

Adapting IT Governance Mechanisms for AI Governance in SMEs

John R. Drake
drakejo@ecu.edu

Jakub Jurasek
jurasekj21@students.ecu.edu

East Carolina University
Greenville, NC, 27858, USA

Abstract

Most small and medium-sized enterprises (SMEs) do not need an entirely new governance system to manage generative artificial intelligence (GenAI). What they need is a principled extension of the governance mechanisms already in place for information technology, adapted to account for AI-specific risk. Prior literature establishes why AI must be governed and why abstract principles require translation into organizational routines, while COBIT 2019 anchors how governance is expressed through recurring components including roles, processes, policies, information flows, skills, culture, and services. The National Institute of Standards and Technology (NIST) Artificial Intelligence Risk Management Framework (AI RMF), the NIST GenAI Profile, and the AI RMF Playbook supply the risk-management structure for AI-specific concerns such as unreliable outputs, data handling, and integration exposure. Using Design Science Research Methodology (DSRM), this research develops a deployer-oriented governance design for SMEs organized around four AI operating contexts: standalone GenAI tools and assistants, embedded AI in business software, connected or agent-like AI systems, and model adaptation or development. The artifact centers on Table 1 as the primary design summary, with Table 2 providing a control catalog logic derived from that summary and supporting templates in a supplementary role. The contribution is a NIST-first governance design that demonstrates how existing IT governance mechanisms can be simplified appropriately for SMEs without sacrificing accountability or traceability.

Keywords: AI Governance, SMEs, IT Governance, COBIT, NIST AI RMF, design science research

Adapting IT Governance Mechanisms for AI Governance in SMEs

John R. Drake and Jakub Jurasek

1. INTRODUCTION

Generative AI has entered organizational work not through formal adoption programs but through the tools people already use. In many SMEs, this means drafting assistants, summarization features, meeting tools, and AI functions embedded in existing productivity and business software. The governance question this creates is therefore not whether AI matters but how existing governance mechanisms must change as these tools become part of routine operations (National Institute of Standards and Technology, 2024; Taeihagh, 2025).

For the purposes of this study, SMEs are defined as organizations with fewer than 250 employees, consistent with the European Commission's widely adopted benchmark for medium-sized enterprises, though this paper treats employee count primarily as a boundary condition for a resource-constrained organizational context (European Commission, n.d.). Within that context, SMEs are understood as organizations that lack dedicated AI governance capacity but nonetheless require governance arrangements concrete enough to guide daily use, support accountability, manage vendor relationships, and handle incidents.

Two bodies of prior work provide the necessary foundations. The first is organizational IT governance research, which establishes that AI governance involves rules, practices, processes, and accountability arrangements that translate abstract principles into operational routines (Mäntymäki et al., 2022; Papagiannidis et al., 2023, 2025; Mökander & Floridi, 2023). The second is governance framework logic. COBIT 2019 distinguishes governance from management and defines governance-system components—processes, organizational structures, policies and procedures, information flows, people and skills, culture and behavior, and services and infrastructure—that give governance expression across the enterprise (ISACA, 2018a, 2018b). NIST complements this by providing the AI-specific risk structure through the AI RMF, the GenAI Profile, and the AI RMF Playbook (National Institute of Standards and Technology, 2023, 2024, 2025).

Together, these foundations expose the gap this paper addresses. The literature explains why AI must be governed; NIST explains how AI risks are organized; but no account yet exists of how existing governance mechanisms should be simplified appropriately for SMEs operating primarily as deployers of GenAI. That gap is consequential precisely because SMEs generally do not build AI governance from scratch. They adapt what they already have (National Institute of Standards and Technology, 2023, 2024, 2025; Papagiannidis et al., 2025; Mökander & Floridi, 2023). To that end, this paper asks: How can existing IT governance mechanisms be adapted to govern generative AI use in SMEs?

To sharpen the problem further, this paper distinguishes four AI operating contexts relevant to SMEs. First, standalone GenAI tools and assistants cover direct use of external or desktop tools for drafting, summarization, analysis, translation, and question answering, often described as AI Chatbots. Second, embedded AI in business software covers AI features integrated into productivity, customer support, marketing, analytics, and enterprise applications. Third, connected or agent-like AI systems are internally built systems that connect to enterprise knowledge, call external tools, retrieve information, or trigger actions across systems with some degree of workflow autonomy. The fourth context, model adaptation or development, includes fine-tuning, customization beyond ordinary prompting, and in-house model development. The first three contexts are the primary focus here, as they align most directly with the deployer-oriented concerns the NIST GenAI Profile emphasizes; the fourth imposes a substantially heavier burden around data, expertise, infrastructure, and evaluation that places it outside the typical SME deployer case (National Institute of Standards and Technology, 2024; Ji et al., 2023; Liu et al., 2024; Mitchell et al., 2019; Gebru et al., 2021; Bender & Friedman, 2018).

The paper makes four contributions. First, it identifies a small set of IT governance mechanisms that carry the most weight for AI governance in SMEs. Second, it combines COBIT as the governance framework anchor with NIST as the AI risk-management structure, a pairing that, to the best of our knowledge, has not been developed in this way for the SME deployer context. Third, it makes the SME difference explicit by showing how large-enterprise governance expressions can be compressed without

forfeiting accountability. Fourth, it delivers the result as a design artifact centered on Table 1 and the control catalog logic in Table 2.

The design draws on three distinct bodies of knowledge, each informing a different layer of the artifact. IT governance and IT operations literature shapes the treatment of roles, decision rights, monitoring, and incident response (ISACA, 2018a, 2018b; Nelson et al., 2025). AI governance research informs the approach to human oversight, documentation, auditability, and GenAI-specific risk (Mäntymäki et al., 2022; Papagiannidis et al., 2025; Mökander & Floridi, 2023). NIST guidance provides the risk-first structure and GenAI risk vocabulary that run throughout (National Institute of Standards and Technology, 2023, 2024, 2025). Design science literature determines how the artifact is constructed, justified, and positioned as a research contribution rather than a practitioner toolkit (Peppers et al., 2007; Wieringa, 2009; Gregor & Hevner, 2013; Mandviwalla, 2015). This article is structured as follows. First, it reviews the IT governance framework and AI governance literatures that motivate the design. Second, it describes the DSRM methodology that guides artifact development. Third, it presents the proposed governance design, including Table 1 and the control catalog logic. Last, it discusses contributions, implications, and limitations before concluding.

2. LITERATURE REVIEW

2.1 Existing Governance Frameworks as the Starting Point

COBIT 2019 distinguishes governance from management and organizes enterprise governance of information and technology through recurring components. By considering processes, organizational structures, principles, policies and procedures, information flows, people and skills, culture and behavior, and services and infrastructure, COBIT gives governance expression across the organization (ISACA, 2018a, 2018b). It also groups governance and management objectives into recognizable domains: Evaluate, Direct and Monitor; Align, Plan and Organize; Build, Acquire and Implement; Deliver, Service and Support; and Monitor, Evaluate and Assess (ISACA, 2018a, 2018b).

This paper uses COBIT as the foundational logic for understanding what governance mechanisms are and how they function. In operational terms, this research restates COBIT's seven governance-system components as seven practical questions that any working governance arrangement must answer: who decides (organizational structures), what is documented (principles, policies, and procedures), how work is approved (processes), how information moves (information flows), how staff are prepared (people and skills together with culture and behavior), how systems are monitored (services and infrastructure, exercised through the Monitor, Evaluate and Assess domain), and how incidents are handled (processes within the Deliver, Service and Support domain) (ISACA, 2018a, 2018b). The restatement is deliberate: operators recognize questions more readily than component labels.

It is important to note that COBIT alone did not produce this list. Each question was cross-checked against the NIST AI RMF Govern function, whose categories address documented policies (Govern 1), accountability, roles, and workforce training (Govern 2), risk culture (Govern 4), and third-party oversight (Govern 6), while the Measure and Manage functions cover monitoring and incident response (National Institute of Standards and Technology, 2023). The questions therefore sit at the intersection of the two frameworks: COBIT supplies the organizational vocabulary, and NIST confirms that each question remains load-bearing when the system being governed is AI. Section 5 traces each Table 1 mechanism back to this intersection.

It is important to note that COBIT is explicit on the question of scale: governance systems should be tailored to enterprise context (ISACA, 2018a). Smaller organizations may require reduced roles, structures, and governance components relative to large enterprises. That point is central to this research. The objective is not to transpose large-enterprise governance architecture into SMEs but to take existing governance framework logic and reduce it to a form SMEs can realistically operate.

2.2 AI and GenAI Governance Literature

The organizational AI governance literature establishes why this adaptation is necessary. Mäntymäki et al. (2022) define organizational AI governance in terms of rules, practices, processes, and tools that must be aligned with organizational strategy, values, and legal requirements. Papagiannidis et al. (2023, 2025) extend this view by identifying governance practices, barriers, and outcomes, and by framing responsible AI governance as an organizational capability rather than a narrow compliance function. Mökander and Floridi (2023) further demonstrate why evidence-oriented auditing and operational

routines are indispensable when organizations attempt the transition from abstract ethics principles to actual governance practice.

Generative AI intensifies these concerns in identifiable ways. The NIST GenAI Profile catalogs cross-sector risks that are either unique to or substantially worsened by generative systems, including confabulation, data privacy exposure, information integrity threats, information security risks, intellectual property complications, and value-chain integration problems (National Institute of Standards and Technology, 2024). Ji et al. (2023) reinforce the governance significance of unreliable but confident outputs, while Liu et al. (2024) show why prompt injection and untrusted content manipulation become pressing concerns when AI is connected to external systems or enterprise data. In combination, these findings establish that GenAI introduces not merely incremental risk but qualitatively distinct governance obligations for deployers.

2.3 Research Gap

Taken together, the literature supplies three of the four pieces needed for the design developed here. COBIT explains how governance mechanisms are structured within organizations (ISACA, 2018a, 2018b). AI governance research explains why AI requires explicit accountability arrangements, operational routines, and documentary evidence (Mäntymäki et al., 2022; Papagiannidis et al., 2023, 2025; Mökander & Floridi, 2023). NIST explains how AI and GenAI risks should be organized and managed across the deployment lifecycle (National Institute of Standards and Technology, 2023, 2024, 2025). What is missing is an account of how these bodies of knowledge are combined and simplified for the SME deployer context.

While ISO/IEC 42001, ISO/IEC 23894, ISO/IEC 38507, and the EU AI Act provide important governance and compliance context (International Organization for Standardization & International Electrotechnical Commission, 2022, 2023a, 2023b; European Parliament and Council of the European Union, 2024), COBIT and NIST are more prevalent in the United States context within which this study takes place. Furthermore, this research uses NIST as the primary AI risk structure because NIST offers deployer-accessible risk functions and GenAI-specific profiles that are not only technically rigorous but practically oriented toward organizations that did not design the systems they are deploying. COBIT provides the organizational governance-mechanism vocabulary that gives those risk functions a home inside the enterprise.

3. RESEARCH METHODOLOGY

This study uses Design Science Research Methodology (DSRM) because the problem is prescriptive rather than descriptive: the goal is to design a governance intervention, not to explain an existing phenomenon. DSRM is appropriate when research seeks to create and justify an artifact that addresses an identified organizational problem and when the contribution resides in the artifact and the knowledge embedded in it (Peppers et al., 2007). Following Peppers et al. (2007), DSRM supplies a nominal methodology running from problem identification through communication. Gregor and Hevner (2013) complement that methodology by clarifying how artifact contributions can be positioned to maximize scholarly impact.

Wieringa's nested problem-solving account is also useful here because it separates the practical problem from the knowledge questions nested within it (Wieringa, 2009). The top-level problem in this study is practical: SME deployers lack a workable governance design for GenAI. Literature synthesis, mapping-table construction, and artifact justification are knowledge tasks nested within that practical problem, not separate research objectives.

The study maps to the six DSRM activities in a compressed but complete way. Problem identification and motivation address the absence of a deployer-oriented governance design suited to SME capacity. Solution objectives establish a lightweight, NIST-aligned design that can be operated without a dedicated AI governance function. Design and development yield the mechanism summary in Table 1, the control catalog logic derived from it, and supporting materials. Demonstration proceeds through representative SME operating contexts and common deployer use cases. Evaluation is analytical and formative, focusing on coherence, coverage, and traceability. Communication presents the core artifact in the body of the paper while keeping operational materials in supplementary form (Peppers et al., 2007; Gregor & Hevner, 2013).

Following Gregor and Hevner’s (2013) knowledge contribution framework, this study is best characterized as an improvement contribution: it develops a more effective governance solution for a recognized organizational problem within a familiar application domain, building on existing frameworks.

4. PROPOSED GOVERNANCE DESIGN

4.1 From Governance Frameworks to an SME AI Governance Design

The proposed artifact follows a two-step structure. The first step simplifies large-enterprise governance expressions into an SME form that can be operated with fewer roles, lighter documentation requirements, and fewer approval layers. The second step overlays NIST’s AI and GenAI risk-management concerns onto those simplified mechanisms. The result is a deployer-oriented governance design summarized in Table 1 and expressed operationally through the control catalog logic.

Governance mechanism	Typical large-enterprise expression	Simplified SME expression
Decision rights and accountability	Formal board oversight, AI committee, enterprise architecture board, layered RACI, multiple approval bodies	Executive sponsor, use-case owner, IT lead, and one shared review point
Policies and acceptable use	Multi-layer policy stack across IT, legal, compliance, privacy, security, and human resources	One short AI policy plus a practical workforce acceptable use policy
Use-case intake and inventory	Portfolio office, central architecture repository, formal application inventory	Simple intake form and living use-case register
Risk classification and approval gates	Multi-stage risk committee review, formal control gates, enterprise risk scoring	Three-tier rubric with one clear go/no-go decision
Documentation and evidence	Extensive project files, assurance documents, architecture records, audit artifacts	Short system card, evaluation record, and evidence log
Vendor and integration oversight	Dedicated procurement, vendor-risk office, contract negotiation teams, ongoing third-party assurance	Short due diligence checklist and periodic vendor-change review
Monitoring, incident response, and learning	Formal dashboards, management reports, internal audit review, enterprise incident playbooks	Small monthly or quarterly review plus AI incident types added to the existing incident process

Table 1. Governance mechanisms expressed in SME environments.

Table 1 is the paper’s primary design summary. It maps governance mechanisms commonly expressed in large enterprises to simplified SME equivalents, with the explicit purpose of preserving accountability, review, and traceability without demanding governance infrastructure that SMEs do not have.

The seven mechanisms in Table 1 operationalize the seven questions developed in Section 2.1, though not one-to-one. Who decides becomes decision rights and accountability; what is documented becomes documentation and evidence; how work is approved is split between use-case intake and risk classification, because approval in SMEs functions best as registration followed by a tier-scaled decision; how information moves is carried by the register and evidence log; how staff are prepared is embedded in the policies and acceptable-use mechanism, the SME vehicle for workforce expectations and training; and how systems are monitored and how incidents are handled combine into the final mechanism. Vendor and integration oversight is the one mechanism NIST contributes directly, responding to Govern 6 and the GenAI Profile’s value-chain risks, which COBIT’s enterprise-oriented components do not isolate (National Institute of Standards and Technology, 2023, 2024). In essence, every row in Table 1 traces to a named question, component, or NIST function.

The table is organized to improve readability rather than to enumerate an exhaustive framework. In essence, AI governance in SMEs can be built from a small set of familiar mechanisms, and each mechanism earns its place for a specific reason. Decision rights matter because AI use becomes ungovernable when no one clearly owns approval, oversight, exceptions, or escalation. Policies matter because GenAI use is predicated upon clear rules for prompts, uploads, outputs, disclosure, and prohibited use. Intake and risk classification matter because not every AI use warrants the same level of scrutiny. Documentation matters because AI governance requires traceability for decisions, operational limits, settings, and review outcomes. Vendor and integration oversight matter because the majority of SME AI capabilities arrive through external providers or connected systems rather than in-house development. Monitoring and incident response matter because AI governance does not end at deployment—tool behavior, vendor behavior, and user behavior can all change in ways that create new risk (ISACA, 2018a, 2018b; National Institute of Standards and Technology, 2023, 2024, 2025; Ji et al., 2023; Liu et al., 2024; Nelson et al., 2025).

The simplified SME expressions are justified by three design principles applied to every row. The first design principle is functional preservation. Every large-enterprise expression performs three governance functions, 1) someone is accountable for a decision, 2) a record of the decision exists, and 3) a review point can stop or revise it. Table 1 expresses a simplified expression of all three functions. Replacing layered committees with an executive sponsor, a use-case owner, and one shared review point removes coordination structure while preserving ownership, evidence, and veto, eliminating coordination redundancy across business units of the SMEs. The second design principle is role consolidation rather than elimination: responsibilities combine into fewer named people, without neglecting any governance function. The third design principle is proportionality: documentation and review effort scale with the risk tier assigned at intake, so low-risk uses consume little governance capacity. These principles are predicated upon the source frameworks themselves. COBIT directs that governance systems be tailored to enterprise size and context through its design factors (ISACA, 2018a), and NIST presents the AI RMF as voluntary, risk-based, and usable by organizations of all sizes (National Institute of Standards and Technology, 2023). The simplifications thus apply the frameworks' own tailoring logic.

The second step in the design, the control catalog logic, translates Table 1's governance expressions into an operational governance structure. Each entry in the catalog should, at minimum, specify the governance mechanism, the particular control, the trigger condition that activates it, and the repository where evidence of the control is recorded. Table 2 provides representative examples of how this logic might be applied across the governance mechanisms identified in Table 1.

Mechanism	Example control	Trigger	Evidence
Decision rights and accountability	Assign a named use-case owner for every approved AI use case.	Any recurring business use of AI.	Use-case register
Decision rights and accountability	Require executive approval for moderate- or high-risk AI use.	AI use involving customers, employees, regulated data, financial decisions, or external outputs.	Approval record
Policy and acceptable use	Prohibit confidential, regulated, customer, employee, financial, credential, or proprietary data in unapproved AI tools.	Any prompt, upload, attachment, or data connection.	Acceptable-use policy
Policy and acceptable use	Require human review before AI outputs are used externally or in consequential decisions.	AI output informs customer communication, HR, finance, compliance, legal, operations, or public-facing content	Review note
Use-case intake and inventory	Review the use-case register periodically to remove inactive or unauthorized uses.	Monthly or quarterly review cycle.	Use-case register
Use-case intake and inventory	Capture the business purpose for each AI use case.	New AI use is proposed.	Intake form
Risk classification	Assign low/moderate/high tier	Before deployment	Risk-tiering rubric

Risk classification	Allow low-risk uses through simple owner approval.	AI use involves general productivity, drafting, brainstorming, or internal summarization with non-sensitive data.	Owner approval note
Documentation	Keep a simple evaluation record for AI outputs used in important decisions.	AI output supports decisions, recommendations, analysis, or external deliverables.	Evaluation record
Vendor oversight	Review vendor AI terms before approving a tool.	New AI tool, AI-enabled SaaS feature, or vendor-provided AI capability is proposed.	Vendor checklist
Monitoring, incident response, and learning	Review active AI use cases on a monthly or quarterly basis.	Scheduled governance review	Review log
Monitoring, incident response, and learning	Update policy, training, or controls after AI incidents or near misses.	AI incident, user confusion, recurring output problem, or control failure.	Lessons-learned note

Table 2. Examples of Control Catalog Logic for SMEs

5. DISCUSSION

5.1 Identifying the Governance Mechanisms That Matter Most

The first contribution is conceptual and practical. This paper identifies a small set of existing IT governance mechanisms that carry disproportionate weight for AI governance in SMEs: decision rights, policy, intake, risk classification, documentation, vendor and integration oversight, and monitoring plus incident response. This selection not only gives the paper a more tractable and readable structure than one organized around abstract AI principles, but also reflects a deliberate argument about where governance effort in resource-constrained organizations should be concentrated. What distinguishes this selection from a checklist is its derivation. Each mechanism traces to a named COBIT component through the seven questions of Section 2.1 and to a named NIST function or Govern category, so anyone can ask where a mechanism came from and receive a specific answer. That traceability is what the formative evaluation assessed, and it separates a governance design from an opinion about good practice.

5.2 Combining COBIT with NIST

The second contribution is architectural. The paper treats COBIT and NIST as complementary structures that serve different functions in the same design. COBIT anchors governance through recurring components and decision structures that give the enterprise consistent control over its operations. NIST specifies what those components must address once AI enters scope. Combining them keeps the governance design grounded in established organizational logic while ensuring that AI-specific risk is explicitly and traceably managed.

5.3 Making the SME Difference Concrete

The third contribution is the explicit SME adaptation. This paper does not assume that large-enterprise governance forms can be transposed into SMEs without modification, nor that simplification is self-justifying. The three design principles articulated in Section 4.1—functional preservation, role consolidation, and proportionality—specify what a simplification must retain for it to remain governance rather than mere convenience. Risk classification illustrates the point: a multi-stage risk committee and a three-tier rubric with one go/no-go decision differ enormously in cost, but both answer the same question—does this use warrant deployment under these conditions—and both leave evidence of the answer. This analysis extends COBIT’s tailoring guidance by showing that scale-appropriate governance is a matter of preserving governance functions, not governance structures.

5.4 Implications and Boundaries

For practice, this paper argues that SMEs should not begin AI governance by attempting to mirror provider-side AI assurance programs or full regulatory compliance architectures. The more productive

starting point is to adapt the governance mechanisms they already recognize and operate. This is predicated upon the observation that SME governance is largely informal and distributed. It exists in roles, habits, and relationships rather than in formal structures. Any viable governance design must meet organizations where they are. This orientation is especially important in the first three AI operating contexts addressed here, where deployer concerns about prompt handling, output reliability, embedded functions, connector security, data management, vendor change, and incident response dominate day-to-day governance demands.

The design also implies a concrete adoption sequence. An SME applying this artifact would begin with the one-page AI policy and acceptable-use rules, because policy governs behavior already occurring. It would then stand up the use-case register, assign an owner and sponsor to each active use, and apply the three-tier rubric. Documentation, vendor review, and the monitoring cycle follow once the register exists, since each is predicated upon knowing what AI is in use. In essence, Table 1 names the mechanisms, Table 2 supplies their operating logic, and this sequence orders their adoption, giving an SME a path from unmanaged GenAI use to an operating governance system within existing capacity.

The paper also has clear limits that deserve acknowledgment. It is NIST-first by design, with COBIT as the governance framework anchor; organizations operating under EU AI Act requirements or using ISO/IEC frameworks as their primary governance vocabulary will need to adapt the design accordingly. The scope is bounded to the first three AI operating contexts, excluding model development environments where data governance, evaluation infrastructure, and computational resource considerations introduce substantially different governance demands. Evaluation is analytical and formative rather than field-based, which means the design has not been tested against the friction of actual SME implementation. Future work should examine how SMEs adapt this governance design in practice and whether the mechanism-based framing demonstrably improves adoption rates, accountability clarity, and incident handling outcomes.

6. CONCLUSION

This research addressed the problem of how SMEs can govern generative AI when they operate primarily as deployers rather than model developers. The central claim is direct: SMEs do not need a separate governance framework for AI. They need a simplified and extended version of the IT governance mechanisms they already have, adapted to account for the particular risks that generative AI introduces at the deployer level.

Using COBIT as the governance framework anchor and NIST as the AI risk-management spine, the research developed a governance design expressed through Table 1 and the control catalog logic through Table 2. These contributions show how IT governance can be adapted to govern AI use in smaller organizations that lack the capacity to build AI governance from the ground up.

7. REFERENCES

- Bender, E. M., & Friedman, B. (2018). Data statements for natural language processing: Toward mitigating system bias and enabling better science. *Transactions of the Association for Computational Linguistics*, 6, 587-604. https://doi.org/10.1162/tacl_a_00041
- European Commission. (n.d.). *SME definition*. Directorate-General for Internal Market, Industry, Entrepreneurship and SMEs. Retrieved April 16, 2026, from https://single-market-economy.ec.europa.eu/smes/sme-fundamentals/sme-definition_en
- European Parliament and Council of the European Union. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)* (OJ L, 2024/1689, 12 July 2024). <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>
- 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>
- Gregor, S., & Hevner, A. R. (2013). Positioning and presenting design science research for maximum impact. *MIS Quarterly*, 37(2), 337-355. <https://doi.org/10.25300/MISQ/2013/37.2.01>

- ISACA. (2018a). COBIT 2019 framework: Introduction and methodology. ISACA.
- ISACA. (2018b). COBIT 2019 framework: Governance and management objectives. ISACA.
- International Organization for Standardization & International Electrotechnical Commission. (2022). *ISO/IEC 38507:2022 Information technology - Governance of IT - Governance implications of the use of artificial intelligence by organizations* (ISO/IEC 38507:2022). <https://www.iso.org/standard/56641.html>
- International Organization for Standardization & International Electrotechnical Commission. (2023a). *ISO/IEC 23894:2023 Information technology - Artificial intelligence - Guidance on risk management* (ISO/IEC 23894:2023). <https://www.iso.org/standard/77304.html>
- International Organization for Standardization & International Electrotechnical Commission. (2023b). *ISO/IEC 42001:2023 Information technology - Artificial intelligence - Management system* (ISO/IEC 42001:2023). <https://www.iso.org/standard/42001>
- Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y., Chen, D., Dai, W., Chan, H. S., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12), Article 248. <https://doi.org/10.1145/3571730>
- Liu, Y., Jia, Y., Geng, R., Jia, J., & Gong, N. Z. (2024). Formalizing and benchmarking prompt injection attacks and defenses. In *Proceedings of the 33rd USENIX Security Symposium (USENIX Security '24)* (pp. 1831-1847). USENIX Association. <https://www.usenix.org/conference/usenixsecurity24/presentation/liu-yupei>
- Mandviwalla, M. (2015). Generating and justifying design theory. *Journal of the Association for Information Systems*, 16(5), 314-344. <https://doi.org/10.17705/1jais.00397>
- Mäntymäki, M., Minkinen, M., Birkstedt, T., & Viljanen, M. (2022). Defining organizational AI governance. *AI and Ethics*, 2, 603-609. <https://doi.org/10.1007/s43681-022-00143-x>
- 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. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* '19)* (pp. 220-229). Association for Computing Machinery. <https://doi.org/10.1145/3287560.3287596>
- Mökander, J., & Floridi, L. (2023). Operationalising AI governance through ethics-based auditing: An industry case study. *AI and Ethics*, 3, 451-468. <https://doi.org/10.1007/s43681-022-00171-7>
- 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>
- National Institute of Standards and Technology. (2024). *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile* (NIST AI 600-1). <https://doi.org/10.6028/NIST.AI.600-1>
- National Institute of Standards and Technology. (2025, February 6). *NIST AI RMF Playbook*. Retrieved March 1, 2026, from <https://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook>
- Nelson, A., Rekhi, S., Souppaya, M., & Scarfone, K. (2025, April). *Incident response recommendations and considerations for cybersecurity risk management: A CSF 2.0 community profile* (NIST Special Publication 800-61 Rev. 3). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-61r3>
- Papagiannidis, E., Enholm, I. M., Dremel, C., Mikalef, P., & Krogstie, J. (2023). Toward AI governance: Identifying best practices and potential barriers and outcomes. *Information Systems Frontiers*, 25, 123-141. <https://doi.org/10.1007/s10796-022-10251-y>
- Papagiannidis, E., Mikalef, P., & Conboy, K. (2025). Responsible artificial intelligence governance: A review and research framework. *The Journal of Strategic Information Systems*, 34(2), 101885. <https://doi.org/10.1016/j.jsis.2024.101885>
- Peppers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, 24(3), 45-77. <https://doi.org/10.2753/MIS0742-1222240302>

- Taeihagh, A. (2025). Governance of generative AI. *Policy and Society*, 44(1), 1-22.
<https://doi.org/10.1093/polsoc/puaf001>
- Wieringa, R. J. (2009). Design science as nested problem solving. In *Proceedings of the 4th International Conference on Design Science Research in Information Systems and Technology (DESRIST '09)* (Article 8, pp. 1-12). Association for Computing Machinery.
<https://doi.org/10.1145/1555619.1555630>