Decision-Centric Process Mining and Automation Prioritisation for Saudi Public and Financial Services under Vision 2030: Integrating Business Process Management, Robotic Process Automation, and Benefit–Risk Analysis.

Abstract

Digital transformation programmes increasingly possess event traces but lack defensible methods for deciding which workflows should be automated, redesigned or retained under human control. This critical review examines how business process management, process mining, robotic process automation and benefit–risk analysis can jointly improve prioritisation in Saudi public administration and financial services, while extending the analysis to electricity utility interfaces. Thirty DOI-identified scholarly works published between 2021 and 2025 are synthesised through an explicit thematic appraisal covering data quality, decision discovery, technical suitability, value measurement, governance and contextual transferability. The proposed decision-centric framework separates evidence about process behaviour from normative decisions about public value, legal authority and operational risk. It introduces auditable eligibility gates, scenario-sensitive prioritisation, human decision rights and post-deployment monitoring. Application pathways are developed for permits, banking onboarding, electricity connections, metering disputes and grid-related service requests. The analysis identifies tensions between automation efficiency and procedural fairness, brittle interface automation and architectural renewal, and cross-agency visibility and data minimisation. Its principal contribution is a testable governance architecture rather than an unsupported prediction of savings. Saudi institutions can use the framework to document investment rationales, prohibit unsafe automation and measure realised outcomes against service quality, resilience and citizen protection.

Keywords: process mining; decision mining; robotic process automation; business process management; benefit–risk prioritisation; Saudi Vision 2030; electricity services; public governance

UNDER PEER REVIEW IN IJAR

1. Introduction

Saudi Arabia's Vision 2030 positions digital public services, financial-sector development and dependable electrical infrastructure within a common modernisation agenda. Digital portals and bank applications are increasingly the visible surface of processes that traverse agencies, identity providers, payment gateways, utilities and legacy records. Their effectiveness depends less on the number of

automated tasks than on the quality of decisions about exceptions, authorisations and handovers. A licensing request may be completed rapidly yet remain unjust if an opaque rule rejects an eligible applicant; an electricity connection may be processed cheaply yet expose users to safety risks if technical approval is bypassed. These distinctions make automation prioritisation a governance problem rather than a software procurement exercise.

Process mining reconstructs executed workflows from timestamped events, whereas business process management defines, redesigns and governs the intended process. Robotic process automation (RPA) performs specified interactions, often through user interfaces. Existing evidence finds complementarities between mining and RPA but also unresolved data preparation, task granularity and tool-integration problems [1]. Selection studies consequently emphasise feasibility and stable inputs, although operational simplicity does not establish social desirability [2,3]. Bot-level analysis adds visibility into whether automation actually changes waiting times, control breaks and downstream work [4]. These contributions create an opportunity to connect factual process evidence to explicit investment decisions.

The paper asks: how can Saudi public, financial and electricity-related service organisations identify automation candidates without mistaking high transaction volumes for public value? What controls should separate routine execution from consequential judgement? How should benefits be discounted for uncertainty, operational dependence and risk? Its aim is to develop an evidence-grounded, decision-centric prioritisation architecture suitable for critical assessment and local testing. The objectives are to synthesise complementary technical and organisational research; specify a defensible event-to-decision pipeline; design a transparent benefit–risk selection rule; examine domain-specific constraints; and identify falsifiable propositions for future studies. No Saudi deployment or causal benefit is claimed.

2. Conceptual foundations and critical tensions

The traditional process perspective uses cases, activities, timestamps and resources to identify variants, throughput and conformance deviations. Decision-centric analysis shifts the analytical unit to a consequential choice within a process: who authorised it, under which evidence and rule, with what alternatives and downstream effects? Decision discovery from object-centric data illustrates how decisions involve interacting business entities rather than isolated case histories [5]. Research on robotic process mining similarly shows that desktop actions can illuminate previously invisible work, but screen traces introduce surveillance, privacy and interpretability concerns [6]. A decision model is

UNDER PEER REVIEW IN IJAR

therefore neither equivalent to a process map nor sufficient evidence that its rules ought to be automated.

A further tension concerns organisational boundaries. Public services rely on more than one registry, while banks connect compliance, customer, payments and third-party platforms. Cross-organisational mining reveals coordination opportunities but raises challenges involving heterogeneous identifiers, shared semantics and data-access rights [7]. Event-stream conformance research demonstrates that perpetual retention is not technically necessary for all monitoring tasks, offering an important counterweight to indiscriminate data capture [8]. Uncertain traces complicate conclusions about non-compliance when event times, labels or attributes may be incomplete [9]. In Saudi settings, a missing event must not automatically be interpreted as administrative misconduct or customer fault.

The process-mining literature is also moving beyond a single case identifier. An electricity service journey may link one customer, several premises, meters, work orders, contractors and invoices. Object-centric monitoring can represent operational problems across these related entities [10]. New approaches incorporating unstructured sensor or documentary information may improve coverage but increase provenance and validation requirements [11]. These developments favour hybrid evidence models over a simplistic assumption that a single, complete event log exists. They also make architecture choices material: integration at the service or application-programming-interface level may ultimately be more resilient than automating screen clicks.

RPA governance scholarship identifies project phases, operating models and implementation mechanisms, rather than treating bot deployment as a single technical decision. Implementation frameworks emphasise discovery, design, construction, deployment and maintenance [12]; organisational studies distinguish different purposes and mechanisms of automation [13]. Choices between internal and external development, cloud and local hosting, and proprietary or open tools determine long-term accountability [14]. Process audits demonstrate the usefulness of traces for internal control, but an audit trail cannot prove the substantive legitimacy of a policy or the completeness of source events [15]. These insights support a layered model in which process redesign, automation mode and decision authority are selected together.

An additional conceptual distinction is needed between detecting an inefficient path and estimating the effect of changing it. High waiting time may coincide with a particular team because that team handles more difficult cases, not because its staffing causes delay. Nor does observing an automated approval route establish what would have happened under human review. Counterfactual claims

UNDER PEER REVIEW IN IJAR

therefore require comparable cases, defensible outcome definitions and awareness of selection bias. Decision discovery research makes this problem especially salient: inferred branch conditions may reproduce historical choices and their inequities rather than identify an independently justified eligibility rule [18]. The proposed architecture treats mining outputs as diagnostic evidence and requires separate evaluation before a causal or normative claim is made.

3. Review aim, design and evidence appraisal

This article uses a structured critical integrative review, not a census of all published studies or a statistical meta-analysis. The review window was restricted to 2020–2025, English-language research with a traceable digital object identifier and substantive relevance to process discovery, automation selection, governance, Saudi digital services or electricity-system information security. Searches combined phrases equivalent to process mining, decision mining, automation suitability, RPA, public governance, financial services and smart grids across scholarly publisher records, institutional repositories and bibliographic discovery sources. Citation chaining was used selectively to connect technical methods with organisational and regional evidence. Publication metadata and DOI strings were checked against identifiable publisher or scholarly records.

Thirty relevant publications constitute the purposively assembled evidence set. Selection favoured work with explicit method descriptions, implementation detail or analytically distinctive findings, while excluding undated commentary, unsupported vendor savings claims, and sources outside the publication window. Evidence was coded under six themes: event reliability, process and decision discovery, technical feasibility, value and risk evaluation, institutional accountability, and Saudi transferability. Research designs were interpreted separately: conceptual frameworks propose mechanisms; case studies establish situated feasibility; systematic reviews synthesise heterogeneous findings; experiments provide bounded technical evidence. Their claims were not pooled into a fictitious common effect size.

The appraisal used five questions: Is the construct defined clearly? Are data provenance and analytical assumptions transparent? Does the design identify limitations or alternatives? Are reported outcomes distinguished from predicted benefits? Can the result plausibly transfer to high-stakes Saudi services? This approach follows the organisational process-management emphasis on context, engineering fit and continuous improvement [16]. It is also consistent with work arguing that process analytics should enable augmented execution rather than terminate at a static dashboard [17]. The principal limitations are deliberate selectivity, reliance on published evidence, uneven access to full

UNDER PEER REVIEW IN IJAR

texts and the absence of an independently replicated database screening census. Accordingly, the synthesis supports conceptual design and testable recommendations, not prevalence estimates or claims of universal optimality.

The regional literature was deliberately used as contextual evidence rather than as proof that a particular technology has been adopted in Saudi agencies. For instance, survey measures of trust or acceptance cannot verify achieved reductions in processing time, while conceptual smart-grid architectures cannot establish operating performance at a named utility. The appraisal also separated ex ante projected benefits from post-deployment measurements and recorded whether studies acknowledged failure cases. Differences in publication venues and research methods prevented a meaningful pooled ranking of technologies. This interpretive discipline matters because confident, cross-sector recommendations can otherwise rest on incomparable denominators, undisclosed baselines or enthusiasm for an implementation method.

4. Evidence architecture: from events to defensible decisions

Figure 1 specifies six sequential but revisable layers. First, an agency defines the service outcome and identifies which decisions have legal, financial, safety or citizen-rights consequences. Second, it constructs a minimal event schema joining case and object identifiers, activity, timestamp, actor class, system, decision outcome and reason code. Logs should distinguish customer waiting from organisational processing, authorised from unauthorised rework, and failed technical actions from rejected business rules. Third, mining tools reconstruct variants, bottlenecks, handovers and conformance patterns. Fourth, decision discovery associates branch choices with documented policy conditions. Fifth, a cross-functional committee evaluates the candidate against value, risk and authority gates. Sixth, a controlled intervention is monitored, with an explicit ability to suspend or redesign automation.

UNDER PEER REVIEW IN IJAR

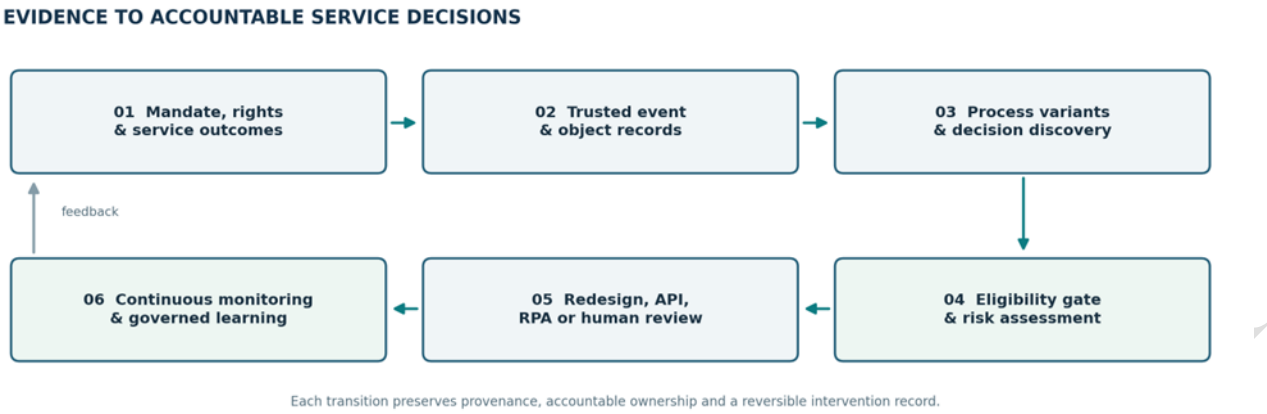


Figure 1. Decision-centric event-to-intervention architecture with governed feedback (original conceptual synthesis).

The architecture must not equate statistically frequent paths with correct paths. A common approval route could reflect routine compliance or an entrenched inequity; a rare exception could be a legitimate accessibility accommodation. Process-aware decision discovery offers a means of making rule branches visible, but the inferred rules remain hypotheses requiring policy-owner verification [5]. Monitoring research highlights the importance of detecting emerging deviations through related objects [10]. The crucial distinction is that event logs answer what happened and sometimes why a particular technical branch fired; authorised decision makers determine whether that branch should remain permissible.

Event quality controls should include stable clocks, versioned reason codes, unique correlation identifiers, access logging and provenance at extraction. Data joins spanning government departments should use lawful minimisation and role-specific permissions. Saudi personal-data requirements and national cybersecurity controls create relevant governance boundaries, while banking controls add sector-specific confidentiality and continuity duties. Electricity operational-technology environments require separation between business-process analytics and commands affecting physical equipment. A mined pattern can trigger a recommendation to review a connection backlog, but it must not directly issue an energisation instruction without engineering authority and independent safeguards.

Automation modes should be ordered by durability rather than fashionable terminology. Simplifying or eliminating a redundant step often dominates bot replication; a validated service integration can be preferable to interface-level scripting. RPA remains appropriate where systems expose no usable integration, tasks are stable, inputs are structured and exceptions can be routed safely. Intelligent document extraction may support staff but should not make unreviewable eligibility

UNDER PEER REVIEW IN J.

judgements. The intelligent-automation literature describes opportunities for more adaptive workflows, yet its broad conceptual claims do not establish that increasingly autonomous agents are proportionate in every service [19]. Selection must therefore consider a deliberate no-automation option.

The decision register associated with Figure 1 should capture each decision point's accountable owner, policy instrument, permitted inputs, approval threshold, explanation requirement, appeal path and retention rule. It should also preserve policy-version history so that an apparent drift in routing can be distinguished from a legitimate regulatory revision. Where a case depends on third-party evidence, the register records whether the evidence is authoritative, inferred or supplied by the applicant. This granular design supports review of individual adverse outcomes without giving process analysts unrestricted access to personal files. Importantly, a decision register is an organisational control: machine-generated decision rules cannot replace formal delegations of authority.

5. Benefit–risk prioritisation model

The proposed method treats each candidate as a portfolio decision with three stages: non-negotiable eligibility, comparative scoring and observed-benefit review. In the first stage, process owners document statutory authority, safety criticality, personal-data purpose, cybersecurity classification, input provenance, accessibility obligations and escalation routes. Candidates fail the gate if adequate decision rights, safe rollback or a legally defensible data basis are absent. Such failures cannot be offset by an attractive labour-saving forecast. They instead trigger process redesign, data remediation or supervised decision support.

Candidates passing the gates receive a transparent, context-sensitive evaluation. A benefit vector includes expected reduction in user waiting, improved first-time-right completion, staff capacity released, fraud or error prevention, accessibility and reliability gains. An implementation vector covers event completeness, task stability, interface fragility, cross-entity dependencies and maintenance demands. A risk vector covers wrongful decisions, privacy exposure, attack surface, service interruption, dependence on suppliers and loss of meaningful review. Rather than compressing fundamentally different harms into a single opaque grade, decision makers first assign thresholds to each critical risk dimension and then compare remaining options using stakeholder-approved weights. Costa and colleagues demonstrate that multi-criteria techniques can support automation selection [3]; their contribution does not remove the need for local value judgements.

UNDER PEER REVIEW IN IJAR

A defensible scoring expression is P(i) = sum of weighted, normalised benefit measures minus sum of weighted, normalised implementation burdens and residual risks, conditional on all mandatory controls being satisfied. Weights are published within the institution, set before inspecting candidate rankings and tested under alternative scenarios. Benefits are expressed against a recorded baseline and corrected for displaced work, extra review, bot licences, maintenance and exceptions. Consequences with fundamentally non-compensable safety or legal implications are treated as gates rather than negative points. Farinha and colleagues demonstrate how automation criteria can be assembled and evaluated across organisational processes [2], while wider RPA reviews caution that reported gains depend on implementation conditions [20,21].

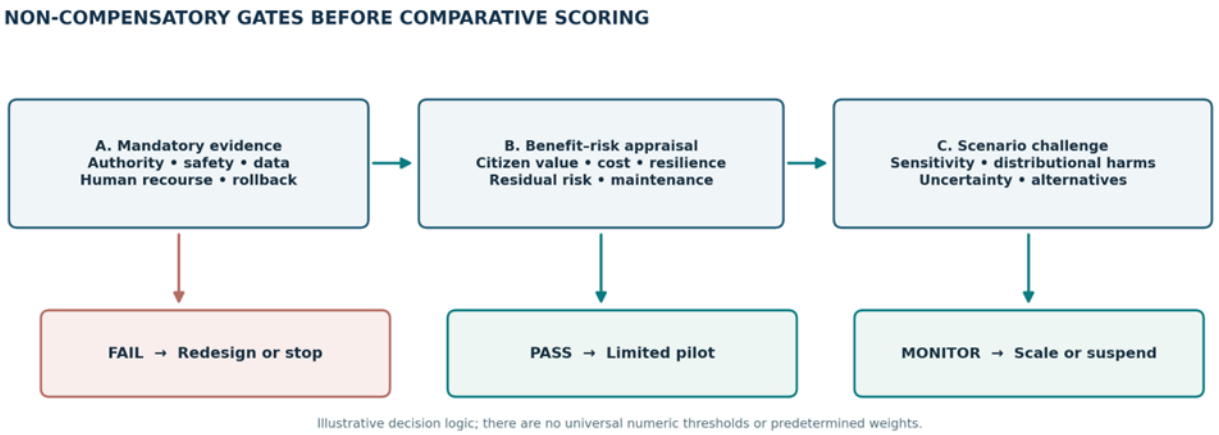


Figure 2. Non-compensatory eligibility gate and scenario-sensitive automation prioritisation (original conceptual synthesis).

Sensitivity analysis is essential because a single weight set may conceal disagreements between finance, compliance, accessibility, cybersecurity and operational staff. The committee should publish whether a ranking changes when privacy risk or service continuity receives greater emphasis, and explicitly document dissent. A candidate with modest but robust public value may be preferable to a high-volume activity whose potential saving depends on perfect document classification. Decision models must include a not-ready category, not merely a ranked queue. Figure 2 shows how the three-stage structure can separate attractive, uncertain and prohibited opportunities without inventing numerical cut-offs.

Three separate financial viewpoints should accompany the score. The transaction-level view measures touch time, rework and exception frequency; the programme view measures integration effort, governance, support, training and licensing; and the public-value view considers citizen effort, avoidable delay and service availability. Discounting projected benefits for uncertain event coverage

UNDER PEER REVIEW IJAR

is preferable to treating missing observations as zero-cost cases. The committee should distinguish one-off migration expenditure from recurrent operating expense and specify the period over which benefits are expected. It should also identify changes in demand: an increase in applications can obscure real efficiency gains unless productivity and user outcomes are normalised appropriately. None of these adjustments licenses an invented return-on-investment figure.

6. Saudi public, banking and electricity applications

Saudi digital government research provides evidence that technology adoption depends on institutional readiness, coordination and stakeholder trust. A multi-case study of emerging technologies across government sectors identifies context-specific drivers and barriers [22]. The transition from electronic transactions to joined-up government requires organisational change and inclusive design, not only new interfaces [23]. Studies of smart-government adoption and e-government participation similarly identify infrastructure, perceived usefulness, security and human factors as material determinants [24,25]. These findings support extending process mining beyond back-office efficiency towards verified service outcomes.

For public permits, the relevant unit is the complete applicant journey, including identity verification, document requests, inter-agency approvals and appeals. A candidate for unattended RPA might be deterministic status synchronisation between authorised systems. Conversely, decisions involving eligibility interpretation, adverse outcomes or contested evidence should remain under accountable review until decision rules and appeal mechanisms are tested. A baseline should capture median and tail waiting times, repeat submissions, incorrect rejection rates and accessibility performance. A shortened process is not successful if it shifts effort to citizens or suppresses legitimate appeals.

Table 1. Cross-sector evidence-to-action matrix for Saudi service workflows.

Public permit or licence	Journey variants; repeated data requests; appeal reversals	Remove duplicate checks; automate authorised status exchange	Retain appeal rights and legal interpretation; measure rejections
Bank	Reconciliation	API validation and	Preserve due diligence,

UNDER PEER REVIEW IN IJAR

onboarding	queues; exception patterns; document rework	deterministic reconciliation	segregation of duties and audit trail
Electricity connection	Linked applicant, site, work order and inspection events	Automate routing, document completeness and notifications	Engineer approves safety, capacity and energisation
Meter or billing dispute	Reading provenance; adjustments; complaint sequences	Automatic evidence collation and case assignment	Human review contested charges and attribution
Outage service communication	Ticket chronology; dispatch handovers; status corrections	Synchronise approved customer-facing updates	No bot access to switching or protection controls

Note. Analytical applications are proposals, not observed Saudi deployments; each pathway requires local legal and technical validation.

Banking onboarding and transaction servicing offer high event volumes but stringent consequences. Mining may reveal repeated customer-document collection, unnecessary handovers, sanction-screening queues or reconciliation delays. Deterministic reconciliation can be prioritised when segregation of duties and exception management survive automation; customer due-diligence exceptions and suspicious-activity decisions require distinct oversight and evidence retention. Studies of RPA implementation show that operating-model choice affects accountability [14], while Saudi e-government cybersecurity research reinforces the importance of organisational security practices [26]. Human approval should be risk-specific rather than a nominal click that routinely ratifies bot outputs.

Electricity provides a strategically important bridge between public-service delivery and physical infrastructure. Saudi research describes smart grids as an enabler of renewable integration, demand management and more capable electricity services [27]. A Saudi-focused security model identifies requirements for protecting connected grid information flows [28]. Process mining could reconstruct connection applications, meter installation, billing disputes, outage communications, procurement approvals and maintenance work-order administration. These are plausible analytical use cases, not

UNDER PEER REVIEW IN IJAR

claims that a Saudi utility already operates the proposed system. Repetitive customer updates or standard document transfer may be suitable for automation; network switching, protection settings and safety-critical energisation should remain outside ordinary office-bot privileges.

A practical electricity example involves a connection request linking the applicant, plot, transformer capacity assessment, contractor visit, payment record and commissioning certificate. One-case mining can wrongly attribute delays or duplicate shared events; object-aware methods help preserve the relationships. A false-positive compliance alert may originate in mismatched timestamps rather than an actual missed inspection. The appropriate intervention may therefore be event standardisation before automation. Process-aware architectures developed in Saudi research also suggest the importance of tracking evidence across distributed digital services [29]. Strong data lineage and operational-technology segregation are prerequisites for moving from service-workflow automation towards dependable utility-wide intelligence.

In the electricity context, a specific governance hazard is the conflation of administrative closure with technical completion. A work order may be marked closed after a contractor upload even though an independent inspection or safety certificate is still pending. Process mining can reveal mismatched sequences, but the correct response is to verify the engineering evidence and correct the workflow before accelerating closure. Likewise, meter-event anomalies may arise from network outages, synchronisation failures or lawful adjustments, not necessarily customer misconduct. A decision-centric approach permits automatic notification and evidence collation while reserving disputed attribution and financial adjustment for a controlled review path. This retains the productivity value of automation without turning technical uncertainty into unchallengeable charges.

7. Critical discussion: accountability, equity and resilience

The framework challenges the assumption that a favourable cost-benefit calculation settles an automation decision. Benefits and harms are distributed unevenly. A public agency may save staff hours while applicants face repeated digital identity failures; a bank may reduce case handling time while vulnerable customers encounter inexplicable exceptions. Saudi research on the perceived security and privacy of digital government services underlines the relevance of user trust in this context [30]. Distributional assessment must therefore record outcomes by service channel and appropriately governed user group, without generating unnecessary sensitive profiles.

The second tension is between visibility and data minimisation. Event enrichment improves model interpretability, yet indiscriminate linking can expose banking, welfare or electricity-consumption

UNDER PEER REVIEW IN IJAR

information beyond its original purpose. Privacy-preserving extraction, pseudonymised identifiers, retention schedules and federated analysis should be evaluated alongside analytic accuracy. Zaman and colleagues show that conformance monitoring can be designed under retention constraints [8]. The proper response to incomplete observation is uncertainty reporting, not quiet enlargement of surveillance. Human reviewers need both a reasoned explanation of the model and the capacity to access authoritative primary evidence.

Third, automation changes the risk landscape rather than merely eliminating manual error. Interface bots can fail silently after software updates, credentials may be overprivileged, and incident response may depend on vendor expertise. Monitoring bot actions and whole-process effects can detect rework and handover displacement [4], but resilience additionally requires fallbacks, independent reconciliation, credential isolation, controlled change management and rehearsed restoration. In electricity services, the boundary between corporate information systems and operational equipment should be explicit; no service-level performance target justifies bypassing a safety interlock.

Fourth, governance can become ceremonial. A risk committee may approve scores without examining evidence quality, while staff nominally assigned to oversight face intolerable workloads or lack override authority. The governance design therefore allocates named decision ownership, conflict-resolution rules, escalation service levels and a documented power to stop a bot. Auditors should sample rejected and appealed cases, inspect dataset representativeness and compare expected with realised benefits. The committee should reassess the entire workflow after major policy, platform or threat changes, not simply renew a licence. These obligations distinguish accountable automation from dashboard-centric managerial optimism.

The model also exposes a conflict between operational transparency and security disclosure. Citizens and customers should receive intelligible reasons for adverse outcomes, yet revealing fraud thresholds or sensitive infrastructure details may enable circumvention. Layered explanations can reconcile these demands: an affected person receives the decisive rule category, supporting evidence and route to challenge; supervisors receive operational diagnostics; authorised security teams examine sensitive detection parameters. Roles, not a vague promise of explainability, determine which information may be disclosed. A meaningful appeal requires a reviewer who can change the outcome, not simply restate the automated result. This principle is particularly important where service access affects household electricity, financial inclusion or essential public entitlements.

UNDER PEER REVIEW IN IJAR

8. Research gaps and empirical validation agenda

Five gaps follow from the synthesis. First, available RPA criteria often treat suitability as a stable property of an activity, whereas Saudi services change with policy, digital identity, payment rules and infrastructure investment. Longitudinal studies should examine ranking drift before and after such changes. Second, evidence on cross-organisational mining seldom resolves which institution has authority to approve or appeal a decision. Comparative governance experiments could measure coordination quality under centralised, federated and service-owner models.

Third, outcome measurement remains inconsistent. Studies should prospectively register baseline definitions, costs and adverse outcomes; evaluate both throughput and public value; and report uncertainty rather than point estimates alone. A stepped implementation or matched comparison design could test whether the proposed prioritisation method reduces errors and appeals relative to existing selection practice. Fourth, research must evaluate the usability of human oversight. Reviewers should be assessed for understanding, override frequency, workload, response times and independent judgement, not merely the presence of an approval field.

Fifth, electricity-linked public services warrant explicit analysis of cyber-physical boundaries. Research could combine synthetic, properly anonymised connection and meter-workflow data with structured safety scenarios to test the effect of object-centric reconstruction and risk gates. The evaluation should report false conformance alerts, missed critical exceptions, data linkage failures and recovery times. None of these measures should be inferred from simulated examples as evidence of actual Saudi service performance. A public evaluation protocol and transparent documentation of excluded cases would allow results to be challenged and replicated.

A realistic evaluation could recruit multiple service teams, collect a preregistered baseline, and assign comparable workflows to conventional selection or the proposed governed selection process. The primary outcomes would include corrected end-to-end turnaround time, materially wrong outcomes, appeal reversals and service disruptions; secondary outcomes would include staff workload, operational cost and perceived legitimacy. Analyses should examine process complexity and baseline data completeness as potential moderators. For security-sensitive systems, published datasets may need synthetic substitutes, but independent reviewers should inspect the generation assumptions and confirm that confidentiality constraints do not conceal unfavourable results. Researchers must report implementation failures and withdrawals, not only successful candidates.

UNDER PEER REVIEW IN IJAR

Evaluation should also consider opportunity costs: selecting an administratively convenient bot may divert scarce engineering and governance capacity from a more valuable integration project. Prioritisation records should therefore display deferred alternatives, their readiness dependencies and the expected consequences of postponement. This portfolio perspective provides an explicit counterfactual for management decisions and helps explain why an eligible automation was not funded. It also discourages the misleading equation of a longer bot inventory with greater institutional maturity.

9. Practical implementation roadmap

Implementation should begin with a bounded service inventory and a shared glossary of decisions, activities, objects and outcomes. A mixed team comprising service owners, business architects, data stewards, finance, legal, cyber specialists, accessibility representatives and operational engineers should choose one low-risk process for detailed mapping. The team should validate sampled event histories against authoritative files and establish baseline performance before recommending automation. Table 1 translates the literature into sector-specific diagnostic questions; Table 2 provides an auditable measurement catalogue.

Table 2. Operational rubric for transparent benefit–risk appraisal.

Public/service value	End-to-end waiting; accessibility; first-time resolution	Normalise to pre-intervention baseline; stratify by channel	Do not trade away appeal or legal access
Financial benefit	Net operating cost; rework; maintenance; displaced tasks	Separate one-off from recurrent costs; challenge assumptions	No claimed saving without attributable evidence
Technical readiness	Event coverage; stable rules; API/interface durability	Check source provenance and version sensitivity	Remediate unreliable event data first
Decision quality	Error rates; explanations; appeal	Sample consequential cases and monitor drift	Manual decision authority for disputed

UNDER PEER REVIEW IN IJAR

	reversals		outcomes
Privacy and security	Data purposes; privileged access; attack surface	Data minimisation, role controls, threat assessment	Block unlawful access or unacceptable exposure
Operational resilience	Fallback time; incident recovery; third-party dependencies	Test recovery and independent reconciliation	Prohibit unsafe grid-control integration
Portfolio uncertainty	Sensitivity; alternative investments; policy changes	Record weight changes, variance and deferred options	Return unstable rankings for further evidence

Note. Weights and thresholds must be selected by accountable Saudi service owners; no universal scoring coefficients are asserted.

The subsequent pilot should compare at least three alternatives: remove or simplify the step, integrate systems through governed interfaces, or automate a stable task with supervised exceptions. The decision record should explain why one route was chosen and preserve evidence of security, data protection, service equality and recovery testing. During deployment, outcome monitoring must cover the full user journey rather than only bot execution success. Persistent deviations, unexpected adverse decisions, unexplained data access or unresolved safety exceptions should suspend the intervention pending investigation.

A scaled programme requires portfolio governance, not a growing collection of isolated scripts. Shared event definitions, reusable controls, versioned decision tables and consistent benefit baselines permit cross-service comparison. Nevertheless, national standardisation should not erase differences between public entitlements, bank risk controls and electricity engineering rules. Each institution remains accountable for the lawful purpose of its automation, while an enterprise-level board can govern architectural coherence, assurance evidence and aggregate investment decisions.

Before scaling, managers should establish explicit stop conditions and a quarterly review cadence tailored to the service's risk and change rate. Example triggers include unexplained growth in adverse

UNDER PEER REVIEW IN IJAR

decisions, a sudden change in exception mix, repeated access-control violations, missed safety attestations or persistent disagreement between process logs and authoritative records. An incident review should ask whether the defect arose in source data, process design, decision rules, bot execution or monitoring. Each cause requires a different remedy; replacing a bot cannot fix an ambiguous legal rule. Reusable incident taxonomies could strengthen learning across Saudi organisations without requiring uncontrolled sharing of case-level personal data.

10. Conclusion

Decision-centric process mining offers a more defensible basis for Saudi automation investment than transaction volume or predicted labour savings alone. The review connects process evidence, decision discovery, BPM redesign, RPA feasibility and explicit benefit–risk governance in a single analytical architecture. Its central principle is asymmetry: efficiency benefits can be traded against ordinary costs, but missing legal authority, inadequate human recourse and unacceptable safety exposures cannot be compensated by higher scores. A Saudi service portfolio can consequently prioritise reliable administrative tasks while reserving sensitive banking and electricity decisions for accountable judgement. The framework remains a proposition requiring longitudinal evaluation, independently checked baselines and stakeholder scrutiny. Its value will be demonstrated only if institutions can show better service outcomes, preserved rights and resilience after deployment.

References

[1] El-Gharib, N. M. and Amyot, D. (2023). Robotic process automation using process mining—A systematic literature review. Data & Knowledge Engineering, 148, 102229. https://doi.org/10.1016/j.datak.2023.102229

[2] Farinha, D., Pereira, R. and Almeida, R. (2024). A framework to support robotic process automation. Journal of Information Technology, 39(1). https://doi.org/10.1177/02683962231165066

[3] Costa, D. S., Mamede, H. S. and da Silva, M. M. (2023). A method for selecting processes for automation with AHP and TOPSIS. Heliyon, 9(3), e13683. https://doi.org/10.1016/j.heliyon.2023.e13683

[4] Egger, A., ter Hofstede, A. H. M., Kratsch, W., Leemans, S. J. J., Röglinger, M. and Wynn, M. T. (2024). Bot log mining: An approach to the integrated analysis of robotic process automation and process mining. Information Systems, 126, 102431. https://doi.org/10.1016/j.is.2024.102431

UNDER PEER REVIEW IN IJAR

[5] Goossens, A., De Smedt, J. and Vanthienen, J. (2024). Extracting process-aware decision models from object-centric process data. Information Sciences, 682, 121263. https://doi.org/10.1016/j.ins.2024.121263

[6] Leno, V., Polyvyanyy, A., Dumas, M., La Rosa, M. and Maggi, F. M. (2021). Robotic process mining: Vision and challenges. Business & Information Systems Engineering, 63, 301–314. https://doi.org/10.1007/s12599-020-00641-4

[7] Rott, J., Böhm, M. and Krcmar, H. (2024). Laying the ground for future cross-organizational process mining research and application: A literature review. Business Process Management Journal, 30(8), 144–206. https://doi.org/10.1108/BPMJ-04-2023-0296

[8] Zaman, R., Hassani, M. and van Dongen, B. F. (2023). Conformance checking of process event streams with constraints on data retention. Information Systems, 117, 102228. https://doi.org/10.1016/j.is.2023.102228

[9] Felli, P., Gianola, A., Montali, M., Rivkin, A. and Winkler, S. (2023). Multi-perspective conformance checking of uncertain process traces: An SMT-based approach. Engineering Applications of Artificial Intelligence, 126, 106895. https://doi.org/10.1016/j.engappai.2023.106895

[10] Park, G. and van der Aalst, W. M. P. (2025). Operational process monitoring: An object-centric approach. Computers in Industry, 164, 104170. https://doi.org/10.1016/j.compind.2024.104170

[11] Egger, A., Fehrer, T., Kratsch, W., Wördehoff, N., König, F. and Röglinger, M. (2025). Refining the process picture: Unstructured data in object-centric process mining. Information Systems, 134, 102582. https://doi.org/10.1016/j.is.2025.102582

[12] Herm, L.-V., Janiesch, C., Helm, A., Imgrund, F., Hofmann, A. and Winkelmann, A. (2023). A framework for implementing robotic process automation projects. Information Systems and e-Business Management, 21, 1–35. https://doi.org/10.1007/s10257-022-00553-8

[13] Crisan, E. L., Chis, D. M., Bodea, E. E. and Buchmann, R. (2023). Mechanisms for robotic process automation implementation in organizations: A systematic literature review. Journal of Advances in Management Research, 20(5), 920–946. https://doi.org/10.1108/JAMR-04-2023-0100

[14] Asatiani, A., Copeland, O. and Penttinen, E. (2023). Deciding on the robotic process automation operating model: A checklist for RPA managers. Business Horizons, 66(1), 109–121. https://doi.org/10.1016/j.bushor.2022.03.004

UNDER PEER REVIEW IN IJAR

[15] Imran, M., Hamid, S. and Ismail, M. A. (2023). Advancing process audits with process mining: A systematic review of trends, challenges, and opportunities. IEEE Access, 11, 68340–68357. https://doi.org/10.1109/ACCESS.2023.3292117

[16] Monti, F., Mathew, J. G., Leotta, F. et al. (2024). On the application of process management and process mining to Industry 4.0. Software and Systems Modeling, 23, 1407–1419. https://doi.org/10.1007/s10270-024-01175-z

[17] Chapela-Campa, D. and Dumas, M. (2023). From process mining to augmented process execution. Software and Systems Modeling, 22, 1977–1986. https://doi.org/10.1007/s10270-023-01132-2

[18] Leewis, S., Smit, K. and Versendaal, J. (2024). Discovering operational decisions from data—A framework supporting decision discovery from data. Decision, 51, 417–436. https://doi.org/10.1007/s40622-024-00402-2

[19] Ng, K. K. H., Chen, C.-H., Lee, C. K. M., Jiao, J. and Yang, Z.-X. (2021). A systematic literature review on intelligent automation: Aligning concepts from theory, practice, and future perspectives. Advanced Engineering Informatics, 47, 101246. https://doi.org/10.1016/j.aei.2021.101246

[20] Moreira, S., Mamede, H. S. and Santos, A. (2023). Process automation using RPA—A literature review. Procedia Computer Science, 219, 244–254. https://doi.org/10.1016/j.procs.2023.01.287

[21] Pramod, D. (2022). Robotic process automation for industry: Adoption status, benefits, challenges and research agenda. Benchmarking: An International Journal, 29(5), 1562–1586. https://doi.org/10.1108/BIJ-01-2021-0033

[22] Alshahrani, A. (2025). Adopting emerging technologies in digital government: A multi-case analysis of drivers, enablers, and challenges in Saudi Arabia. Digital Government: Research and Practice, 6(3), article 36. https://doi.org/10.1145/3719297

[23] Zygiaris, S. and Maamari, B. E. (2023). The journey from e-government to digital transformation: The case of Saudi Arabia. Electronic Government, an International Journal, 19(1), 95–111. https://doi.org/10.1504/EG.2023.127578

[24] Alajmi, M., Mohammadian, M. and Talukder, M. (2023). The determinants of smart government systems adoption by public sector organizations in Saudi Arabia. Heliyon, 9(10), e20394. https://doi.org/10.1016/j.heliyon.2023.e20394

UNDER PEER REVIEW IN IJAR

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

[26] Al-Hawamleh, A. M. (2024). Investigating the multifaceted dynamics of cybersecurity practices and their impact on the quality of e-government services: Evidence from the KSA. Digital Policy, Regulation and Governance, 26(3), 317–336. https://doi.org/10.1108/DPRG-11-2023-0168

[27] Khan, K. A., Quamar, M. M., Al-Qahtani, F. H., Asif, M., Alqahtani, M. and Khalid, M. (2023). Smart grid infrastructure and renewable energy deployment: A conceptual review of Saudi Arabia. Energy Strategy Reviews, 50, 101247. https://doi.org/10.1016/j.esr.2023.101247

[28] Akkad, A., Wills, G. and Rezazadeh, A. (2023). An information security model for an IoT-enabled smart grid in the Saudi energy sector. Computers and Electrical Engineering, 105, 108491. https://doi.org/10.1016/j.compeleceng.2022.108491

[29] Alzhrani, F., Saeedi, K. and Zhao, L. (2024). A process-aware framework to support process mining from blockchain applications. Journal of King Saud University – Computer and Information Sciences, 36, 101956. https://doi.org/10.1016/j.jksuci.2024.101956

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

UNDER PEER REVIEW IN IJAR