Global Academic Journal of Economics and Business

Available online at https://www.gajrc.com
DOI: https://doi.org/10.36348/gajeb.2026.v08i05.037

ISSN: 2706-9001 (P)
ISSN: 2707-2584 (O)

Original Research Article

Intelligent Decision Automation for Saudi Digital Government: Integrating AI, Business Rules and Case Management for Vision 2030

Arshad Pasha Mohammed[1]*
Business Development Manager (Non Fuel Retail), Shell Wafi Saudi: Riyadh, SA

***Corresponding Author**
Arshad Pasha Mohammed
Business Development Manager (Non Fuel Retail), Shell Wafi Saudi: Riyadh, SA

Article History
Received: 29.07.2026
Accepted: 24.09.2026
Published: 28.09.2026

Abstract: Saudi Arabia's digital-government agenda increasingly depends on decisions that must be fast, consistent, explainable and adaptable rather than merely digitised. This review examines how artificial intelligence, business rules and case management can be combined into an intelligent decision-automation architecture suitable for public services operating under Vision 2030. The paper synthesises peer-reviewed literature published from 2020 to 2025 across five intersecting streams: artificial intelligence in public administration, Saudi digital transformation, explainable and responsible artificial intelligence, hybrid decision management, and data-driven case management. The evidence indicates that artificial intelligence is strongest when used for prediction, classification, prioritisation and recommendation; business rules are strongest when policy, eligibility, thresholds and legal constraints require deterministic control; and case management is strongest when services involve incomplete information, exceptions, appeals and human judgement. Treating these capabilities as substitutes creates avoidable governance and operational risks. The review therefore proposes a layered framework in which trusted data and case context feed a decision-orchestration layer that coordinates artificial intelligence models, executable rules and case logic under human oversight. A seven-stage lifecycle—sense, validate, decide, explain, authorise, execute and learn—connects model outputs to accountable service actions and continuous improvement. The framework emphasises versioned policies, traceable evidence, appealability, bias monitoring, privacy controls and public-value metrics. The study contributes a practical architecture and governance model for Saudi public entities seeking to automate high-volume decisions without sacrificing legal compliance, institutional accountability or citizen trust.

Keywords: Artificial Intelligence, Digital Government, Decision Automation, Business Rules, Case Management, Explainable AI, Saudi Arabia, Vision 2030.

Copyright © 2026 The Author(s): This is an open-access article distributed under the terms of the Creative Commons Attribution **4.0 International License (CC BY-NC 4.0)** which permits unrestricted use, distribution, and reproduction in any medium for non-commercial use provided the original author and source are credited.

1. INTRODUCTION

Saudi Vision 2030 has placed digital transformation within a wider programme of institutional modernisation, economic diversification and service improvement. The transformation is not limited to moving forms and transactions online; it increasingly concerns how government entities use data, algorithms and connected platforms to make and implement decisions at scale. Earlier analysis of Vision 2030 emphasised the institutional significance of coordinated reform, while subsequent Saudi initiatives have positioned data and artificial intelligence as national capabilities rather than isolated technical projects [1, 2]. This shift matters because a digital service becomes genuinely intelligent only when it can interpret context, apply policy, manage uncertainty and act consistently across channels.

Research on Saudi public-sector artificial intelligence shows that assimilation is influenced by management attention, data quality, infrastructure, culture, governance and the alignment between technical capabilities and decision-making needs [3]. More recent case evidence also indicates that leadership attention and communication are decisive in translating artificial-intelligence initiatives into operational public-service change [4]. At the citizen interface, adoption of e-government services depends on trust, usability, motivation and perceptions of value [5]. At the infrastructure interface, application programming interfaces increasingly enable interoperable services, platform integration and automated transactions across organisational boundaries [6]. These developments suggest that the next stage of digital government is not a single artificial-intelligence application; it is an operational decision system connecting models, policies, workflows, data and accountable human intervention.

The international literature reinforces this interpretation. Reviews of artificial intelligence in public governance identify opportunities for efficiency, responsiveness and new service forms, but also unresolved questions about risk, accountability, performance measurement and the governance of scaling [7]. Research on adoption and diffusion in public administration similarly shows that public value depends on the full innovation lifecycle, not only on technical acquisition [8]. Empirical work on public-service delivery finds that artificial intelligence can improve efficiency when embedded in organisational processes rather than treated as a detached analytical tool [9]. Consequently, the central design problem is how to convert algorithmic capability into decisions that are lawful, understandable, operationally executable and adaptable to exceptional cases.

This paper addresses that problem through the concept of intelligent decision automation. The concept differs from conventional workflow automation in three ways. First, it treats a decision as a governed asset with explicit inputs, logic, evidence and outcomes. Second, it separates probabilistic intelligence from deterministic policy so that a model can estimate risk or recommend an action without silently rewriting eligibility rules. Third, it recognises that many government services are case-based: documents arrive asynchronously, facts change, exceptions occur, and citizens may appeal. In such settings, a rigid end-to-end workflow is insufficient, while unconstrained artificial intelligence is inappropriate.

The aim of this review is therefore to develop a publishable, evidence-based framework for integrating artificial intelligence, business rules and case management in Saudi digital government. Four objectives guide the analysis: to clarify the distinct role of each decision capability; to identify architectural and governance requirements for their integration; to propose a lifecycle that links automated decisions to explanations, authorisation, execution and learning; and to derive implementation priorities relevant to Vision 2030. The paper does not claim that every public decision should be automated. Instead, it asks where automation can safely increase speed and consistency, where human judgement remains necessary, and how both forms of decision-making can operate within one auditable service architecture.

2. Conceptual Foundations

Intelligent decision automation sits at the intersection of digital government, decision science and process management. Public-sector decisions differ from many commercial decisions because they are exercised under statutory authority, affect rights or access to services, and are expected to be reviewable. A useful architecture must therefore optimise more than predictive accuracy. It must preserve procedural legitimacy, policy consistency, traceability and the possibility of correction. Research on explainable artificial intelligence for government demonstrates that the type of explanation can influence perceived accuracy, fairness and trustworthiness, especially when an algorithmic decision directly affects an individual [10]. Work on public-sector decision toolkits further shows that machine-learning assumptions can become misaligned with real administrative decision conditions if technical design ignores institutional context [11]. Citizens also tend to prefer an advisory or bounded role for artificial intelligence rather than unrestricted autonomous authority in sensitive public decisions [12].

These findings give explainability a functional rather than cosmetic role. Explainability is not merely a visual layer placed over a completed prediction. It links a decision to the reasons, evidence and policy basis that justify action. Broad reviews of explainable artificial intelligence distinguish between transparent models, post-hoc explanations and different levels of interpretability [13]. Analytical work has further emphasised that explainability should be evaluated against the needs of the user receiving the explanation, not only against technical fidelity [14]. Evaluation research similarly warns that explainability lacks a single universal measure; usefulness depends on whether the explanation supports understanding, contestability and correct action [15]. For government, this implies at least three explanation audiences: the citizen, the case officer and the auditor. Each requires different detail while relying on the same underlying evidence.

A second foundation is hybrid decision management. Business rules encode deterministic logic such as eligibility thresholds, mandatory exclusions, document requirements, escalation conditions and legal constraints. Artificial intelligence, by contrast, is well suited to inference under uncertainty: predicting likely demand, classifying documents, detecting anomalies, estimating risk or prioritising cases. Hybrid decision-management research shows that combining machine-learning capabilities with structured decision models can improve analysis without surrendering the explicit logic that business users need to inspect and maintain [16]. Decision Model and Notation research likewise demonstrates that formal decision models can describe variability through understandable rules, separating decision logic from process sequencing [17]. The conceptual implication is important: artificial intelligence should enrich a governed decision model, not obscure the policy logic inside opaque code.

The third foundation is case management. Many public services are knowledge-intensive rather than fully predictable. Licensing, social support, complaints, regulatory investigations, complex permits and cross-agency requests often involve uncertain sequences, missing documents, exceptional facts and discretionary review. Research on data-driven case management shows that compliance can be checked against rules both at design time and while a case is running [18]. Research on exception handling further demonstrates that case-management approaches can support unforeseeable events without requiring every path to be modelled in advance [19]. This flexibility is essential for government because automation must handle the common path efficiently while allowing lawful deviation when facts justify it.

The Saudi context adds a fourth foundation: organisational and sectoral readiness. Evidence from Saudi organisations indicates that artificial intelligence can improve decision speed and quality when training, suitability and effectiveness are addressed [20]. Saudi applications of explainable artificial intelligence also show growing attention to transparency alongside technical performance [21]. Algorithmic decision systems are appearing in strategic organisational functions such as human-resource management, illustrating how automated recommendations increasingly affect consequential choices [22]. Sectoral reviews of digital finance similarly show that Saudi institutions are experimenting with artificial-intelligence-enabled decision capabilities under broader transformation objectives [23]. These developments strengthen the case for a common government architecture that can support different service domains without forcing each entity to reinvent controls.

Finally, intelligent decision automation must address risks that become more visible as generative and predictive systems expand. Bias can enter through data, modelling choices and contextual assumptions, creating unequal outcomes if monitoring is weak [24]. Generative systems also create privacy and security concerns when organisational rules, confidential data and externally generated content interact [25]. Managerial acceptance studies suggest that decision-makers become more cautious when artificial intelligence crosses from support into autonomous implementation, particularly in sensitive functional areas [26]. Together, these findings support a bounded-autonomy principle: models may generate evidence and recommendations, but policy rules, risk thresholds and human authorisation should determine when the system may act automatically.

3. REVIEW METHODOLOGY

The study uses a structured integrative review designed to combine conceptual synthesis with architecture development. The approach was selected because the research question spans several literatures that rarely appear in one search vocabulary: public-sector artificial intelligence, digital government, explainability, business-rule management, decision modelling, case management and Saudi digital transformation. Review-method guidance stresses transparent boundary setting, reproducible source selection and explicit synthesis rather than an unsystematic narrative summary [27–30]. Consistent with that guidance, the review was organised around a defined period, relevance criteria, source-quality requirements and a coding framework linked directly to the study objectives.

The review corpus contains thirty peer-reviewed sources published between 2020 and 2025. Inclusion required a persistent DOI and a direct contribution to at least one of six analytical domains: Saudi digital transformation; artificial intelligence in public administration; public-service decision quality and acceptance; explainable or responsible artificial intelligence; hybrid decision or business-rule management; and flexible case management or compliance. Sources were excluded when they were outside the publication window, lacked a verifiable DOI, focused only on model accuracy without decision or governance implications, or duplicated an argument already represented by a stronger source. The final set deliberately combines Saudi-specific evidence with international studies because an architecture for Saudi government requires local institutional relevance and transferable technical concepts.

Analysis proceeded in three stages. First, each paper was coded for its primary contribution, decision layer, governance concern, and implications for automation. Second, evidence was synthesised into recurring design themes: data and context, artificial-intelligence inference, deterministic policy rules, case progression, explainability, human authorisation, integration, monitoring and learning. Third, these themes were translated into a reference architecture and lifecycle. The synthesis is therefore interpretive rather than statistical; it does not estimate pooled effects. Instead, it asks which capabilities repeatedly appear necessary for trustworthy decision automation and how their interfaces should be designed.

Methodological quality was strengthened through source triangulation and reference integrity checks. No single study was used to justify the complete framework. Saudi case evidence was paired with international public-administration research, while technical design claims about rules and case management were supported by specialised information-systems literature. All thirty references were checked for DOI consistency and publication-year compliance before document preparation. The resulting method is appropriate for a review whose contribution is architectural synthesis rather than bibliometric ranking or meta-analysis.

4. Evidence Synthesis: From AI Tools to Decision Systems

The literature converges on a central distinction between automating tasks and automating decisions. A chatbot may answer a question, a classifier may tag a document, and a predictive model may estimate risk, but a government decision additionally requires authority, policy logic, evidence, timing and a mechanism for review. Public-governance research repeatedly warns that benefits from artificial intelligence are contingent on implementation arrangements, performance evaluation and governance [7, 8]. Saudi public-sector studies add that attention, infrastructure, data integrity and leadership determine whether artificial intelligence becomes embedded in operational routines [3, 4]. Intelligent decision automation should therefore be designed as a system of coordinated capabilities rather than as a collection of models.

The first capability is contextual intelligence. Government decisions draw on identity, applications, prior interactions, documents, sensor or transactional data, and information from other entities. Artificial intelligence can reduce the cognitive burden of interpreting these inputs by extracting entities, detecting anomalies, estimating probabilities and ranking cases. Evidence on public-service efficiency supports the value of these capabilities when they are integrated into service operations [9]. However, model outputs should be treated as decision evidence rather than self-executing truth. A risk score, for example, should carry its model version, input timestamp, confidence and intended use so that later reviewers can reconstruct why it influenced an outcome. This evidence discipline is particularly important when models change after deployment.

The second capability is executable policy. Government entities routinely implement laws, regulations, delegated authorities, internal policies and service criteria. Encoding these constraints in business rules creates a separation between policy logic and application code. Hybrid decision management shows why this separation matters: explicit rules remain understandable and testable, while artificial intelligence can assist with complex pattern recognition around them [16]. Formal decision modelling also enables policy variants to be represented without duplicating entire workflows [17]. In practice, a rule service might determine whether a citizen meets statutory eligibility, whether a request exceeds a financial threshold, or whether an automated outcome requires human approval. Rules should be versioned with effective dates because a lawful decision depends on the policy in force when the case was decided, not simply on the current rule set.

The third capability is case management. A large share of public work does not follow a perfectly predictable sequence. A case may pause while evidence is requested, branch because a condition changes, or require escalation after an exception. Data-driven case-management research demonstrates that flexible cases can still be checked against compliance constraints [18]. Exception-handling research shows that case systems can support unforeseen events while maintaining structured control [19]. This is a better fit than forcing complex services into one rigid workflow. The case becomes the durable container for facts, documents, decisions, communications, deadlines and appeals, while rules and models provide services to the case at appropriate moments.

Table 1: Evidence streams and design implications for intelligent decision automation

Evidence stream	Synthesis	Design implication	Key references
Public-sector AI assimilation	Value depends on organizational attention, data readiness, governance and process integration.	Treat AI as an operational capability with named service owners and governance forums.	[3–9]
Explainability and legitimacy	Affected users and officials need explanations that support fairness, trust and review.	Generate role-specific reasons for citizens, officers and auditors from shared evidence.	[10–15]
Hybrid decision management	Machine learning and explicit decision logic are complementary rather than interchangeable.	Keep eligibility and legal constraints in versioned rules; use AI for inference and prioritisation.	[16,17]
Case management and compliance	Knowledge-intensive services require runtime flexibility without losing compliance control.	Use a case record for documents, exceptions, human tasks, deadlines and appeals.	[18,19]
Saudi decision context	Saudi studies show rapid AI expansion across public services and regulated organisational domains.	Adopt common reusable decision services while allowing sector-specific risk controls.	[2–23]
Bias, privacy and acceptance	Autonomy increases risk when bias, privacy or human-control requirements are weak.	Apply bounded autonomy, fairness monitoring, privacy controls and risk-based authorisation.	[24–26]

These three capabilities interact most effectively through orchestration rather than direct coupling. A case-management engine should not embed a machine-learning model inside every task, and an artificial-intelligence service should not independently determine legal eligibility. Instead, a decision-orchestration layer should call the required services, preserve evidence and return a structured decision result. The result can include recommended action, applicable rules, confidence, explanation, mandatory approvals and next case state. This pattern allows one model to support several services, one policy rule set to be reused across channels, and one case to survive technology changes without losing its history.

Explainability is the fourth capability and acts across the entire architecture. Government-facing explainability must answer different questions at different moments. Before a decision, officers need to understand whether the model is suitable for the case. At the decision point, citizens need intelligible reasons and the policy basis for the outcome. After the decision, auditors need technical evidence sufficient to reconstruct the process. Research on explainable artificial intelligence shows that explanations can affect perceived fairness and trust [10], while public-sector studies caution that technical assumptions must match administrative realities [11]. Citizen studies further indicate discomfort with fully autonomous artificial intelligence in consequential decisions [12]. The system should therefore distinguish an explanation from a justification: the model explains its contribution, while the organisation justifies the final action through policy, evidence and authority.

The fifth capability is bounded human oversight. Human review is often discussed as a generic safeguard, but an effective architecture specifies when and why it is required. Human authorisation should be triggered by defined conditions: low model confidence, conflicting evidence, exceptional hardship, high financial or legal impact, suspected bias, an unresolved rule conflict, or a citizen appeal. Explainability research supports this role by showing that explanations must be usable for evaluation rather than merely technically available [13–15]. Acceptance evidence also suggests that managers are less comfortable when artificial intelligence autonomously implements decisions without prior human control [26]. Human intervention should therefore be risk-based, not universal; requiring manual approval for every low-risk decision would eliminate the efficiency gains that automation is intended to create.

The sixth capability is integration. Saudi digital government increasingly operates through shared platforms, digital identity, mobile channels and interconnected services. API research in the Saudi context highlights the economic and operational importance of programmable interfaces [6]. For decision automation, APIs should expose well-defined services such as identity verification, document retrieval, rule evaluation, model scoring, payment status and notification. Integration contracts should specify data ownership, permitted purposes, latency, security, versioning and failure behaviour. A resilient decision system must also define what happens when an upstream service is unavailable: whether the case is queued, routed to manual review, or allowed to proceed using cached evidence.

The final capability is learning without uncontrolled drift. Automated decisions generate valuable operational evidence: processing time, override rate, appeal rate, false-positive patterns, service outcomes and citizen feedback. That evidence should improve rules, models and case design, but changes should pass through governance. A model that retrains silently can invalidate previously approved risk thresholds; a rule changed without effective dating can make two identical cases produce unexplained differences. Bias research emphasises the need to detect systematic disparities [24], while privacy and security research warns that more powerful generative capabilities can create new organisational rule violations [25]. Continuous learning must therefore be controlled learning: monitored, tested, approved, versioned and reversible.

Saudi-specific evidence reinforces the practical need for this integrated model. Studies of decision-making quality show that artificial intelligence can improve organisational decisions when users are trained and systems fit the task [20]. Explainable Saudi applications demonstrate that transparency can be engineered alongside performance [21]. Algorithmic decision tools in organisational functions illustrate the increasing reach of automated recommendations [22], while financial-sector reviews show how intelligent automation is expanding into regulated domains [23]. These examples suggest that public entities will increasingly face the same architectural question across multiple services: how to scale intelligent decisions while keeping policy ownership, accountability and citizen rights visible.

5. Proposed Framework for Saudi Digital Government

The proposed framework translates the evidence into a layered architecture for intelligent decision automation. At the top, citizen and business channels capture requests through portals, mobile applications, service centres and machine-to-machine interfaces. The next layer assembles trusted data and case context: verified identity, submitted evidence, relevant history, current service status and authorised external data. These inputs feed a decision-orchestration layer containing three distinct engines—artificial-intelligence models, business rules and case management—surrounded by policy controls and human oversight. The output is a government service action such as approval, rejection, a request for information, escalation or notification. Outcomes, appeals and monitoring then feed controlled improvement. Figure 1 expresses this architecture as a connected decision system rather than a linear workflow.

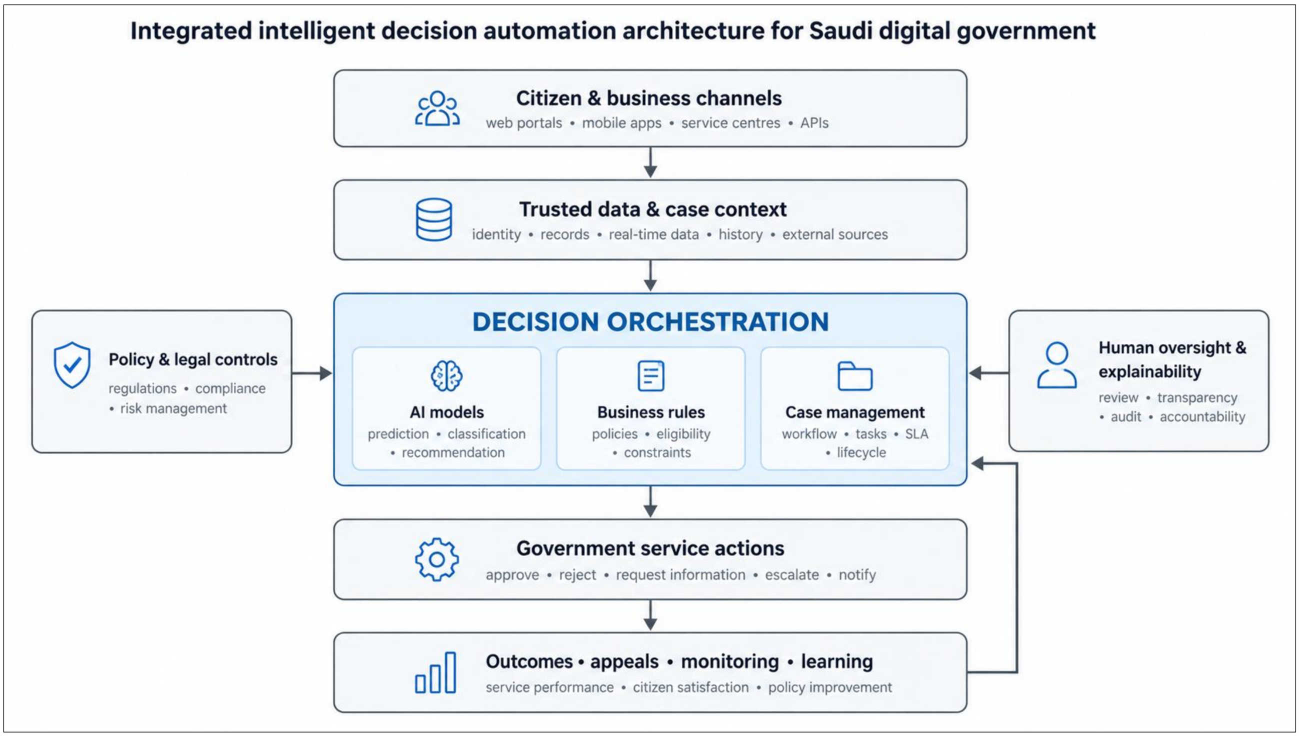

Figure 1: Integrated intelligent decision-automation architecture for Saudi digital government.

The architecture is intentionally modular. Artificial-intelligence services perform inference; business rules implement explicit policy; case management controls the evolving context and work state. The orchestrator determines which combination is appropriate for the current case. For a simple renewal, verified data and deterministic rules may be enough for straight-through processing. For a fraud-sensitive request, a model may provide a risk estimate that influences routing but not legal eligibility. For a complex permit, the case may move through several evidence-gathering and review stages before a final rule evaluation. This modularity allows automation depth to vary with risk while preserving a common governance foundation.

Figure 2 converts the architecture into a seven-stage decision lifecycle. Sense captures the event, request and contextual need. Validate checks identity, data quality, document completeness and applicable rules. Decide combines model inference and deterministic logic to produce a proposed action. Explain produces the human-readable reasons and auditable evidence needed for review. Authorise applies delegated authority, risk thresholds and any mandatory human approval. Execute performs the service action and updates the case record. Learn measures outcomes and proposes improvements to models, rules or process design. The feedback loop is deliberately separated from direct execution so that learning cannot change production behaviour without controlled governance.

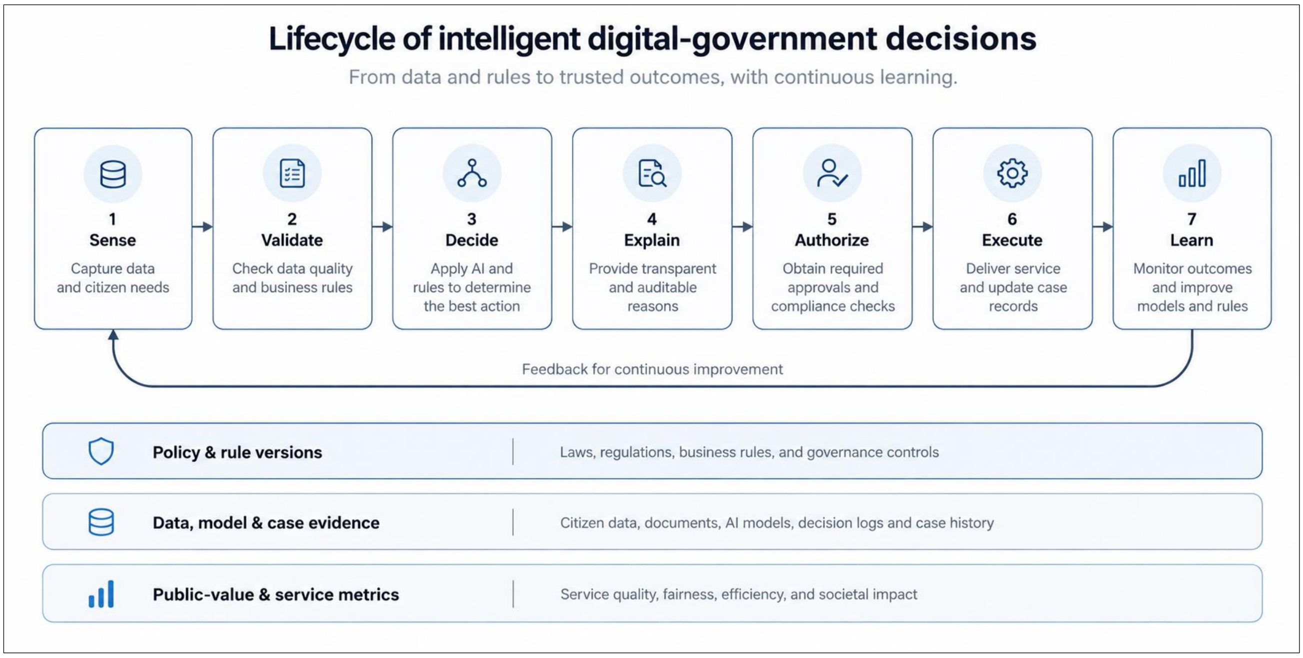

Figure 2: Governed lifecycle for intelligent digital-government decisions

Table 2: Decision-lifecycle controls, evidence and performance measures

Stage	Primary control	Minimum audit evidence	Illustrative measure
Sense	Purpose limitation; authorised data capture	Request ID, channel, consent or legal basis, timestamp	Input completeness; duplicate rate
Validate	Identity, data quality and rule prerequisites	Verification results, missing evidence, data lineage	Validation failure rate; rework rate
Decide	Model suitability plus deterministic policy constraints	Model version, score, rule set, confidence, case facts	Decision latency; exception rate
Explain	Reason generation and evidence linking	Citizen reason, officer rationale, feature or rule evidence	Explanation completeness; complaint rate
Authorise	Delegated authority and escalation thresholds	Approval identity, risk class, override rationale	Manual-review rate; override rate
Execute	Controlled action and transactional integrity	Action response, downstream confirmations, notifications	Completion rate; failed execution rate
Learn	Governed change management and outcome monitoring	Appeals, reversals, fairness metrics, model/rule changes	Appeal rate; reversal rate; fairness drift

Three evidence foundations should persist across all seven stages. The first is policy and rule provenance: every decision should identify the relevant rule version, effective date, authority and

© 2026: Global Academic Journal's Research Consortium (GAJRC)

exceptions applied. The second is data, model and case evidence: source records, model version, input features where appropriate, confidence measures, documents, officer actions and timestamps should be retained according to lawful retention rules. The third is public-value and service evidence: processing time, completion rate, appeal rate, reversal rate, fairness indicators, citizen effort, satisfaction and operational cost. Together, these foundations allow a public entity to answer not only "What did the system decide?" but also "Was the decision lawful, understandable, efficient and beneficial?"

A practical governance model can be organised around four decision classes. Class A covers low-risk, deterministic services with high-quality data and explicit rules; these are candidates for straight-through automation. Class B covers decisions where artificial intelligence provides a recommendation but rules determine the allowable action; automated execution is possible when confidence and risk thresholds are satisfied. Class C covers high-impact or ambiguous cases requiring human authorisation, with artificial intelligence used as decision support. Class D covers novel, contested or exceptional cases that should remain primarily case-managed until evidence justifies further automation. This classification prevents a common failure mode in digital transformation: assuming that the most technologically advanced service is necessarily the most automated one.

Implementation should proceed through decision inventory rather than technology procurement. Each entity should identify high-volume decisions, map their legal basis, inputs, exceptions, current delays and appeal patterns, then classify them by automation suitability. Business owners and legal teams should define policy rules before model deployment, while data teams establish quality and lineage requirements. Artificial-intelligence models can then be introduced only where inference adds value beyond deterministic logic. Case management should provide the operational container for documents, tasks, deadlines and exceptions. Leadership evidence from Saudi public-sector organisations indicates that successful assimilation depends on sustained managerial attention and communication, so governance forums must connect service owners, legal specialists, data teams and operational managers rather than leaving automation to technical units alone [3, 4].

The framework is also compatible with gradual modernisation. Legacy systems do not need to be replaced simultaneously. Rules, model scoring and case services can be exposed through APIs and integrated incrementally. This is especially valuable in government environments where core registries and sector systems have different renewal cycles. The API economy evidence from Saudi Arabia supports the broader importance of interoperable digital infrastructure [6]. A modular decision architecture turns that infrastructure into reusable decision capabilities, allowing new services to consume trusted components instead of duplicating identity checks, policy logic or audit mechanisms.

6. DISCUSSION

The review suggests that the strategic value of intelligent decision automation lies in institutionalising good decision practice, not simply reducing headcount or processing time. Vision 2030 requires public services that are responsive, data-driven and capable of operating at national scale, but public legitimacy depends on consistency and accountability as well as speed. The proposed framework addresses this tension by separating probabilistic inference from deterministic policy and by using case management to preserve flexibility. It therefore treats automation as a governed continuum rather than a binary choice between human and machine decision-making.

This design also changes how performance should be measured. A system that approves applications faster but produces more appeals is not necessarily better. Likewise, a highly accurate model may create little public value if its recommendations cannot be operationalised or explained. Public-administration research stresses the need to evaluate implementation and public value across the innovation lifecycle [7, 8]. For Saudi entities, performance dashboards should therefore combine operational measures with governance outcomes: processing time, straight-through rate, model override rate, rule exceptions, appeal volume, reversal rate, fairness indicators and citizen satisfaction.

The framework further implies that governance should be built into architecture rather than added through policy documents after deployment. Versioned rules, decision logs, explanation records and escalation conditions are technical controls that make accountability executable. This approach is consistent with the growing Saudi emphasis on data and artificial-intelligence capability [2] and with evidence that public-sector assimilation is shaped by governance, infrastructure and organisational attention [3, 4]. The result is a model of digital government in which technology supports institutional responsibility instead of displacing it.

7. Research Agenda and Limitations

Future research should evaluate the framework through longitudinal case studies in Saudi government services with different risk profiles. Priority questions include how citizens respond to different explanation formats, how rule and model conflicts should be resolved, which decisions benefit most from case-based flexibility, and how appeal outcomes can be used to recalibrate automation safely. Arabic-language explainability, cross-agency data provenance, generative-artificial-intelligence controls and measurable fairness across demographic groups also require dedicated investigation.

This review has limitations. The final corpus was intentionally selective rather than exhaustive and was designed for conceptual integration. It includes thirty DOI-bearing peer-reviewed sources from 2020–2025, but relevant standards, legislation, official strategy documents and practitioner reports were not counted in the evidence corpus because the study imposed a DOI requirement. The synthesis also combines evidence from different countries where public-law conditions vary. Accordingly, the proposed architecture should be treated as a design framework that requires service-specific legal validation, not as a universal operating procedure.

8. CONCLUSION

Intelligent government requires more than artificial intelligence. It requires a disciplined way to combine prediction, policy, case context, explanation, authority and learning. The literature reviewed in this paper shows that artificial intelligence can strengthen public-service decisions, but its value is constrained when models are detached from institutional processes, legal rules and human accountability. Business rules provide explicit policy control; case management provides flexibility for incomplete and exceptional situations; and artificial intelligence provides inference where uncertainty or scale makes manual analysis inefficient.

For Saudi digital government, the proposed architecture offers a practical route from isolated artificial-intelligence projects to reusable decision capabilities aligned with Vision 2030. Its central principle is bounded automation: automate when evidence, policy clarity and risk permit; require human judgement when consequences or uncertainty demand it; and preserve an auditable record in both cases. The seven-stage lifecycle and layered architecture make this principle operational by connecting sensing, validation, decision, explanation, authorisation, execution and controlled learning.

The resulting model is not technology-neutral in the sense of ignoring technical choices, but it is technology-resilient: models, rules and systems can change while decision evidence and governance obligations remain stable. That property is essential for a government environment in which policies evolve, technologies mature and citizens expect services to become faster without becoming less fair or less accountable. This supports scalable, trustworthy, citizen-centred public administration across diverse services.

REFERENCES

1. Moshashai D, Leber AM, Savage JD. Saudi Arabia plans for its economic future: Vision 2030, the National Transformation Plan and Saudi fiscal reform. British Journal of Middle Eastern Studies. 2020;47(3):381–401. https://doi.org/10.1080/13530194.2018.1500269
2. Memish ZA, Altuwaijri MM, Almoeen AH, Enani SM. The Saudi Data & Artificial Intelligence Authority (SDAIA) Vision: Leading the Kingdom's Journey toward Global Leadership. Journal of Epidemiology and Global Health. 2021;11(2):140–142. https://doi.org/10.2991/jegh.k.210405.001
3. Alshahrani A, Dennehy D, Mäntymäki M. An attention-based view of AI assimilation in public sector organizations: The case of Saudi Arabia. Government Information Quarterly. 2022;39(4):101617. https://doi.org/10.1016/j.giq.2021.101617
4. Alshahrani A, Griva A, Dennehy D, Mäntymäki M. The role of leadership and communication in AI assimilation: case studies from Saudi Arabia's public sector organizations. Transforming Government: People, Process and Policy. 2025;19(4):875–894. https://doi.org/10.1108/TG-06-2025-0187
5. Hasan A, Alenazy AA, Habib S, Husain S. Examining the drivers and barriers to adoption of e-government services in Saudi Arabia. Journal of Innovative Digital Transformation. 2024;1(2):139–157. https://doi.org/10.1108/JIDT-09-2023-0019
6. Ali MA, Salih SM. Impact of Application Programming Interfaces (APIs) Economy on Digital Economics in Saudi Arabia. Sustainability. 2025;17(9):4104. https://doi.org/10.3390/su17094104
7. Zuiderwijk A, Chen Y-C, Salem F. Implications of the use of artificial intelligence in public governance: A systematic literature review and a research agenda. Government Information Quarterly. 2021;38(3):101577. https://doi.org/10.1016/j.giq.2021.101577
8. Madan R, Ashok M. AI adoption and diffusion in public administration: A systematic literature review and future research agenda. Government Information Quarterly. 2023;40(1):101774. https://doi.org/10.1016/j.giq.2022.101774

9. Kulal A, Rahiman HU, Suvarna H, Abhishek N, Dinesh S. Enhancing public service delivery efficiency: Exploring the impact of AI. Journal of Open Innovation: Technology, Market, and Complexity. 2024;10(3):100329. https://doi.org/10.1016/j.joitmc.2024.100329
10. Aoki N, Tatsumi T, Naruse G, Maeda K. Explainable AI for government: Does the type of explanation matter to the accuracy, fairness, and trustworthiness of an algorithmic decision as perceived by those who are affected? Government Information Quarterly. 2024;41(4):101965. https://doi.org/10.1016/j.giq.2024.101965
11. Fischer-Abaigar U, Kern C, Barda N, Kreuter F. Bridging the gap: Towards an expanded toolkit for AI-driven decision-making in the public sector. Government Information Quarterly. 2024;41(4):101976. https://doi.org/10.1016/j.giq.2024.101976
12. Haesevoets T, Verschuere B, Van Severen R, Roets A. How do citizens perceive the use of Artificial Intelligence in public sector decisions? Government Information Quarterly. 2024;41(1):101906. https://doi.org/10.1016/j.giq.2023.101906
13. Barredo Arrieta A, Díaz-Rodríguez N, Del Ser J, Bennetot A, Tabik S, Barbado A, Garcia S, Gil-Lopez S, Molina D, Benjamins R, Chatila R, Herrera F. Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion. 2020;58:82–115. https://doi.org/10.1016/j.inffus.2019.12.012
14. Angelov PP, Soares EA, Jiang R, Arnold NI, Atkinson PM. Explainable artificial intelligence: an analytical review. WIREs Data Mining and Knowledge Discovery. 2021;11(5):e1424. https://doi.org/10.1002/widm.1424
15. Vilone G, Longo L. Notions of explainability and evaluation approaches for explainable artificial intelligence. Information Fusion. 2021;76:89–106. https://doi.org/10.1016/j.inffus.2021.05.009
16. Bork D, Ali SJ, Dinev GM. AI-Enhanced Hybrid Decision Management. Business & Information Systems Engineering. 2023;65(2):179–199. https://doi.org/10.1007/s12599-023-00790-2
17. Speck A, Jagenow A, Windrich M. Decision Model and Notation for Describing Variability in Business Process Product Lines. In: Proceedings of the 16th International Conference on Evaluation of Novel Approaches to Software Engineering. 2021:445–452. https://doi.org/10.5220/0010497204450452
18. Haarmann S, Holfter A, Pufahl L, Weske M. Formal Framework for Checking Compliance of Data-Driven Case Management. Journal on Data Semantics. 2021;10(1–2):143–163. https://doi.org/10.1007/s13740-021-00120-3
19. Andree K, Ihde S, Weske M, Pufahl L. An exception handling framework for case management. Software and Systems Modeling. 2022;21:939–962. https://doi.org/10.1007/s10270-022-00993-3
20. Aljohani NB, Albliwi S. Impacts of Applying Artificial Intelligence on Decision-Making Quality: A Descriptive Study in Saudi Arabian Private Sector Organizations. International Transaction Journal of Engineering, Management, & Applied Sciences & Technologies. 2022;13(5):1–14. https://doi.org/10.14456/ITJEMAST.2022.104
21. Al Mazroa A, Maray M, Alashjaee AM, Alotaibi FA, Alzahrani AA, Alkharashi A, Alotaibi SD, Alnfiai MM. Computer Vision With Explainable Artificial Intelligence for Visual Pollution Detection in the Kingdom of Saudi Arabia. IEEE Access. 2024;12:193014–193027. https://doi.org/10.1109/ACCESS.2024.3513696
22. Alabdali MA, Khan SA, Yaqub MZ, Alshahrani MA. Harnessing the Power of Algorithmic Human Resource Management and Human Resource Strategic Decision-Making for Achieving Organizational Success: An Empirical Analysis. Sustainability. 2024;16(11):4854. https://doi.org/10.3390/su16114854
23. Al-Baity HH. The Artificial Intelligence Revolution in Digital Finance in Saudi Arabia: A Comprehensive Review and Proposed Framework. Sustainability. 2023;15(18):13725. https://doi.org/10.3390/su151813725
24. Wei X, Kumar N, Zhang H. Addressing bias in generative AI: Challenges and research opportunities in information management. Information & Management. 2025;62(2):104103. https://doi.org/10.1016/j.im.2025.104103
25. Choudhary SK, Kar AK. Modeling the role of generative AI in organizational privacy and security. Decision Support Systems. 2025;196:114500. https://doi.org/10.1016/j.dss.2025.114500
26. Gieselmann M, Erdsiek D, Rost V, Sassenberg K. Do managers accept artificial intelligence? Insights into the role of business area and AI functionality. Journal of Economic Psychology. 2025;108:102804. https://doi.org/10.1016/j.joep.2025.102804
27. Donthu N, Kumar S, Mukherjee D, Pandey N, Lim WM. How to conduct a bibliometric analysis: An overview and guidelines. Journal of Business Research. 2021;133:285–296. https://doi.org/10.1016/j.jbusres.2021.04.070
28. Kraus S, Breier M, Dasí-Rodríguez S. The art of crafting a systematic literature review in

entrepreneurship research. International Entrepreneurship and Management Journal. 2020;16(3):1023–1042. https://doi.org/10.1007/s11365-020-00635-4
29. Linnenluecke MK, Marrone M, Singh AK. Conducting systematic literature reviews and bibliometric analyses. Australian Journal of Management. 2020;45(2):175–194. https://doi.org/10.1177/0312896219877678
30. Marzi G, Balzano M, Caputo A, Pellegrini MM. Guidelines for Bibliometric-Systematic Literature Reviews: 10 steps to combine analysis, synthesis and theory development. International Journal of Management Reviews. 2025;27(1):81–103. https://doi.org/10.1111/ijmr.12381.