

Beyond Authorized Adoption: Shadow AI, Governance Visibility, and Information Security Exposure in Knowledge Work

Ade Ayuni Syam¹

¹ Universitas Terbuka, Indonesia

Correspondence: y.ayuni12@gmail.com

Article Info

Article History:

Received Aug 15th, 2026

Revised Aug 21th, 2026

Accepted Aug 30th, 2026

Keyword:

AI governance; Governance visibility; Information security exposure; Knowledge work; Shadow AI.

ABSTRACT

This study examined why formal authorization did not ensure that organizational use of generative artificial intelligence was visible enough to be governed. It used a theory-driven integrative literature synthesis across Shadow IT, generative artificial intelligence governance, information security governance, and AI-enabled knowledge work. The synthesis distinguished formal authorization as institutional status, governance visibility as epistemic capacity, and governability as organizational action capacity. Crossing authorization with visibility produced four ideal-type conditions consisting of Governed AI, Opaque Authorized AI, Detectable Shadow AI, and Deep Shadow AI. The analysis also separated information security exposure from vulnerability, incident, and realized harm. Opaque Authorized AI emerged as the most consequential conceptual correction because an approved AI-use practice could remain difficult to observe, reconstruct, and govern. The study concluded that approval alone was an insufficient basis for organizational AI governance and that governance analysis needed to focus on the visibility of actual AI-use practices in knowledge work.



© 2026 The Authors. Published by Punggawa Legacy Center. This is an open access article under the CC BY license (<https://creativecommons.org/licenses/by/4.0/>)

INTRODUCTION

Formal approval does not guarantee that organizational AI use is visible to the people expected to govern it. Generative AI now enters knowledge work through ordinary browsers, embedded software, personal accounts, and enterprise platforms. Much of this use can be productive, and field and experimental studies show meaningful gains in speed and output quality for some tasks (Noy & Zhang, 2023) (Brynjolfsson et al., 2025). Yet the same ease of access weakens a familiar governance assumption. Organizations may know which tools they approve while knowing much less about how those tools are actually used, what information enters them, or where generated material later travels. The governance problem therefore begins inside authorized environments as well as outside them.

Shadow IT research provides an obvious starting point because it has long examined technology that escapes formal approval. Unsanctioned systems often grow from workarounds, unmet local needs, or tension between formal procedures and the demands of work (Ferneley & Sobrepez, 2006) (Silic & Back, 2014). Later research treated shadow systems as arrangements that can be detected, tolerated, absorbed, or brought under business management (Haag & Eckhardt, 2017) (Kopper et al., 2018) (Fürstenau et al., 2021). Hiddenness and authorization were never quite the same thing.

GenAI makes the separation harder to ignore. Shadow AI research describes unsanctioned organizational use and the governance drift that appears when employee practice develops faster than formal controls (Silic et al., 2025). The difficulty is not confined to unauthorized tools. An approved service may be well known at the application level while the organizationally relevant use remains obscure. Sensitive material can enter a prompt. Generated text can be copied into another system. A recommendation can influence a decision without the AI interaction remaining attached to the final workflow. Approval tells the organization that a category of use is permitted. It does not tell the organization what happened in a particular AI-enabled work practice.

Visibility is not new to AI governance. Research on AI agents treats visibility as information that allows oversight bodies to understand where systems are deployed and how they are used (Chan et al., 2024). Internal auditing work also stresses records that allow organizations to reconstruct decisions and assign responsibility (Raji et al., 2020). Reviewability research treats automated decision making as a sociotechnical process that needs suitable records before meaningful review is possible (Cobbe & Singh, 2020). The unresolved issue is narrower. Shadow AI discussion still leans heavily on the authorized and unauthorized divide, while visibility is rarely theorized as an independent property of organizational AI-use practice.

Knowledge work makes the distinction especially important because GenAI is often used as a cognitive resource rather than as a discrete enterprise system. Workers draft, interpret, compare, and explore through tools that can sit beside the formal workflow. Learning algorithms already alter expertise and coordination in ways that are difficult to reduce to conventional software use (Faraj et al., 2018). Agentic information systems also make delegation central to the user-system relation (Baird & Maruping, 2021). Recent GenAI studies show movement from experimentation toward routine professional practice (Retkowsky et al., 2024) (Butollo et al., 2025). The organizational AI-use practice, not the software asset alone, is therefore the more useful unit of analysis.

Two questions follow from this problem. The first asks how governance visibility should be distinguished from formal authorization in organizational AI use. The second asks how their interaction shapes an organization's capacity to recognize and govern information security exposure in knowledge work. The article develops a conceptual account through an integrative literature synthesis. Its central claim is modest but consequential. Authorization is an institutional status. Governance visibility is an epistemic capacity. Governability is an organizational action capacity. When those properties diverge, an approved AI-use practice can remain difficult to govern, while an unauthorized practice can sometimes be visible enough to investigate and contain.

RESEARCH METHODS

This study uses conceptual research supported by a theory-driven integrative literature synthesis. The design fits a problem concerned with construct boundaries rather than prevalence, prediction, or causal estimation. Conceptual work can create value by separating constructs that have been treated as equivalent and then rebuilding their relations in a more useful theoretical form (Jaakkola, 2020). Integrative review methods are suitable when relevant ideas are distributed across adjacent literatures and the purpose is interpretive synthesis rather than exhaustive aggregation (Whitemore & Knafl, 2005) (Snyder, 2019).

Literature was selected purposively from four domains tied directly to the research questions. Shadow IT supplied the lineage of unauthorized technology use. Shadow AI and GenAI governance supplied the contemporary setting. Information security and AI governance contributed work on oversight and organizational control, while knowledge-work scholarship explained everyday professional use. Peer-reviewed research with verifiable DOI records was prioritized. Targeted searching was supplemented by backward and forward citation tracing when a paper exposed a construct needed for the framework.

Rather than treating the literature as a dataset for counting studies, the analysis used it as theoretical evidence. Analysis proceeded through conceptual comparison, construct differentiation, and relational synthesis. Each source was read for how it described authorization, what information about use could become known, what organizational response became possible once use was known, and how information moved beyond an intended governance boundary. Terms were retained only when they performed a distinct analytical function. This process separated formal authorization from governance visibility, moved reviewability into governability, and defined the organizational AI-use practice as the unit to which the framework applies.

Claims produced by the synthesis remain conceptual and within that evidentiary limit. The synthesis can distinguish governance conditions and explain why they create different capacities for recognition and response. It cannot show how common those conditions are, estimate effect sizes, or establish that low visibility causes a security incident. Information security exposure is therefore treated as susceptibility within an information-handling arrangement, not as proof of breach or loss. Questions about frequency, magnitude, and causal effects require later empirical designs with different evidence.

RESULTS AND DISCUSSION

Formal authorization and governance visibility solve different organizational problems. Authorization indicates whether a particular AI-enabled practice is permitted under organizational rules. It is a status produced by policy, approval, procurement, or another recognized decision process. Visibility concerns what the organization can know about the practice after or while it occurs. Treating the two as one construct overstates what approval can achieve. Shadow IT studies already show that technology can be identified before it is formally accepted and can remain locally managed even after detection (Fürstenau et al., 2017) (Kopper et al., 2018). GenAI widens that gap because approval may attach to the platform while opacity emerges later in the interaction.

An organizational AI-use practice is the relevant unit. It is narrower than the tool and broader than a single prompt. The term refers to work in which AI processes organizationally relevant information or shapes an organizational output. This prevents a common classification error. The same approved platform can support one practice that fits policy and another that does not. It can also support a permitted practice whose traces disappear once output moves into a downstream workflow. Authorization should therefore attach to the practice that matters, not merely to the application on an approved-tool register.

Governance visibility is best understood as an epistemic organizational capacity. Three elements are sufficient for the present framework. Observability concerns whether the organization can know that a consequential AI-use practice occurred. Traceability concerns whether the relevant interaction or information flow can be reconstructed with enough detail for governance purposes. Attributability concerns whether that practice can be connected to a responsible actor, role, or workflow. Chan et al. (2024) show why deployment visibility matters for oversight, while Raji et al. (2020) demonstrate the organizational value of records that support auditing and accountability. Visibility does not require exhaustive surveillance. It requires enough knowledge for the governance purpose at stake.

Reviewability does not belong inside visibility. A practice can be visible and still remain beyond meaningful review because no role has authority to challenge it, no procedure exists for reconsideration, or the organization lacks the competence to assess what occurred. Cobbe and Singh (2020) treat reviewability as a broader accountability mechanism grounded in records and organizational process. Novelli et al. (2024) likewise place accountability around identifiable responsibilities and answerability rather than transparency alone. Governability is therefore defined here as the practical capacity to review, constrain, correct, escalate, or otherwise intervene in an AI-use practice. Visibility can enable that capacity, but it does not guarantee it.

Table 1 Construct Boundaries in the Revised Framework

Construct	Working definition	Primary question	Analytical status	Boundary
Formal authorization	Institutional permission for a particular organizational AI-use practice	Is the practice permitted	Institutional status	Not visibility
Governance visibility	Capacity to observe, trace, and attribute consequential AI use	What can the organization know about the practice	Epistemic capacity	Not reviewability
Governability	Capacity to review, constrain, correct, escalate, or intervene in AI use	What can the organization do once the practice is known	Organizational action capacity	Not visibility alone
Information security exposure	Susceptibility created when organizational information is handled partly outside its intended governance boundary	What information-handling condition needs governance attention	Exposure condition	Not vulnerability, incident, or harm
Shadow AI	Organizationally relevant AI-use practice occurring outside formal authorization	Is the practice outside sanctioned arrangements	Governance phenomenon	Not synonymous with invisibility

This separation changes the logic of Shadow AI. A simple authorized and unauthorized divide remains useful for identifying compliance status, yet it cannot describe whether the organization knows enough to act. Organizational AI governance research increasingly recognizes that policies require

operational translation into structures and practices (Mäntymäki et al., 2022) (Birkstedt et al., 2023) (Papagiannidis et al., 2025). Research on sociotechnical guardrails and user involvement makes the same point from a different direction (Jung & Winter, 2025) (Sturm et al., 2026). Formal permission can exist without the information needed for review. Detection can occur before permission. The two properties therefore vary independently enough to justify a two-dimensional account.

AI Governance Visibility Gap names the discrepancy between an organizationally consequential AI-use practice and the extent to which that practice can be observed, traced, and attributed through available governance arrangements. The concept does not claim that visibility has been absent from earlier AI governance research. Its contribution is to detach visibility from authorization within the Shadow AI problem. That move matters because the governance failure can occur inside an approved environment. A platform may be sanctioned while prompt content, downstream reuse, or responsibility for an AI-assisted output remains difficult to reconstruct. The gap is therefore about what the organization can know about use, not about whether the user has broken a rule.

Table 2 Literature Streams Used in the Conceptual Synthesis

Literature stream	Main concern	What it contributes	Limit that matters here	Use in this article
Shadow IT	Unauthorized and locally managed technology	Authorization boundaries and post-detection governance	Developed largely before general-purpose GenAI	Shows that detection and authorization can diverge
Shadow AI and GenAI governance	Unsanctioned AI use and operational governance	AI-specific diffusion and governance drift	Often retains an authorized and unauthorized framing	Locates the current organizational problem
AI and information security governance	Oversight, accountability, records, and organizational control	Visibility, auditability, responsibility, and action capacity	Can assume the governed object is already known	Supports visibility and governability as separate constructs
AI-enabled knowledge work	AI embedded in professional activity	Delegation, workflow diffusion, and changing expertise	Governance is often secondary to work outcomes	Explains why use-level opacity can arise inside ordinary work

Crossing authorization with visibility produces four ideal-type conditions. Governed AI describes an authorized practice with comparatively high visibility. Detectable Shadow AI is unauthorized but visible enough for the organization to recognize and investigate it. Deep Shadow AI combines unauthorized status with low visibility, which limits both recognition and the possibility of timely intervention. Opaque Authorized AI is the more revealing case. The practice is permitted, yet the organization lacks enough knowledge about consequential use to govern it confidently. These are analytical types rather than measured categories. High and low visibility should be understood relative to the governance purpose and the information needed to meet it.

Opaque Authorized AI exposes a weakness in tool-centred governance. Organizations often approve products, vendors, or technical environments. Knowledge work occurs through practices that do not stay inside those boundaries. An employee can use an approved assistant to transform material and then move the output into another document, message, model, or decision process. At that point the tool remains authorized while the use becomes harder to reconstruct. Tool visibility therefore does not imply use visibility. Berente et al. (2021) describe AI as difficult to manage partly because autonomy, learning, and inscrutability complicate conventional control. In everyday GenAI work, the added difficulty is that the relevant object of governance can dissolve into ordinary cognitive and documentary activity.

Table 3 Authorization and Governance Visibility Matrix

Condition	Authorization	Governance visibility	Governance capacity	Exposure legibility
Governed AI	Authorized	Relatively high	Stronger basis for review and intervention	Exposure is more observable and reconstructable
Opaque Authorized AI	Authorized	Relatively low	Constrained despite formal permission	Exposure can remain difficult to reconstruct
Detectable Shadow AI	Unauthorized	Relatively high	Potential for investigation and remediation	Exposure can be recognized despite noncompliance

Table 3 Continuation

Condition	Authorization	Governance visibility	Governance capacity	Exposure legibility
Deep Shadow AI	Unauthorized	Relatively low	Recognition and intervention are both constrained	Exposure can remain poorly legible

Knowledge work is not decorative in this argument. Discretion is often built into professional work, and the information being handled is less standardized than a transaction recorded in a core system. AI can be consulted privately, briefly, and iteratively while the formal output appears in a conventional document or communication channel. Faraj et al. (2018) connect learning algorithms to changing expertise and organizational boundaries. Baird and Maruping (2021) show why delegation to agentic artifacts changes the user-system relation. GenAI extends both dynamics because conversational interaction can become part of reasoning without becoming a durable organizational record. Visibility may therefore erode at the exact point where AI becomes ordinary. A broader sociotechnical account of GenAI points in the same direction (Storey et al., 2025).

Information security exposure should also be kept narrower than the earlier version of this framework implied. Exposure exists when organizational information is handled in a way that places it outside, or partly outside, the governance boundary intended for that information. It is not the same as a vulnerability in the technical system. Nor is it an incident or realized harm. Shadow IT research explains why unauthorized technology can create security concerns when information handling escapes formal control (Silic & Back, 2014) (Silic et al., 2017). With GenAI, the exposure can arise through the information flow itself even when the application is legitimate. The analytical focus should remain on susceptibility and governance knowledge rather than presumed damage.

Low visibility does not necessarily create the exposure. It changes how legible that exposure is to the organization and how readily the organization can respond. Sensitive-data transfer may occur in a visible environment or an opaque one. The difference is that the first can leave enough evidence for attribution and review, while the second may remain difficult to reconstruct. Table 3 therefore avoids ranking the configurations by security risk. Such a ranking would require empirical evidence about the information involved, technical safeguards, and surrounding controls. The framework supports a claim about exposure legibility and response capacity, not breach probability.

Proposition 1 follows directly from this distinction. Formal authorization and governance visibility are analytically independent properties of an organizational AI-use practice. This does not mean they are statistically unrelated. Organizations may increase visibility when they authorize a system. The point is that one cannot be inferred from the other. An approved practice may remain opaque, while an unauthorized practice may still be detected through network, identity, endpoint, or workflow evidence.

A second proposition concerns governability. Greater governance visibility expands the information available for organizational review and intervention, but governability also depends on authority and procedure. A visible practice can remain poorly governable when no actor owns the issue or when review rights are unclear. This boundary condition matters because it prevents visibility from becoming another oversized construct. Monitoring alone is not governance. Organizational AI governance depends on roles, decision rights, and processes that can translate knowledge into action (de Almeida & dos Santos Júnior, 2025) (Papagiannidis et al., 2025). The framework therefore treats visibility as an enabling epistemic condition and governability as a separate action capacity.

Proposition 3 is narrower. Authorization and visibility together condition the organization's capacity to recognize and govern information security exposure. Deep Shadow AI is difficult because the organization has neither formal permission nor adequate knowledge of use. Detectable Shadow AI preserves more room for response because the practice is known even though it is not permitted. Opaque Authorized AI presents a different failure. The organization has granted permission but may lack the evidentiary basis needed to reconstruct the use. None of these conditions determines whether harm occurs. They change what the organization can know, attribute, and do when exposure needs to be assessed.




Figure 1 Configurational Governance Framework for Organizational AI-Use Practice

A normative limit is also clearly necessary. Maximum visibility is not an obvious governance objective. Continuous capture of prompts, identities, and work products can conflict with privacy, professional autonomy, proportionality, or data-minimization duties. Visibility should therefore be purpose-relative. Organizations need enough observability, traceability, and attributability to support the governance task, not an unlimited record of employee cognition. This qualification also explains why high visibility is not defined as complete surveillance.

Several empirical paths remain open. Later studies could operationalize authorization and the three visibility dimensions separately, then classify AI-use practices rather than organizations as a whole. Longitudinal research could examine movement between configurations. Outcome studies would need another evidentiary layer. Measures of unreviewable information transfer or missing lineage could capture exposure conditions, while breach frequency or realized loss would require incident data. Keeping these levels separate is essential if later testing is to preserve the conceptual boundaries developed here.

CONCLUSION

Shadow AI becomes harder to understand when authorization is asked to carry the whole governance problem. Formal permission tells an organization whether an AI-use practice is sanctioned. It does not establish whether the practice can be observed, reconstructed, attributed, or meaningfully reviewed. The distinction is especially important in knowledge work, where GenAI can enter reasoning and document production without remaining attached to a durable system record. Treating the organizational AI-use practice as the unit of analysis makes that gap visible. Three properties structure the revised framework. Authorization is institutional status. Governance visibility is epistemic capacity. Governability is organizational action capacity. Their interaction produces four ideal-type conditions, including Opaque Authorized AI, where permission exists but use remains difficult to govern. Information security exposure sits alongside these governance conditions rather than downstream from governability. Visibility affects how exposure can be recognized and handled, not whether harm has already occurred. This narrower formulation preserves the article's theoretical claim without borrowing causal force from evidence that the study does not possess.

What changes here is the organization of the Shadow AI problem, not a claim that visibility itself is new. Existing work already shows the importance of monitoring, records, and review in AI oversight. The present synthesis explains why those capacities must be separated from authorization in organizational knowledge work. Empirical research can now test how often the four configurations occur, how practices move between them, and whether particular visibility arrangements improve response. Until that evidence exists, the defensible conclusion is simpler. An AI practice can be approved and still remain opaque enough to challenge governance.

REFERENCES

- Baird, A., & Maruping, L. M. (2021). The next generation of research on IS use. A theoretical framework of delegation to and from agentic IS artifacts. *MIS Quarterly*, 45(1), 315–341. <https://doi.org/10.25300/MISQ/2021/15882>
- Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Managing artificial intelligence. *MIS Quarterly*, 45(3), 1433–1450. <https://doi.org/10.25300/MISQ/2021/16274>
- Birkstedt, T., Minkkinen, M., Tandon, A., & Mäntymäki, M. (2023). AI governance. Themes, knowledge gaps and future agendas. *Internet Research*, 33(7), 133–167. <https://doi.org/10.1108/INTR-01-2022-0042>
- Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–942. <https://doi.org/10.1093/qje/qjae044>
- Butollo, F., Haase, J., Katzinski, A.-K., & Krüger, A. K. (2025). Generative AI and changes to knowledge work. *Weizenbaum Journal of the Digital Society*, 5(4). <https://doi.org/10.34669/WI.WJDS/5.4.1>
- Chan, A., Ezell, C., Kaufmann, M., Wei, K., Hammond, L., Bradley, H., Bluemke, E., Rajkumar, N., Krueger, D., Kolt, N., Heim, L., & Anderljung, M. (2024). Visibility into AI agents. *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, 958–973. <https://doi.org/10.1145/3630106.3658948>
- Cobbe, J., & Singh, J. (2020). Reviewable automated decision-making. *Computer Law & Security Review*, 39, 105475. <https://doi.org/10.1016/j.clsr.2020.105475>
- de Almeida, P. G. R., & dos Santos Júnior, C. D. (2025). Artificial intelligence governance. Understanding how public organizations implement it. *Government Information Quarterly*, 42(1), 102003. <https://doi.org/10.1016/j.giq.2024.102003>
- Faraj, S., Pachidi, S., & Sayegh, K. (2018). Working and organizing in the age of the learning algorithm. *Information and Organization*, 28(1), 62–70. <https://doi.org/10.1016/j.infoandorg.2018.02.005>
- Ferneley, E. H., & Sobrepey, P. (2006). Resist, comply or workaround? An examination of different facets of user engagement with information systems. *European Journal of Information Systems*, 15(4), 345–356. <https://doi.org/10.1057/palgrave.ejis.3000629>
- Fürstenau, D., Rothe, H., & Sandner, M. (2017). Shadow systems, risk, and shifting power relations in organizations. *Communications of the Association for Information Systems*, 41, 43–61. <https://doi.org/10.17705/1CAIS.04103>
- Fürstenau, D., Rothe, H., & Sandner, M. (2021). Leaving the shadow. A configurational approach to explain post-identification outcomes of shadow IT systems. *Business & Information Systems Engineering*, 63, 97–111. <https://doi.org/10.1007/s12599-020-00635-2>
- Haag, S., & Eckhardt, A. (2017). Shadow IT. *Business & Information Systems Engineering*, 59, 469–473. <https://doi.org/10.1007/s12599-017-0497-x>
- Jaakkola, E. (2020). Designing conceptual articles. Four approaches. *AMS Review*, 10, 18–26. <https://doi.org/10.1007/s13162-020-00161-0>
- Jung, J. Y., & Winter, S. J. (2025). Policy for sociotechnical Gen AI assessment. Leveraging end users. *Journal of the Association for Information Systems*, 26(2), 287–293. <https://doi.org/10.17705/1jais.00934>
- Kopper, A., Fürstenau, D., Zimmermann, S., Klotz, S., Rentrop, C., Rothe, H., Strahringer, S., & Westner, M. (2018). Shadow IT and business-managed IT. A conceptual framework and empirical illustration. *International Journal of IT/Business Alignment and Governance*, 9(2), 53–71. <https://doi.org/10.4018/IJITBAG.2018070104>
- Mäntymäki, M., Minkkinen, 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>
- Novelli, C., Taddeo, M., & Floridi, L. (2024). Accountability in artificial intelligence. What it is and how it works. *AI & Society*, 39, 1871–1882. <https://doi.org/10.1007/s00146-023-01635-y>
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>
- 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>

- Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). 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*, 33–44. <https://doi.org/10.1145/3351095.3372873>
- Retkowsky, J., Hafermalz, E., & Huysman, M. (2024). Managing a ChatGPT-empowered workforce. Understanding its affordances and side effects. *Business Horizons*, 67(5), 511–523. <https://doi.org/10.1016/j.bushor.2024.04.009>
- Silic, M., & Back, A. (2014). Shadow IT. A view from behind the curtain. *Computers & Security*, 45, 274–283. <https://doi.org/10.1016/j.cose.2014.06.007>
- Silic, M., Barlow, J. B., & Back, A. (2017). A new perspective on neutralization and deterrence. Predicting shadow IT usage. *Information & Management*, 54(8), 1023–1037. <https://doi.org/10.1016/j.im.2017.02.007>
- Silic, M., Silic, D., & Kind-Trüller, K. (2025). From Shadow IT to Shadow AI. Threats, risks and opportunities for organizations. *Strategic Change*. Advance online publication. <https://doi.org/10.1002/jsc.2682>
- Snyder, H. (2019). Literature review as a research methodology. An overview and guidelines. *Journal of Business Research*, 104, 333–339. <https://doi.org/10.1016/j.jbusres.2019.07.039>
- Storey, V. C., Yue, W. T., Zhao, J. L., & Lukyanenko, R. (2025). Generative artificial intelligence. Evolving technology, growing societal impact, and opportunities for information systems research. *Information Systems Frontiers*, 27(5), 2081–2102. <https://doi.org/10.1007/s10796-025-10581-7>
- Sturm, S., Krause, F., & van Giffen, B. (2026). Beyond control? Governing AI in organizations for responsible use. *European Management Journal*. Advance online publication. <https://doi.org/10.1016/j.emj.2026.06.003>
- Whittemore, R., & Knafl, K. (2005). The integrative review. Updated methodology. *Journal of Advanced Nursing*, 52(5), 546–553. <https://doi.org/10.1111/j.1365-2648.2005.03621.x>