



## **Governing AI in Practice: A Typology of Four Risk-Driven Archetypes**

Downloaded from: <https://research.chalmers.se>, 2026-10-29 00:02 UTC

Citation for the original published paper (version of record):

Shteriyarov, V., Dzhusupova, R., Bosch, J. et al (2026). Governing AI in Practice: A Typology of Four Risk-Driven Archetypes. Proceedings 2026 IEEE ACM 5th International Conference on AI Engineering Software Engineering for AI Cain 2026: 210-215.  
<http://dx.doi.org/10.1145/3793653.3793787>

N.B. When citing this work, cite the original published paper.

© 2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, or reuse of any copyrighted component of this work in other works.

# Governing AI in Practice: A Typology of Four Risk-Driven Archetypes

Vasil Shteriyarov  
Mathematics and Computer Science  
Eindhoven University of Technology  
Eindhoven, Netherlands  
v.shteriyarov1@tue.nl

Jan Bosch  
Chalmers University of Technology  
Gothenburg, Sweden  
jan.bosch@chalmers.se

Rimma Dzhusupova  
Eindhoven University of Technology  
Eindhoven, Netherlands  
r.m.dzhusupova@tue.nl

Helena Holmström Olsson  
Malmö University  
Malmö, Sweden  
helenaholmstrom.olsson@mau.se

## Abstract

Organizations are rapidly adopting Artificial Intelligence (AI), but the governance frameworks needed to manage its risks remain poorly understood, creating a gap between high-level policy and real-world practice. Addressing this, our paper presents findings from an in-depth case study of a large Engineering, Procurement, and Construction (EPC) company, supplemented by five comparative interviews with leaders from other engineering and technology organizations. Our primary contribution is the identification of four distinct governance archetypes: High-Assurance, Cautious, Enterprise-Driven, and Developer-Led. We argue that their adoption is based on an organization's specific risk profile, defined by operational domain, AI application, and organizational scale. We further show how these governance archetypes directly shape technical architectures, from mandated hybrid systems in high-assurance environments to emergent designs in developer-led settings. The resulting typology provides a novel diagnostic framework for practitioners and an explanatory tool for researchers, bridging the crucial gap between regulation and implementation.

## CCS Concepts

• **Software and its engineering** → **Risk management; Software architectures**; • **Computing methodologies** → **Artificial intelligence**; • **Social and professional topics** → **Management of computing and information systems**.

## Keywords

AI Governance, Responsible AI, Governance Archetypes, Software Engineering for AI (SE4AI), AI Risk Management

## ACM Reference Format:

Vasil Shteriyarov, Rimma Dzhusupova, Jan Bosch, and Helena Holmström Olsson. 2026. Governing AI in Practice: A Typology of Four Risk-Driven Archetypes. In *2026 IEEE/ACM 5th International Conference on AI Engineering - Software Engineering for AI (CAIN '26)*, April 12–13, 2026, Rio de Janeiro, Brazil.



This work is licensed under a Creative Commons Attribution 4.0 International License. CAIN '26, Rio de Janeiro, Brazil

© 2026 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2475-6/26/04  
<https://doi.org/10.1145/3793653.3793787>

Brazil. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3793653.3793787>

## 1 Introduction

The adoption of Artificial Intelligence (AI) and automation systems is accelerating in the industry, promising to improve efficiency and drive innovation [20]. However, this shift creates an organizational problem: while AI's capabilities are advancing rapidly, the organizational practices for managing its risks and ensuring accountability are lagging behind [15]. In high-stakes domains like critical infrastructure, where errors can have catastrophic consequences, this is a critical safety and emerging legal imperative [5].

Extensive research has addressed AI governance from two primary angles. At the macro level, scholars have focused on high-level ethical principles and regulatory frameworks [6]. At the micro level, the focus has been on prescriptive solutions [10, 21]. What is missing is the crucial bridge connecting these two ends: the empirical understanding of the emergent governance frameworks that organizations are creating to connect principles to practice. There is a significant "governance gap" regarding the concrete rules, roles, and processes that organizations are implementing. It remains unclear how these frameworks are shaped by an organization's specific context, such as its risk environment and scale, and crucially, how these organizational choices constrain the resulting technical systems.

This paper identifies how organizational risk orientation and domain context shape AI governance. We adopt the definition of AI governance from [14] as a system of rules, practices, processes, and technological tools employed to ensure that an organization's use of AI aligns with its strategies, values, and legal requirements. To explore this, we conducted an exploratory qualitative study centered on an in-depth case study of McDermott, a major global Engineering, Procurement, and Construction (EPC) firm, supplemented by expert interviews from five other organizations. While we acknowledge the limited sample size characteristic of early-stage empirical work, this design provides a rich, contextualized understanding of how governance and engineering co-evolve.

We address the following research questions:

- (1) What governance frameworks for AI and automation systems are adopted by engineering organizations, and what factors shape the choice of framework?

- (2) What solution approaches and architectures emerge in practice based on these governance frameworks?

The paper makes the following contributions:

- **A Typology of Four Governance Archetypes:** We identify and define the *High-Assurance*, *Cautious*, *Enterprise-Driven*, and *Developer-Led* models. We demonstrate that these emerge predictably from the interplay of application risk, domain risk, and organizational scale.
- **The Governance-Architecture Link:** We establish a direct link between governance and software engineering, showing how archetypes mandate specific architectural patterns, including the use of hybrid systems for reliability and the formalization of Human-in-the-Loop (HITL) as an auditable component in high-risk settings.

## 2 Related Work

AI governance research spans high-level principles, prescriptive frameworks, and literature syntheses, yet a persistent gap remains between prescribed principles and organizational practice. At the macro level, scholars have articulated core ethical principles [6] that are increasingly codified into legal frameworks like the EU AI Act [5], regional initiatives [11, 23], and international standards such as ISO 42001 [9]. Notably, Ebers [3] points out that legal ambiguity complicates the translation of these high-level requirements into the auditable outcomes required for high-risk systems, such as risk management and data quality assurance.

At the micro level, research proposes prescriptive frameworks and tools to guide organizational implementation. These include conceptual models for responsible AI [17] and guidance frameworks like the NIST AI Risk Management Framework (RMF) [16]. While valuable, these are often disconnected from specific engineering approaches that translate requirements into verifiable technical properties [10]. Furthermore, practitioner guides for sectors like finance [8] and construction [21] often rely on single-domain case studies, frequently in healthcare [13, 24], which limits their explanatory power across diverse industrial contexts.

Systematic reviews confirm this empirical gap, highlighting the challenge of operationalizing theoretical principles [14]. Recent syntheses identify a scarcity of research on high-risk AI governance [11] and find that most existing solutions focus on organizational policy but rarely address the practical “Who, What, When, and How” of implementation [1]. Our study addresses this by providing an explanatory framework that links organizational context to four distinct governance archetypes and their resulting architectures.

## 3 Research Methodology

We employ a mixed-method qualitative design centered on an in-depth, embedded case study of McDermott, a major EPC company (Company 1) [22], supplemented by five comparative expert interviews to establish the boundaries of our framework [12] through analytical generalization [18].

### 3.1 Participant Selection and Characteristics

Our selection captures variation across risk profiles and organizational scales (Table 1). Company 1 (McDermott) is the primary case, and the lead author’s role as an embedded member of its AI team

enabled the in-depth participant observation. Companies 2–6 provide a comparative sample of diverse engineering and technology providers.

**Table 1: Participating Organizations’ Characteristics**

| Company       | Sector / Domain        | Size  | Primary AI Use Type           | Informant Role            |
|---------------|------------------------|-------|-------------------------------|---------------------------|
| 1 (McDermott) | EPC                    | Large | GenAI & Automated Engineering | Senior IT Specialist      |
| 2             | EPC                    | Large | Automated Engineering         | Software Product Manager  |
| 3             | EPC                    | Large | Automated Engineering         | Digital Director          |
| 4             | Tech Conglomerate      | Large | Enterprise GenAI              | R&T Manager               |
| 5             | Eng. Software Provider | SME   | Platform Automation           | Product Manager           |
| 6             | Energy Systems         | SME   | Industrial Control Systems    | Dev Team (6 participants) |

## 3.2 Data Collection and Analysis

Data for Company 1 (McDermott) was collected (March–July 2025) via participant observation in AI steering committees, analysis of internal strategy and policy documents, and semi-structured interviews. For Companies 2–6, we conducted 60-minute semi-structured anonymous interviews with senior leaders.

Analysis followed an iterative thematic approach [2] using a systematic coding process. We first employed **open coding** to identify emergent practices (e.g., “mandatory review gates”, “human-in-the-loop audit”, “off-the-shelf vs. proprietary”). This was followed by **axial coding** to group patterns into three primary coding categories: *Formality*, *Risk Strategy*, and *Technology Adoption*. To ensure rigor and credibility, we applied two key techniques. First, we used **data triangulation** within our primary case by comparing findings from observations, internal documents, and interviews. Second, throughout the analysis process, we performed informal **member checking** by communicating with interviewees to clarify points of ambiguity and ensure our interpretations were accurate.

Our analysis is based on existing literature by structuring our risk assessment around three dimensions derived from established enterprise governance and regulatory logic [5, 19]:

- **Domain Risk:** Inherent risk of the business sector (e.g., critical infrastructure vs. software services).
- **Application Risk:** Determined by the system’s function and potential for harm, aligning with the EU AI Act’s distinction between high-risk and accessory functions as defined in [4].
- **Organizational Risk:** Pressures related to scale, reputation, and standardization common in large enterprises [7].

## 4 The Four Risk-Driven Archetypes of AI Governance

Our analysis of the six organizations reveals four distinct archetypes of AI governance, each emerging predictably from an organization’s specific risk profile. These archetypes should be understood not as static “types” but as dynamic “states” because an organization may transition between them as its risk posture, regulatory environment, or maturity evolves. Furthermore, while presented as distinct for clarity, a single large organization may exhibit mixed characteristics, applying different archetypes to different organizational divisions based on their specific risk context. The four archetypes are themselves a synthesis of three central themes identified in our data: Governance Formality, Risk Management Strategy, and Technology Adoption Approach.

#### 4.1 Archetype 1: High-Assurance Governance

This archetype is characteristic of any organization deploying an application with **High Application Risk**, where errors can have catastrophic or mission-critical consequences. It relies on strict documentation, continuous internal audits, and the development of a formal “responsible AI” framework with pillars for accountability and transparency.

**Governance Formality:** The framework is centralized, formal, and non-negotiable. As observed in our primary case study, McDermott made a deliberate shift to consolidate control, creating a single, authoritative “AI solution team” to “oversee all the AI use cases and projects across the company.” This central team, along with a formal “steering committee” and “AI Council,” forms the undisputed backbone of their framework. Governance is executed through mandatory, rigid IT tollgates where certain risk levels require “audit, AI Council approval.” Company 2 employs a similar multi-stage approval process, beginning with a mandatory business case gate where “The first question they ask is what problem are you solving?” followed by a formal risk and impact analysis.

**Risk Management Strategy:** The strategy is proactive and quantitative, designed to produce auditable evidence of compliance. McDermott’s internal documentation revealed a sophisticated diagnostic tool that moves beyond subjective assessment to a classified system, evaluating projects across six critical dimensions: **Data Sensitivity, Access, Capabilities (Autonomy), Output Quality & Accuracy, Legal, and Fairness**. This rigor is a direct response to unforgiving domain requirements where verification is paramount: “the AI solution needs to be verified by various standardization bodies. Otherwise, the clients won’t trust it.” The framework is explicitly aligned with emerging regulations, with checklist questions addressing accountability, oversight, and documentation. This directly creates the demand for governance artifacts like assurance logs and audit trails.

**Technology Adoption and Architectural Consequences:** The stringent governance directly mandates a constrained, verifiable, and human-centric architecture. This is not merely a preference but a required outcome of the risk management strategy.

- **Hybrid and Constrained Systems:** The primary architectural feature is the use of hybrid systems where the non-deterministic power of AI is explicitly constrained by deterministic logic. Development at Company 2 is confined to a secure “Azure bubble,” and prototypes favor linking an LLM to a formal ontology or using rule-based validation to ensure engineering outputs are predictable and trustworthy.
- **Mandatory Human-in-the-Loop:** The architecture must include formally designed and auditable points for human intervention. McDermott’s internal framework moves HITL from an informal concept to a required architectural component, prescribing specific patterns based on assessed autonomy levels, such as requiring an explicit approval gate where a “System acts; but human explicitly approves first.” This shows how governance rules are translated directly into non-negotiable components of the system design.

#### 4.2 Archetype 2: Cautious Governance

This archetype emerges in organizations that manage their **High Domain Risk** by strategically pursuing only applications with **Low Application Risk**. Governance is designed to prevent risk contagion by strictly limiting AI to accessory functions. This allows the organization to experiment with AI without incurring the full compliance overhead of high-risk systems, though it presents a potential scaling challenge, as more autonomous solutions would trigger a shift in governance.

**Governance Formality:** In stark contrast to the High-Assurance model, Company 3’s governance is less formal and more iterative. They manage projects through lightweight, time-boxed reviews focused on empirical results over extensive documentation: “We will inspect after 6 weeks and 12 weeks and look at what you’ve been able to solve, and we’ll use that as a real-world litmus test to... pause that idea for now and try something else.”

**Risk Management Strategy:** The core risk strategy is one of deliberate containment. This is achieved through a “fast follower” technology posture—“We only want to be involved with technology when it’s on that slope of enlightenment”—and a strict focus on low-impact applications. Company 3 manages its high-stakes environment as a government contractor, where they “have to be absolutely right,” by limiting AI to accessory functions like “fairly simple robotic automations” that do not influence critical decisions.

**Technology Adoption and Architectural Consequences:** This risk-averse strategy fosters an architecture of isolated, low-risk solutions. The landscape is dominated by **standalone, off-the-shelf tools**, with a strategic **avoidance of deep integration** with critical systems. This choice is strongly informed by a past failure with an ambitious proprietary LLM that “did not go well,” which now acts as a powerful architectural constraint against using generative AI for safety-critical tasks.

#### 4.3 Archetype 3: Enterprise-Driven Governance

This archetype is characteristic of organizations with **Low Application and Domain Risk** but **High Organizational Risk**, where governance is not motivated by the immediate correctness of applications but by the need for enterprise-wide standardization, control, and predictability. It includes formal stages, a central unit that vets all models, and a mandate for human accountability, which are hallmarks of a mature organization maintaining control across a complex enterprise.

**Governance Formality:** Company 4’s governance is defined by a formal, centralized process that consciously evolved from informal methods to a structured, gated process to gain control over spending and technology proliferation. Control is enforced by a central unit with absolute veto power over technology choices: “Internally we have a central unit approving the usage of... basically anything, such as pretrained models.”

**Risk Management Strategy:** Risk is managed through a universally applied corporate principle. A prime example is Company 4’s non-negotiable rule of human accountability, which functions as a core governance pillar that shapes all AI development: “We can never trust a model. We always have to have a human taking responsibility for the output of the model.”

**Technology Adoption and Architectural Consequences:** This governance posture funnels all development onto a **standardized, platform-based foundation**. This creates a pragmatic architecture where teams often prefer classical, deterministic algorithms over AI for reliability or resource constraints. As one informant noted, “We relatively often discover that developing classical algorithms might be better.” This platform model, while enabling scale, necessarily **constrains developer agility** and tooling freedom, creating a trade-off between enterprise control and local innovation.

#### 4.4 Archetype 4: Developer-Led Minimal Governance

Found in organizations with **Low Domain, Application, and Organizational Risk**, this archetype explicitly lacks the documented policies and overarching processes of formal governance. Its minimalist controls are proportionate to the risk profile, enabling fast experimentation with pilots launched directly with customers or through internal developer trials.

**Governance Formality:** Governance is explicitly informal and minimalist. Company 5’s informant stated directly, “We don’t have that... formally established,” relying instead on informal knowledge sharing. At Company 6, the entire framework was reduced to a single, simple guideline: “a warning against putting internal data on public AI,” with no formal committees or stage-gates.

**Risk Management Strategy:** The minimalist framework implicitly **transfers the burden of managing complex risks**—such as code quality, correctness, and long-term maintainability—from the organization to the individual developer. This creates systemic challenges, such as the risk of skill erosion, where a developer might “introduce complex code without fully understanding it.” At Company 5, risk is managed externally by product-market feedback and internally by pragmatic constraints, such as choosing not to train their own models because they “don’t have the capacity to outperform the established players.”

**Technology Adoption and Architectural Consequences:** The architecture is **decentralized and emergent**, shaped by the ad-hoc integration of third-party tools into existing workflows. With minimal built-in controls, the **human developer becomes the primary safeguard**. Interestingly, developers at Company 6 are also discovering the need for hybrid solutions through practical experience, recognizing that “Combining AI with some rule-based system would actually be the way forward.” Company 5 constrains its architecture through technical guardrails, creating a “safe sandbox” with **strict scope limitation** and no external web access.

#### 4.5 A Diagnostic Framework for AI Governance

To synthesize these findings into a practical tool and show the analytical move from data to archetypes, we present a diagnostic framework (Figure 1) and a combined summary of the archetypes (Table 2). This framework provides a logical path for an organization to identify its governance posture based on its risk profile.

### 5 Discussion

This paper presents the findings of a preliminary study on AI governance in practice. Based on an in-depth case study and supplemented by comparative interviews, our work identified four distinct

AI governance archetypes based on an organization’s risk profile. This finding empirically grounds the risk-based logic of regulations like the EU AI Act [5]. This section discusses the implications of our findings by connecting them to key literature gaps.

Our primary contribution as a preliminary study is to provide initial empirical evidence that recent literature reviews have identified as critically missing. While scholars have mapped the landscape of published academic solutions, they have confirmed that our understanding of how high-level principles are operationalized within organizations is limited [14]. Our study, combining a deep case with cross-industry interviews, addresses this gap, moving beyond prescription to explain what is actually practiced across diverse contexts. We show that governance is not a monolithic concept but a spectrum of distinct, context-driven archetypes.

#### 5.1 Implications for Theory and Practice

**Extending Governance Models.** Our typology complements established governance models like the NIST AI Risk Management Framework (RMF) [16] and ISO 42001 [9]. These standards provide a comprehensive set of prescriptive functions for a robust AI management system. Our archetypes, in contrast, provide an explanatory framework that connects an organization’s risk context to its capacity and motivation for implementing those functions. Our findings challenge a linear view of governance maturity, suggesting instead that effective governance is about achieving a state that is *appropriate* for the organization’s risk profile.

**Explaining Organizational Readiness for Compliance.** Our archetypes provide crucial organizational context for technical compliance. The formal processes required by normative guides [8, 10] are a natural fit for a High-Assurance archetype but would likely fail in a Developer-Led one, which lacks the necessary formal structures. This highlights that technical compliance is an organizational challenge, and our typology can help assess an organization’s readiness for such frameworks.

**Dynamic States and Parallels to Technology Adoption.** The archetypes are best understood as dynamic “states,” not static “types.” An organization may transition between them, for instance, moving from Developer-Led to Enterprise-Driven as it scales, or a division may adopt a High-Assurance posture when a product enters a safety-critical market. This dynamic view aligns our findings with established literature on technology adoption. The Enterprise-Driven archetype, for instance, mirrors traditional, centralized IT governance designed to manage costs and risks at scale, while the Developer-Led archetype is analogous to the rise of bottom-up adoption of new tools and open-source software.

**Nuancing the Governance-Architecture Link.** A key nuance is that a single, large organization can exhibit multiple archetypes simultaneously. For instance, a conglomerate might apply High-Assurance governance to a safety-critical division but use a Developer-Led approach for internal tools. This shows that an organization may apply different governance archetypes to different systems based on their specific risk contexts.

**Making the Governance-Architecture Link Concrete.** A key contribution is linking these organizational archetypes to concrete software engineering outcomes. This link is most evident in two cross-archetype architectural trends: the normalization of **hybrid**




Figure 1: A Risk-based Diagnostic Framework for Identifying AI Governance Archetypes.

Table 2: Summary of AI Governance Archetypes: Risk Profiles, Mechanisms, and Architectural Consequences

| Archetype (Companies)                    | Risk Profile (based on Diagnostic Framework)                          | Key Governance Mechanisms & Artifacts                                                                                         | Resulting Architectural Patterns                                                                                         |
|------------------------------------------|-----------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| <b>High-Assurance</b><br>(McDermott, C2) | High Application Risk,<br>Any Domain/Organizational Risk              | Central steering committee; mandatory, auditable stage-gates;<br>formal risk checklists; detailed documentation requirements. | Mandated hybrid (probabilistic/deterministic) systems;<br>formal, auditable Human-in-the-Loop (HITL) components.         |
| <b>Cautious</b><br>(C3)                  | Low Application Risk,<br>High Domain Risk,<br>Any Organizational Risk | Iterative, time-boxed reviews; problem-driven prioritization;<br>strategic avoidance of high-risk applications.               | Dominated by isolated, off-the-shelf tools;<br>strategic avoidance of deep system integration.                           |
| <b>Enterprise-Driven</b><br>(C4)         | Low Application & Domain Risk,<br>High Organizational Risk            | Centralized technology vetting unit; enterprise-wide policies<br>(e.g., mandatory human accountability).                      | Standardized, centrally managed platform foundation;<br>pragmatic use of classical algorithms over AI where appropriate. |
| <b>Developer-Led</b><br>(C5, C6)         | Low Application, Domain,<br>& Organizational Risk                     | Minimalist rules (e.g., IP protection); informal knowledge sharing;<br>no formal committees or stage-gates.                   | Decentralized, emergent architecture; ad-hoc integration<br>of third-party tools; human developer as primary safeguard.  |

systems (probabilistic/deterministic) for reliability, and the evolution of **Human-in-the-Loop (HITL)** from an informal concept into a formal, auditable component in high-risk settings.

## 5.2 Limitations

The findings of our early-stage research have several limitations. Most notably, our archetypes are derived from a sample of only six organizations, offering **analytical generalization** rather than statistical generalization [18]. The analysis is also interpretive, influenced by the varying data depth between our in-depth primary case and the supplementary interviews. This introduces potential bias. We sought to mitigate this through data triangulation in the primary case and a systematic coding process grounded in established literature. Finally, our study provides a snapshot in time of a rapidly evolving field, and the archetypes observed will likely adapt as new technologies and regulations emerge.

## 5.3 Future Work

As an early-stage study, this work provides a foundational framework that requires further validation. Future work should first validate this framework with a larger, more diverse sample of organizations to test its generalizability and refine the archetype boundaries. Second, longitudinal studies are needed to explore how organizations transition between archetypes, which would clarify how governance and technical architectures co-evolve. Finally, future work could investigate the technical artifacts for enacting governance in each archetype, such as the formal assurance logs and monitoring systems for High-Assurance settings and the lightweight security guardrails and developer tooling sufficient

for Developer-Led environments. This would strengthen the link between organizational policy and software engineering practice.

## 6 Conclusion

This study was motivated by the critical gap between high-level AI principles and the real-world organizational practices needed to implement them, a challenge amplified by risk-based regulations like the EU AI Act. We identified and defined four distinct AI governance archetypes: High-Assurance, Cautious, Enterprise-Driven, and Developer-Led Minimal Governance. Our primary contribution is demonstrating that these archetypes are not arbitrary but emerge predictably from an organization's specific risk profile, defined by the interplay of application, domain, and organizational scale. In doing so, this paper provides an initial explanatory tool for understanding how and why AI governance varies in practice. This typology, along with its identified link to technical architecture, offers a structured lens for practitioners and scholars to bridge the crucial gap between high-level policy and actual implementation.

## Acknowledgments

This work is supported by McDermott, Software Center (Gothenburg, Sweden), and Eindhoven University of Technology.

## References

- [1] Amna Batool, Didar Zowghi, and Muneera Bano. 2025. AI governance: a systematic literature review. *AI and Ethics* (2025), 1–15.
- [2] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (2006), 77–101. doi:10.1191/1478088706qp0630a
- [3] Martin Ebers. 2025. Truly risk-based regulation of artificial intelligence how to implement the EU's AI Act. *European Journal of Risk Regulation* 16, 2 (2025),

- 684–703.
- [4] Arnoud Engelfriet. 2025. *AI and Algorithms: Mastering Legal and Ethical Compliance*. ICTRecht, Netherlands.
  - [5] European Union. 2024. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 on laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>. Official Journal of the European Union, L 168, 12 July 2024, pp. 1–254.
  - [6] Luciano Floridi and Josh Cowls. 2022. *A Unified Framework of Five Principles for AI in Society*. John Wiley & Sons, Ltd, Chapter 22, 535–545. doi:10.1002/9781119815075.ch45
  - [7] Charles J Fombrun and Cees BM Van Riel. 2004. *Fame & fortune: How successful companies build winning reputations*. FT press.
  - [8] Sergio Genovesi. 2025. Introducing an AI Governance Framework in Financial Organizations. Best Practices in Implementing the EU AI Act (Practitioner Track). In *Symposium on Scaling AI Assessments (SAIA 2024)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 9–1.
  - [9] International Organization for Standardization. 2023. Information technology – Artificial intelligence – Management system. <https://www.iso.org/standard/81230.html>
  - [10] Jessica Kelly, Shanza Ali Zafar, Lena Heidemann, João-Vitor Zacchi, Delfina Espinoza, and Núria Mata. 2024. Navigating the EU AI Act: A Methodological Approach to Compliance for Safety-critical Products. In *2024 IEEE Conference on Artificial Intelligence (CAI)*. 979–984. doi:10.1109/CAI59869.2024.00179
  - [11] Byeong-Je Kim, Seunghoo Jeong, Bong-Kyung Cho, and Ji-Bum Chung. 2025. AI Governance in the Context of the EU AI Act. *IEEE Access* 13 (2025), 144126–144142. doi:10.1109/ACCESS.2025.3598023
  - [12] Steinar Kvale. 2009. *Interviews: Learning the craft of qualitative research interviewing*. Sage.
  - [13] Frank Liao, Sabrina Adelaine, Majid Afshar, and Brian W Patterson. 2022. Governance of clinical AI applications to facilitate safe and equitable deployment in a large health system: key elements and early successes. *Frontiers in digital health* 4 (2022), 931439.
  - [14] Matti Mäntymäki, Matti Minkkinen, Teemu Birkstedt, and Mika Viljanen. 2022. Defining organizational AI governance. *AI and Ethics* 2, 4 (2022), 603–609.
  - [15] Jessica Morley, Luciano Floridi, Libby Kinsey, and Anat Elhalal. 2020. From what to how: an initial review of publicly available AI ethics tools, methods and research to translate principles into practices. *Science and engineering ethics* 26, 4 (2020), 2141–2168.
  - [16] National Institute of Standards and Technology. 2023. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. Technical Report NIST AI 100-1. U.S. Department of Commerce, National Institute of Standards and Technology. doi:10.6028/NIST.AI.100-1
  - [17] Emmanouil Papagiannidis, Patrick Mikalef, and Kieran Conboy. 2025. Responsible artificial intelligence governance: A review and research framework. *The Journal of Strategic Information Systems* 34, 2 (2025), 101885.
  - [18] Geoff Walsham. 1995. Interpretive case studies in IS research: nature and method. *European Journal of information systems* 4, 2 (1995), 74–81.
  - [19] Peter Weill and Jeanne W Ross. 2004. *IT governance: How top performers manage IT decision rights for superior results*. Harvard Business Press.
  - [20] World Economic Forum. 2023. *The Future of Jobs Report 2023*. Technical Report. World Economic Forum. <https://www.weforum.org/publications/the-future-of-jobs-report-2023/>
  - [21] Lichao Yang, Gavin Allen, Zichao Zhang, and Yifan Zhao. 2024. Achieving on-site trustworthy AI implementation in the construction industry: A framework across the AI lifecycle. *Buildings* 15, 1 (2024), 21.
  - [22] Robert K Yin. 2018. *Case study research and applications*. Vol. 6. Sage Thousand Oaks, CA.
  - [23] Esmat Zaidan and Imad Antoine Ibrahim. 2024. AI governance in a complex and rapidly changing regulatory landscape: A global perspective. *Humanities and Social Sciences Communications* 11, 1 (2024).
  - [24] Jie Zhang and Zong-ming Zhang. 2023. Ethics and governance of trustworthy medical artificial intelligence. *BMC medical informatics and decision making* 23, 1 (2023), 7.