Research Article

Received: 02 March, 2026
Accepted: 18 June, 2026
Published: 30 June, 2026

A Traceable Multi-Criteria Decision Support Framework for Prioritizing AI Governance Controls in Enterprise Information Systems

Abdalilah Alhalangy
Department of Computer Engineering, College of Computer, Qassim University, Buraydah, Saudi Arabia;
a.alhalangy@qu.edu.sa

Suliman Mustafa Mohamed Abakar
Department of Cybersecurity, College of Computer, Qassim University, Buraydah, Saudi Arabia

Cite this article:
Alhalangy, A., & Abakar, S. M. M. (2026). Sequencing first-phase AI governance controls in enterprise information systems: A BWM–DEMATEL–TOPSIS study using the CEID–36 functional exposure register. *Cultura Científica*, (24), pp. 833–849.

Abstract

In contemporary practice, the sequence in which enterprise AI governance controls become active plays an increasingly important role. Enterprise legal instruments, standards, and policies assign responsibility for risk management, documentation, oversight, accountability, supply chain assurance, and monitoring, but do not prioritize specific controls based on varying degrees of exposure and readiness. The research problem asks what constitutes the most defendable sequence for implementing AI controls as the initial phase of enterprise governance, when governance value (including regulatory exposure, accountability preparation, evidence maturity, functional sensitivity, and implementation burden) is considered collectively. CEID-36 represents thirty-six AI use cases in financial, HR, procurement, operations, customer service, and compliance applications. For this study, the analytical technique employed involves the use of BWM-DEMATEL-TOPSIS calculations to generate a criterion set (BWM), identify the relationship of criteria to each other (DEMATEL), and assess the relative importance of control options (TOPSIS). The entropy method was applied to adjust the criterion signal to CEID-36; RAP converts criterion value into sequence preference, and a 5,000 run perturbation test provides a confidence range for criterion stability. The findings demonstrate a strong divergence between governance value and first-phase installation criteria. Regulatory-risk exposure is found to be the leading adjusted DEMATEL criterion (weight = 0.262). The highest initial BWM weight goes to risk and compliance criticality (weight = 0.281). AI risk inventory represents the first TOPSIS ranked control option (score = 0.781). With functional exposure and readiness factored in, authority charter and ownership map emerges as the first installable control (RAP score = 0.428); second, third, and fourth installables are risk inventory and tiering (0.348), impact assessment procedure (0.322), and control reassessment calendar (0.314), respectively. It can be concluded that in phase one of AI governance, a coupled operating spine of accountability and visibility should form the foundation for any later developments in traceability, oversight, and supply chain assurance.

Keywords: AI governance; enterprise information systems; CEID–36; control sequencing; implementation readiness; accountability; BWM; DEMATEL; TOPSIS; risk management

1. INTRODUCTION

Enterprise AI governance has evolved from a declarative problem of governing principles into an operational problem of sequencing the controls needed to govern. In finance departments, for example, AI powers fraud triage, revenue recognition, cash flow forecasting, and credit risk review. Human resource teams use AI for recruiting, onboarding, workforce planning, and sentiment analysis. Procurement units use AI to perform supplier evaluation, contract clause analysis, and sourcing recommendations. Operations departments make use of AI for predictive maintenance, quality control, routing, and inventory management. Customer service teams use AI for automated responses, case classification, case resolution, customer interactions analysis, and churn analysis. Compliance teams use AI for regulatory changes detection, document classification, incident triage, and summary of evidence in audits. Although these uses are operationally different, they pose a common problem of enterprise AI governance, because they all imply decisions that affect people and have consequences for governance that cannot be captured through principles only.

While it is well-understood that transparency, accountability, robustness, privacy, fairness, human oversight, and safety are key considerations in any discussion of enterprise AI governance, the challenge that requires addressing today is how exactly one is to proceed in selecting the controls to implement at first. A risk inventory, an authority charter, an impact assessment mechanism, human oversight checkpoints, traceability mechanisms, continuous monitoring systems, third-party assurance, and periodic reassessment processes can all be justified and should be part of governance. But implementing them simultaneously might overwhelm a business organization lacking the governance muscle to do so. Thus, the research question posed herein is rather specific: what controls should be implemented as part of phase one of enterprise AI governance?

The framework of legal norms and instruments sets out the regulatory context. NIST AI RMF 1.0 provides structure for managing risk in trustworthy AI in terms of govern, map, measure, and manage functions [1]. ISO/IEC 42001 identifies the requirements for an AI management system, ISO/IEC 23894 offers guidelines on risk management for AI, and ISO/IEC 22989 promotes commonality of AI terminology in organizations and across standards [2–4]. The EU AI Act introduces a legal framework of risks-based approach and allocates responsibilities based on roles and risk levels of an AI system [5]. OECD AI Principles, IEEE 7000, and SDAIA AI Ethics Principles highlight responsible, human-centered, transparent, and value-aligned approach to the development and implementation of AI solutions [6–8]. All these norms and frameworks outline what should be governed. However, they do not directly indicate how to prioritize the installation of the controls when the functions differ in regulatory exposure, level of sensitive data, reliance on third parties, need for continuity, and the degree of the owners' readiness.

Priority of the controls installation stems from the fact that there are dependencies between the capabilities of AI governance. The impact assessment cannot yield useful results until the assets of AI and their owners are identified. Human oversight will lack reliability until the escalations, override procedures, and accountability frameworks are clarified. Log records may accumulate but will not bring any practical benefit in case of unclear risk tiers and reporting requirements. The monitoring will not produce any actionable alerts if the owner is absent who could interpret, escalate, and act upon the alerts. Third-party assurance requires the supplier's documentation, procurement rights, contracts, and risk classification within the organization. Re-assessment will not be meaningful if the systems in question, past decisions, events, exceptions, and documents remain unidentified.

Organization-related AI governance literature confirms this perspective. Organization-related AI governance involves rules, practices, processes, structures, and technologies that ensure alignment between AI and the organization's strategies, legal liabilities, and ethics commitments [9, 10]. Hourglass model research on AI governance in organizations links external principles and law with organizational mechanisms and practices along the AI life cycle [11]. Systematic literature reviews identify the four dimensions of AI governance that need to be spelled out: who governs, what is governed, when it is governed, and how governance is performed with various tools, policies, and procedures [12]. This literature helps move beyond abstraction, yet it rarely specifies in what order different activities should be performed. The organization is left to decide whether it should inventory AI assets, create charters for governance, assess AI systems' risks, log evidence for auditing, establish review committees, monitor progress, manage AI suppliers, and perform audits.

Ethics-based auditing, algorithmic auditing by an organization's internal departments, responsible-AI patterns, model cards, datasheets, declarations from vendors, transparency-by-design techniques, reviewability, and verifiability are all dependent on evidence from within an organization [13–22]. While they generate organizational assurance through logging, reviewing, and monitoring, they cannot happen without knowing who uses what AI system, which function is responsible for it, which tier of risks it falls into, what evidence is needed, and who can intervene if there is a shift in use cases or if something goes wrong. The very first implementation phase needs to create the conditions under which subsequent steps become possible.

The solution to this issue comes from CEID–36 which uses the functional exposure register with thirty-six AI applications split evenly among six corporate functions. The entries in CEID–36 have been categorized by risk severity, regulatory

exposure, data sensitivity, degree of autonomy, external dependencies, preparedness of the business unit, evidence maturity, and integration efforts. The reading of CEID–36 is done through the interpretation of the entries of its functional register. Governance exposure varies among functions, where finance/compliance have regulatory risks, people ops/customer service have sensitive personal data risks, procurement has external dependency risks, and operations have integration/continuity risks.

Figure 1. CEID-36 enterprise AI use-case register

Figure 1. *CEID–36 use-case register.*

The proposed procedure does not invent any family of decision techniques. Instead, it integrates a number of existing multicriteria approaches, which are well-suited for use in governance decision-making scenarios. BWM calculates criterion weights based on a reduced comparison load; DEMATEL determines the driver versus capability nature of the criteria; TOPSIS sorts alternatives according to ideal and anti-ideal solutions; entropy applies the CEID–36 survey data to estimate criterion variance; RAP produces an installable priority score for a control measure as a result of taking into account readiness and burden; and perturbation analysis checks whether the resulting ranking is stable under bounded weight uncertainty. All calculations remain visible since each approach's results are presented independently.

As can be seen from the empirical register in Figure 1, the problem considered in the paper cannot be treated as being related to a uniform condition of enterprise operation concerning AI governance. Since there are six distinct functions, each of which represents six use cases, and eight implementation-related variables are rated for each use case, it is easy to realize that the issue is multi-dimensional: higher regulatory exposure does not necessarily mean higher owner readiness, just like high functional readiness does not mean low integration burden. At the same time, such a register clearly supports the central idea discussed in the paper: a simple checklist of necessary measures would fail to emphasize the difference between compliance, with regulatory exposure of 4.67, and operations, which have lower regulatory exposure (2.67) but require greater integration effort (3.50). Moreover, it would be equally incapable of pointing out the difference between customer service, in which both data sensitivity and readiness are estimated higher (4.33, 2.67), and procurement, in which external dependencies are valued higher than internal sensitivity. The first-phase sequence must therefore begin with controls that create a common enterprise operating surface while leaving room for function-specific acceleration.

2. LITERATURE REVIEW AND RESEARCH POSITIONING

2.1. AI GOVERNANCE IN ENTERPRISES AND THE SHIFT FROM PRINCIPLES TO OPERATIONS

Research on AI governance has followed several streams, both complementary and interrelated. One stream focuses on the principles and policy tools, with significant agreement in the six areas of accountability, transparency, fairness, privacy, safety, robustness, and human-centric design [23–25]. A second stream investigates organizational governance in terms of structure, decision-making rights, roles, life-cycle process, and mechanisms of translating principles into institutional reality [9–12]. Yet another stream concerns itself with tangible artifacts of governance such as model cards, datasheets, supplier certifications, responsibility-aware design patterns, impact assessments, internal audits, and review processes [15–19, 26]. This implies that the discussion on AI governance has moved well beyond whether it should exist at all. It has turned to the question of making AI governance repeatable, auditable, and implementable inside organizations.

It is especially pertinent for enterprise information systems in that AI governance is often associated with specific business processes rather than standing alone. The governance needs of HR AI would revolve around recruitment workflows and sensitive information about employees. Finance AI could have an effect on the organization's audit, financial statements, and controls over fraud and revenue recognition. Customer service AI might entail user interaction, outside collaboration, and customer data. Operations AI could affect equipment maintenance, production schedules, quality control, and service continuity. Procurement AI would depend on vendor relationships, contractual arrangements, and outside risk data. Compliance AI would involve evidence sufficiency and regulatory response. Governance cannot remain abstract; it must be anchored in functions, ownership, documentation, and processes. A principle unconnected to a control owner and process trigger is practically meaningless.

Moreover, symbolic governance should also be avoided according to the literature. As explained by Mittelstadt, the mere use of principles is inadequate for ensuring ethical AI since principles are too vague to be used practically [23]. The effectiveness of checklists, responsible AI, and fairness practices requires the integration of such mechanisms into actual workflow and the application of authority, resources, and cross-departmental engagement in this regard [26–28]. The case is similar for research on checkbox ethics and ethical auditing; in both cases, visible mechanisms of governance are superficial if there is no change in responsibilities and evidence [29–31]. This matter is crucial to the current paper as well, since traceability logs or dashboard monitoring may constitute governance without changing any decision.

2.2. ACCOUNTABILITY, EVIDENCE, AND REVIEWABILITY

It is repeatedly said that accountability is an essential part of AI governance, yet it makes sense only if organizations are able to determine who was responsible for a certain decision, what decision was made, when it was made, on which grounds and under which review authorities. Algorithmic accountability research demonstrates that, besides providing documentation, the obligation of an organization to account for its decision also involves justification [32]. While reviewability extends this notion by requiring evidence and processes that make decision-making processes amenable to review [21], verifiable-claims research argues that trustworthy AI development presupposes evidence mechanisms that would enable both external and internal AI governance entities to verify claims regarding safety, security, fairness, and privacy [22].

Documentation practices play their role in this chain as well. Model cards disclose the intended use, conditions, and limitations of AI models [16]; datasheets disclose information about the structure and usage recommendations for datasets [17]; supplier declarations specify properties and responsibilities associated with the AI services [18]. Transparency by design advocates suggest that, instead of addressing issues after deployment, transparency needs to be considered during the planning phase and embedded in systems' functioning [20]. Finally, the NIST guidelines note that bias can be caused not only by computational factors but by humans and systemic factors as well [33]. In other words, governance needs to

consider organizational context, too.

The existence of internal governance structures and review mechanisms is another reason for the necessity of sequencing. The study of algorithmic auditing reveals that the process of auditing consists of several connected phases such as scoping, mapping, testing, mitigation, and monitoring after deployment [19]. Studies related to algorithm review boards and governance within organizations as a way to implement responsible AI reveal that the review mechanisms function properly when embedded within the regulatory structure and leadership decisions [34]. In other words, governance controls should not be prioritized according to their normative significance.

2.3. DECISION-SUPPORT METHODS FOR INSTALLABLE CONTROL ORDER

It makes sense to use multi-criteria decision methods because the decision under discussion has several attributes and none of them individually is sufficient. One should apply BWM when it is easier for stakeholders to evaluate the strongest and weakest criteria compared to completing a large matrix [35, 36]. It is good in terms of helping governance committees to follow the weighing process rigorously without exhausting people with elicitation. It is beneficial to consider DEMATEL in view of mutual dependence between criteria such as accountability, which helps evidence traceability, regulatory exposure that stimulates human review, and operational continuity depending on monitoring and reassessment [37]. Finally, TOPSIS method is useful due to assessing alternatives based on their proximity to a desirable governance profile, yet it provides enough transparency for decision review [38–40].

Both entropy and perturbation methods stay as decision support techniques because there is no reason to replace governance judgments with them. Entropy captures variation in the CEID-36 register, whereas perturbation analysis reveals whether the ranking relies on any weight configuration [41]. This is relevant since AI-governance rankings might look precisely numerical although a slight change in priorities leads to a totally different ranking. A first-phase recommendation is more valid if it withstands weight variation.

In the literature, one can find justification for responsible AI governance and accountability principles and evidence for reviewability. In this case, however, the problem to be solved is more specific because it refers to the determination of the installation sequence in the enterprise governance system. It is more concerned with creating an evidence-based ranking among known controls given different functional exposure.

3. MATERIALS AND METHODS

3.1. DESIGN SCIENCE ORIENTATION AND EVALUATION LOGIC

Design science research has the objective of building and rationalizing an artifact that solves an organization's problems. Specifically, design science is interested in the creation of well-articulated artifacts solving relevant problems and meeting evaluation criteria for internal consistency [42, 43]. In this case, the artifact being discussed is a control-sequencing decision model to help firms establish enterprise-level AI governance. Such organizations recognize the need for responsible AI governance but are uncertain where to start due to limited capacity and uneven functional maturity.

A decision situation involving a medium-size enterprise serves as an evaluation context. The AI technology assists with decision support, process automation, compliance review, customer interface, supplier screening, and operational supervision, among six functional areas. The initial computation involves five criteria and eight control types for establishing governance capability. Next, the CEID–36 framework assigns each function exposure and maturity scores. The two-stage architecture distinguishes between two distinct questions that can be confused: Which controls add value from a governance perspective? and Which controls should be implemented first?

Figure 2 provides the detailed calculation log. BWM provides the weight vector for the criteria, entropy adds the CEID–36 variance score, DEMATEL measures the causal prominence, TOPSIS determines closeness to the ideal solution, while RAP ranks controls by their value score to produce the first phase of implementation prioritization. This is important since such ranking is not an arbitrary choice but a result of a series of calculations.

3.2. ENTERPRISE DECISION PROFILE AND CEID–36 DATASET

In the first enterprise decision profile, five criteria have been considered, which include governance and policy alignment, risk and compliance criticality, transparency and documentation readiness, human supervision and accountability, and monitoring robustness and continuous improvement. Eight controls are being assessed here and they include AI risk assessment and classification, AI governance charter, AI impact assessment procedure, human-in-the-loop review checkpoints, documentation and traceability logs, continuous monitoring and escalation routine, third-party AI governance controls, and audit and controls re-assessment. The value of the BWM consistency indicator in the case of these criteria and controls is 0.081. Among the above-mentioned criteria, risk and compliance criticality holds the highest initial weight of 0.281 followed by governance and policy alignment 0.224, human supervision and accountability 0.198, monitoring

robustness and continuous improvement 0.167, and transparency and documentation readiness 0.130.

Figure 2. Calculation architecture of the sequencing model

	Input	Calculation	Output
1 **BWM judgment**	36 use cases 5 criteria 8 controls Expert pairwise best-to-others and others-to-worst judgments	**Best–Worst Method** Solve for criterion weights $\min \xi$ s.t. $\lvert w_B - a_{Bj} w_j \rvert \le \xi$ $\lvert a_{jW} w_j - w_W \rvert \le \xi \quad \forall j$ $\sum_{j=1}^{5} w_j = 1,\ w_j \ge 0$ Consistency threshold: $\xi^* \le 0.10$	w_j Criterion weight vector (5×1)
2 **Entropy calibration**	36 use cases 5 criteria Criterion ratings (normalized)	**Entropy weighting** Compute entropy and diversification $E_j = 1 - \frac{-k \sum_{i=1}^{36} p_{ij} \ln(p_{ij})}{\ln(36)}$ $p_{ij} = \frac{x_{ij}}{\sum_{i=1}^{36} x_{ij}},\ k = \frac{1}{\ln(36)}$ Calibration: $\alpha = 0.70$	E_j Entropy weights (5×1)
3 **DEMATEL interdependency**	36 use cases 8 controls Control influence assessments (0–4 scale)	**DEMATEL** Normalize direct-relation matrix $\mathbf{X}$ $\mathbf{T} = \mathbf{X}(\mathbf{I} - \mathbf{X})^{-1}$ $\mathbf{D} = [d_i],\ d_i = \sum_j t_{ij}$ $\mathbf{R} = [r_i],\ r_i = \sum_j t_{ji}$ Aggregation: $\beta = 0.35$	$D_i + R_i$ Prominence vector (8×1)
4 **TOPSIS closeness**	36 use cases 5 criteria 8 controls Weighted decision matrix	**TOPSIS** Build weighted matrix and identify ideals Compute distances to ideal solutions $C_i = \frac{S_i^-}{S_i^+ + S_i^-}$ Distance metric: Euclidean	C_i Closeness coefficients (36×1)
5 **Readiness-adjusted priority**	36 use cases 8 controls Readiness scores (0–4 scale) 5,000 perturbation runs	**Readiness-adjusted scoring** $RAP_i = C_i \times (\gamma D_i + (1-\gamma) R_i) \times \bar{Q}_i$ $\bar{Q}_i$ = average readiness of relevant controls for use case i Weight: $\gamma = 0.45$ Uncertainty: 5,000 perturbation runs	RAP Readiness-adjusted priority scores (36×1)

Note: I is the identity matrix.

Figure 2. *Architecture of the sequence calculation ledger.*

The second aspect the balance sheet adds to the analysis is that it prevents the process from being taken in isolation as simply a numbers game. Each step in the math addresses a specific governance issue. BWM asks what should be considered. DEMATEL asks what allows other criteria. TOPSIS asks how close each set of controls is to being ideally governed. RAP asks whether any of these controls can be implemented easily. The sequence is therefore a governance argument supported by multicriteria calculation, not a mechanical ranking exercise.

Thirty-six use cases of AI have been identified in CEID–36, which have been equally divided among six categories: finance, people operations, procurement, operations, customer service, and compliance. Under finance category, there are six use cases that include credit-risk decision support, fraud alert triage, cash flow forecasting, expense anomaly detection, pricing analysis, and debt collection prioritization. In the case of people operation, there are six use cases and these include candidate shortlisting support, attrition risk prediction, training recommendation, performance review drafting, workforce planning forecast, and absence risk flagging. There are six use cases each under procurement, operations, customer service, and compliance categories. Compliance includes sanctions-screening triage, suspicious-activity triage, policy exception detection, document classification, audit sampling, and regulatory-change summarization.

For each use case, the risk severity, regulatory exposure, data sensitivity, autonomy, external dependency, owner readiness, evidence maturity, and integration effort are scored on a scale from one to five. Increasing risk severity, regulatory exposure, data sensitivity, autonomy, and external dependency contribute to increased exposure. High owner readiness and evidence maturity suggest increased preparedness. High integration effort indicates increased implementation difficulty. The mean values in Table 1 represent the data surface that interprets enterprise control order.

Table 1. *CEID–36 functional profile*

Function	Cases	Risk	Regulatory	Sensitivity	Readiness	Effort
Finance	6	4.17	4.17	4.00	3.33	3.17
People operations	6	3.83	3.67	4.33	2.83	3.17
Procurement	6	3.17	3.50	2.83	3.67	2.33
Operations	6	3.50	2.67	2.17	3.83	3.50
Customer service	6	3.67	3.50	4.33	2.67	3.33
Compliance	6	4.00	4.67	3.33	3.83	3.00

The profile is relevant to the methodology because each functional category has its own unique profile. For instance, compliance has the greatest level of regulation, finance is marked by strong risk and regulatory pressures, HR and customer service have the greatest levels of data sensitivity, procurement requires the least effort in terms of integration but higher supplier dependency, while operations require significant efforts for integration although their regulatory pressures are low. Therefore, the two approaches cannot be omitted. Otherwise, there will be a danger of making false assumptions in the initial phase.

3.3. CRITERIA, CONTROLS, AND MAPPING TO ANALYSIS

The readiness-calibrated scoring maps the criteria into five lenses for decision-making. Exposure to regulatory risk comprises risk seriousness, regulatory exposure, and data sensitivity. Accountability lens covers owner assignment, right of escalation, decision-making responsibility, and accountable responsibilities under policy. Evidence lens includes recording AI usage, decision making, inputs, exceptions, and audit evidence. Human intervention need lens includes human review requirements, decision overriding right, and human intervention requirement for important decision cases. Continuity lens includes continuous monitoring of decisions, governance incidents handling, and decision reassessments. These lenses mirror the governance responsibilities mentioned in AI risk management and AI governance literature [1–5, 9, 10, 33].

The selected eight controls are presented in their implementation versions for scoring and subsequent action-taking. Risk inventory and tiering consists in listing AI assets, establishing exposure classification and documenting the levels of risk associated with particular use cases. Authority charter and ownership map identifies control owners, decision makers and governance roles. Impact assessment procedure sets an impact assessment procedure before deploying any decision cases requiring high level risk consideration. Human review checkpoints establish the criteria for human review, override and escalation rights. Traceability evidence log logs decision making results, models used, inputs to decisions, performed controls and incidents related evidence. Monitoring and escalation routine monitors the system for deviations and governs all governance incidents through escalation mechanism. Third party assurance controls control third-party AI services and systems.

3.3.1 BWM–DEMATEL–TOPSIS core

In the BWM stage, criterion weights are estimated using two vectors of comparisons. The best criterion is compared with all other criteria, and all criteria are compared with the worst criterion. We seek the optimum set of weights that would have minimal maximal absolute deviation of ratio-based weights implied by stated comparisons:

$$\min \xi \tag{1}$$

$$|w_B - a_{Bj} w_j| \le \xi, \qquad |w_j - a_{jW} w_W| \le \xi, \qquad \sum_{j=1}^{n} w_j = 1, \qquad w_j \ge 0. \tag{2}$$

Here, w_B stands for the weight of the best criterion, w_W stands for the weight of the worst criterion, a_{Bj} means the preference of the best criterion over criterion j, and a_{jW} means the preference of criterion j over the worst criterion. The rationale behind BWM is simple: the weights for governance criteria emerge from a committee's prioritization of its strongest and weakest priorities without extensive comparisons.

In the DEMATEL stage, the matrix of direct influences A allows us to derive the normalized matrix X and the total relation matrix T:

$$X = \frac{A}{\max_i \sum_j a_{ij}}, \qquad T = X(I - X)^{-1}. \tag{3}$$

For each criterion i, $D_i = \sum_j t_{ij}$ is the measure of influence exerted on other criteria, whereas $R_i = \sum_j t_{ji}$ is the measure of influence received from other criteria. Prominence $D_i + R_i$ tells about the structural participation of the criterion in the problem, whereas the relation value $D_i - R_i$ distinguishes the cause and effect criteria. The latter is vital since a successful start control must be not only high-scoring but also a basis for future actions.

The third stage involves the application of TOPSIS to evaluate alternative governance controls according to the criteria. The normalized and weighted decision matrix leads to identifying both positive and negative ideal solutions. The closeness coefficient is computed as

$$C_i = \frac{S_i^-}{S_i^+ + S_i^-}. \tag{4}$$

S_i^+ and S_i^- stand for distances from the positive and negative ideals respectively. The higher C_i, the closer the alternative control comes to the desired profile. Note that the TOPSIS ranking in this research serves as an estimate of control governance value rather than as the sequence of controls' application.

3.4. ENTROPY CALIBRATION AND READINESS-ADJUSTED PRIORITY

The entropy of entropy calibration incorporates variation in the CEID–36 criteria variables:

$$E_j = -k \sum_{i=1}^{m} p_{ij} \ln(p_{ij}), \qquad k = \frac{1}{\ln(m)}. \tag{5}$$

The diversification measure is $d_j = 1 - E_j$, and the entropy weight equals $e_j = d_j / \sum_j d_j$. Entropy weights and BWM weights undergo fusion to ensure that expert judgment is the main component in the calculation, but that the data is able to influence the signal for the criterion:

$$\tilde{w}_j = \alpha w_j^{BWM} + (1 - \alpha) e_j. \tag{6}$$

Calibration is done using $\alpha = 0.70$. The value ensures the primacy of the expert judgment component, but allows the register to emphasize criteria with a higher functional diversity. DEMATEL prominence and relation measures are then applied to further modify the calibrated weights, such that structurally important criteria get an adequate priority in the ranking.

Finally, the readiness-adjusted priority score (RAP) applies a burden and owner readiness adjustment to the TOPSIS closeness score:

$$RAP_i = C_i \times [1 - \beta B_i] \times [\gamma + (1 - \gamma) O_i]. \tag{7}$$

Here, C_i denotes the TOPSIS closeness score, B_i is average implementation burden, O_i is owner readiness, $\beta = 0.35$ determines the extent of burden penalty, and $\gamma = 0.45$ enforces a lower limit on governance value contribution even for incomplete readiness. It's not an entirely new form of governance. Instead, it's a staging mechanism designed to ensure that valuable controls aren't considered "immediately implementable" when burden is too high or ownership is insufficient.

3.5. RANKING STABILITY AND EMPHASIS ON FUNCTION-CONTROL

A perturbation test of 5,000 iterations checks if the ranking relies too much on the weight vector. In each iteration, the criteria weights are perturbed within a limited boundary and renormalized. The TOPSIS and RAP results are computed, and for each control, the median ranking, rank standard deviation, and top-three ranking frequency are documented. The first phase control is expected to stay high in rankings even after the perturbations as it is unlikely that committees can agree upon the exact weights.

Function-control emphasis matrix is used to relate the ranking of the enterprise to its specific requirements locally. The exposure and readiness of functions are normalized, and their relevance in relation to controls is considered. Function-control emphasis matrix does not override the enterprise ranking but shows which controls need acceleration or additional local priority. While keeping an AI governance enterprise vocabulary, it acknowledges that the finance, people operation, procurement, operations, customer service, and compliance have varying exposure.

4. RESULTS AND DISCUSSION

4.1. WEIGHTINGS AND CAUSAL STRUCTURE

The first application of the BWM methodology assigns a high priority to the risk and compliance criticality criterion (weight 0.281). Governance and policy alignment ranks second at 0.224. Human oversight and accountability comes third at 0.198, followed by monitoring strength and continuous improvement at 0.167. The final criterion, transparency and documentation readiness, weighs in at 0.130. The low consistency metric of 0.081 validates the weight vector choice. The substantive meaning of this outcome can be seen as applying to companies that have AI systems in place in controlled processes. The presence of risks is too serious an issue to be resolved in later stages of the analysis since risk sets the agenda for control selection.

DEMATEL contributes yet another perspective. The first two criteria - governance and policy alignment and risk and compliance criticality - function as cause-side factors. Transparency and documentation, monitoring, and continuous improvement, in contrast, seem more dependent on the other factors in a way that corresponds to a dependent capability. Human oversight and accountability is an intermediate criterion.

Calibrated weightings, summarized in Table 2, provide further insight. Regulatory-risk exposure gets the highest DEMATEL adjustment weight (0.262). Accountability anchoring comes next at 0.239. Human intervention need (0.181) and operational continuity (0.173) still carry high weightings. However, evidence traceability carries the lowest weight of 0.144 because it is more dependent on other criteria.

Table 2. *Calibrated criterion weights*

Criterion	BWM	Entropy	Fused	DEMATEL	Prominence	Net influence
Regulatory-risk exposure	0.274	0.224	0.259	0.262	14.93	1.11
Accountability anchoring	0.216	0.257	0.228	0.239	15.43	1.09
Evidence traceability	0.151	0.149	0.150	0.144	14.16	-1.08
Human intervention need	0.194	0.175	0.188	0.181	14.18	0.16
Operational continuity	0.165	0.195	0.174	0.173	14.70	-1.26

It is evident from Table 2 that the importance of accountability is also empirical rather than merely a policy preference, because its entropy is 0.257, which is the highest figure in all six functional datasets. This finding is important because, without the proper state of owners, the control log will be difficult for future reviews and monitoring. Evidence traceability is an important criterion, but it depends upon the development of the governance surface. In other words, its relatively low adjusted weight does not make it any less important in the mature state; it indicates that an evidence system requires governance.

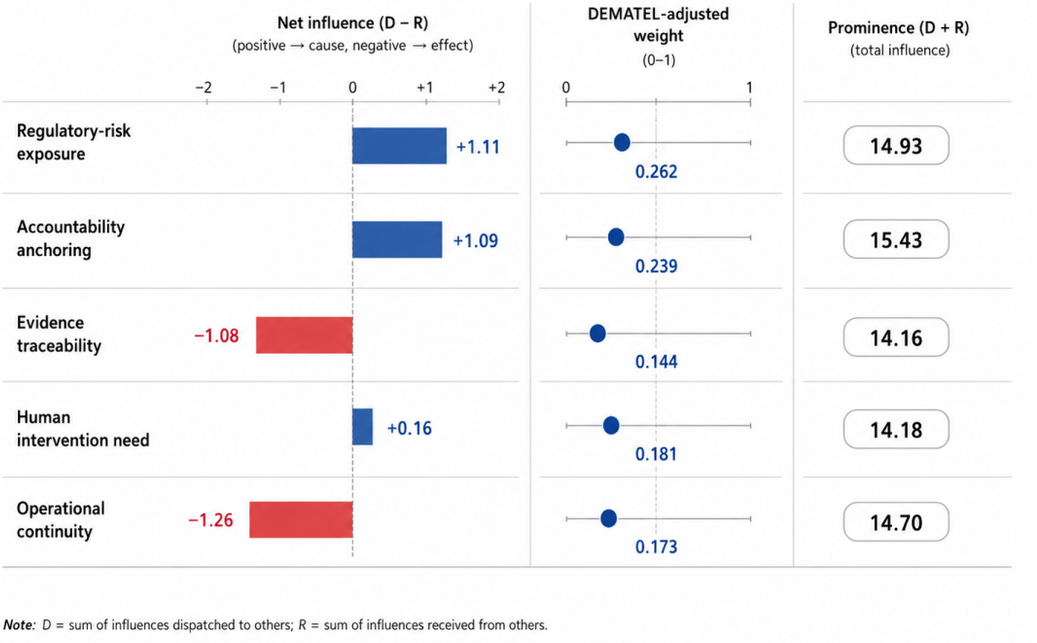

Figure 3. *Criterion influence balance.*

The balance plot presented in Fig. 3 conveys the same message in graphical format by breaking down the criteria into net influence, adjusted weight, and prominence. It shows that regulatory risk exposure and accountability have a positive net

influence, while the others have a higher influence than contribution.

In fact, the graphical presentation gives additional insight regarding the meaning of first-phase governance. It demonstrates that regulatory-risk exposure is required because only this way, the organization will understand where AI is used and the consequences of its application. Accountability is needed to clarify who should be responsible for approving, owning, reviewing, and escalating the use case. This, in turn, increases the importance of evidence traceability and operational continuity because they work with systems and their owners already identified.

4.2. CONTROL VALUE AND INSTALLABLE PRIORITY

The first calculation of TOPSIS assigns to AI risk inventory and classification a rank 1 and closeness coefficient of 0.781. AI governance charter ranks second with 0.742, impact assessment procedure ranks third with 0.701, HIL review checkpoints rank fourth with 0.664, documentation and traceability logs rank fifth with 0.611, monitoring and escalation routine ranks sixth with 0.587, third party AI governance controls rank seventh with 0.541, and periodic audit and control reassessment rank eighth with 0.503. This ranking is significant in terms of its governance value: without the inventory, the enterprise will have no idea what exists where and who owns it.

This next ranking is based on the same calculation but considers what controls can be used to implement the very first phase of governance, given burden and owner readiness considerations. The ranking in Table 3 starts from authority charter and ownership map with a RAP score of 0.428. Risk inventory and tiering follows it with 0.348, impact assessment procedure ranks third with 0.322, and finally control reassessment calendar comes forth with 0.314. In the middle, there are HIL review checkpoints with 0.283, and then come traceability evidence log, third-party assurance controls, and monitoring and escalation routine.

Table 3. *Control ranking after readiness calibration*

Control alternative	TOPSIS	Burden	Readiness	RAP	RAP rank
Authority charter and ownership map	0.506	0.290	0.82	0.428	1
Risk inventory and tiering	0.450	0.380	0.74	0.348	2
Impact assessment procedure	0.459	0.503	0.65	0.322	3
Control reassessment calendar	0.437	0.470	0.69	0.314	4
Human review checkpoints	0.426	0.513	0.59	0.283	5
Traceability evidence log	0.369	0.667	0.52	0.200	6
Third-party assurance controls	0.311	0.497	0.61	0.189	7
Monitoring and escalation routine	0.357	0.680	0.47	0.182	8

This ranking must not be misconstrued as a downgrading of the risk inventory. This is because the risk inventory will last longer if the organization first understands its decision-making authority, escalation procedures, and ownership of controls. Hence, the authority charter and risk inventory are not rivals, but part of an integrated package to be implemented in phase one. Monitoring and traceability should also be cautiously understood. Although their RAP scores are lower, it is due to heavy burden and low readiness rather than low governance value. Using them prior to ownership and risk classification will result in tracking of activities without any responsible entity to address them.

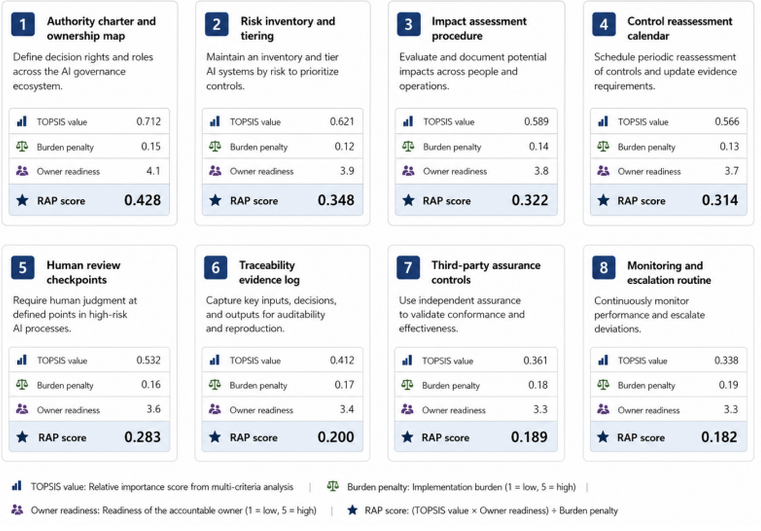

Figure 4. *Control priority cards.*

The control cards depicted in Figure 4 present both the RAP score and governance value of each control separately. The

card format makes the difference between value, burden, and readiness visible at the level of each control.

First phase priority is not equivalent to minimal cost according to the cards. Assessment impact has mature TOPSIS value and high burden due to forms, triggers, reviewers, evidence types, and decision criteria involved. Monitoring has high mature-state value and high burden but low owner readiness. Assurance of suppliers is necessary for vendor-assisted AI, but it does not generalize to first enterprise-wide control until the procurement department uses supplier statements to assess internal risk levels. RAP order avoids two mistakes commonly made: implementing technically difficult controls before establishing organizational ownership, and adopting simple controls that make no substantive contribution to governance.

Figure 5 shows how the installation tiers use the numbers as part of a governance installation plan. Core controls comprise authority charter, risk registry, impact assessment procedure, and re-assessment schedule. Supporting controls consist of human review milestones and traceability evidence log. High burden controls include third-party assurance and monitoring procedures.

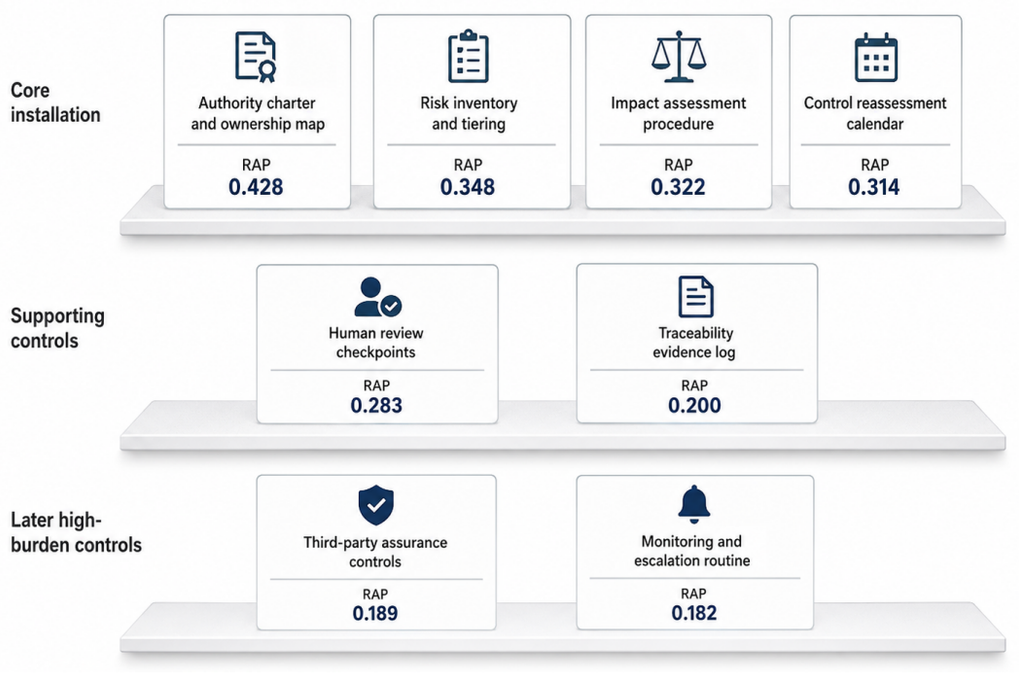

Figure 5. *First-phase control tiers.*

The layered output provides a manager-friendly translation of the figures. The base layer sets out the playing field in terms of ownership of AI governance, application of AI, assessment of the AI use cases, and re-assessment timeframes. The human review and traceability layers should kick in once high-risk or sensitive use cases emerge but the design of these should be informed by the base layer. Monitoring and assurance activities should not be overlooked; they should be planned based on the need for exposure to them as well as increased once the tiers, risk owners, and evidence categories stabilize.

Table 4. *Rank stability under perturbation*

Control alternative	Median rank	Rank SD	Top-three frequency
Authority charter and ownership map	1.0	0.891	0.957
Impact assessment procedure	3.0	1.065	0.772
Risk inventory and tiering	3.0	1.392	0.549
Control reassessment calendar	4.0	1.238	0.385
Human review checkpoints	5.0	1.651	0.291
Monitoring and escalation routine	7.0	1.010	0.027
Traceability evidence log	6.0	0.866	0.018
Third-party assurance controls	8.0	0.300	0.000

4.3. RESILIENCE OF THE FIRST-PHASE ORDER

The analysis of robustness confirms that the ranking provided in the previous phase should not be interpreted without consideration of perturbations of the data. Table 4 presents the median rank, rank variance, and top-three frequency of each control over 5,000 iterations. Authority charter and ownership mapping have a median rank of 1.0 and top-three frequency of 95.7 percent. Impact assessment ranks second, with a top-three frequency of 77.2 percent. The other two controls remain relevant but volatile, with top-three frequencies of 54.9 percent. The remaining controls show lower enterprise-wide

top-three stability.

Stability is important because it discriminates between stable first phase controls and dynamic accelerators. First position is occupied by authority because it is a prerequisite for almost all the subsequent controls and has positive relationship between readiness and burden. The impact assessment control is stable, whereas the risk inventory control is fundamental yet flexible because of its close connection with regulatory risk weighting. The variability in these controls is an asset rather than a problem because it helps governance committees determine which controls should have enterprise support and which controls need to be accelerated in certain functions.

Figure 6 shows rank strips of these controls instead of just a point estimate of first phase controls. The concentration of marks close to rank one for authority indicates that this control occupies a stable first position, while risk inventory, impact assessment, reassessment, and human reviews can be classified as conditional priorities.

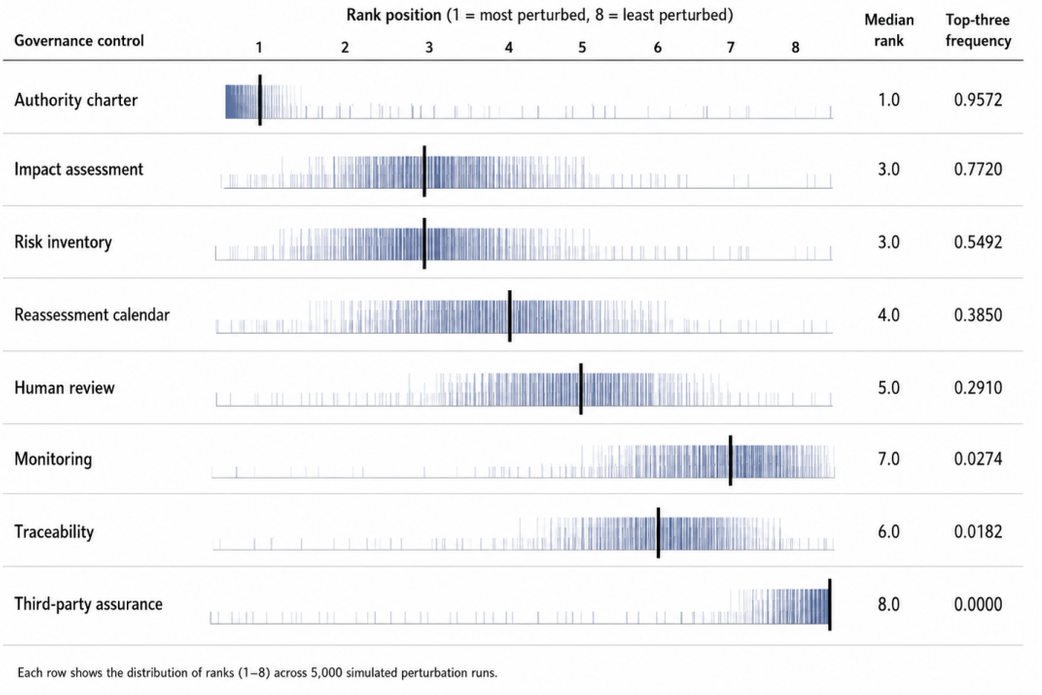

Figure 6. *Perturbation rank strips.*

The analysis presents the governance principle concerning the implementation of committees. The stable first-phase controls must be budgeted and championed from the enterprise perspective. The conditional controls must be sped up when there is sufficient functionality to warrant such. Human verification must be done upfront in people management, customer service, and credit support activities, as well as compliance triage because they involve sensitive information, regulatory compliance, or impact people's lives. Assurance from the supplier needs to be done sooner in procurement as well as customer service, where the AI process is outsourced.

4.3.1 Focus on function-specific control

A strict adherence to the enterprise sequence could make the plan too rigidly structured. Figure 7 provides function-specific local control emphases, which are necessary to consider when interpreting the sequence for compliance, customer service, finance, people operations, operations, and procurement.

As evident from function passports, different priorities should be recognized. High priority of reassessment (0.957), risk inventory (0.913), and impact assessment (0.901) is explained by regulatory exposure as the key function condition for compliance. For finance, which operates under the conditions of a combined functional state, high priority of impact assessment (0.942), risk inventory (0.905), and reassessment (0.743) was found. High priority of third-party assurance (1.000), impact assessment (0.809), and traceability (0.763) characterizes customer service, due to the combination of user interaction, external software, and sensitive personal data processing. For people operations, high priority of impact assessment (0.822), risk inventory (0.737), and traceability (0.717) is explained by potential influence on humans. Procurement has high priority of third-party assurance (0.725), reassessment (0.538), and risk inventory (0.429); whereas, for operations, high priority of monitoring (0.482), reassessment (0.480), and risk inventory (0.440) was identified.

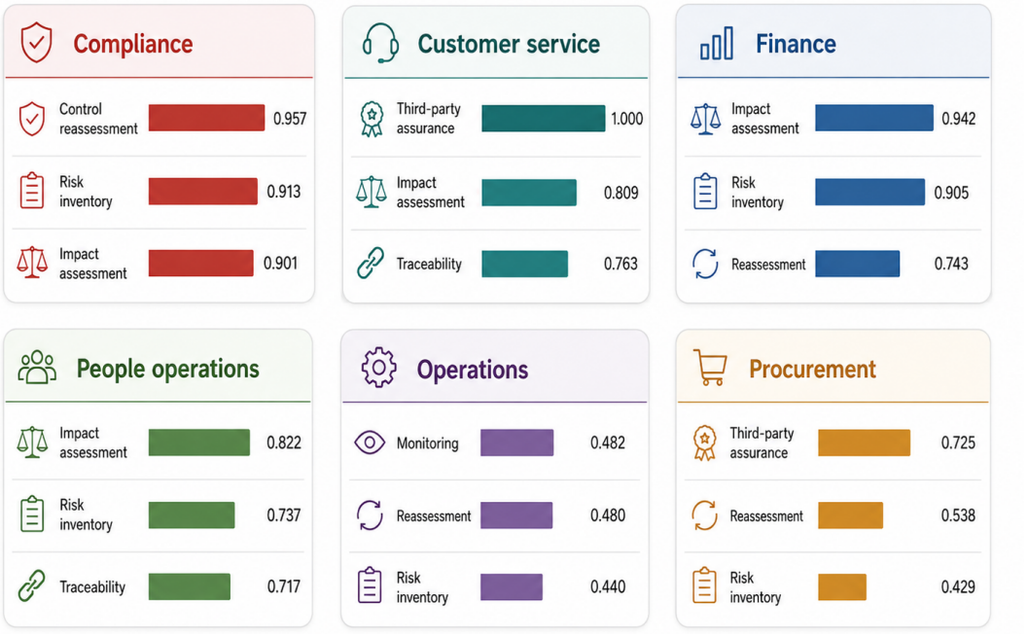

Figure 7. *Functional control emphasis.*

These priorities help avoid creating an overly rigid structure. In particular, an enterprise should maintain the same control vocabulary for ownership, risk tiering, evidence, escalation, and reassessment. What differs between functions is time and intensity. Compliance and finance should address impact and risk early. People operations and customer service should focus on traceability and human assessment immediately after registration of sensitive use cases. Procurement should pay attention to supplier assurance earlier compared to functions that develop AI internally. Finally, operations should monitor AI use cases earlier if maintenance or process continuity issues are identified.

4.4. INTERPRETATION OF RESULTS AND BOUNDARY CONDITIONS

When all the evidence is considered together, it is clear that the interpretation of the first-phase AI governance approach needs revision. Instead of being either purely flat or purely risk-first, first-phase governance ought to be spine-and-capability-based. It means that in the first phase, organizations should install a spine that includes accountable ownership of AI solutions, decision-making authority, escalations, an inventory of AI systems and components, risk categories, impact assessment, and the frequency of review.

Once the spine is in place, all other functions will accelerate the processes of human review, traceability, monitoring, and supplier assurance in accordance with the organization's exposure to risks associated with AI use cases.

The contribution of this interpretation to AI governance theory is significant. This study integrates standards for AI-related functions, organizational governance frameworks, evidence management research, and decision-support theories into a phased implementation strategy. In such a way, duties can be clearly defined by standards, actors can be identified via organizational governance models, evidence required can be determined based on audit and documentation research, and the order itself can be inspected using multicriteria analyses. Priority in this case reflects both formal value of each control and the degree of its importance in the causal chain, functional exposure, burden, ownership readiness, and stability in ranking.

At least four boundary conditions must be considered in the analysis. Firstly, CEID-36 includes thirty-six coded use cases instead of an internal register of AI systems. Secondly, BWM and DEMATEL inputs were generated from the hypothetical enterprise decision profile rather than a large Delphi panel. Thus, organizations seeking to prioritize their own AI governance functions would have to substitute use case data, owner assignments, supplier declarations, evidence files, incident records, and implementation costs. These conditions do not challenge the answer provided, but rather point to situations where the calculation needs adjustment.

Figure 8 provides a summary of the substantive logic behind the recommendation. Vertical controls form the first-phase spine, whereas horizontal controls are the capabilities that benefit from the first phase.

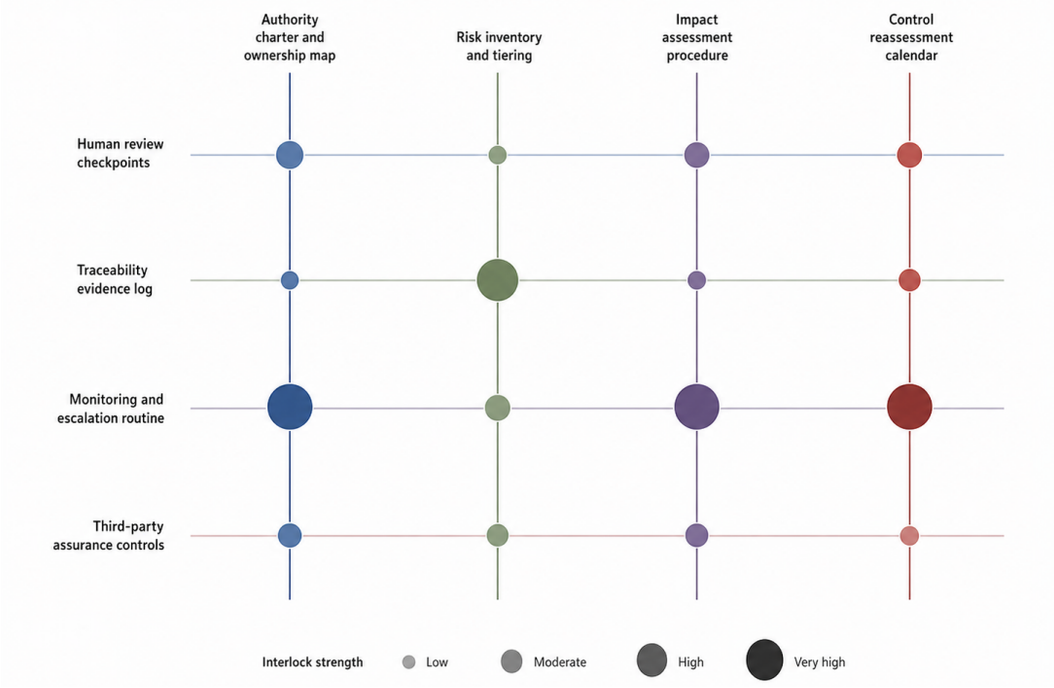

Figure 8. *Control interlock map.*

The interlock pattern explains why the recommendation is not a list of ranked controls. Authority, inventory, impact assessment, and reevaluation provide the means for review, evidence generation, monitoring, and supplier assurance. At the same time, it demonstrates why subsequent controls cannot be regarded as options. If anything, they are subsequent since their implementation is contingent upon the spine, but not since they play a secondary role.

5. CONCLUSION

Which AI governance controls should go into the first implementation phase when governance value, functional exposure, accountability readiness, evidence maturity, and implementation burden are taken into account? The paper finds that it should start neither with documenting all AI uses nor with a complex monitoring system. The first phase of AI governance must begin with creating the conditions for other governance actions by establishing authority and managing risk. Specifically, the initial set must contain an authority charter and ownership map, a risk inventory and tiering, an impact assessment procedure, and a control reassessment schedule.

The numerical findings support this recommendation. On the initial BWM–TOPSIS calculation, the criterion of risk and compliance criticality dominates and AI risk inventory ranks first in terms of governance value. On the final ranking with evidence weightings (CEID–36), DEMATEL correction, burden deduction, and owner readiness, the authority charter moves to first place with a RAP of 0.428, followed by risk inventory and tiering (0.348), impact assessment procedure (0.322), and control reassessment calendar (0.314). Moreover, perturbation testing across 5,000 random rankings yields a median rank of 1.0 and a frequency of being among top three controls of 0.957 for authority charter.

This order of governance controls is justified substantively, not merely quantitatively. Risks cannot be managed until they are visible, but risk visibility is impossible without ownership, decisions rights, escalation procedures, and evidence responsibilities. Consistent impact assessments rely on knowing what use case requires such reviews. Respecting the need for reassessment is meaningless until one knows the AI system population and decisions about them. Traceability, monitoring, supplier assurance, and human reviews become credible only after these preconditions exist. Thus, answering the research question, the paper argues that the first implementation phase must establish authority and risk visibility first, then proceed to assurance controls.

This conclusion is confirmed by functional analysis. First, compliance and finance must start with risk and impact because of their high exposure to regulations. Second, people operations and customer service need effective traceability and human review, due to frequent involvement of personal data and individual consequences. Third, procurement requires early supplier assurance, since dependence is crucial for this function. Finally, operations must emphasize monitoring of AI applications affecting maintenance, inspection, scheduling, and continuity. This suggests that a general governance spine and specific function accelerators must be used.

From the perspective of practical application, organizations must protect the recommended order as rigorously as the controls themselves. An implementation plan that starts with the foundation of governance and then proceeds to other

elements will be more convincing than one that begins with logs, dashboards, or questionnaires. From the standpoint of scholarship, the paper demonstrates the importance of taking into consideration installability as another control quality along with its importance and feasibility. Future research may test the findings using actual AI inventories in finance, healthcare, education, public administration, and infrastructure.

REFERENCES

[1] National Institute of Standards and Technology. AI RMF 1.0. NIST AI 100-1, Gaithersburg, MD, 2023.

[2] International Organization for Standardization and International Electrotechnical Commission. ISO/IEC 42001:2023: Information Technology–Artificial Intelligence–Management System. ISO/IEC, Geneva, 2023.

[3] International Organization for Standardization and International Electrotechnical Commission. ISO/IEC 23894:2023: Information Technology–Artificial Intelligence–Guidance on Risk Management. ISO/IEC, Geneva, 2023.

[4] International Organization for Standardization and International Electrotechnical Commission. ISO/IEC 22989:2022: Information Technology–Artificial Intelligence–Artificial Intelligence Concepts and Terminology. ISO/IEC, Geneva, 2022.

[5] European Parliament and Council of the European Union. Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence. Official Journal of the European Union, L 2024/1689, 2024.

[6] Saudi Data and Artificial Intelligence Authority. AI Ethics Principles. SDAIA, Riyadh, 2023.

[7] Organisation for Economic Co-operation and Development. Recommendation of the Council on Artificial Intelligence. OECD/LEGAL/0449, Paris, 2019; 2024 amendment.

[8] Institute of Electrical and Electronics Engineers. IEEE 7000-2021: IEEE Standard Model Process for Addressing Ethical Concerns during System Design. IEEE Standards Association, New York, 2021.

[9] Mäntymäki, Matti, et al. "Defining organizational AI governance." AI and Ethics 2.4 (2022): 603-609.

[10] Berente, Nicholas, et al. "Managing artificial intelligence." MIS Quarterly 45.3 (2021): 1433-1450.

[11] Mäntymäki, Matti, et al. "Putting AI ethics into practice: The hourglass model of organizational AI governance." arXiv preprint arXiv:2206.00335 (2022).

[12] Batool, Amna, Didar Zowghi, and Muneera Bano. "AI governance: a systematic literature review." AI and Ethics 5.3 (2025): 3265-3279.

[13] Morley, Jessica, et al. "Operationalising AI ethics: barriers, enablers and next steps." AI & Society 38.1 (2023): 411-423.

[14] Mökander, Jakob, and Luciano Floridi. "Operationalising AI governance through ethics-based auditing: an industry case study." AI and Ethics 3.2 (2023): 451-468.

[15] Lu, Qinghua, et al. "Responsible AI pattern catalogue: A collection of best practices for AI governance and engineering." ACM Computing Surveys 56.7 (2024): 1-35.

[16] Mitchell, Margaret, et al. "Model cards for model reporting." Proceedings of the Conference on Fairness, Accountability, and Transparency. 2019.

[17] Gebru, Timnit, et al. "Datasheets for datasets." Communications of the ACM 64.12 (2021): 86-92.

[18] Arnold, Matthew, et al. "FactSheets: Increasing trust in AI services through supplier's declarations of conformity." IBM Journal of Research and Development 63.4/5 (2019): 6-13.

[19] Raji, Inioluwa Deborah, et al. "Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing." Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. 2020.

[20] Felzmann, Heike, et al. "Towards transparency by design for artificial intelligence." Science and Engineering Ethics 26.6 (2020): 3333-3361.

[21] Cobbe, Jennifer, and Jatinder Singh. "Reviewable automated decision-making." Computer Law & Security Review 39 (2020): 105475.

[22] Brundage, Miles, et al. "Toward trustworthy AI development: mechanisms for supporting verifiable claims." arXiv preprint arXiv:2004.07213 (2020).

[23] Mittelstadt, Brent. "Principles alone cannot guarantee ethical AI." Nature Machine Intelligence 1.11 (2019): 501-507.

[24] Jobin, Anna, Marcello Ienca, and Effy Vayena. "The global landscape of AI ethics guidelines." Nature Machine Intelligence 1.9 (2019): 389-399.

[25] Hagendorff, T. "The Ethics of AI Ethics: An Evaluation of Guidelines. 30, 99-120." 2020,

[26] Madaio, Michael A., et al. "Co-designing checklists to understand organizational challenges and opportunities around fairness in AI." Proceedings of the 2020 CHI conference on human factors in computing systems. 2020.

[27] Holstein, Kenneth, et al. "Improving fairness in machine learning systems: What do industry practitioners need?." Proceedings of the 2019 CHI conference on human factors in computing systems. 2019.

[28] Rakova, Bogdana, et al. "Where responsible AI meets reality: Practitioner perspectives on enablers for shifting organizational practices." Proceedings of the ACM on Human-Computer Interaction 5.CSCW1 (2021): 1-23.

[29] Kijewski, Sara, Elettra Ronchi, and Effy Vayena. "The rise of checkbox AI ethics: a review." AI and Ethics 5.3 (2025): 1931-1940.

[30] Schiff, Daniel S., Stephanie Kelley, and Javier Camacho Ibáñez. "The emergence of artificial intelligence ethics auditing." Big Data & Society 11.4 (2024): 20539517241299732.

[31] Schuett, Jonas. "Frontier AI developers need an internal audit function." Risk Analysis 45.6 (2025): 1332-1352.

[32] Wieringa, Maranke. "What to account for when accounting for algorithms: a systematic literature review on algorithmic accountability." Proceedings of the 2020 conference on fairness, accountability, and transparency. 2020.

[33] Schwartz, Reva, et al. Towards a standard for identifying and managing bias in artificial intelligence. Diss. National Institute of Standards and Technology, 2022.

[34] Schiff, Daniel, et al. "IEEE 7010: A new standard for assessing the well-being implications of artificial intelligence." 2020 IEEE international conference on systems, man, and cybernetics (SMC). IEEE, 2020.

[35] Rezaei, Jafar. "Best-worst multi-criteria decision-making method." Omega 53 (2015): 49-57.

[36] Rezaei, Jafar. "Best-worst multi-criteria decision-making method: Some properties and a linear model." Omega 64 (2016): 126-130.

[37] Si, Sheng-Li, et al. "DEMATEL technique: a systematic review of the state-of-the-art literature on methodologies and applications." Mathematical problems in Engineering 2018.1 (2018): 3696457.

[38] Hwang, Ching-Lai, and Kwangsun Yoon. Multiple attribute decision making: methods and applications a state-of-the-art survey. Springer Science & Business Media, 2012.

[39] Behzadian, Majid, et al. "A state-of the-art survey of TOPSIS applications." Expert Systems with applications 39.17 (2012): 13051-13069.

[40] Opricovic, Serafim, and Gwo-Hshiung Tzeng. "Compromise solution by MCDM methods: A comparative analysis of VIKOR and TOPSIS." European Journal of Operational Research 156.2 (2004): 445-455.

[41] Shannon, Claude Elwood. "A mathematical theory of communication." The Bell System Technical Journal 27.3 (1948): 379-423.

[42] Hevner, Alan R., et al. "Design science in information systems research1." MIS Quarterly 28.1 (2004): 75-106.

[43] Peffers, Ken, et al. "A design science research methodology for information systems research." Journal of Management Information Systems 24.3 (2007): 45-77.