



Beyond control? Governing AI in organizations for responsible use

Simon Sturm^{a,*}, Florian Krause^{a,b}, Benjamin van Giffen^c

^a University of St.Gallen, Switzerland

^b Leibniz University of Hanover, Germany

^c University of Liechtenstein, Liechtenstein

ARTICLE INFO

Keywords:

Artificial Intelligence
AI Governance
Responsible Use of AI
Structuration Theory
Ethics
Case Study

ABSTRACT

Ensuring the responsible use of artificial intelligence (AI) has become a central concern for researchers, business leaders, and policymakers. There is a broad consensus that governing AI at the organizational level is crucial. While the existing literature acknowledges that responsible AI use depends on institutional structures and human agency, it is divided over which factor is decisive. We take an integrative perspective and examine how structure and agency interact when organizations govern AI for responsible use. To this end, we enrich Giddens' structuration theory with ethics (i.e., deontology, consequentialism, virtue ethics) and conceptualize AI governance as reflexive organizational praxis. We illustrate our analytical framework with the case of OTTO, Germany's largest e-commerce company. We conclude that while AI governance cannot define the responsible use of AI, organizations can scaffold how their members reason about whether and how AI should be used in a specific social context.

1. Introduction

Recent advances in machine learning-based artificial intelligence (AI) technologies have created major opportunities and risks for individuals, organizations, and societies (Davenport, 2018; Taddeo & Floridi, 2018); for example, the same machine learning models that help discover lifesaving therapeutics could be (mis)used to design deadly biological weapons (Urbina et al., 2022). Therefore, ensuring the responsible use of AI has become a central concern for researchers, business leaders, and policymakers (Floridi et al., 2018; Garibay et al., 2023; Jobin et al., 2019; Mikalef et al., 2022; Ruster & Daniell, 2025).

The general consensus is that responsible use requires governance of AI technologies at the organizational level (Birkstedt et al., 2023; Cihon et al., 2021; Mäntymäki et al., 2022; Mökander et al., 2022; Schuett, 2023). The organizational use of AI is shaped by institutional arrangements (*structure*) (Birkstedt et al., 2023; Schneider et al., 2023) and by the capacity of organization members to deliberate on decisions regarding whether and how AI is applied in a given situation (*agency*) (Griffin et al., 2024, 2025). Nonetheless, much of the AI governance literature implicitly assumes that structure and agency are distinct ontological entities, exemplified by the frequent use of gap metaphors (Mökander & Floridi, 2023; Morley et al., 2020; Schiff et al., 2021;

Shneiderman, 2020b). This implicit *dualism* has tended to divide the literature into camps that regard either structure (Black & Murray, 2019; Fukuda-Parr & Gibbons, 2021; Koniakou, 2023; Maclure & Morin-Martel, 2025; Schuett, 2023; Schuett et al., 2025) or agency (Bezuidenhout & Ratti, 2021; Constantinescu et al., 2021; Griffin et al., 2024; Hagendorff, 2020, 2022; Hayes et al., 2025; Ratti & Graves, 2021) as the *prima causa* of responsible AI use. If we accept that the behavior of organizations and their members is neither fully determined by institutional forces nor completely free from these forces (Walgenbach, 1999), as long as one perspective is privileged over the other, we will only partially capture how organizations govern AI for responsible use. The our study aims to address this issue, asking: *what roles do structure and agency play when organizations govern AI for responsible use?*

We address this research question by enriching Giddens' (1984) structuration theory with ethics (i.e., consequentialism, deontology, and virtue ethics) to conceptualize how the responsible use of AI emerges from the recursive interplay of structure and agency across social, organizational, and individual levels. Our conceptual work yields an analytical framework, which we illustrate with the case of OTTO, Germany's largest e-commerce company.

We identify three modalities (den Hond et al., 2012) through which organizations can govern AI for responsible use: not by defining what

This article is part of a special issue entitled: (Ir)responsible uses of technologies published in European Management Journal.

* Corresponding author.

E-mail addresses: simongabriel.sturm@unisg.ch (S. Sturm), florian.krause@unisg.ch (F. Krause), benjamin.vangiffen@uni.li (B. van Giffen).

<https://doi.org/10.1016/j.emj.2026.06.003>

Received 4 April 2025; Received in revised form 13 June 2026; Accepted 23 June 2026

Available online 26 June 2026

0263-2373/© 2026 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

responsible use of AI is (for we argue that organizations cannot do so), but rather by *scaffolding* how their members deliberate about whether and how AI should be used in a given context. We conclude that responsible AI governance is best understood and enacted as a *reflexive organizational praxis*.

Our findings support previous research highlighting the inherent complexity of governing AI for responsible use (Mittelstadt, 2019; Moats & Trifuljesko, 2025; Orr & Davis, 2020; Winfield & Jirotko, 2018). We add to the existing knowledge by explaining this complexity through the recursive interplay of structure and agency. Thus, we provide a theoretical foundation for future integrative AI governance research and offer actionable recommendations for practitioners in building effective AI governance systems.

The paper is organized as follows. We first provide the necessary background (Section 2), then develop our analytical framework (Section 3) and apply it to the case of OTTO for illustration (Section 4). Finally, we relate our research to existing knowledge and offer concluding remarks (Sections 5 and 6).

2. Theoretical background

2.1. Why organizational AI governance is needed and how to define it

In contemporary understanding, AI refers not to specific technologies but rather to “a continually evolving frontier of emerging computing capabilities” (Berente et al., 2021, p. 1433) that enables machines to act without user intervention and learn through data and experience, often in ways that humans cannot fully understand or explain (p. 1438). Recent advances include generative AI technologies that can produce text, images, audio, or code on demand (Feuerriegel et al., 2023) and agentic AI systems that can pursue complex goals with minimal user intervention (Acharya et al., 2025).

AI technologies offer organizations numerous opportunities to create value, but their use also creates new risks. A life-threatening example is the dual use of AI in drug discovery (Urbina et al., 2022); a less life-threatening example is Microsoft's chatbot Tay, which quickly became a platform for racist hate speech (Wolf et al., 2017). Ensuring the responsible use of AI thus has become a central concern for researchers, business leaders, and policymakers (Floridi et al., 2018; Garibay et al., 2023; Jobin et al., 2019; Mikalef et al., 2022; Ruster & Daniell, 2025).

Organizations play a pivotal role in deciding which AI technologies are developed and how they are used (Cihon et al., 2021). Consequently, there has been growing interest in understanding AI governance at the organizational level. The literature on organizational AI governance, an emerging field of research, remains largely conceptual and builds on insights from adjacent governance domains (Birkstedt et al., 2023; Mäntymäki et al., 2022). Much of this work focuses on defining and distinguishing AI governance from related forms of technology governance (Birkstedt et al., 2023).

Mäntymäki et al. (2022, p. 604) define AI governance as “a system of rules, practices, processes, and technological tools that are employed to ensure an organization's use of AI technologies aligns with the organization's strategies, objectives, and values; fulfills legal requirements; and meets principles of ethical AI followed by the organization.” This definition highlights the inherent complexity of governing AI at the organizational level: external expectations must be translated into internal structures and practices while reconciling different, and sometimes conflicting, objectives (Mäntymäki et al., 2022). Among these objectives, the responsible use of AI is particularly important (Mikalef et al., 2022; Taddeo & Floridi, 2018). Yet there is little agreement on how organizations can (and should) govern AI for responsible use.

2.2. The Structure–Agency controversy in current AI governance research

As indicated in the Introduction, the responsible use of AI depends on

both institutional arrangements (i.e., structure) and the capacity of organization members to deliberate about whether and how AI should be used in a given situation (i.e., agency) (Birkstedt et al., 2023; Griffin et al., 2024, 2025; Schneider et al., 2023). What remains contested, however, is their relative importance, which has tended to divide the literature on responsible AI governance into opposing camps.

The structure camp (Black & Murray, 2019; Fukuda-Parr & Gibbons, 2021; Koniakou, 2023; Maclure & Morin-Martel, 2025; Schuett, 2023; Schuett et al., 2025) views human agency as a potential threat to responsible AI use that must be disciplined through external constraints. Maclure and Morin-Martel (2025, p. 18), for example, argue that “focusing on individual virtue in AI development is, at best, insufficient and, at worst, integral to an ethics-washing strategy.” Consequently, the structure camp suggests that AI should be governed through oversight functions analogous to data protection officers (Birkstedt et al., 2023), institutionalized risk management frameworks such as the three lines of defense model (Schuett, 2023), and binding legal frameworks (Fukuda-Parr & Gibbons, 2021; Koniakou, 2023; Maclure & Morin-Martel, 2025).

The agency camp (Bezuidenhout & Ratti, 2021; Constantinescu et al., 2021; Griffin et al., 2024; Hagendorff, 2020, 2022; Hayes et al., 2025; Ratti & Graves, 2021) starts from the inverse assumption that responsible AI use emerges from the dispositions of individuals and cannot be secured through external constraints alone. Ratti and Graves (2021, p. 1822), for example, warn that formulating ever more specific rules leads to “mindless compliance” that neither provides “a robust ‘infrastructure’ for ethical action, nor ... leads to autonomous ethical agents.” Similarly, Hagendorff (2022, p. 4) argues that AI ethics needs to shift from an external “ethics assessment” to an internal “ethics for design.” Hence, the agency camp suggests that AI should be governed by cultivating professional dispositions that enable individuals to respond appropriately to the situation at hand (Bezuidenhout & Ratti, 2021; Hagendorff, 2020, 2022; Hayes et al., 2025; Ratti & Graves, 2021).

The ongoing debate between structure and agency has led to much of the literature adopting one-sided positions. However, both perspectives capture essential aspects of organizational reality but have characteristic blind spots in isolation. Structural approaches tend to treat organizations as static arrangements that determine behavior through external constraints, whereas agency-oriented approaches risk overemphasizing voluntarism and overlooking the enabling and constraining effects of institutional contexts (Walgenbach, 1999). For a deeper understanding of behavior by and within organizations, it is therefore necessary to take both structure and agency into account (Walgenbach, 1999).

The purpose of this study is not to resolve the debate between structure and agency in the current AI governance literature by privileging one perspective over the other. Instead, we aim to bridge them and challenge the underlying assumption that structure and agency are two distinct ontological entities. To this end, we build on Giddens' (1984) theory of structuration.

2.3. Structuration theory as an integrative perspective for AI governance research

Giddens developed structuration theory to reconcile “empire-building endeavours” between functionalist/structuralist and interpretive traditions in social science (Giddens, 1984, p. 2). He argued that the conflict surfaces epistemologically but is rooted in different beliefs about the nature of the social world. These beliefs tend to either reduce structure to an aggregated residue of individual actions or reduce human agency to the execution of structural imperatives (Giddens, 1984, pp. 1–2). Therefore, Giddens suggested that social sciences should focus on practices and resolve the ontological conflict of structure and agency in a *duality*: “structure is both medium and outcome of the reproduction of practices. Structure enters simultaneously in the constitution of the agent and social practices, and ‘exists’ in the generating moments of this constitution” (Giddens, 1979, p. 5).

Structuration theory proposes that the stratification of human agency lets social life unfold in “essential recursiveness” (Giddens, 1979, p. 5). Individuals are presumed to monitor their own actions and physical and social contexts (including the conduct of others), and to expect everyone else to do the same (i.e., reflexive monitoring of action; Giddens, 1984, p. 5). Individuals further maintain a “theoretical understanding” of the grounds of their actions and are expected to explain them if asked (i.e., rationalization of action; Giddens, 1984, p. 5). Actions themselves, however, are prompted by wants that individuals cannot necessarily report on (i.e., motivation of action; Giddens, 1984, p. 5). Hence, individuals only ever possess partial knowledge of the context in which they act (Giddens, 1984, p. 7). Therefore, purposive action produces unintended consequences that become conditions of future action (Giddens, 1984, p. 8).

Structuration theory thus provides a processual ontology of social reality (Fig. 1) that allows us to reframe central issues in the current AI governance discourse. From a structurationist standpoint, for example, there can be no “gap” between principles and practice that AI governance could close (Bach et al., 2025; Birkstedt et al., 2023; Schiff et al., 2021). After all, rules and norms exist only at the moment they are brought into existence through social practice (Giddens, 1984, p. 25). Gaps may, however, lie between the principles that organization members supposedly enact and those they actually draw upon.

Giddens's theory of structuration is a general theory of social conduct (den Hond et al., 2012; Walgenbach, 1999; Whittington, 2025) that requires specification to capture the nuances of governing AI for responsible use. We undertake this specification in the next section.

3. Conceptualizing responsible AI governance in organizations

3.1. Theorizing by multidimensional institutional analysis

If structure and agency are viewed as mutually constitutive facets, they cannot be distinguished within specific social practices. Therefore, applying structuration theory to AI governance requires us to separate analytically what cannot be separated phenomenologically. Giddens proposed two types of analytic bracketing that hold one facet of social practice constant to examine the other (Whittington, 2025). *Strategic conduct analysis* brackets institutional features to examine how knowledgeable actors enact rules and resources in situated action (Whittington, 2025), while *institutional analysis* brackets strategic conduct to foreground “rules and resources as chronically reproduced features of social systems” (Giddens, 1981, p. 28).

Since AI governance is concerned with creating structural conditions that enable the expected form of AI use (Mäntymäki et al., 2022), institutional analysis is the appropriate analytical orientation for this study. We first conduct a *vertical* institutional analysis to identify the *levels* across which the recursive interplay of structure and agency unfolds. Then we conduct a *horizontal* institutional analysis to specify the *modalities* (den Hond et al., 2012) through which structure and agency

bring each other into being when organizations govern AI for responsible use.

3.2. Vertical institutional analysis: levels in responsible AI governance

The recursive interplay of structure and agency is by nature a multilevel phenomenon (Giddens, 1984). The AI governance literature typically distinguishes three levels of analysis: society, organization, and individual (Birkstedt et al., 2023; Shneiderman, 2020a, 2020b).

The overarching societal concern is to ensure that AI technologies are used for the common good (Floridi et al., 2018; Shneiderman, 2020b); however, at the societal level, AI use can be monitored but not directly steered. Therefore, expectations regarding how AI is to be used are codified and assumed to be enacted through subordinate levels, most explicitly in the form of (supra)national regulation, such as the EU AI Act (Floridi et al., 2018; Tursunbayeva et al., 2025).

At the organizational level, regulatory requirements and societal expectations are translated into AI governance structures. Again, the organizational level has no immediate influence on the use of AI. Instead, AI governance structures constitute the rules and resources that enable or constrain the conditions of action at the individual level.

It follows that the recursive interplay of structure and agency is most visible at the individual level, where members of an organization determine, monitor, and adjust how AI is used in a given context. They do so by drawing on the conditions of action shaped by organizational AI governance structures. But because knowledge is always bounded (Giddens, 1984), whether and how AI is used in a given situation may alter the conditions of future action and thereby (re)produce governance structures. Instances of AI use accumulate to shape the organizational and ultimately the societal use of AI and close the recursive loop (Fig. 2).

3.3. Horizontal institutional analysis: modalities of responsible AI governance

We turn to the anthropologic premise at the heart of structuration theory to further clarify how structure and agency bring each other into being when organizations govern AI for responsible use. According to Giddens, “[t]o be a human being is to be a purposive agent, who both has reasons for his or her activities and is able, if asked, to elaborate discursively upon those reasons” (Giddens, 1984, p. 3). Whereas Giddens positions deliberation as a “routine characteristic of human conduct” (p. 4), he leaves open which reasons human agents can draw on to justify their conduct. Ethical theories offer clarity.

We must first clarify what an action is to understand how it can be justified. However, the notion of action and the distinction between action and behavior are philosophically contested (Piñeros Glasscock & Tenenbaum, 2023). Intuitively, we speak of an action if there is an *actor*, an *outcome*, and an *intention*. We typically do not speak of an action if any of these elements is missing. For example, we would not call plans that are never realized an action, as they lack both outcome and




Fig. 1. Theory of structuration.




Fig. 2. Levels in responsible AI governance.

intentional follow-through. Similarly, we do not typically refer to something as an action if the actor lacks intentional control, such as sleepwalking.

Just as the notion of action is contested, so too is the question of how it can be justified. Over the history of human thought, three major ethical theories have emerged: consequentialism, deontology, and virtue ethics (Gal et al., 2022). All theories recognize actor, outcome, and intention as constitutive elements of action but disagree regarding which of these elements determines its moral quality. Deontology suggests evaluating action by intention (e.g., Kant, 1797/1996), consequentialism by outcome (e.g., Bentham, 1996; Mill, 2009), and virtue ethics by personal disposition of the actor (e.g., Aristotle, 2007).

The fundamental disagreement among the “big three” (Gal et al., 2022) ethical theories regarding which action element holds primacy has both negative and positive implications for this study. The negative implication is that the responsible use of AI cannot be defined in *substance*. Consider cases in which AI technologies were used to track contacts during the COVID-19 pandemic with questionable user consent (Anshari et al., 2023). Deontologists like Kant (1797/1996) would reject such use, because a lack of consent risks treating humans solely as means. In contrast, consequentialists like Bentham (1996) and Mill (2009) would accept the same use of AI if the outcomes were beneficial (e.g., reduced infection rates and saved lives). Therefore, responsible use of AI is best understood as a context-bound interpretation that can always be challenged from alternative ethical perspectives.

The positive implication is that although responsible use of AI cannot be defined in substance, human reasoning about it can be defined in *form*. If we assume that humans reason about responsible AI use no differently from any other action, and that the three major ethical theories cover the basic forms by which action is justified, individuals will refer to intention, outcome, or actor disposition when they deliberate whether and how AI should be used in a specific social context.

We argue that it is in these basic forms that organizations can govern AI for responsible use, by scaffolding their members toward a specific type of normative reasoning. Each of the “big three” ethical theories comes with an implicit theory of action motivation that when mapped to Giddens’ stratification model (1984, pp. 5–7) can establish analytic connections between AI governance structures and human agency.

Deontologists like Kant (1797/1996) believe that humans possess a natural sense of duty that when engaged makes them want to align with overarching norms. Consequentialists like Bentham (1996) and Mill

(2009) believe that humans possess instrumental rationality that when engaged makes them act in ways that optimize the consequences of action. Finally, virtue ethicists like Aristotle (2007) believe that humans cultivate practical wisdom that when engaged allows them to make an appropriate decision in a given situation.

Enriching Giddens’s structuration theory with ethics results in three modalities through which organizations can govern AI for responsible use (Fig. 3). AI governance structures target a component of action, engage the corresponding motivation, and thereby scaffold how organization members rationalize and reflexively monitor the use of AI.

First, AI governance structures can target intention, engage the sense of duty, and scaffold individuals toward *compliant* use of AI in ways that respect overarching norms, such as regulatory requirements and organizational values, regardless of potential outcomes. Second, AI governance structures can target outcomes, engage instrumental rationality, and scaffold individuals toward *calculated* use of AI in ways that maximize positive and minimize negative consequences. Third, AI governance structures can target actor disposition, engage practical wisdom, and scaffold individuals toward *competent* use of AI in ways that reflect the judgment of virtuous practitioners.

Importantly, scaffolding does not determine normative reasoning; responsible use of AI cannot be defined in substance, so organization members must decide what counts as responsible for the situation at hand. That said, the orientation that each scaffold can provide is limited by the normative theory on which it draws. This is well documented in the literature. Governing AI use by intention inherits the deontological problem of duty conflicts (Mittelstadt, 2019; Seger, 2022; Shneiderman, 2020b; Whittlestone et al., 2019, pp. 195–200); governing AI use by outcome faces the consequentialist issue of outcome uncertainty (Mökander & Floridi, 2023); and governing AI use by actor disposition inherits the virtue-ethical problem of guiding specific action (Hagendorff, 2022; Hayes et al., 2025).

Therefore, individuals will reason beyond the scaffolds that prove to be insufficient and in doing so (re)produce AI governance structures that scaffold future reasoning. The three modalities thus constitute a bidirectional relationship through which the recursive interplay of organizational AI governance structures and human agency unfolds over time.

3.4. Responsible AI governance as reflexive organizational praxis

Combining vertical and horizontal institutional analysis yields our




Fig. 3. Modalities of responsible AI governance.

conceptualization of AI governance as a *reflexive organizational praxis*. Our analytical framework (Fig. 4) explains how responsible AI use emerges through the recursive interplay of structure and agency across societal, organizational, and individual levels.

At the societal level, (supra)national regulation codifies normative expectations for the acceptable use of AI and monitors its effects across society. The organizational level translates these expectations into AI governance structures that enable or constrain the conditions of action.

At the individual level, organization members interpret, monitor, and adjust whether and how AI should be used in a given situation. Their normative reasoning is influenced by AI governance structures that (1) target intention, outcome, or actor disposition; (2) engage sense of duty, instrumental rationality, or practical wisdom; and (3) scaffold compliant, calculated, or competent use of AI. Since responsible AI use cannot be defined in substance, and the orientation that each scaffold can provide is inherently limited, organization members draw on the available scaffolds but are not confined to them. In doing so, they reflexively (re)produce the AI governance structures that scaffold future conduct. The recursive interplay of structure and agency becomes most visible in moments when the “rules, practices, processes, and technological tools” (Mäntymäki et al., 2022, p. 604) that compose an AI governance system are adapted. The resulting instances of responsible AI use accumulate to the organizational and, ultimately, societal use of AI (Fig. 4).

A brief note on the epistemological status of the proposed framework is warranted. Since it builds on Giddens's processual ontology of the social world, our framework is not a model in the representationalist sense (i.e., a “map” of AI governance structures that exist independent of situated practice). Rather, it is an analytical device designed to foreground “chronically reproduced features” (Giddens, 1981, p. 28) of organizational AI governance as a social system. Thus, it should be evaluated not by its correspondence to a fixed social reality but rather by its coherence and practical usefulness. We therefore invite readers to consider whether our framework helps illuminate what otherwise would remain opaque, and whether it does so without generating obvious absurdities (Polanyi, 1962).

4. Illustrative case study

We adopt Yin's (1994) case study research method to illustrate our analytical framework in a real-world context. A single case study design is appropriate because the recursive interplay of structure and agency requires close attention to contextual dynamics (den Hond et al., 2012).

4.1. Case selection, data collection, and data analysis

We selected the case of OTTO, Germany's largest e-commerce company, for illustration. The company offers a rich context in which examine how organizations govern AI for responsible use. AI is crucial to OTTO's business success; the company has been developing AI for more than a decade and actively positions itself as using AI responsibly (Table 1).

Research access was secured through a multiyear research

Table 1

Characteristics of OTTO and their relevance for empirical illustration.

#	Characteristic	Relevance for Empirical Illustration
1	Has been developing AI for more than 10 years (Christophersen & Pärn, 2021)	Increases the likelihood that AI governance practices have matured and can be observed
2	Employs over 6000 employees in Germany (Otto, 2026)	Indicates that the organizational setting is sufficiently complex
3	AI is central to competitive position (Christophersen & Pärn, 2021; Müller-Wünsch et al., 2025)	Increases the likelihood that AI governance is a strategic concern and not treated as a peripheral issue
4	Positions itself as using AI responsibly (Otto, 2021)	Ensures that responsible AI use is an explicit organizational concern

collaboration that allowed us to collect data from interviews, observation, and company documents (Table 2). We triangulated complementary perspectives (Cañibano et al., 2025) on AI governance at OTTO through multiple data sources. Interviews provided insights into the perceptions, challenges, and expectations of key informants; observation allowed us to contrast stated intentions with enacted conduct; and internal documents and archival records provided historical and strategic context.

We conducted 12 semistructured interviews with key informants. We selected the interviewees by balancing background, function, and seniority. All interviews were conducted in German, recorded with explicit consent, and lasted between 30 and 78 minutes. The Whisper speech recognition system (Radford et al., 2022) was used for transcription. The research team reviewed all transcripts and corrected them

Table 2

Case data by type.

Data Type	Σ	Duration
Interviews	12	605 min
(ID-01) Chief AI Officer	1	50 min
(ID-02) Senior Expert, AI Management	1	78 min
(ID-03) Senior Human Resources Manager	1	55 min
(ID-04) Business Innovation Manager	1	65 min
(ID-05) Digital Responsibility Manager	1	34 min
(ID-06) Department Manager, Business Intelligence	1	45 min
(ID-07) Senior Expert, Business Analytics	1	47 min
(ID-08) GenAI Expert & Product Manager	1	45 min
(ID-09) Senior Business Intelligence Product Manager	1	30 min
(ID-10) Business Intelligence Analyst	1	65 min
(ID-11) Information Security Officer	1	30 min
(ID-12) Digital Content Management Expert	1	61 min
Observation	9	555 min
Meetings	8	465 min
Workshops	1	90 min
Documents & Archival Records	22	310 min
Audio Transcripts (Podcasts/Videos)	10	310 min
Documents, Websites & Practitioner Publications	14	N/A
Total Contact Time (Interviews + Observation)		19 h, 20 min




Fig. 4. AI governance as reflexive organizational praxis.

as needed.

We observed eight internal meetings and one AI governance workshop and collected data through field notes. Again, participant consent was explicitly obtained. We also collected data from internal materials, such as company presentations, policy documents, and digital whiteboards. We also reviewed publications by OTTO practitioners (Christophersen & Pärn, 2021; Müller-Wünsch et al., 2025) and archival records, such as podcasts and videos. We focused on archival records in which OTTO executives addressed issues related to governance and responsible AI use; these records were retrieved from public sources such as YouTube.

After collection, we stored all data in a shared ATLAS.ti database and analyzed it in four steps. First, we marked all data relating to AI governance and responsible AI use, resulting in more than 290 excerpts. Second, we deductively coded the excerpts using the core constructs of our analytical framework and combined case evidence from multiple sources (Yin, 1981, 1994). Third, we captured and abstracted additional insights from the empirical data through inductive coding. Finally, we integrated the results from deductive and inductive coding and drew on our case knowledge for interpretation. Table 3 illustrates our data analysis procedure.

For readability, we described the data analysis as a linear process; in practice, however, this sensemaking was iterative. We repeatedly revisited and (re)coded the data, derived new insights, and refined earlier insights (Cañibano et al., 2025). In line with the illustrative purpose of the case study, we continued our analysis until all components of our analytical framework were substantiated by converging evidence from multiple data sources (Cañibano et al., 2025). Specifically, all authors agreed that explaining AI governance at OTTO as reflexive organizational praxis was free from obvious absurdities (Polanyi, 1962), and that additional data no longer altered our interpretation. Finally, we discussed our findings with OTTO practitioners for external validation.

4.2. Introducing OTTO, Germany's largest e-commerce company

OTTO is a family-owned e-commerce company based in Hamburg, Germany. Werner Otto founded the company in 1949 as a catalog-based mail-order business. Since then, OTTO has undergone a digital transformation. Today, the core of its operations is otto.de, with the company acting as both a platform operator and a retailer. The platform serves 12.6 million active users and offers more than 19 million products from OTTO and more than 6100 partners (Otto, 2026). This makes OTTO the largest player in the German e-commerce market after Amazon, with a gross revenue of 7 billion euros.

OTTO has been using AI technologies for longer than a decade, and AI has been crucial to the company's digital platform transformation. Currently, OTTO employs more than 100 experts who develop AI solutions in-house or in collaboration with partners such as Google and Microsoft. OTTO uses AI, for example, to interpret customer search queries and optimize

warehouse logistics, and it has launched its own version of ChatGPT to increase employee productivity. The CIO predicts that up to 50 percent of OTTO's technology portfolio will include AI components by 2030 (Müller-Wünsch et al., 2025).

4.3. The need for AI governance at OTTO

Given the centrality of AI to its operations, a top priority for OTTO is to ensure that AI is used in ways that align with both business needs and social expectations. As the CIO stated:

We are very much determined by the values of our owners. They have given us a Code of Ethics ... One sub-category ... is the responsible use of technology. ... We are already making extensive use of AI and ... are giving this a lot of thought. (Otto, 2021).

In recent years, the need for AI governance at OTTO has become

increasingly urgent because of emerging (supra)national regulations, such as the EU AI Act. Yet the company's ambition reaches beyond regulatory compliance. Leadership seeks to position OTTO as:

the trustworthy platform that does more than just comply with the law. We want to make the idea that Europe stands for a tangible system of values. We want to demonstrate that technology, when used responsibly, can bring real positive benefits. (Otto, 2021).

Accordingly, OTTO has invested in building appropriate governance structures so that it is not a matter of chance that developers deal with them [issues of responsible AI use], but that they are dealt with systematically (ID-02). In what follows, we apply our analytical framework (Fig. 5) to explain how OTTO governs AI for responsible use.

4.4. How OTTO governs AI use by intention: responsible AI guidelines

OTTO's leadership outlined the company's AI governance approach in a recent press conference. The first goal is to align intention behind AI use with legal aspects to consider. But there are also values that our company stands for. And it is very important to us to reflect those as well (Otto, 2024).

Accordingly, OTTO Group's central strategy consulting and technology unit has specified the normative expectations in responsible AI guidelines (one.O, 2026). The guidelines stipulate, for example, that human judgment takes precedence over machine decisions, that machine-generated content is proactively disclosed as such, and that AI will not be used to exploit vulnerabilities.

The responsible AI guidelines are based on the deontological idea that certain forms of AI use are (ir)responsible per se and can be expressed as general rules to follow (ID-06). Once formulated, these rules engage an inner sense of duty and scaffold individuals to use AI accordingly, irrespective of the outcome (i.e., compliant use of AI). This means that if a customer said, "I want an AI model that acts in a racist manner," then we will not work on that project. Period (ID-04).

The responsible AI guidelines provide clear normative directives for individuals to draw on:

Everyone is glad that the [responsible AI] guidelines exist. ... We had ... colleagues who approached us and said that they thought it was great that we were thinking about such topics... That even made them want to join us in the first place. (ID-04)

However, individuals found abstract principles like "trust through transparency" (one.O, 2026) difficult to operationalize: *If you just write that down and say we can always explain how an AI makes a decision, then everyone ... knows that this is not possible in many cases. (ID-06).* Furthermore, different principles of responsible AI guidelines may conflict, and additional complications may arise when collaborating with AI vendors that bring their own normative understandings to the table. For example, when OTTO Group's central technology unit developed an AI shopping assistant,

[t]hings like clothes and furniture worked, but sex toys didn't because the AI said that they didn't meet the safety level and filtered them all out. We then talked to [the AI vendor], and they said it was an American model that needed to be adapted. (ID-04)

The responsible AI guidelines thus provide valuable structural resources to draw on, but they do not confine human agency. As one interviewee put it:

They are not rules that you can follow slavishly. You always have to evaluate ... do they make sense or not? And if they don't make sense, you may have to violate them, but you also have to explain why. (ID-06)

Should the responsible AI guidelines repeatedly fail to provide adequate guidance in organizational practice, they may be adapted:

We have excellent specialists ... who ... would give us, the management, a heads-up and tell us that there is a new problem or a new solution. ... Then we would adjust the AI guidelines if necessary. That has not been the case yet, so everything seems to be okay for the time being. (ID-06)

At moments like these, the recursive interplay of the AI governance structure and human agency becomes most visible.

Table 3
Examples of case data analysis.

Source	Data Excerpt	Deductive Code(s)	Inductive Code(s)	Interpretation
Interview (ID-04)	<i>The [AI] guidelines apply to all developers, they apply to us as a company ..., and they are also relevant to our customers for the projects we carry out with them. This means that if a customer says, “No, I don't want that, I want an AI model that acts in a racist manner,” then we will not work on that project. Period.</i>	Component of Action: Intention Motivation of Action: Sense of Duty Rationalization & Reflexive Monitoring of Action: Compliant Use of AI	Scaffold: Responsible AI Guidelines Applicability of Responsible AI Guidelines	Responsible AI guidelines scaffold compliant use of AI at OTTO. They specify normative expectations as general rules that engage individuals' sense of duty, irrespective of outcome. While they provide clear orientation to draw upon, abstract principles can be difficult to operationalize.
Interview (ID-06)	<i>They are guidelines, not laws. There are issues such as data protection and resource consumption that need to be contextualized. ... They are not rules that you can follow slavishly. You always have to evaluate do they make sense or not? And if they don't make sense, you may have to violate them, but you also have to explain why. ... Ultimately, it is humans who make the decision and bear the responsibility for it.</i>	Component of Action: Intention Motivation of Action: Sense of Duty Rationalization & Reflexive Monitoring of Action: Compliant Use of AI	Scaffold: Responsible AI Guidelines Limitation(s) of Responsible AI Guidelines	Should they repeatedly fail to provide adequate guidance, the guidelines may be adapted over time.
Interview (ID-06)	<i>We have excellent specialists in our company who think about these things and would give us, the management, a heads-up and tell us that there is a new problem or a new solution. We rely on our specialists to raise their hands and let us know. Then we would adjust the AI guidelines if necessary. That has not been the case yet, so everything seems to be okay for the time being.</i>	Component of Action: Intention Motivation of Action: Sense of Duty Rationalization & Reflexive Monitoring of Action: Compliant Use of AI	Scaffold: Responsible AI Guidelines (Future) Adaptation of Responsible AI Guidelines	
Interview (ID-10)	<i>In our utility analysis for AI use cases, there are some points that are difficult to quantify and perhaps cannot be determined directly. It could also be, of course, that the assumed values in the utility analysis will change in the future. ... So there is a certain risk involved.</i>	Component of Action: Outcome Motivation of Action: Instrumental Rationality Rationalization & Reflexive Monitoring of Action: Calculated Use of AI	Scaffold: Utility Analysis for AI Use Cases Limitation(s) of Utility Analysis for AI Use Cases	A utility analysis for AI use cases scaffolds calculated use of AI at OTTO. It engages decision-makers' instrumental rationality by weighing potential outcomes into a single score. While it brings transparency to AI investment decisions, it cannot remove outcome uncertainty and depends on contextual judgment. The utility analysis may be adapted when it no longer fits the business context.
Interview (ID-02)	<i>We have seen that it is difficult to define the evaluation criteria [for the utility analysis of AI use cases]. How many points does a criterion receive for a specific case? ... What does economic benefit mean? Is it a million more in annual revenue? Two million? Ten million? ... We need to provide general guidance, otherwise the evaluation process will be very difficult. This is where we see the tool reaching its limits. The advantage is that it gets us talking.</i>	Component of Action: Outcome Motivation of Action: Instrumental Rationality Rationalization & Reflexive Monitoring of Action: Calculated Use of AI	Scaffold: Utility Analysis for AI Use Cases Limitation(s) of Utility Analysis for AI Use Case	
Observation [Company Workshop]	<i>During a workshop, we observed practitioners applying OTTO's utility analysis to evaluate and select AI use cases. However, the predefined scale (s) for assessing business value did not allow meaningful differentiation between AI use cases. The participants adapted the scale (s) to better reflect the specific business context.</i>	Component of Action: Outcome Motivation of Action: Instrumental Rationality Rationalization & Reflexive Monitoring of Action: Calculated Use of AI	Scaffold: Utility Analysis for AI Use Cases (Future) Adaptation of Utility Analysis for AI Use Cases	
Archival Record (Otto, 2024) [Press Conference]	<i>The second is internal knowledge building and professional education. We all need to keep learning. I've mentioned this several times already. And here, for example, with our internal learning tool, Masterplan, we already have two AI learning paths where you can really learn the basics of AI, where you can watch the videos and view the training courses to learn.</i>	Component of Action: Actor Disposition Motivation of Action: Practical Wisdom Rationalization & Reflexive Monitoring of Action: Competent Use of AI	Scaffold: AI Trainings for Employees Focus of AI Trainings for Employees	AI trainings for employees scaffold competent use of AI at OTTO. They cultivate practical wisdom by combining knowledge transfer with hands-on exercises. The mindset focus makes the AI trainings adaptable, but they offer little guidance for daily conduct. OTTO evaluates whether the AI trainings still help to cultivate practical wisdom and may adapt them in the future.
Interview (ID-05)	<i>Masterplan is also an important building block. However, I believe that we need to go into much more detail, especially when it comes to responsible AI. What is on Masterplan is not necessarily sufficient. There are courses on ethics and AI, but these tend to be at a meta level and not so application-oriented that a data scientist can look at them and immediately know what to do.</i>	Component of Action: Actor Disposition Motivation of Action: Practical Wisdom Rationalization & Reflexive Monitoring of Action: Competent Use of AI	Scaffold: AI Trainings for Employees Limitation(s) of AI Trainings for Employees	
Interview (ID-03)	<i>Let's see if, in two or three years' time, we'll say, “Okay, we need to take another look [at the learning paths].” If we actually roll out something like Microsoft Copilot at a large scale, I imagine that something new could come up. But that's just a hypothesis for now.</i>	Component of Action: Actor Disposition Motivation of Action: Practical Wisdom Rationalization & Reflexive Monitoring of Action: Competent Use of AI	Scaffold: AI Trainings for Employees (Future) Adaptation of AI Trainings for Employees	




Fig. 5. AI governance at OTTO in the analytical framework.

4.5. How OTTO governs AI use by outcome: utility analysis for AI use cases

OTTO's AI governance is also concerned with the outcome of AI use: *We don't want to go around with a watering can and sprinkle AI all over the place. We look very carefully at where it makes sense to use AI ... to get the most out of it.* (Otto, 2024).

To this end, OTTO has developed a utility analysis instrument for AI use cases. Originally, the instrument was developed to decide which generative AI use cases proposed by consultancies and technology partners should be implemented. The utility analysis uses a qualitative evaluation method that combines 16 weighted criteria across three dimensions—opportunities, challenges, and strategy—into a single utility score. AI use cases that exceed a defined threshold are implemented, while those that do not are placed in a backlog.

The utility analysis rests on the consequentialist idea that responsible AI use is a matter of outcome and can be assessed by calculation. It engages the instrumental rationality of decision-makers and scaffolds them toward calculated use of AI, so that OTTO focuses its scarce resources on AI use cases *with the best ratio of benefit to effort* (ID-07).

The utility analysis brings transparency and comprehensibility to AI investment decisions. Employees find it intuitive and appreciate that it makes decision-making criteria visible to everyone:

Our experience shows that, once you have such a tool, many questions get triggered ... What do you expect from what you want to do with AI? How much value does the AI actually contribute? Do I need AI for this, or can I use other methods? Do I need the additional effort of AI with retraining, monitoring, compliance requirements, or are other approaches sufficient? These discussions come to the table within this standardized utility analysis. (ID-02)

However, interviewees reported that it is very difficult to anticipate and quantify outcomes, or to attribute actual benefits to AI, since it is often part of a larger IT product. During a workshop, OTTO legal experts pointed to similar issues with regard to risks:

We typically calculate risk as probability of occurrence times severity of consequences. But what if we had an AI chatbot that provides false product information? For example, it says a bed can safely support X kilograms, but then the bed collapses. In such product liability cases, it is almost impossible to assess the severity of consequences, because potential compensations will be determined in individual lawsuits. And only God knows what the outcome will be.

While the utility analysis can support the demanding task of outcome assessment, it is not independent of human agency. What can and should be weighed, and how, remains a matter of contextual judgment:

There are aspects ..., such as availability of suitable data, where we need people who can make that judgment. These are usually not the same people who can judge the business value. (ID-07)

As an interviewee noted, *This is where we see the tool reaching its limits. The advantage is that it gets us talking* (ID-02).

We observed the exercise of contextual judgment and reflexive (re) production of AI governance structure during an internal workshop when OTTO practitioners applied the utility analysis to AI use cases in a new business domain. The practitioners found that the existing scales did not allow for meaningful differentiation between AI use cases and

decided to adapt them. As an interviewee explained afterwards:

We adjusted the "savings" criteria. We also adjusted two or three weightings, because the biggest challenge for us tends to be obtaining data in the first place. We may need to revisit these adjustments when we apply the tool to other areas. (ID-12)

4.6. How OTTO governs AI use by actor disposition: AI trainings for employees

Ensuring that AI is used with appropriate intentions and outcomes also requires OTTO to take *our employees with us on this journey. We depend on their curiosity, knowledge, and motivation. We want to engage everyone at all levels [and] foster curiosity and continuous learning* (Otto, 2024).

To engage all employees, OTTO has introduced AI trainings on Masterplan, its e-learning platform. The platform provides Netflix-like learning content for over 42,000 employees across more than 50 companies of OTTO Group. Currently, Masterplan offers three learning paths through which employees can build a foundational understanding of AI and generative AI, as well as familiarize themselves with how to introduce the technologies in business practice.

The AI trainings are based on the virtue-ethical idea that responsible AI use requires practical wisdom that individuals build and cultivate through practice. Once cultivated, practical wisdom scaffolds individuals to use AI in ways that are appropriate to the situation at hand (i.e., competent use of AI). Thus, these trainings combine education with hands-on exercises and a clear call to action: *please try it out, don't be intimidated by it, and just give it a try!* (ID-03)

The AI trainings are designed to equip employees with a digital mindset rather than to instruct them on specific AI technologies. This is due in part to the rapid technological development:

[Y]ou're barely out of the door and already three versions behind. What is much more sustainable is to teach people a general mindset [b]ecause the mindset that you need to deal with AI will not change much in the next five to ten years. (ID-03)

While mindset-oriented AI trainings come with the virtue-ethical strength of adaptability, they also inherit the virtue-ethical issue of guiding specific action. As an interviewee observed:

I believe that we need to go into much more detail ... There are courses on ethics and AI, but these tend to be at a meta level and not so application-oriented that a data scientist can look at them and immediately know what to do. (ID-05)

Therefore, OTTO continuously monitors whether the existing learning paths still cultivate the practical wisdom needed in organizational practice. The current understanding is:

The story has been told for now. We've covered the fundamentals of AI, GenAI, as well as the steps you need to take when introducing AI. Now, we need to consider: Which questions remain unanswered? Where is it worth making adjustments or going deeper? But to do that, we first need to bring the existing learning paths into the organization and get as many people as possible to engage with them. After that, we will see what the next logical step is. ... If we roll out something like Microsoft Copilot on a large scale, I can imagine new questions will arise. But that's just a hypothesis for now. (ID-03)

Again, the recursive interplay of AI governance structure and human agency will become most visible when OTTO adapts the AI trainings for employees in the future.

4.7. Conclusions from the illustrative case study

Taken together, the OTTO case illustrates that no single AI governance modality suffices. Governing AI by intention (responsible AI guidelines), by outcome (utility analysis), or by actor disposition (AI trainings for employees) scaffolds the normative reasoning of individuals in a distinct and valuable way, yet each also has limitations that human agency must compensate for. Accordingly, the OTTO case illustrates that AI governance is a matter not only of using “rules, practices, processes, and technological tools” (Mäntymäki et al., 2022, p. 604), but equally of creating the conditions for organization members to enact them in daily conduct, exercise situational judgment when the scaffolds reach their limits, and adapt AI governance structures over time. This case study thus supports our structurationist proposition that responsible AI governance is best understood and enacted as reflexive organizational praxis.

5. Discussion

We conducted this study to clarify how organizations can govern AI for responsible use. To this end, we enriched Giddens's (1984) structuration theory with ethics (i.e., deontology, consequentialism, and virtue ethics) and conceptualized AI governance as a *reflexive organizational praxis*. We then applied our conceptual framework to Germany's largest e-commerce company (OTTO) for illustration. Here we relate our research to existing knowledge, state the limitations of this study, and discuss implications for practice.

5.1. Theoretical contributions

We make four interrelated contributions to the existing knowledge on AI governance and responsible AI use in organizations. First, we offer an ontological reorientation for the AI governance discourse. Much of the existing literature implicitly assumes a dualism between structure and agency, which has tended to divide the field into competing camps. Structure-oriented scholars argue that responsible AI use requires binding regulation, formalized oversight roles, and risk management frameworks (Maclure & Morin-Martel, 2025; Schuett, 2023; Schuett et al., 2025). Agency-oriented scholars counter that responsible use cannot be imposed through external constraints but must emerge from the situational judgment and professional disposition of those who develop and apply AI technologies (Hagendorff, 2022; Ratti & Graves, 2021; Seger, 2022). Although both camps correctly describe one facet of the phenomenon, they share the assumption that structure and agency are distinct ontological entities whose relative importance can, and thus must, be adjudicated. Building on Giddens' (1984) concept of the duality of structure, we reconceptualize responsible AI governance as a recursive process in which structure and human agency mutually constitute each other through social practices. When viewed through this lens, the question of the relative importance of structure and agency loses its meaning. This ontological reorientation also resolves a tension in the most widely cited definition of organizational AI governance as a “system of rules, practices, processes, and technological tools” (Mäntymäki et al., 2022, p. 604). This definition identifies the constituent elements of AI governance but leaves open how they affect how AI is used. From our structurationist perspective, rules, practices, processes, and technological tools are not governance per se, they only *become* governance at the moment and to the extent that human agents enact them in daily conduct. This turn to social practice also offers the necessary theoretical explanation for why AI governance efforts frequently fall short of their stated aims (Mittelstadt, 2019; Moats & Trifuljesko, 2025; Orr & Davis, 2020; Winfield & Jirotko, 2018). We

thus provide a foundation for future integrative research that asks how structure and agency interact when organizations govern AI for responsible use.

Second, we demonstrate that responsible AI use is a normatively contested concept. Much of the literature treats responsible AI use as a stable reference point that AI governance can be designed to reach (Batool et al., 2025; Birkstedt et al., 2023). Our research challenges this assumption and shows that governing AI for responsible use requires organizations to navigate multiple, and sometimes conflicting, interpretations of responsibility. We trace Mittelstadt's (2019, p. 1) observation that many AI ethics principles “fail to address fundamental normative and political tensions embedded in key concepts” to the deeper issue that the three major ethical theories—deontology, consequentialism, and virtue ethics—disagree about whether the moral quality of an action is determined by intention, outcome, or actor disposition. We argue that the normative underdetermination of responsibility is not something that AI governance can fix but something it must work within. For conceptual clarity, we suggest distinguishing responsibility in terms of compliant, calculated, and competent use of AI.

Third, we reframe what responsible AI governance does and how it works. We show that organizations can govern AI for responsible use by scaffolding how their members deliberate about whether and how AI should be used in a specific social context, rather than by prescribing what responsibility is. We have demonstrated that normative scaffolding operates through three modalities and targets either intention, outcome, or actor disposition. Our research thus complements Seger's (2022) defense of principlism by specifying the mechanism through which principles act as cultural scaffolds. Our findings also extend Hagendorff's (2022) virtue-based framework and Ratti and Graves's (2021) concept of moral attention by situating virtue-ethical cultivation within a broader governance architecture that operates simultaneously through deontological and consequentialist modalities. Our research further offers a classificatory framework that moves beyond the structural-procedural-relational schema (e.g., Papagiannidis et al., 2023) and allows differentiation of AI governance mechanisms by the normative orientation they provide.

Fourth, we explain why AI governance fails in characteristic and predictable ways. Prior research has documented that ethical principles are difficult to translate into actionable guidance (Mittelstadt, 2019; Seger, 2022), that impact assessments face inherent epistemic limitations (Mökander et al., 2022), and that cultivating practical wisdom among AI practitioners remains a persistent challenge (Hagendorff, 2022; Ratti & Graves, 2021). We explain these issues as inherent limitations that AI governance modalities inherit from the ethical theories on which they implicitly draw. Governing the use of AI by intention inherits deontological problems like duty conflicts, governing AI by outcome inherits the consequentialist challenge of outcome uncertainty, and governing AI by actor disposition inherits the virtue-ethical issue of providing specific action guidance. We argue that because these issues are tied to the basic forms of normative reasoning, they will surface in any sufficiently complex organizational context and are not “implementation errors” that sophisticated designs can eliminate. This conclusion is not grounds for pessimism, however; as we have shown, organizations can address these inherent limitations if AI governance is understood and enacted as reflexive praxis.

5.2. Practical implications

Our research also has actionable implications for practitioners. Crucially, our proposition that responsible AI use emerges from situated judgment does not absolve managerial accountability but transforms it from exercising unilateral control to stewarding an evolving social system. Three managerial recommendations follow.

First, start to deliberately design how your organization scaffolds the responsible use of AI. As with any management task, improving

responsible AI governance requires an understanding of the status quo. Cross-functional working groups (e.g., including legal, HR, and AI leadership) can establish whether and how an organization currently scaffolds compliant, calculated, and competent use of AI. The resulting map reveals potential blind spots and provides the foundation for continuous improvement. The three modalities of responsible AI governance provide a vocabulary for aligning leadership and communicating governance priorities to internal and external stakeholders.

Second, actively monitor how your organization enacts AI governance. Employees are typically the first to notice where rules conflict, where impact assessments fall short, or where training leaves them without sufficient guidance. Without dedicated channels, these signals never reach leadership. Measures could include appointing digital responsibility roles with a direct line to senior leadership and establishing lightweight incident reporting. The raising of AI governance issues should be explicitly valued and incentivized.

Third, place continuous improvement of AI governance on your corporate agenda. Governance structures and feedback channels are insufficient if leadership does not act on the issues that surface. Thus, natural points in the corporate calendar should be used to improve the AI governance system, including monthly team check-ins, quarterly leadership reviews to assess whether AI governance structures remain suitable, and annual benchmarking against regulatory and industry developments. Here, format is less important than regularity.

5.3. Limitations and opportunities for future research

Because we developed our analytical framework from both general theories of social conduct and theories of the good, we argue that it is transferable beyond the illustrative context examined in the present study. Testing and challenging the explanatory power of our framework in other settings may identify valuable directions for future research. The framework could be applied across different regional and industry contexts, such as to compare which ethical modalities are most salient, where reflexivity is weak, or how structuration is constrained or enabled under varying regulatory and competitive conditions. Moreover, our analytical framework could be applied to high-stakes domains of AI use. While we expect the fundamental dynamics of the recursive interaction between structure and agency to hold, the ways in which ethical orientations manifest may differ as the severity of unintended consequences increases (e.g., in military or medical contexts).

Additional research opportunities arise from the limitations of our study. First, our analytical framework was developed through institutional analysis, which brackets the strategic conduct of situated actors to foreground AI governance structures. This bracketing is provisional (Giddens, 1979). Our framework illuminates how organizations can scaffold normative reasoning but leaves open how individual members enact AI governance structures in daily conduct. Future empirical research examining this strategic conduct dimension would complement our institutional analysis and provide a fuller account of responsible AI governance as reflexive praxis. Second, while our case data illustrate how AI governance mechanisms are implemented, we did not explore how the practical challenges that we identified can be addressed. Future design-oriented research could focus on developing tools and artifacts to help organizations navigate the complexities of governing AI for responsible use.

6. Conclusion

Our research demonstrates that responsible AI use emerges from the recursive interplay among formal structures, situational judgment, and continuous feedback. This finding underscores the importance of understanding AI governance not as a static system of unidirectional control but rather as a dynamic and dialogical process that evolves alongside organizational practice and technological change. By clarifying the ontological and normative foundations of responsible AI

governance, we contribute to the joint endeavor of research and practice to create desirable futures in which AI technologies are used for good.

CRediT authorship contribution statement

Simon Sturm: Writing – review & editing, Writing – original draft, Visualization, Methodology, Data curation, Conceptualization. **Florian Krause:** Writing – review & editing, Writing – original draft, Methodology, Conceptualization. **Benjamin van Giffen:** Writing – review & editing, Conceptualization.

Declarations

As required by publisher Elsevier, the author(s) declare the following:

During the preparation of this work, the author(s) used DeepL, ChatGPT, and Claude for translation and to improve language and readability. After using the tools/services, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the publication.

References

- Acharya, D., Kuppan, K., & Bhaskaracharya, D. (2025). Agentic AI: Autonomous intelligence for complex goals—A comprehensive survey. *IEEE Access: Practical Innovations, Open Solutions*, 13, 18912–18936.
- Anshari, M., Hamdan, M., Ahmad, N., Ali, E., & Haidi, H. (2023). COVID-19, artificial intelligence, ethical challenges and policy implications. *AI & Society*, 38(2), 707–720.
- Aristotle. (2007). *Die Nikomachische Ethik: Griechisch - Deutsch [Nicomachean Ethics: Greek - German]*. De Gruyter.
- Bach, T. A., Kaarstad, M., Solberg, E., & Babic, A. (2025). Insights into suggested responsible AI (RAI) practices in real-world settings: A systematic literature review. *AI and Ethics*, 5(3), 3185–3232.
- Batool, A., Zowghi, D., & Bano, M. (2025). AI governance: A systematic literature review. *AI and Ethics*, 5(3), 3265–3279.
- Bentham, J. (1996). *The collected works of Jeremy Bentham: An introduction to the principles of morals and legislation*. Clarendon Press.
- Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Managing artificial intelligence. *MIS Quarterly*, 45(3), 1433–1450.
- Bezuidenhout, L., & Ratti, E. (2021). What does it mean to embed ethics in data science? An integrative approach based on microethics and virtues. *AI & Society*, 36(3), 939–953.
- Birkstedt, T., Minkinen, M., Tandon, A., & Mäntymäki, M. (2023). AI governance: Themes, knowledge gaps and future agendas. *Internet Research*, 33(7), 133–167.
- Black, J., & Murray, A. (2019). Regulating AI and machine learning: Setting the regulatory agenda. *European Journal of Law and Technology*, 10(3). <https://www.ejlt.org/index.php/ejlt/article/view/722>.
- Canibano, A., Dudau, A., & Muratbekova-Touron, M. (2025). Publishing qualitative research in EMJ: Editorial reflections. *European Management Journal*, 43(3), 344–350.
- Christophersen, T., & Pärn, J. (2021). Data Science bei OTTO. In P. Buxmann, & H. Schmidt (Eds.), *Künstliche Intelligenz* (pp. 101–115). Springer Berlin Heidelberg.
- Cihon, P., Schuett, J., & Baum, S. D. (2021). Corporate governance of artificial intelligence in the public interest. *Information*, 12(7), 275.
- Constantinescu, M., Voinea, C., Uszkai, R., & Vică, C. (2021). Understanding responsibility in responsible AI. *Dianoetic virtues and the hard problem of context. Ethics and Information Technology*, 23(4), 803–814.
- Davenport, T. (2018). *The AI advantage: How to put the artificial intelligence revolution to work*. MIT Press.
- den Hond, F., Boersma, F. K., Heres, L., Kroes, E. H. J., & van Oirschot, E. (2012). Giddens à la Carte? Appraising empirical applications of Structuration Theory in management and organization studies. *Journal of Political Power*, 5(2), 239–264.
- Feuerriegel, S., Hartmann, J., Janiesch, C., & Zschech, P. (2023). Generative AI. *Business & Information Systems Engineering*. <https://doi.org/10.1007/s12599-023-00834-7>
- Floridi, L., Cows, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People-An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689–707.
- Fukuda-Parr, S., & Gibbons, E. (2021). Emerging consensus on 'ethical AI': Human rights critique of stakeholder guidelines. *Global Policy*, 12(S6), 32–44.
- Gal, U., Hansen, S., & Lee, A. S. (2022). Research perspectives: Toward theoretical rigor in ethical analysis: The case of algorithmic decision-making systems. *Journal of the Association for Information Systems*, 23(6), 1634–1661.
- Garibay, O. O., Winslow, B., Andolina, S., Antona, M., Bodenschatz, A., Coursaris, C., Falco, G., Fiore, S. M., Garibay, A., Grieman, K., Havens, J. C., Jirotko, M., Kacorri, H., Karwowski, W., Kider, J., Konstan, J., Koon, S., Lopez-Gonzalez, M., Maifeld-Carucci, I., ... Xu, W. (2023). Six human-centered artificial intelligence

- grand challenges. *International Journal of Human–Computer Interaction*, 39(3), 391–437.
- Giddens, A. (1979). *Central problems in social theory: Action, structure and contradiction in social analysis*. Red Globe Press.
- Giddens, A. (1981). *Contemporary critique of historical materialism*. University of California Press.
- Hagendorff, T. (1984). *The constitution of society: Outline of the theory of structuration*. University of California Press.
- Griffin, T. A., Green, B. P., & Welie, J. V. M. (2024). The ethical agency of AI developers. *AI and Ethics*, 4(2), 179–188.
- Griffin, T. A., Green, B. P., & Welie, J. V. M. (2025). The ethical wisdom of AI developers. *AI and Ethics*, 5(2), 1087–1097.
- Hagendorff, T. (2023). The ethics of AI ethics: An evaluation of guidelines. *Minds and Machines*, 30(1), 99–120.
- Hagendorff, T. (2022). A virtue-based framework to support putting AI ethics into practice. *Philosophy & Technology*, 35(3), 1–24.
- Hayes, P., Fitzpatrick, N., & Ferrández, J. M. (2025). From applied ethics and ethical principles to virtue and narrative in AI practices. *AI and Ethics*, 5(2), 1217–1239.
- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399.
- Kant, I. (1797/1996). *The metaphysics of morals*. Cambridge University Press.
- Koniakou, V. (2023). From the “rush to ethics” to the “race for governance” in artificial intelligence. *Information Systems Frontiers: A Journal of Research and Innovation*, 25(1), 71–102.
- Maclure, J., & Morin-Martel, A. (2025). AI ethics’ institutional turn. *Digital Society: Ethics, Socio-Legal and Governance of Digital Technology*, 4(1). <https://doi.org/10.1007/s44206-025-00174-x>
- Mäntymäki, M., Minkkinen, M., Birkstedt, T., & Viljanen, M. (2022). Defining organizational AI governance. *AI and Ethics*, 2(4), 603–609.
- Mikalef, P., Conboy, K., Lundström, J. E., & Popović, A. (2022). Thinking responsibly about responsible AI and “the dark side” of AI. *European Journal of Information Systems: An Official Journal of the Operational Research Society*, 31(3), 257–268.
- Mill, J. S. (2009). *Utilitarianism [Utilitarianism]*. Felix Meiner Verlag.
- Mittelstadt, B. (2019). Principles alone cannot guarantee ethical AI. *Nature Machine Intelligence*, 1(11), 501–507.
- Moats, D., & Trifunjesko, S. (2025). Valuable actions and actionable values: Tinkering with principles and practices in AI ethics. *The Sociological Review*. , Article 00380261251316661. <https://doi.org/10.1177/00380261251316661>
- Mökander, J., & Floridi, L. (2023). Operationalising AI governance through ethics-based auditing: An industry case study. *AI and Ethics*, 3(2), 451–468.
- Mökander, J., Sheth, M., Gersbro-Sundler, M., Blomgren, P., & Floridi, L. (2022). Challenges and best practices in corporate AI governance: Lessons from the biopharmaceutical industry. *Frontiers in Computer Science*, 4, Article 1068361.
- Morley, J., Floridi, L., Kinsey, L., & Elhalal, A. (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), 2141–2168.
- Müller-Wünsch, M., Brenner, W., & Brenner, B. (2025). “(Gen)AI First” at OTTO, Germany’s largest online store. *MIS Quarterly Executive*, 24(3), 7.
- one.O. (2026). Responsible commerce. <https://www.oglo.com/en/about-us/ai-guidelines>.
- Orr, W., & Davis, J. L. (2020). Attributions of ethical responsibility by Artificial intelligence practitioners. *Information, Communication and Society*, 23(5), 719–735.
- Otto. (2021). *Ethik vs. Wirtschaftlichkeit, können wir KI vertrauen? | MAIN Session – OTTO [Ethics vs. economics, can we trust AI?]*. Youtube [Video] <https://www.youtube.com/watch?v=aiaj2hfDiOg>.
- Otto. (2024). *Bilanz-Presskonferenz 2024 [Annual Results Press Conference 2024]*. <https://www.ottogroup.com/de/bpk/2024/index.php>.
- Otto. (2026). *Zahlen, Daten, Fakten: Die wichtigsten Kennzahlen rund um OTTO [Facts, figures and data: The most important key figures about OTTO]*. In . <https://www.otto.de/unternehmen/de/unternehmen/auf-einen-blick>.
- 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: A Journal of Research and Innovation*, 25(1), 123–141.
- Piñeros Glasscock, J. S., & Tenenbaum, S. (2023). Action. In *The Stanford encyclopedia of philosophy (Spring 2023 edition)*.
- Polanyi, M. (1962). The Republic of science: Its political and economic theory. *Minerva*, 1(1), 54–73.
- Radford, A., Kim, J. W., Xu, T., Brockman, G., McLeavey, C., & Sutskever, I. (2022). Robust speech recognition via large-scale weak supervision. *arXiv [eess.AS]*. [arXiv. http://arxiv.org/abs/2212.04356](http://arxiv.org/abs/2212.04356).
- Ratti, E., & Graves, M. (2021). Cultivating moral attention: A virtue-oriented approach to responsible data science in healthcare. *Philosophy & Technology*, 34(4), 1819–1846.
- Ruster, L. P., & Daniell, K. (2025). How to operationalize responsible use of artificial intelligence. *MIS Quarterly Executive*, 24(2), 6.
- Schiff, D., Rakova, B., Ayes, A., Fanti, A., & Lennon, M. (2021). Explaining the principles to practices gap in AI. *IEEE Technology and Society Magazine*, 40(2), 81–94.
- Schneider, J., Abraham, R., Meske, C., & Vom Brocke, J. (2023). Artificial intelligence governance for businesses. *Information Systems Management*, 40(3), 229–249.
- Schuett, J. (2023). *Three lines of defense against risks from AI*. AI & Society. <https://doi.org/10.1007/s00146-023-01811-0>
- Schuett, J., Reuel, A.-K., & Carlier, A. (2025). How to design an AI ethics board. *AI and Ethics*, 5(2), 863–881.
- Seger, E. (2022). In defence of principlism in AI ethics and governance. *Philosophy & Technology*, 35(2), 1–7.
- Shneiderman, B. (2020a). Human-centered artificial intelligence: Three fresh ideas. *AIS Transactions on Human-Computer Interaction*, 12(3), 109–124.
- Shneiderman, B. (2020b). Bridging the gap between ethics and practice: Guidelines for reliable, safe, and trustworthy human-centered AI systems. *ACM Trans. Interact. Intell. Syst.*, 10(4), 1–31.
- Taddeo, M., & Floridi, L. (2018). How AI can be a force for good. *Science*, 361(6404), 751–752.
- Tursunbayeva, A., Fernandez, V., Gallardo-Gallardo, E., & Moschera, L. (2025). Artificial intelligence and digital data in recruitment. Exploring business and engineering candidates’ perceptions of organizational attractiveness. *European Management Journal*. <https://doi.org/10.1016/j.emj.2025.03.002>
- Urbina, F., Lentzos, F., Invernizzi, C., & Ekins, S. (2022). Dual use of artificial intelligence-powered drug discovery. *Nature Machine Intelligence*, 4(3), 189–191.
- Walgenbach, P. (1999). Giddens’ Theorie der Strukturierung [Giddens’ Theory of Structuration]. In A. Kieser (Ed.), *Organisationstheorien [Organisation Theories]* (pp. 355–375). Kohlhammer.
- Whittington, R. (2025). Structuration theory: Giddens and beyond. In *Cambridge handbook of strategy as practice* (pp. 113–129). Cambridge University Press.
- Whittlestone, J., Nyrup, R., Alexandrova, A., & Cave, S. (2019). The role and limits of principles in AI ethics: Towards a focus on tensions. *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*.
- Winfield, A. F. T., & Jirotko, M. (2018). Ethical governance is essential to building trust in robotics and artificial intelligence systems. *Philosophical Transactions. Series A, Mathematical, Physical, and Engineering Sciences*, 376(2133), Article 20180085.
- Wolf, M. J., Miller, K., & Grodzinsky, F. S. (2017). Why we should have seen that coming. *ACM SIGCAS Computers and Society*, 47(3), 54–64.
- Yin, R. K. (1981). The case study crisis: Some answers. *Administrative Science Quarterly*, 26(1), 58–65.
- Yin, R. K. (1994). *Case study research: Design and methods*. SAGE Publications.