Data Governance and Quality Management in AI-Driven Enterprise Information Systems

Abstract

Enterprise information systems increasingly combine transactional databases, cloud services, analytical platforms and artificial intelligence. The value of these capabilities depends on whether the underlying data are accurate, contextualised, timely, lawfully used and accountable to named owners. In Saudi Arabia, public-sector modernisation and enterprise digital transformation make data governance and quality management central to trustworthy AI adoption. This narrative review examines governance responsibilities, data-quality dimensions, metadata and lineage, model reliability, security, privacy and operational assurance. It develops an eight-stage data-to-decision framework linking business purpose, ownership, ingestion, quality controls, authorised access, AI assessment, human action and feedback. The synthesis indicates that governance and quality must be designed as continuous operational practices rather than one-time cleansing projects. Illustrative applications cover government services, universities, logistics and cross-enterprise reporting. Two tables translate the review into control responsibilities and performance indicators, while two conceptual figures show the proposed architecture and the closed-loop assurance cycle. The paper contributes a practical evaluation model for organisations implementing AI-enabled enterprise information systems in Saudi Arabia. It makes no claim of primary empirical testing, and it identifies research priorities for prospective local validation.

Keywords

data governance; data quality; enterprise information systems; trustworthy artificial intelligence; Saudi Arabia; data lineage; metadata; model assurance; data stewardship; digital transformation.

UNDER PEER REVIEW IN IJAR

1 Introduction

Enterprise systems rarely operate as isolated applications. A case management platform may exchange identity details with a registry, a document repository and a dashboard; a logistics system may combine delivery records, warehouse updates and vendor information; and a university may need to reconcile student and financial records across several platforms. The resulting information is valuable only when users can establish which source is authoritative, what a field means and whether its value remains current. When AI is added to these environments, inconsistencies can be magnified because automated classification, forecasting and retrieval operate at scales beyond routine manual review.

Data governance defines how decisions concerning data are made and by whom. Data quality management establishes whether information is fit for particular uses and how defects are detected, corrected and prevented. These disciplines overlap but are not interchangeable. Governance can assign responsibility for a dataset that still contains incorrect entries; a cleansing routine can repair entries without deciding who may disclose them or how long they may be retained. Enterprise readiness requires both. The central managerial challenge is to coordinate ownership, technology, process design and organisational behaviour rather than treat data quality as the responsibility of analysts alone.

Saudi Arabia provides a consequential context for this challenge. National strategies encourage digital public services and data-enabled innovation, while the National Data Management Office and personal-data protection framework establish expectations for responsible handling. Institutions undertaking AI adoption must therefore connect technological experimentation with accountable policies, purpose limitation, metadata, information security and documented rights of access. The existence of strong governance policies, however, does not itself prove their effectiveness in operational workflows.

This review investigates how Saudi organisations can turn governance principles into measurable controls across the enterprise data lifecycle. It focuses on AI-assisted decision support, not fully autonomous administrative authority. The contribution is an integrated and testable operating model connecting source data, quality evidence, model validation and responsible business decisions.

2 Aim, Objectives, and Research Questions

The aim is to develop a practical, evidence-informed data governance and quality management framework for AI-driven enterprise information systems in Saudi Arabia, balancing operational value with accountability, security and lawful processing.

The five objectives are to identify the leading causes and consequences of enterprise data defects; compare governance responsibilities and quality-assurance mechanisms; examine the ways data quality affects conventional analytics and AI models; design a scalable decision-and-control architecture compatible with Saudi institutional requirements; and propose implementation metrics and research priorities for evaluating real operational outcomes.

UNDER PEER REVIEW IN IJAR

Four questions guide the paper. How should organisations allocate responsibility for data as it crosses departmental and vendor boundaries? Which quality dimensions and testing approaches are most meaningful for AI-assisted enterprise decisions? How can metadata, lineage, access controls and human verification prevent misleading outputs or improper processing? Which implementation and assurance practices allow governance maturity to be assessed beyond policy documents or technology adoption counts?

The study treats an enterprise information system as a socio-technical arrangement encompassing information assets, processes, people, platforms and external interfaces. It therefore assesses operational and organisational quality as well as algorithmic performance.

3 Review Methodology

This narrative review develops a conceptual synthesis from the supplied manuscript bibliography and a targeted editorial source check on 24 November 2026. The original database search dates, exported results, duplicate counts and screening decisions were not supplied; consequently, exhaustive database coverage, systematic selection and independent screening cannot be claimed. No pooled effect estimate is attempted. The original bibliography is retained as a provisional source inventory pending full-text verification. Sources are distinguished as empirical studies, conceptual works, standards and official guidance. The editorial checks used public web discovery and publisher or government pages; the accompanying register records the checks and unresolved items. The proposed framework and measurement protocols are design propositions, not observations from a Saudi deployment.

4 Enterprise Data Governance in Saudi Arabia

A governance programme should start with business decisions and legal responsibilities rather than the purchase of a catalogue or data platform. Saudi organisations commonly operate heterogeneous systems and must distinguish personal information from other business records, public information from restricted datasets, and authoritative data from convenience copies. Data strategies should therefore identify the information used to deliver public services, approve expenditure, monitor contracts, assess performance or inform automated recommendations. Every critical dataset requires a documented purpose, owner, permitted uses and decision authority.

The Saudi Data and Artificial Intelligence Authority and its National Data Management Office provide national policy context for data governance, management and protection. The Saudi Personal Data Protection Law and implementing framework require organisations to examine the lawful basis and conditions of personal data processing. Practical governance should therefore address collection minimisation, handling of sensitive data, access and disclosure, retention, data-subject requests and cross-border arrangements as relevant to the organisation and specific activity. The technical design must not substitute for a case-specific legal or compliance assessment.

UNDER PEER REVIEW IN IJAR

In enterprise operations, governance becomes visible at interfaces. A municipality may rely on a shared identifier transferred from an external service; a university may need to merge duplicated staff records; and an enterprise analytics team may receive departmental extracts with incompatible definitions. A governance forum can settle definitions and ownership, but its decisions must be implemented through contracts, mappings, validation checks and escalation procedures. A policy that exists only as a file in a shared folder has little effect on the reliability of the next transaction.

Cross-functional accountability is essential. Business owners define meaning and accept risk, data stewards maintain standards, technology teams implement controls, information-security personnel define security requirements, privacy professionals review personal-data processing and internal assurance evaluates compliance. Responsibilities may vary by institution; what matters is that ambiguities are resolved before a defect affects an official decision or model output.

Table 3. Framework comparison and scope

Framework	Scope and limitation	Proposed addition
ISO/IEC 5259-1:2024	Data quality for analytics and ML: overview and terminology; not a local threshold schedule	Decision-specific quality gates
ISO/IEC 42001:2023	AI management system; not a task-level data test	Named data owner and approval evidence
ISO/IEC 23894:2023	AI risk management guidance; not a performance benchmark	Risk-tiered exception and rollback protocol

Applicable privacy duties must be separated from technical recommendations. Saudi personal-data processing and international transfers require case-specific assessment under the applicable PDPL and implementing framework. SDAIA's guidance on binding common rules identifies the separate transfer regulation and its conditions; it does not provide blanket permission for offshore processing. ISO standards are voluntary technical frameworks unless made binding by applicable obligations or contracts. The proposed quality gates are enterprise design recommendations, not legal requirements.

5 Dimensions of Data Quality and Fitness for Purpose

Accuracy, completeness, consistency, timeliness, validity and uniqueness are familiar quality dimensions, but their interpretation depends on purpose. A blank optional field may be acceptable in a draft case yet unacceptable in a regulatory approval. A customer address may be accurate when recorded and outdated when reused months later. Quality targets must therefore specify the decision, context, unit of analysis and acceptable consequences of error, not simply a universal threshold for every dataset.

Completeness should be assessed against required attributes, not just total non-null cells. Consistency tests compare values across authorised systems and within related records; validity checks assess allowable formats and business rules. Timeliness considers whether information arrives before a decision deadline. Uniqueness concerns unintended duplicates, while integrity addresses the preservation of valid relationships between entities. Representation and coverage

UNDER PEER REVIEW IN IJAR

also matter for AI: a dataset may be internally clean while systematically missing transactions from a particular branch, service channel or period.

Enterprise quality programmes should combine preventive and detective controls. Prevention includes controlled input forms, mandatory identifiers, referential constraints, standard vocabularies and training. Detection includes automated profiling, anomaly checks, reconciliation and exception sampling. Correction requires ownership and an audit trail. Root-cause analysis asks why a defect entered the system and whether a redesign can stop recurrence. Repeated manual cleansing is an expensive sign that a process defect remains unresolved.

Quality is not a property established once at initial migration. Changes to workflows, external APIs, definitions and user incentives can degrade formerly reliable information. A continuously monitored quality register should therefore describe tested rules, responsible owners, failures, remediation priority, affected consumers and closure evidence. This register connects operational hygiene to business impact.

Table 1. Enterprise data quality dimensions and management controls

Domain / responsibility	Focus	Control / evaluation
Accuracy	Validated source values	Field reconciliation and correction
Completeness	Required attributes present	Missing-field monitoring
Consistency	Compatible cross-system definitions	Reconciliation and mapping rules
Timeliness	Information current for decision	Freshness thresholds and alerts
Uniqueness	No unintended entity duplicates	Matching and merge review
Integrity	Valid relationships preserved	Referential checks and audit logs

Table 4. Proposed metric dictionary

Dimension	**Calculation**	**Owner / failure action**
Validity	Passing values / tested values	Steward; quarantine invalid inputs
Completeness	Present required fields / required fields	Business owner; request correction
Consistency	Matching reconciled pairs / comparable pairs	Integration owner; resolve mismatch
Freshness	Records within decision-specific age / tested records	Source owner; refresh or defer
Uniqueness	Distinct valid entity IDs / eligible records	Master-data steward; review duplicates
Integrity	Valid relationships / tested relationships	Database owner; block orphan records
Representativeness	Coverage by relevant cohort versus target population	Model owner; restrict unsupported cohorts

For each metric in Table 4, thresholds must reflect consequence of failure, lawful purpose, baseline capability and validation evidence. A mandatory eligibility identifier may require complete verification, whereas a nonessential reporting field may tolerate documented missingness. These are illustrative rules rather than universal numerical requirements.

UNDER PEER REVIEW IN IJAR

6 Metadata, Master Data and Lineage

Metadata explain a dataset's meaning, source, owner, update schedule, sensitivity and permitted usage. Without these descriptions, analysts may treat two visually similar columns as equivalent despite different calculation rules. A catalogue is useful only if its metadata reflect current systems and are maintained by accountable parties. Definitions should be written for users outside the originating department, and they should distinguish approved metrics from preliminary or exploratory measures.

Master data management addresses the problem of multiple representations of core entities such as organisations, suppliers, locations or citizens. Master data processes may include identifier reconciliation, survivorship rules and controlled changes. These functions demand caution because incorrect merges can create privacy risks and make multiple unrelated cases appear to belong to one person. Matching rules should be evaluated with sampled false positives and false negatives, with human review for consequential decisions.

Data lineage traces information from original source through transformations to a report, model feature or generated recommendation. It allows teams to investigate how a dashboard value was calculated and whether a model used a superseded field. Good lineage links technical processing steps to business definitions, preserving versioned transformations and evidence of authorised changes. In AI settings, lineage also includes training datasets, evaluation splits, model releases, feature definitions and inference logs within proportionate retention limits.

Organisations should avoid treating the existence of a graph of systems as proof of complete lineage. Some business transformations occur in spreadsheets, manual approvals or unlogged scripts. Discovery, interviews and routine reconciliation are needed to identify these missing steps. Lineage is strongest when a discrepancy can be followed back to an accountable source and corrected before downstream decisions are made.

7 Data Quality Risks in AI-Driven Systems

AI systems inherit errors from the data used for development and operation. Incorrect labels can teach a classifier the wrong target; duplicate records can leak between training and evaluation datasets; and outdated features can misrepresent current circumstances. Missingness may be systematically related to branches, service channels or administrative practices. A model can therefore achieve apparently good aggregate results while performing badly on a subgroup that is under-represented in historical records.

Different AI tasks expose different quality risks. Forecasting requires consistent time periods, valid denominators and avoidance of information not available at prediction time. Entity matching must balance missed links against false merges. Document intelligence faces scanning errors, language variation and ambiguous field boundaries. Generative systems can produce plausible narratives that misstate values or cite the wrong record. Reliable practice requires task-specific tests, traceable evidence and an explicit rule for when the system must defer to a reviewer.

AI assurance should begin with the data pipeline, not after a model has been trained. Dataset documentation should describe collection purpose, coverage, known limitations, sensitive features,

UNDER PEER REVIEW IN IJAR

label provenance and version history. Tests should separate development data from independent or later records. For risk scores, accuracy must be interpreted alongside calibration, false-positive and false-negative consequences, model drift and operational response capacity. A model that generates too many unreviewable alerts may reduce service performance despite respectable benchmark scores.

Data governance should also specify prohibited or restricted secondary use. A historical employee performance dataset gathered for operational management should not automatically be repurposed for sensitive personnel scoring. Purpose compatibility, fair treatment and access boundaries require documented review. AI changes the scale of possible inference, so traditional permissions at file level may be insufficient when combined outputs reveal information not obvious in the source records.

8 Architecture for Governed Enterprise AI

The proposed architecture has eight connected stages. The first defines the business purpose and responsible decision owner; the second inventories authoritative sources; the third validates ingestion and transformations; the fourth maintains metadata, lineage and quality rules; the fifth enforces authorised access; the sixth evaluates model or analytical outputs; the seventh requires proportionate human decision-making; and the eighth records outcomes for assurance and correction. Security, privacy, records requirements and internal audit operate across the entire design.

A practical implementation separates immutable or controlled source evidence from derived analytical assets. Raw records should not be silently overwritten by generated classifications. Transformations create versions with identifiers, run dates and exception logs. A quality service can quarantine inputs when critical rules fail and place lower-consequence anomalies into review queues. The AI layer should report confidence and limitations rather than hide uncertainty behind a single green status indicator.

Operational decision rights remain with designated officials. A dashboard can flag a suspicious variance, but the relevant manager should confirm whether the variation reflects missing data, a genuine business issue or a changed definition. An automated case classifier may recommend a routing category while preserving the ability to review the source evidence and correct an error. High-impact decisions should have clear escalation, documented overrides and arrangements for recovering from unavailable models or data services.

The architecture is modular, not a demand for one expensive platform. Smaller organisations may implement the same principles through well-designed database constraints, catalogues, controlled interfaces and sampled reviews. Larger institutions may require central metadata services and distributed stewardship. The defining feature is a continuous connection between data meaning, observed quality, permitted use and the real decision that follows.

UNDER PEER REVIEW IN IJAR

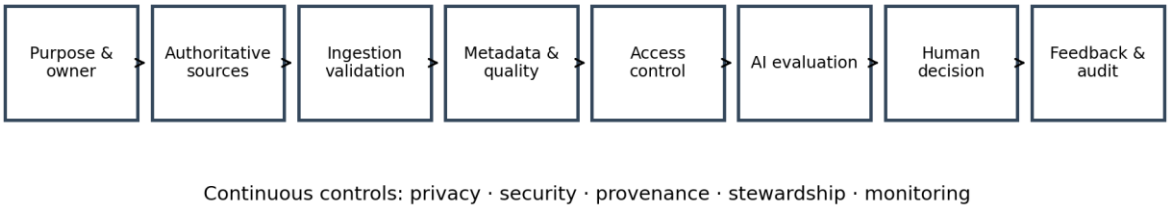

Figure 1. Eight-stage governed enterprise AI data-to-decision architecture

Data provenance records source, extraction time and version; business lineage maps transformations to the decision and accountable owner. The eight-stage proposal requires schema-contract validation before ingestion, separate training and test groups to prevent leakage, and comparison of live drift against an approved baseline. Record quality exceptions, model version, review outcome and overrides. A schema mismatch blocks affected decisions, triggers source-owner remediation and requires revalidation before restoring the approved version.

9 Control Ownership and Quality Scorecards

A mature scorecard measures both defects and their operational consequences. Technical indicators can include required-field completeness, rule failure rate, duplicate entity rate, freshness relative to service needs, reconciliation exceptions and unresolved lineage breaks. Governance indicators can include coverage of documented owners, approved definitions, reviewed access permissions and resolution of high-priority quality issues. Neither set is meaningful without clear denominators, thresholds and time periods.

Scorecard ownership should be distributed. Information stewards investigate recurring attribute defects; system teams resolve interface failures; risk and privacy functions review restricted use; and managers verify whether intervention improved service. A central data office can set methods and publish comparative findings, but local units retain responsibility for context and source accuracy. Escalation is particularly important where an issue crosses two organisations and neither believes it owns correction.

A balanced scorecard avoids incentives to conceal defects. A department that records more exceptions may be improving detection rather than deteriorating. Measures should assess correction quality, repeated causes and user confidence as well as raw incident counts. The aim is dependable decisions rather than an artificially perfect dashboard.

Table 2. Illustrative governance and quality assurance scorecard

Domain / responsibility	Focus	Control / evaluation
Data stewardship	Named owners for critical datasets	Coverage of approved ownership
Source quality	High-priority validation failures	Defect rate and closure age
Data lineage	Traceable transformations	Unresolved lineage exceptions

UNDER PEER REVIEW IN IJAI

Model assurance	Tested model behaviour	Error, drift and reviewer overrides
Access governance	Approved data use	Permission reviews and incidents
Service outcomes	Quality-related rework	Correction time and case delay

Hypothetical logistics scorecard: a carrier sends inconsistent city codes. The logistics data owner opens an exception, the source-system steward corrects the mapping, and integration staff test referential validity before reprocessing. Record affected shipment IDs, baseline and retest invalid-code rates, verification minutes and approval date. Do not populate outcomes until measured. If validation fails, retain the exception and suspend downstream allocation decisions. Recheck after each schema release.

10 Implementation Roadmap for Saudi Enterprises

Implementation should start with a bounded, consequential use case where there is a known data problem and a responsible decision owner. A public-service case-routing dataset, a multi-system university registry or a logistics performance dashboard may provide a practical pilot. The project team first documents the business question, source systems, current data flows, known defects and decision risks. It then establishes a baseline of quality, manual corrections, service delays and unresolved disputes.

The second phase defines ownership, classifications, a minimum dataset, validation rules and access controls. Existing manual workarounds are mapped because they often reveal defects invisible in formal architecture diagrams. The third phase implements automated tests and correction workflows before model deployment. A model should initially operate in a shadow period, comparing recommendations with independently reviewed cases without changing official decisions.

The fourth phase permits controlled deployment only after acceptance criteria are met. Accuracy, coverage, fairness, latency, review burden, security and auditability need separate sign-off. Users should be able to inspect the evidence behind a suggestion, contest outputs and record correction. Internal audits can sample whether controls actually operated, rather than merely verify that policies are documented. Periodic reviews should address model updates, new data feeds and business-rule changes.

Finally, expansion should prioritise reusable standards and trained people over rapid acquisition of additional tools. Common definitions, stewardship networks and controlled APIs make multi-department scaling less fragile. Contracts with technology providers should specify incident reporting, data extraction, model changes, access to audit evidence and exit arrangements. Long-term value depends on ongoing maintenance budgets and governance competence.

UNDER PEER REVIEW IN IJAP

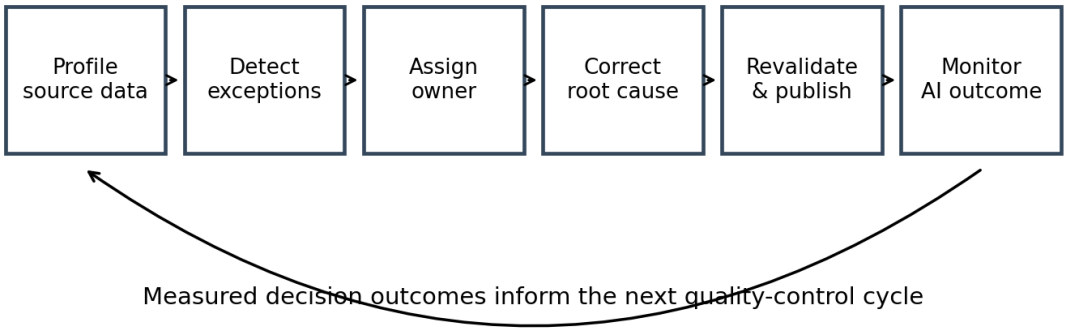

Figure 2. Continuous data-quality assurance and AI outcome feedback cycle

11 Discussion

The synthesis suggests that data governance and data quality function as complementary conditions for trustworthy enterprise AI. Improving a model without fixing broken definitions can automate confusion more efficiently. Conversely, a robust information-control environment enables organisations to choose simpler and more interpretable analytical methods when these are sufficient for the decision. The business objective should therefore determine technology selection rather than the novelty of the algorithm.

The first tension concerns standardisation and local context. Central governance can reduce duplicated rules and conflicting definitions, but a universal measure may obscure legitimate differences in departmental operations. A municipality's service deadline and a university's academic registration cycle require different timeliness thresholds. A federated stewardship model can retain common vocabulary and minimum controls while allowing justified local extensions. Such extensions should be documented to prevent fragmentation from reappearing under a new name.

The second tension concerns control intensity and operational capacity. Centralised approval can improve consistency but create bottlenecks; federated stewardship distributes responsibility but requires mature local teams and shared escalation rules. A hybrid design retains common definitions, minimum controls and central dispute resolution while delegating routine corrections. Its suitability depends on decision risk, staff capacity and cross-system dependencies; comparative evaluation is required.

A third tension concerns transparency and privacy. Extensive lineage and logging improve accountability, yet indiscriminate collection of personal information creates new exposure. Organisations need proportionate logs, restricted access, retention schedules and privacy-aware investigative procedures. Complete openness to every user is not the same as transparent accountability to authorised reviewers. Governance design should make evidence available to those who need it without creating uncontrolled secondary stores.

UNDER PEER REVIEW

For Saudi organisations, the distinctive opportunity is to connect national governance expectations with practical, measurable controls in service operations. Arabic-English datasets, variable legacy-system maturity, external contractors and multi-agency interfaces deserve explicit testing. Cross-institution learning may be valuable, but published success claims must distinguish demonstrations from measured improvements in complete service workflows. Implementation evidence should report what failed as well as what succeeded.

12 Research Gaps and Future Research

First, prospective studies should evaluate governance interventions in Saudi enterprise settings using independently measured before-and-after quality and service indicators. Many conceptual frameworks articulate responsibilities but provide limited evidence about how particular stewardship models change outcomes. Future investigations should compare centralised, federated and hybrid governance arrangements under comparable system complexity and organisational resources.

Second, Arabic-English data interoperability needs task-specific benchmarks. Research should evaluate inconsistent names, mixed-script identifiers, administrative terminology, scanned source documents and translated field definitions. Benchmarks should include errors from routine operational records, not only curated public datasets, with privacy-preserving arrangements for assessment where personal data are involved. Third, the combined consequences of quality failures and model errors require more transparent reporting. Studies should measure the full chain from source defect through model output to human decision and corrective action.

Fourth, causal research should test whether automated validation and stewardship reduce rework, delays or adverse decisions, rather than merely improve a technical quality score. Where randomisation is impractical, interrupted time-series or carefully matched comparisons may be feasible. Economic studies should count technology licensing, stewardship effort, training, human review and incident response alongside savings. Finally, research is needed on the durability of controls during organisational change, platform migration and AI model replacement. Longitudinal evaluation is essential because a high-performing initial system can degrade as definitions, users and source data change.

13 Practical, Managerial and Policy Implications

Executives should regard reliable data as an institutional capability that enables appropriate AI investment. They should require business cases to identify a decision owner, governing policies, expected quality improvements and the costs of exception handling. Funding should cover data stewards, source-system correction and independent assurance, not only cloud capacity and model development. Quarterly management review can examine recurring defects, unresolved cross-system responsibilities and evidence that interventions improved the decisions they were designed to support.

Project managers should include data definitions and governance controls in delivery acceptance criteria. A system should not be declared successfully implemented merely because users can open a dashboard or generate an AI summary. User acceptance must test source reconciliation, incorrect

UNDER PEER REVIEW IN IJAR

identifiers, missing updates, permissions and the ability to challenge outputs. Technology procurement should seek reproducible evidence, documented limits, secure interfaces and supported exit or migration arrangements.

Data and AI professionals should test fairness, leakage, drift and error severity in the context of each use case. Public policymakers and professional bodies can encourage interoperable vocabulary, practical maturity assessment, capability building and privacy-aware benchmarking. Compliance remains organisation-specific and should be assessed against applicable current requirements. The broad policy lesson is that trustworthy innovation depends on accountable, usable information rather than volume of data or number of deployed models.

14 Limitations

This paper is an integrative literature review and a conceptual framework, not an empirical evaluation of deployed Saudi enterprise systems. It does not present original interviews, survey responses, administrative datasets, controlled experiments or independently measured productivity gains. Consequently, it cannot claim that the proposed architecture will reduce costs or improve service outcomes by a stated percentage. The examples of municipal, university and logistics operations illustrate potential applications rather than observations from named organisations.

The reviewed evidence spans different sectors, methods and types of information assets, making statistical aggregation inappropriate. Publication bias may favour successful technology demonstrations, while institutional guidance defines intended governance practice rather than proving implementation effectiveness. Some foundational standards and studies predate the preferred recent literature window because they establish important terminology; rapid changes in AI technologies and Saudi requirements may alter the relevance of particular controls. There is also limited comparable public evidence for bilingual Arabic-English enterprise data and for long-term model governance in operational Saudi settings.

The framework assumes that organisations can identify responsible owners and establish basic access to source-system evidence. In practice, legacy applications, commercial contracts, resource constraints and organisational disagreements may limit those assumptions. Proposed metrics must therefore be adapted to task consequence, data availability and lawful processing conditions. Further validation requires representative datasets, user participation, independent review and longitudinal measurement. The paper should be read as a structured, testable decision-support proposal rather than a certified implementation standard or substitute for legal advice.

15 Conclusion

AI-driven enterprise systems can provide useful classification, forecasting and operational insights, but the credibility of those outputs rests on underlying data and institutional accountability. Saudi organisations seeking sustainable digital transformation should therefore integrate data governance, quality control, metadata, lineage, privacy, security, model assurance and human review into a single operating discipline. Technical tools are valuable when they implement named responsibilities and support decisions whose outcomes can be checked.

UNDER PEER REVIEW IN IJAR

This review proposed an eight-stage data-to-decision architecture, a closed-loop assurance cycle and practical indicators linking data defects to business consequences. It argues for risk-based verification, clearly defined source ownership, controlled transformations and independent evaluation of AI-assisted decisions. The proposed approach supports Saudi digital ambitions while acknowledging that no universal model or quality threshold suits every institution. Future work should test the framework prospectively in multilingual Saudi enterprise environments, report adverse as well as favourable findings, and determine which governance arrangements produce sustained improvements. Trustworthy AI is not merely a model property; it is an achieved and repeatedly demonstrated property of the entire information system.

UNDER PEER REVIEW IN IJAR

Declarations (for author confirmation)

Generative AI assisted editorial rewriting, conceptual control design and preparation of comparison tables in this revision. The author must independently verify the claims, references and figures, confirm the complete history of AI use, and approve the manuscript before submission. No original participant data or empirical intervention is reported. Affiliation, ORCID, corresponding email, funding, conflicts of interest, originality and author contributions remain to be confirmed by the author; no signatures or approvals are asserted.

References

[1] Abraham, R., Schneider, J., & vom Brocke, J. (2019). Data governance: A conceptual framework, structured review, and research agenda. International Journal of Information Management, 49, 424–438. https://doi.org/10.1016/j.ijinfomgt.2019.07.008

[2] Batini, C., Cappiello, C., Francalanci, C., & Maurino, A. (2009). Methodologies for data quality assessment and improvement. ACM Computing Surveys, 41(3), 16. https://doi.org/10.1145/1541880.1541883

[3] Wang, R. Y., & Strong, D. M. (1996). Beyond accuracy: What data quality means to information consumers. Journal of Management Information Systems, 12(4), 5–33.

[4] Khatri, V., & Brown, C. V. (2010). Designing data governance. Communications of the ACM, 53(1), 148–152. https://doi.org/10.1145/1629175.1629210

[5] International Organization for Standardization. (2015). ISO 9001:2015 Quality management systems—Requirements.

[6] International Organization for Standardization. (2022). ISO/IEC 27001:2022 Information security, cybersecurity and privacy protection—Information security management systems—Requirements.

[7] International Organization for Standardization. (2023). ISO/IEC 23894:2023 Information technology—Artificial intelligence—Guidance on risk management.

[8] International Organization for Standardization. (2023). ISO/IEC 5259-1:2023 Artificial intelligence—Data quality for analytics and machine learning—Part 1: Overview, terminology, and examples.

[9] International Organization for Standardization. (2023). ISO/IEC 42001:2023 Information technology—Artificial intelligence—Management system.

[10] National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. https://doi.org/10.6028/NIST.AI.100-1

UNDER PEER REVIEW IN IJAR

[11] National Institute of Standards and Technology. (2024). Artificial Intelligence Risk Management Framework: Generative AI Profile. NIST AI 600-1. https://doi.org/10.6028/NIST.AI.600-1

[12] Saudi Data and Artificial Intelligence Authority, National Data Management Office. National Data Management and Personal Data Protection Standards. https://sdaia.gov.sa/en/SDAIA/about/Documents/PoliciesEN.pdf

[13] Saudi Data and Artificial Intelligence Authority. National Data Management Office: Mission and functions. https://sdaia.gov.sa/en/Sectors/Ndmo/Pages/default.aspx

[14] Saudi Data and Artificial Intelligence Authority. SDAIA Laws and Regulations. https://sdaia.gov.sa/en/SDAIA/about/Pages/RegulationsAndPolicies.aspx

[15] Saudi Data and Artificial Intelligence Authority. Guide to the Saudi Personal Data Protection Law for Controllers and Processors. https://dgp.sdaia.gov.sa/

[16] Saudi Vision 2030. (2016). Vision 2030. https://www.vision2030.gov.sa/

[17] Mirzaie, M., Behkamal, B., Allahbakhsh, M., Paydar, S., & Bertino, E. (2023). State of the art on quality control for data streams: A systematic literature review. Computer Science Review, 48, 100554. https://doi.org/10.1016/j.cosrev.2023.100554

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

[19] Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. Proceedings of FAT* 2019, 220–229. https://doi.org/10.1145/3287560.3287596

[20] Sambasivan, N., Kapania, S., Highfill, H., Akrong, D., Paritosh, P., & Aroyo, L. M. (2021). "Everyone wants to do the model work, not the data work": Data cascades in high-stakes AI. Proceedings of CHI 2021. https://doi.org/10.1145/3411764.3445518

[21] Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J.-F., & Dennison, D. (2015). Hidden technical debt in machine learning systems. Advances in Neural Information Processing Systems, 28.

[22] Breck, E., Cai, S., Nielsen, E., Salib, M., & Sculley, D. (2017). The ML test score: A rubric for ML production readiness and technical debt reduction. IEEE BigData 2017.

[23] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?" Explaining the predictions of any classifier. Proceedings of KDD 2016, 1135–1144.

[24] Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., et al. (2021). On the opportunities and risks of foundation models. arXiv:2108.07258.

[25] European Commission High-Level Expert Group on Artificial Intelligence. (2019). Ethics guidelines for trustworthy AI. European Commission.

UNDER PEER REVIEW IN IJAR