONCLUSION

When generative AI is used as an interpretive layer over controlled evidence and analytics rather than as an independent decision-maker, AI-augmented executive dashboards can become a major decision-support feature for project organisations in Saudi Arabia. The review indicates that visualisation, generative language, predictive analysis, and human judgement each have complementary advantages. Visualisations aid in the quick identification of patterns; analytical models provide quantitative assessments of current situations and forecasts; generative models allow for flexible questioning and synthesis; and it is the executives who offer context, values, accountability, and authority. The degree of value attained in decision-making depends on how well these various elements are coordinated.

The main design challenge is therefore to find a reliable way of getting from evidence to action rather than aiming to maximise the intelligence within the model. Reliable project data, semantic consistency, up-to-date information, cognitive fit, transparent uncertainty, provenance, and human contestability are all fundamental. Furthermore, the implementation in Saudi Arabia also relies on strong leadership attention, organisational readiness, digital infrastructure, governance, and project-delivery practices. It is because of these factors that the same technical dashboard can result in different outcomes in various organisations.

The five-layer framework that has been proposed divides up the components of evidence, analytics, generation, executive experience, and governance and links them together via traceable interfaces. It provides a practical basis for creating dashboards which are able to explain variations, compare different options, retrieve the relevant records, and assist with portfolio decisions all without undermining accountability. Future research should test the framework in actual project environments in Saudi Arabia by using measures of decision quality rather than merely looking at adoption. When generative AI is properly grounded and governed, it can transform executive dashboards from simple reporting into interactive project intelligence while still upholding the principle that significant project decisions remain the work of accountable human beings.

REFERENCES

[1] S. O. Abioye, L. O. Oyedele, L. Akanbi, A. Ajayi, J. M. Davila Delgado, M. Bilal, O. O. Akinade, and A. Ahmed, “Artificial intelligence in the construction industry: A review of present status, opportunities and future challenges,” *Journal of Building Engineering*, vol. 44, Art. no. 103299, 2021, doi: 10.1016/j.jobe.2021.103299. Crossref

[2] T. H. Alaskar, K. Mezghani, and A. K. Alsadi, “Examining the adoption of Big data analytics in supply chain management under competitive pressure: Evidence from Saudi Arabia,” *Journal of Decision Systems*, vol. 30, nos. 2–3, pp. 300–320, 2021, doi: 10.1080/12460125.2020.1859714. Taylor & Francis

[3] A. Alshahrani, D. Dennehy, and M. Mäntymäki, “An attention-based view of AI assimilation in public sector organizations: The case of Saudi Arabia,” *Government Information Quarterly*, vol. 39, no. 4, Art. no. 101617, 2022, doi: 10.1016/j.giq.2021.101617. ScienceDirect

[4] M. Alojail, J. Alshehri, and S. B. Khan, “Critical success factors and challenges in adopting digital transformation in the Saudi Ministry of Education,” *Sustainability*, vol. 15, no. 21, Art. no. 15492, 2023, doi: 10.3390/su152115492. Crossref

[5] W. Alsumari and O. Alrwais, “Artificial intelligence project management methodologies: Insights from field experts in Saudi Arabia,” *International Journal of Informatics, Information System and Computer Engineering*, vol. 6, no. 2, pp. 246–258, 2025, doi: 10.34010/injiiscom.v6i2.16127. INJIISCOM

[6] G. Bansal, T. Wu, J. Zhou, R. Fok, B. Nushi, E. Kamar, M. T. Ribeiro, and D. S. Weld, "Does the whole exceed its parts? The effect of AI explanations on complementary team performance," in *Proc. 2021 CHI Conference on Human Factors in Computing Systems*, 2021, pp. 1–16, doi: 10.1145/3411764.3445717. ACM

[7] S. Bento, L. Pereira, R. Gonçalves, Á. Dias, and R. Lopes da Costa, "Artificial intelligence in project management: Systematic literature review," *International Journal of Technology Intelligence and Planning*, vol. 13, no. 2, pp. 143–163, 2022, doi: 10.1504/IJTIP.2022.126841. Crossref

[8] E. Brynjolfsson, D. Li, and L. R. Raymond, "Generative AI at work," *The Quarterly Journal of Economics*, vol. 140, no. 2, pp. 889–942, 2025, doi: 10.1093/qje/qjae044. Crossref

[9] Z. Buçinca, M. B. Malaya, and K. Z. Gajos, "To trust or to think: Cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making," *Proceedings of the ACM on Human-Computer Interaction*, vol. 5, no. CSCW1, Art. no. 188, pp. 1–21, 2021, doi: 10.1145/3449287. ACM

[10] C. Burnay, M. Lega, and S. Bouraga, "Business intelligence and cognitive loads: Proposition of a dashboard adoption model," *Data & Knowledge Engineering*, vol. 152, Art. no. 102310, 2024, doi: 10.1016/j.datak.2024.102310. Crossref

[11] Y. K. Dwivedi *et al.*, "'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*, vol. 71, Art. no. 102642, 2023, doi: 10.1016/j.ijinfomgt.2023.102642. Crossref

[12] S. Feuerriegel, J. Hartmann, C. Janiesch, and P. Zschech, "Generative AI," *Business & Information Systems Engineering*, vol. 66, pp. 111–126, 2024, doi: 10.1007/s12599-023-00834-7. Springer

[13] A. Fügener, J. Grahl, A. Gupta, and W. Ketter, "Cognitive challenges in human-artificial intelligence collaboration: Investigating the path toward productive delegation," *Information Systems Research*, vol. 33, no. 2, pp. 678–696, 2022, doi: 10.1287/isre.2021.1079. Crossref

[14] E. Glikson and A. W. Woolley, "Human trust in artificial intelligence: Review of empirical research," *Academy of Management Annals*, vol. 14, no. 2, pp. 627–660, 2020, doi: 10.5465/annals.2018.0057. Crossref

[15] W. Han and H.-J. Schulz, "Providing visual analytics guidance through decision support," *Information Visualization*, vol. 22, no. 2, pp. 140–165, 2023, doi: 10.1177/14738716221147289. Crossref

[16] A. Handler, K. R. Larsen, and R. Hackathorn, "Large language models present new questions for decision support," *International Journal of Information Management*, vol. 79, Art. no. 102811, 2024, doi: 10.1016/j.ijinfomgt.2024.102811. Crossref

[17] S. Hjelle, P. Mikalef, N. Altwaijry, and V. Parida, "Organizational decision making and analytics: An experimental study on dashboard visualizations," *Information & Management*, vol. 61, no. 6, Art. no. 104011, 2024, doi: 10.1016/j.im.2024.104011. ScienceDirect

[18] J. Jöhnk, M. Weißert, and K. Wyrtki, "Ready or not, AI comes: An interview study of organizational AI readiness factors," *Business & Information Systems Engineering*, vol. 63, pp. 5–20, 2021, doi: 10.1007/s12599-020-00676-7. Springer

[19] H. Kaur, H. Nori, S. Jenkins, R. Caruana, H. Wallach, and J. W. Vaughan, "Interpreting interpretability: Understanding data scientists' use of interpretability tools for machine learning," in *Proc. 2020 CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–14, doi: 10.1145/3313831.3376219. ACM

[20] A. Mohammad and B. Chirchir, "Challenges of integrating artificial intelligence in software project planning: A systematic literature review," *Digital*, vol. 4, no. 3, pp. 555–571, 2024, doi: 10.3390/digital4030028. Crossref

[21] M. E. Nenni, F. De Felice, C. De Luca, and A. Forcina, "How artificial intelligence will transform project management in the age of digitization: A systematic literature review," *Management Review Quarterly*, vol. 75, pp. 1669–1716, 2025, doi: 10.1007/s11301-024-00418-z. Springer

[22] S. Noy and W. Zhang, "Experimental evidence on the productivity effects of generative artificial

intelligence," *Science*, vol. 381, no. 6654, pp. 187–192, 2023, doi: 10.1126/science.adh2586. Crossref

[23] Y. Pan and L. Zhang, "Roles of artificial intelligence in construction engineering and management: A critical review and future trends," *Automation in Construction*, vol. 122, Art. no. 103517, 2021, doi: 10.1016/j.autcon.2020.103517. Crossref

[24] P. Puranam, "Human-AI collaborative decision-making as an organization design problem," *Journal of Organization Design*, vol. 10, pp. 75–80, 2021, doi: 10.1007/s41469-021-00095-2. Springer

[25] S. Raisch and S. Krakowski, "Artificial intelligence and management: The automation-augmentation paradox," *Academy of Management Review*, vol. 46, no. 1, pp. 192–210, 2021, doi: 10.5465/amr.2018.0072. Crossref

[26] D. Shin, "The effects of explainability and causability on perception, trust, and acceptance: Implications for explainable AI," *International Journal of Human-Computer Studies*, vol. 146, Art. no. 102551, 2021, doi: 10.1016/j.ijhcs.2020.102551. Crossref

[27] B. Shneiderman, "Human-centered artificial intelligence: Reliable, safe & trustworthy," *International Journal of Human-Computer Interaction*, vol. 36, no. 6, pp. 495–504, 2020, doi: 10.1080/10447318.2020.1741118. Taylor & Francis

[28] Taboada, A. Daneshpajouh, N. Toledo, and T. de Vass, "Artificial intelligence enabled project management: A systematic literature review," *Applied Sciences*, vol. 13, no. 8, Art. no. 5014, 2023, doi: 10.3390/app13085014. Crossref

[29] P. C. Verhoef, T. Broekhuizen, Y. Bart, A. Bhattacharya, J. Q. Dong, N. Fabian, and M. Haenlein, "Digital transformation: A multidisciplinary reflection and research agenda," *Journal of Business Research*, vol. 122, pp. 889–901, 2021, doi: 10.1016/j.jbusres.2019.09.022. Crossref

[30] Y. Ye, J. Hao, Y. Hou, Z. Wang, S. Xiao, Y. Luo, and W. Zeng, "Generative AI for visualization: State of the art and future directions," *Visual Informatics*, vol. 8, no. 2, pp. 43–66, 2024, doi: 10.1016/j.visinf.2024.04.003. ScienceDirect