AI-Driven Advanced Decision Management for Intelligent Government Services in Saudi Arabia: A Real-Time, Explainable and Policy-Aware Decisioning Framework for Vision 2030

Abstract:

Saudi Arabia's rapid digital transformation creates an opportunity to move government services beyond online transactions toward intelligent, continuously governed decision management. This review examines how artificial intelligence can support real-time public-service decisions while preserving legal consistency, explainability, human responsibility, security, and public value. The paper synthesises research published from 2020 to 2025 on artificial intelligence in public administration, automated decision-making, explainable artificial intelligence, public-sector governance, data documentation, rules-as-code, human oversight, and Saudi digital-government adoption. Rather than treating artificial intelligence as a stand-alone prediction engine, the review develops a policy-aware decisioning perspective in which models, codified rules, eligibility conditions, service priorities, human authorisation, and evidence are orchestrated as one accountable operational system. The synthesis identifies three recurring weaknesses in current practice: technically accurate models may still be unsuitable for administrative decisions; explanations are often designed after decisions rather than as part of the decision process; and policy controls are commonly separated from model operations, creating risks when rules, data, or models change. In response, the paper proposes a conceptual framework for Saudi intelligent government services based on six interacting capabilities: trusted event and case data, real-time decision orchestration, machine-readable policy controls, contextual explanation, risk-based human oversight, and closed-loop monitoring. The framework links every service outcome to the data, model, rule version, rationale, authorisation path, and subsequent appeal or correction. Its purpose is not maximum automation, but controlled decision velocity: decisions should be fast where rules and evidence are stable, and deliberately escalated where rights, ambiguity, novelty, or material consequences require human judgement. The review concludes that policy-aware advanced decision management can strengthen Vision 2030 service ambitions when operational performance and public-law values are engineered together rather than assessed separately after deployment.

Keywords: artificial intelligence; advanced decision management; digital government; explainable AI; policy-aware systems; public services; Saudi Arabia; Vision 2030

1. Introduction

Digital government increasingly involves more than converting paper procedures into portals. Contemporary public services must detect events, interpret evidence, apply changing rules, prioritise cases, communicate reasons, and learn from outcomes while maintaining accountability. Artificial intelligence can improve classification, forecasting, anomaly detection, language processing, and resource allocation, but public administration cannot evaluate such systems only by predictive accuracy. Government decisions affect entitlements, permissions, payments, inspections, referrals, and access to services; consequently, legality, consistency, contestability, privacy, and human responsibility are part of system quality. Research on public-sector artificial intelligence repeatedly shows that adoption creates tensions between efficiency and values such as fairness, transparency, responsiveness, and accountability [2,5,6,11,14]. Studies of automated administrative decision-making further show that automation can support consistency and speed while simultaneously creating difficult trade-offs when discretion, context, or rights are material [20].

Saudi Arabia is a particularly important setting for this problem because digital transformation and artificial-intelligence assimilation are advancing alongside broad institutional modernisation. Research across Saudi public organisations has identified management alignment, data integrity, infrastructure, cultural and linguistic considerations, and governance as recurring conditions for successful artificial-intelligence assimilation [9]. Saudi managers' willingness to adopt artificial intelligence is also associated with perceived usefulness, impact, trust, security, privacy, transparency, and accountability [24]. At the citizen level, adoption of electronic government services is shaped by digital literacy, infrastructure, policy support, security, privacy, and trust [25], while recent research on Saudi digital public services emphasises usable security and privacy as determinants of confidence in electronic channels [29]. These findings suggest that the next stage of intelligent government cannot be framed simply as adding models to existing workflows. Decision logic, policy authority, explanations, oversight, service operations, and evidence must be integrated.

This paper reviews the literature through the lens of advanced decision management: the coordinated use of data, analytics, machine learning, business rules, policy constraints, workflow orchestration, and feedback to make repeatable operational decisions. The central proposition is that intelligent public services require a policy-aware architecture in which a model is one input to a governed decision rather than the decision itself. This distinction matters because public authorities must be able to identify which rule applied, which evidence was considered, why an outcome followed, who or what authorised it, and how a person can challenge or correct it. Accountability research warns that algorithmic systems can diffuse responsibility when technical complexity obscures who can answer for an outcome [5]. Public-sector cases also show that organisational and legal institutions may fail to interrogate opaque systems even after errors become visible [21].

The study therefore asks how real-time, explainable and policy-aware decision management can be structured for Saudi intelligent government services. Its aim is to synthesise recent evidence and propose an operational framework that reconciles decision speed with administrative legitimacy. Four objectives guide the review: first, to identify design requirements for real-time artificial-intelligence decisioning in public services; second, to examine how explainability, transparency and human oversight influence trust and accountability; third, to analyse how laws, regulations, policies and eligibility rules can remain authoritative when decisions incorporate adaptive models; and fourth, to derive a lifecycle framework that supports continuous monitoring, appeal, correction and controlled learning. The result is a review-based conceptual contribution rather than an empirical evaluation of a deployed Saudi system.

2. Review Methodology

The review used a structured integrative method designed for an interdisciplinary topic that spans information systems, public administration, artificial intelligence, law, digital government, and Saudi public-sector transformation. The methodological organisation follows the logic of defining information sources, search terms, eligibility criteria, analytical categories, and synthesis before interpreting findings. This approach is appropriate because the research question concerns how multiple technical and institutional literatures can be combined into a coherent decision-management architecture rather than the pooled statistical effect of a single intervention.

Searches were designed around five concept groups: public-sector artificial intelligence and automated administrative decision-making; explainability, transparency, fairness and human oversight; artificial-intelligence governance and accountability; machine-readable policy or rules-as-code; and Saudi digital government, artificial-intelligence adoption and public-service transformation. Priority was given to peer-reviewed literature from 2020 through 2025, supplemented by one policy research paper with a persistent DOI on machine-consumable government rules. Studies were included when they contributed evidence, frameworks or empirical findings relevant to operational decision design; publications outside the date window or without a public-service governance implication were excluded.

The synthesis proceeded in three stages. First, each source was coded against the decision lifecycle: sense, validate, decide, explain, authorise, execute and learn. Second, evidence was grouped into four control domains: data and model integrity; policy and legal alignment; human and organisational accountability; and citizen-facing transparency, trust and contestability. Third, these domains were mapped to Saudi public-service conditions identified in the literature. The review intentionally distinguishes explanatory evidence from normative claims. For example, experimental research indicates that explanations can affect perceived trustworthiness [16,19], but this does not imply that any single explanation format is universally sufficient. Similarly, citizen preferences may trade transparency against perceived effectiveness in some contexts [17], yet administrative systems still face independent legal and accountability requirements.

Quality appraisal emphasised publication credibility, methodological clarity, relevance to public-sector decision-making, and DOI traceability. Thirty references were retained. The review targets the architecture and governance problem implied by advanced decision management rather than every government artificial-intelligence application, allowing focused comparison of controls used when algorithms influence administrative outcomes.

Table 1. Review design and analytical scope.

Review element	**Operationalisation**
Time window	2020-2025 publications only for the final reference set.
Core domains	Public-sector AI; automated decisions; explainability; AI governance; rules-as-code; Saudi digital government.
Primary inclusion test	Direct contribution to decision design, public-service governance, accountability, trust, policy execution or Saudi implementation context.
Analytical lens	Decision lifecycle: sense, validate, decide, explain, authorise, execute and learn.
Control domains	Data/model integrity; policy/legal alignment; human accountability; citizen transparency and contestability.
Quality checks	Peer-review credibility, methodological clarity, topical relevance and DOI traceability.

3. From Predictive Models to Advanced Decision Management

Artificial intelligence becomes operationally valuable in government when it is embedded in a decision process. A predictive model may estimate fraud risk, service demand, eligibility likelihood, case urgency, or the probability of an adverse event, but an administrative action normally requires more than a score. The system must identify the relevant service and jurisdiction, verify identity and data sufficiency, apply mandatory rules, recognise exceptions, determine whether discretion is permitted, generate a reason, route higher-risk cases for authorisation, execute the action, and preserve evidence. Studies of public-sector artificial intelligence adoption show that organisational context and absorptive capacity shape whether technical capabilities produce public value [10,14].

A real-time decision-management architecture therefore needs event awareness and orchestration. "Real time" should not be interpreted as an instruction to automate every case instantly. It means that the system can respond to an event with the appropriate decision pathway at operational speed. A low-risk request supported by complete data and deterministic policy may be decided automatically; an ambiguous or materially consequential case may be paused, enriched, or escalated. Human-in-the-loop research in government illustrates the practical importance of designing machine assistance around human review instead of treating manual intervention as a failure of automation [4]. Human-machine interaction experiments also warn of automation bias and selective adherence: officials may follow algorithmic advice too readily in some conditions and discount it selectively in others [13]. Accordingly, oversight must be engineered through thresholds, role definitions and evidence requirements, not left to informal discretion.

Data quality and provenance are the first control layer. Documentation approaches such as datasheets for datasets were developed to expose how data were collected, what populations and purposes they represent, and what limitations accompany reuse [7]. Lifecycle analyses of machine-learning harm show that problems may arise at problem formulation, data collection, measurement, model development, evaluation, deployment and feedback [8]. For government decisioning, these insights imply that the decision record should preserve not only the model version but also the source, timestamp, quality state and permitted use of decisive data. A "trusted data" layer is therefore a governed operational object rather than a generic data lake. Eligibility checks should distinguish verified facts from inferred variables; missingness should be explicit; and sensitive attributes should be controlled according to their legitimate role in the service.

Advanced decision management also requires separation between prediction and authority. Machine learning is useful where relationships are probabilistic, complex or high-dimensional, while deterministic rules remain valuable when policy requires a stable, auditable condition. Research comparing automation approaches in public services argues that expert or rule-based systems can remain appropriate where process transparency is essential, and that hybrid designs may be preferable to using machine learning for every step [15]. The rules-as-code literature extends this principle by examining how legal or policy rules can be expressed in machine-consumable form [3,22]. However, codification creates its own risks: text may be mistranslated into software, exceptions may be oversimplified, and implementation logic may become detached from the authoritative legal source. Policy-aware decisioning must therefore preserve links between executable rules and their source provisions, effective dates, interpretations and approval history.

Figure 1 conceptualises this architecture. Citizen and business channels produce events and cases that enter a trusted data layer. The central orchestration component combines artificial-intelligence models, explicit rules, eligibility conditions, service priorities and risk thresholds. Policy and legal controls constrain the decision from one side; explainability and human oversight constrain it from the other. The result is a service action with an evidence trail, followed by monitoring, appeal and learning. The feedback path does not permit uncontrolled self-modification. Instead, outcomes inform reviewed changes to data, models, policies and thresholds through governed release processes.

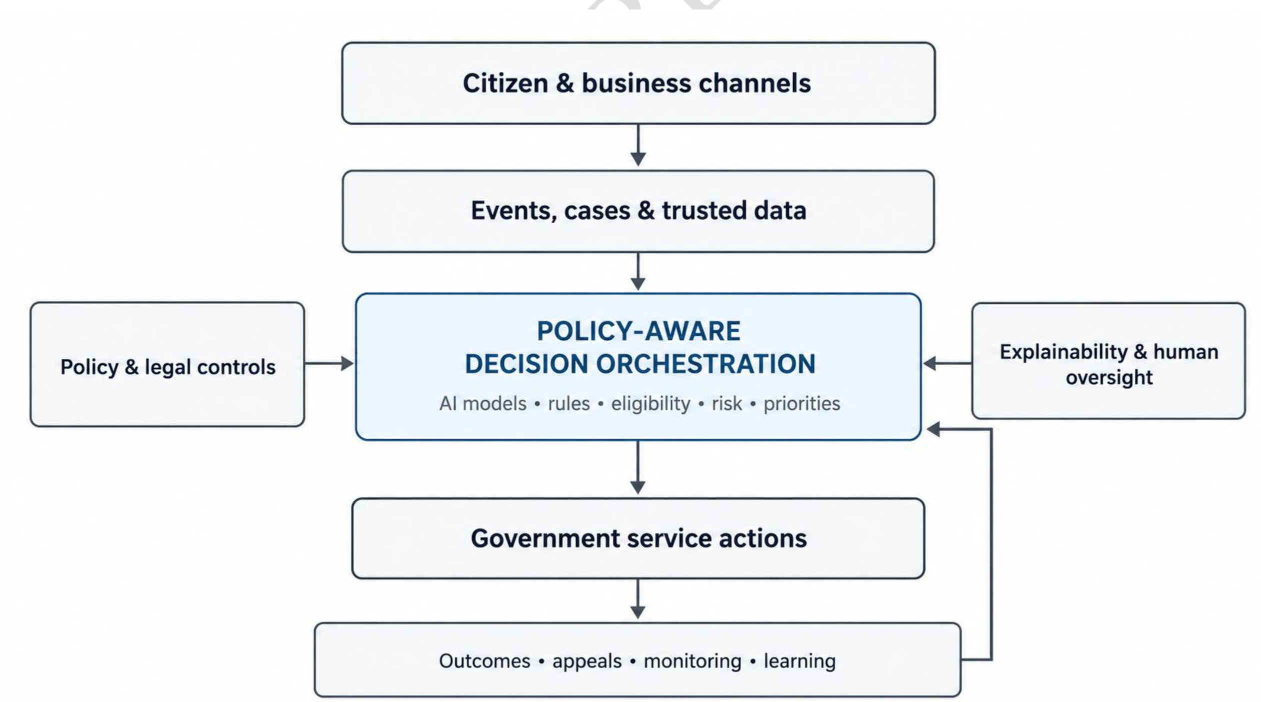

Figure 1. Policy-aware real-time decisioning architecture for intelligent government services.

4. Explainability, Transparency and Human Authorisation

Explainability is often treated as a technical feature that is added after a model produces a result. For administrative decision management, this is insufficient. Explanations must be generated from the actual decision path, including material data, model contribution, policy rules, exceptions, confidence, and human intervention. The broad explainable-artificial-intelligence literature distinguishes multiple forms of explanation and emphasises that the appropriate technique depends on audience, model and purpose [1]. Public-sector evidence makes the audience issue especially important. Experimental work shows that algorithmic explainability can influence perceived trustworthiness more strongly than merely providing access to information about the system [16]. More recent government-focused research finds that different explanation types can affect perceived accuracy, fairness and trustworthiness differently across decision domains [19]. A single generic explanation template is therefore unlikely to meet every service need.

The framework proposed here separates four explanation targets. The citizen-facing explanation states the decision, material reasons, applicable policy basis, important evidence, and route to review. The caseworker explanation provides more diagnostic detail, including confidence, data quality flags, exceptions and comparable rules. The governance explanation supports audit by preserving model, rule, data and workflow versions. The technical explanation supports engineers and validators with feature behaviour, monitoring indicators and error analysis. This layered design avoids two common failures: exposing technical complexity to citizens without meaningful reasons, or giving citizens a simplified message while leaving auditors unable to reconstruct the decision.

Transparency must also be paired with contestability. Accountability cannot be achieved solely by publishing a model description if affected persons cannot correct data, question a rule application, or obtain meaningful review. Public-value research shows that failures of fairness and transparency can reduce citizen evaluations of government even when people are not directly affected by a decision [11]. Survey experiments likewise indicate that citizens may regard artificial-intelligence-led government decisions as less trustworthy than human decisions in some settings [12]. At the same time, research on public-sector algorithms suggests that people sometimes prioritise effectiveness over transparency or participation [17]. These findings should not be read as contradictory. They show that trust is context-dependent and that public authorities cannot rely on one attitudinal measure to define acceptable governance. A robust system must provide efficiency, reasons, correction mechanisms and independent accountability together.

Human authorisation should consequently be risk-based. Not every decision needs manual approval, because excessive review can recreate delays that intelligent services are intended to remove. Yet automation level should vary with consequence, uncertainty, discretion and reversibility. High-volume deterministic transactions can be straight-through processed when policy and data are clear. Medium-risk cases can require exception review only when thresholds are triggered. High-impact or contested decisions should require accountable human authorisation before final adverse action, particularly when the model contributes a probabilistic assessment rather than a legally sufficient fact. Human reviewers should see both the recommendation and the evidence needed to challenge it; otherwise, the interface can convert oversight into ceremonial confirmation.

The authorisation design must also guard against organisational blackboxing. Case research demonstrates that errors can persist when agencies and oversight institutions defer to automated systems or fail to examine how a system actually produced decisions [21]. Accountability scholarship similarly emphasises that complex algorithmic systems can weaken answerability if responsibility is fragmented among vendors, data owners, developers and public officials [5]. The decision framework therefore assigns named ownership to policy rules, data products, models, service workflows and final authorisation. Escalations, overrides and appeals become analysable events. This converts "human in the loop" from a slogan into an observable governance mechanism.

5. Policy-Aware Decisioning and the Saudi Government Context

Policy awareness is the feature that distinguishes a government decision-management system from a generic artificial-intelligence platform. Public services operate within legal and administrative authority. A model may identify patterns, but it should not silently redefine eligibility, create new sanctions, or alter the meaning of a regulation. Machine-readable policy can improve consistency and speed when rules are explicit, versioned and linked to authoritative sources. The rules-as-code literature describes the potential for governments to produce machine-consumable representations of rules [3], while legal analysis cautions that software implementation can threaten reviewability if code obscures interpretation or departs from legal meaning [22]. These perspectives support a dual representation: human-readable policy remains authoritative, while executable rules are controlled artefacts derived from it.

The decision engine should therefore apply policy through a versioned policy service. Each rule object needs an identifier, source reference, effective period, owner, approval record, test cases, dependencies and interpretation notes. When policy changes, affected services can be identified before deployment. Historical decisions remain reproducible because the system retains the version that was active at the decision time. This is particularly important for retrospective audit, appeals and compensation where a later rule must not be projected backward. Policy tests should include ordinary cases, boundary values, exceptions and conflict scenarios. Models should be tested against policy constraints so that a statistically plausible output cannot authorise an impermissible action.

Saudi public-sector evidence indicates why this integration matters. Studies of artificial-intelligence assimilation in Saudi public organisations identify misalignment between artificial intelligence and management decision-making, data integrity and sharing, infrastructure, national culture and language, and ethics and governance as significant challenges [9,18]. These are not separate implementation issues; they interact within the decision pathway. Weak data sharing can reduce model validity, while a model that cannot handle Arabic language variation may distort service triage. Management misalignment can result in technically successful pilots that do not fit administrative authority. Governance concerns intensify if models are acquired or developed without a clear relationship to accountable service owners.

Manager adoption research in Saudi Arabia further shows that perceived impact and usefulness influence intentions to adopt artificial intelligence, while transparency, privacy, security and related ethical considerations shape trust [24]. Citizen-focused e-government research identifies similar conditions around security, infrastructure, literacy and trust [25]. Recent Saudi public-service research specifically emphasises that security controls must remain usable because overly burdensome protections can themselves undermine digital-service adoption [29]. A policy-aware architecture can address these factors by making compliance and trust mechanisms part of the transaction flow rather than additional documents around the system. Identity assurance, consent or lawful-use controls, data minimisation, explanation, security checks and escalation can be executed as decision components.

Saudi governance scholarship also places artificial intelligence within wider institutional change [23]. For this review, the relevant implication is that technological capability must be matched by formal roles, review powers and organisational learning. Vision 2030-oriented service modernisation is strengthened when agencies can reuse common decision controls without centralising all discretion into one opaque platform. Shared components can include identity, data quality, policy registries, model governance, explanation services, audit logging and monitoring, while domain agencies retain responsibility for the substantive rules and decisions within their mandates. This federated pattern supports interoperability without erasing institutional accountability.

Table 2. Decision-management control domains, risks and required evidence.

Control domain	Principal risk	Required evidence/control
Trusted data	Incorrect, stale or unlawfully reused information	Source, timestamp, quality state, permitted use, missingness and provenance.
Models	Drift, bias, miscalibration or use outside approved purpose	Model version, validation scope, confidence, fairness checks, monitoring and change approval.
Policy rules	Outdated or mistranslated requirements	Source provision, rule version, effective dates, owner, tests, exceptions and interpretation notes.
Human oversight	Automation bias or ceremonial review	Risk thresholds, accountable role, escalation triggers, override reason and authorisation record.
Explanation	Opaque or misleading reasons	Audience-specific reason, material evidence, policy basis, uncertainty and review route.
Execution and learning	Irreversible action or uncontrolled adaptation	Reversibility, appeal status, outcome monitoring, controlled retraining and release governance.

6. A Real-Time, Explainable and Policy-Aware Framework

The proposed framework consists of six capabilities that operate as one governed decision system. The first is event and case intelligence. Government interactions arrive through portals, mobile applications, contact centres, sensors, partner systems and scheduled processes. Each event is converted into a case context containing identity, purpose, jurisdiction, service type, verified facts, relevant history and data-quality indicators. The second capability is decision orchestration. Orchestration determines which models, rules, validations and human steps are required for the specific service and risk level. It also controls sequencing.

The third capability is policy execution. This layer contains versioned rules, thresholds, permissions and exceptions derived from approved policy. It must support effective dates and conflict resolution while retaining human-readable links. The fourth is model intelligence: machine-learning or statistical components used for forecasting, classification, ranking, anomaly detection or natural-language processing. Models are treated as governed assets with approved purposes, training lineage, performance bounds, fairness checks and monitoring requirements. The fifth capability is explainability and human authorisation. Depending on risk, the system either completes the service, requests additional evidence, routes to an authorised officer, or issues a provisional outcome pending review. The sixth capability is evidence and learning, which records the full decision trace and monitors outcomes, appeals, overrides, drift, service times and residual errors.

Figure 2 expresses the framework as a lifecycle rather than a one-time model deployment. The system senses an event, validates identity and evidence, decides using models and rules, explains the proposed outcome, authorises according to risk, executes the service action, and learns from subsequent outcomes. Three control layers remain active throughout: policy and rule versions; data, model and decision evidence; and public-value and service metrics. This lifecycle reflects responsible-governance research that emphasises governance across the artificial-intelligence lifecycle rather than only at initial approval [26-28]. Empirical research on how public organisations implement artificial-intelligence governance likewise finds that training, standards, organisational processes and governance practices determine whether ethical principles are translated into operational behaviour [27,30].

Real-time operation requires predefined service-level controls. Latency should be measured separately from deliberation. The objective is not the lowest possible response time but the shortest justified path. A low-risk renewal may have a sub-second machine decision; a complex entitlement case may have real-time triage followed by deliberate review. Decision latency metrics should therefore be paired with outcome quality, explanation completeness, override rate, appeal rate, policy exceptions, drift and fairness indicators. This prevents teams from optimising speed while silently transferring risk to citizens or caseworkers.

The framework also treats monitoring as a decision function. Model drift is only one form of change. Policy amendments, data-source changes, integration failures, new fraud patterns, seasonal demand, demographic shifts and operational workarounds can alter decision behaviour even when a model remains statistically stable. Monitoring must therefore compare technical, legal and service signals. A rise in overrides may indicate poor model calibration, unclear policy or a new case pattern. A rise in appeals may indicate explanation failure, data quality problems or inconsistent application. Decision records should support root-cause analysis across these layers.

Finally, learning is controlled rather than automatic. Production outcomes may improve future models, but retraining should pass through data review, policy checks, validation and authorisation. Direct online learning is inappropriate where a model could change the practical standard applied to citizens without review. Responsible artificial-intelligence governance literature increasingly distinguishes principles from the organisational practices needed to operationalise them [26,28]. The framework follows this distinction: fairness, transparency and accountability are translated into versioning, ownership, testing, escalation, evidence preservation and reviewable change.

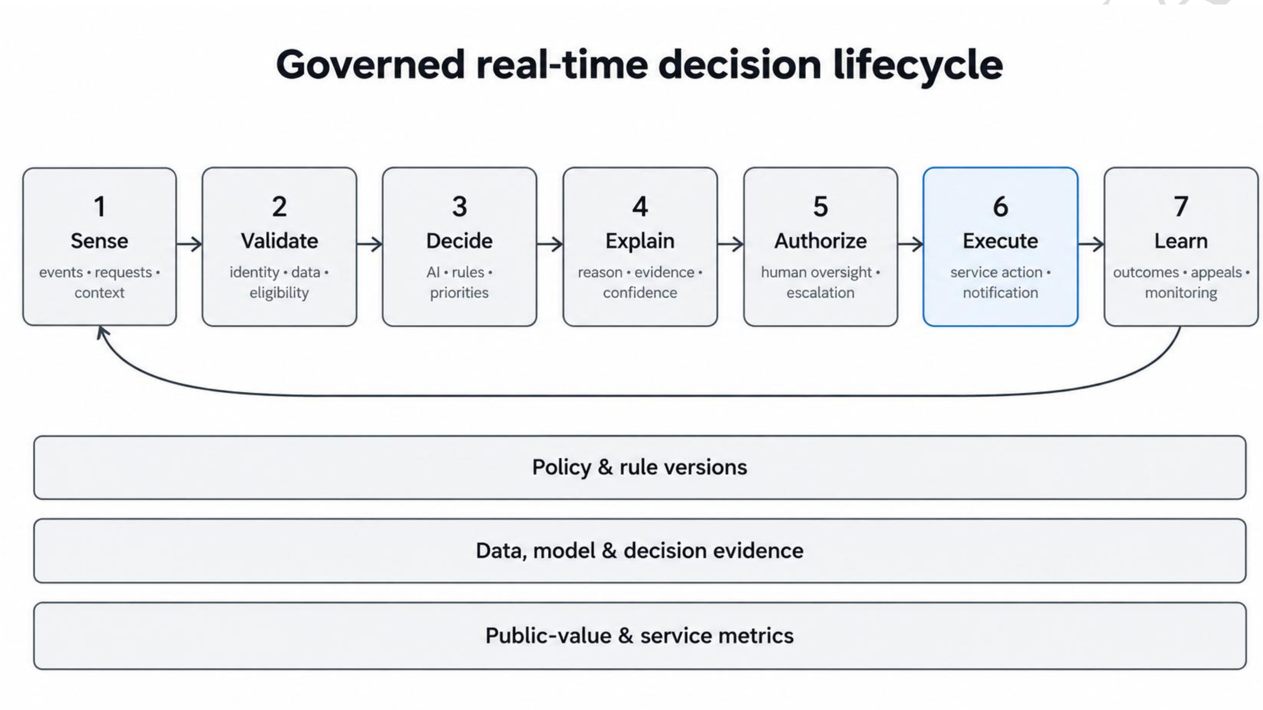

Figure 2. Governed real-time decision lifecycle with continuous evidence and public-value controls.

7. Critical Synthesis and Research Agenda

The literature suggests that intelligent government decisioning should be evaluated as a sociotechnical institution, not a software feature. Governance studies consistently identify the need to specify who governs, what is governed, when governance occurs and how controls are implemented [26]. Public-organisation evidence shows that governance becomes more mature when agencies train decision-makers, developers and users and embed standards into operating processes [27]. In practical terms, an advanced decision platform cannot compensate for unclear authority. Saudi agencies adopting a shared framework would still need explicit owners for service policy, data products, models, citizen communications, security and final administrative responsibility.

A second research priority concerns explanation quality. Current studies demonstrate that explanations can influence trust and fairness perceptions, but effects vary by domain and explanation type [16,19]. Future Saudi research should therefore test explanations in actual service contexts and in Arabic as well as English, distinguishing informational explanations from legally sufficient reasons. Measures should include comprehension, ability to identify an error, perceived fairness, successful appeal navigation, and caseworker diagnostic value. Explanation research should also test adverse and favourable decisions separately because citizens may scrutinise them differently.

Third, policy-aware artificial intelligence needs stronger methods for semantic traceability. Machine-readable rules can make services faster and more consistent, yet legal analysis warns that software can harden one interpretation or obscure how discretionary concepts are operationalised [22]. Research should investigate governance patterns in which policy experts and engineers jointly maintain executable rules, with automated comparison between source text, rule versions and test cases. The objective is not to convert all law into code, but to make the parts that are operationalised in digital services visible, testable and reviewable.

Fourth, human oversight requires empirical calibration. Human involvement can protect against errors but also introduce inconsistency, delay or automation bias [13]. Agencies need evidence about which decisions benefit from mandatory authorisation, which are better managed through exception review, and which can be safely automated. Experiments and field studies should compare thresholds using service quality, error severity, fairness, workload, review time and appeal outcomes. The appropriate oversight model will vary by domain; therefore, a national architecture should support configurable controls rather than a single automation rule.

Fifth, public-value monitoring should be expanded beyond technical performance. Automated decision-making literature shows that good governance involves trade-offs among efficiency, equality, fairness, resilience, responsiveness, privacy, rule of law and transparency [20]. A production dashboard that reports only accuracy and latency is consequently incomplete. Future work should define auditable measures for decision reversibility, explanation coverage, unresolved appeals, distributional error, policy exceptions, human override, and time to correction. These measures can reveal whether performance improvements are accompanied by hidden governance costs.

The resulting research agenda is directly aligned with the paper's title. "Advanced decision management" should be investigated as the integration point between artificial-intelligence engineering and public administration. "Real-time" should be studied as controlled responsiveness, not unqualified automation. "Explainable" should mean reasons that fit the recipient and the administrative purpose. "Policy-aware" should mean that legal and policy authority is explicitly represented and versioned. Saudi research on artificial-intelligence assimilation, manager adoption, citizen e-government use and digital-service security already provides an empirical base [9,24,25,29]. The next step is to connect those studies to operational decision-system design and evaluation.

8. Limitations

This review has four limitations. First, it is an integrative synthesis rather than an exhaustive census of every publication on artificial intelligence in government. The selected literature was intentionally bounded around decision management, explainability, governance, policy representation and Saudi digital government. Second, the evidence base is heterogeneous: it includes conceptual work, reviews, experiments, case studies and technical-governance research, so findings cannot be interpreted as a single causal effect. Third, several influential public-sector studies come from jurisdictions with legal and administrative traditions different from Saudi Arabia. Their mechanisms are useful for design reasoning, but local validation remains necessary. Fourth, the proposed framework is conceptual. It has not yet been tested against service-level outcomes, appeals, caseworker behaviour, Arabic-language explanation quality, or cross-agency interoperability. These limitations define priorities for empirical evaluation rather than weakening the need for integrated governance.

9. Conclusion

Artificial intelligence can make Saudi government services more responsive, predictive and personalised, but only if decision speed is matched by visible authority and reviewable evidence. The literature shows that public-sector artificial intelligence creates value and tension at the same time: automation can increase efficiency and consistency while raising questions about fairness, transparency, responsibility, privacy and the rule of law [5,14,20]. Saudi studies add practical constraints around data integrity, infrastructure, management alignment, security, trust, culture and organisational readiness [9,24,25,29]. These findings support a shift from model-centric projects toward governed advanced decision management.

The framework developed in this review positions trusted data, policy rules, artificial-intelligence models, explanations, human authorisation and monitoring inside one lifecycle. Every outcome should be reconstructable from the evidence, model, rule version, rationale and authorisation path that existed at the decision time. Automation should expand where policy is deterministic and consequences are limited, while uncertainty, discretion and material adverse effects should trigger additional evidence or accountable human review. Appeals and corrections should feed monitored learning, but changes to models or policy logic should remain controlled releases rather than invisible adaptation.

For Vision 2030, the strategic opportunity is therefore not simply to automate more decisions. It is to build a reusable decisioning capability that allows government entities to deliver faster services without separating innovation from administrative legitimacy. A real-time, explainable and policy-aware architecture can provide that bridge. Its success should be measured by service performance and by the ability of citizens, officials, auditors and courts to understand, challenge and correct decisions when necessary. In intelligent government, the strongest system is not the one that removes humans from the process; it is the one that makes the relationship among data, models, policy and accountable human authority explicit enough to support both speed and trust.

Declarations

Funding: No external funding is declared for this review manuscript.

Data availability: No new empirical dataset was created. The review is based on the cited literature.

Conflict of interest: The author(s) should insert the appropriate declaration before submission.

References

1. Arrieta, A.B., Diaz-Rodriguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., Herrera, F.: Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion 58, 82-115 (2020). https://doi.org/10.1016/j.inffus.2019.12.012

2. Kuziemski, M., Misuraca, G.: AI governance in the public sector: Three tales from the frontiers of automated decision-making in democratic settings. Telecommunications Policy 44(6), 101976 (2020). https://doi.org/10.1016/j.telpol.2020.101976

3. Mohun, J., Roberts, A.: Cracking the Code: Rulemaking for Humans and Machines. OECD Working Papers on Public Governance, No. 42 (2020). https://doi.org/10.1787/3afe6ba5-en

4. Benedikt, L., Joshi, C., Nolan, L.A., Henstra-Hill, R., Shaw, L., Hook, S.: Human-in-the-loop AI in government: A case study. Proceedings of the 25th International Conference on Intelligent User Interfaces, 488-497 (2020). https://doi.org/10.1145/3377325.3377489

5. Busuioc, M.: Accountable artificial intelligence: Holding algorithms to account. Public Administration Review 81(5), 825-836 (2021). https://doi.org/10.1111/puar.13293

6. 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 38(3), 101577 (2021). https://doi.org/10.1016/j.giq.2021.101577

7. Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J.W., Wallach, H., Daume III, H., Crawford, K.: Datasheets for datasets. Communications of the ACM 64(12), 86-92 (2021). https://doi.org/10.1145/3458723

8. Suresh, H., Guttag, J.: A framework for understanding sources of harm throughout the machine learning life cycle. Equity and Access in Algorithms, Mechanisms, and Optimization, 1-9 (2021). https://doi.org/10.1145/3465416.3483305

9. Alshahrani, A., Dennehy, D., Mantymaki, M.: An attention-based view of AI assimilation in public sector organizations: The case of Saudi Arabia. Government Information Quarterly 39(4), 101617 (2022). https://doi.org/10.1016/j.giq.2021.101617

10. Valle-Cruz, D., Fernandez-Cortez, V., Gil-Garcia, J.R.: From E-budgeting to smart budgeting: Exploring the potential of artificial intelligence in government decision-making for resource allocation. Government Information Quarterly 39(2), 101644 (2022). https://doi.org/10.1016/j.giq.2021.101644

11. Schiff, D.S., Schiff, K.J., Pierson, P.: Assessing public value failure in government adoption of artificial intelligence. Public Administration 100(3), 653-673 (2022). https://doi.org/10.1111/padm.12742

12. Ingrams, A., Kaufmann, W., Jacobs, D.: In AI we trust? Citizen perceptions of AI in government decision making. Policy & Internet 14, 390-409 (2022). https://doi.org/10.1002/poi3.276

13. Alon-Barkat, S., Busuioc, M.: Human-AI interactions in public sector decision making: Automation bias and selective adherence to algorithmic advice. Journal of Public Administration Research and Theory 33(1), 153-169 (2023). https://doi.org/10.1093/jopart/muac007

14. Madan, R., Ashok, M.: AI adoption and diffusion in public administration: A systematic literature review and future research agenda. Government Information Quarterly 40(1), 101774 (2023). https://doi.org/10.1016/j.giq.2022.101774

15. Long, Y., Gil-Garcia, J.R.: Understanding the extent of automation and process transparency appropriate for public services: AI in Chinese local governments. International Journal of Electronic Government Research 19(1), 1-20 (2023). https://doi.org/10.4018/IJEGR.322550

16. Grimmelikhuijsen, S.: Explaining why the computer says no: Algorithmic transparency affects the perceived trustworthiness of automated decision-making. Public Administration Review 83(2), 241-262 (2023). https://doi.org/10.1111/puar.13483

17. Konig, P.D., Felfeli, J., Achtziger, A., Wenzelburger, G.: The importance of effectiveness versus transparency and stakeholder involvement in citizens' perception of public sector algorithms. Public Management Review 26(4), 1061-1082 (2024). https://doi.org/10.1080/14719037.2022.2144938

18. Kumar, R., Singh, A., Kassar, A.S.A., Humaida, M.I., Joshi, S., Sharma, M.: Leveraging artificial intelligence to achieve sustainable public healthcare services in Saudi Arabia: A systematic literature review of critical success factors. Computer Modeling in Engineering & Sciences 142(2), 1289-1349 (2025). https://doi.org/10.32604/cmes.2025.059152

19. 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 41(4), 101965 (2024). https://doi.org/10.1016/j.giq.2024.101965

20. Roehl, U.B.U., Hansen, M.B.: Automated, administrative decision-making and good governance: Synergies, trade-offs, and limits. Public Administration Review 84(6), 1184-1199 (2024). https://doi.org/10.1111/puar.13799

21. Kronblad, C.: When justice is blind to algorithms: Multilayered blackboxing of algorithmic decision-making in the public sector. MIS Quarterly 48(4), 1637-1662 (2024). https://doi.org/10.25300/MISQ/2024/18251

22. Kennedy, R.: Rules as code and the rule of law: Ensuring effective judicial review of administration by software. Law, Innovation and Technology 16(1), 170-193 (2024). https://doi.org/10.1080/17579961.2024.2313801

23. Ibrahim, N.M.H.: Artificial Intelligence (AI) and Saudi Arabia's governance. Journal of Developing Societies 40(4), 500-530 (2024). https://doi.org/10.1177/0169796X241288590

24. Majrashi, K.: Determinants of public sector managers' intentions to adopt AI in the workplace. International Journal of Public Administration in the Digital Age 11(1), 1-26 (2024). https://doi.org/10.4018/IJPADA.342849

25. Hasan, A., Alenazy, A.A., Habib, S.: Examining the drivers and barriers to adoption of e-government services in Saudi Arabia. Journal of Innovative Digital Transformation 1(2), 139-157 (2024). https://doi.org/10.1108/JIDT-09-2023-0019

26. Batool, A., Zowghi, D., Bano, M.: AI governance: A systematic literature review. AI and Ethics 5, 3265-3279 (2025). https://doi.org/10.1007/s43681-024-00653-w

27. Almeida, P.G.R., Santos Junior, C.D.: Artificial intelligence governance: Understanding how public organizations implement it. Government Information Quarterly 42(1), 102003 (2025). https://doi.org/10.1016/j.giq.2024.102003

28. Papagiannidis, E., Mikalef, P., Conboy, K.: Responsible artificial intelligence governance: A review and research framework. Journal of Strategic Information Systems 34(2), 101885 (2025). https://doi.org/10.1016/j.jsis.2024.101885

29. Saeed, S.: Digital transformation in governmental public service provision and usable security perception in Saudi Arabia. Information 16(3), 247 (2025). https://doi.org/10.3390/info16030247

30. Chen, J., Lou, Z., Chen, G., Zhang, C., Li, Y.: Research on the application of AI-based intelligent decision support systems in e-government services. Journal of Organizational and End User Computing 37(1) (2025). https://doi.org/10.4018/JOEUC.392618