

Activating Employees' AI Literacy for Responsible Generative AI Use: A Sensitization Framework

Completed Research Paper

Michael Dominic Harr

University of Duisburg-Essen
Universitätsstr. 9, 45141 Essen
michael.harr@ris.uni-due.de

Leonardo Banh

University of Duisburg-Essen
Universitätsstr. 9, 45141 Essen
leonardo.banh@uni-due.de

Reinhard Schütte

University of Duisburg-Essen
Universitätsstr. 9, 45141 Essen
reinhard.schuette@ris.uni-due.de

Gero Strobel

University of Duisburg-Essen
Universitätsstr. 9, 45141 Essen
gero.strobel@uni-due.de

Abstract

As generative AI (GenAI) diffuses rapidly into organizational practice, developing employees' AI literacy for responsible use has become a strategic imperative. However, organizations struggle to move beyond ad-hoc training toward sustained, actionable competence. Current research has mapped what AI literacy entails but left the organizational process underexplored through which these competencies become actionable. We address this gap by theorizing GenAI sensitization as an organizationally embedded process operating at the intersection of AI literacy, governance, and organizational practice. Drawing on 20 semi-structured interviews with human resource professionals analyzed through Gioia methodology, we identify four aggregate dimensions: organizational factors, individual factors, sensitization practices, and tensions and barriers. Four structural tensions and barriers characterize conditions that explain why responsible GenAI use fails even among AI-literate employees. The sensitization profile matrix synthesizes these findings into a diagnostic framework for organizational practice, guiding employees toward responsible GenAI use.

Keywords: Generative AI, responsibility, sensitization, AI literacy, shadow AI, human resource management.

Introduction

Generative artificial intelligence (GenAI) – a class of artificial intelligence (AI), built on large-scale foundation models such as large language models, that is capable of producing novel text, images, code, and other content from natural language prompts (Banh & Strobel, 2023; Feuerriegel et al., 2024) – has emerged as one of the most consequential technological developments in recent years. Since OpenAI released ChatGPT in late 2022, a rapidly expanding body of research has investigated how GenAI can be productively used to augment work (e.g., Dell'Acqua et al., 2026), adopted and integrated into organizational processes (e.g., Sharma et al., 2026), and used to collaborate with human teams (e.g., Grange et al., 2026). What distinguishes GenAI from previous waves of digital innovation is the convergence of two dynamics: the democratization of knowledge, which has placed sophisticated AI capabilities within reach of any employee, and the unprecedented velocity of its diffusion across both public and organizational domains (Retkowsky et al., 2024). Yet, the fast-paced diffusion and adoption of GenAI does not come without peril. Across industries, a governance vacuum has opened between the pace of employee adoption

and the ability of organizations to guide, regulate, and support responsible GenAI use (Choudhary & Kar, 2025; Gartner, 2025). For instance, a recent survey found that more than half of GenAI-using employees rely on personal or non-company-governed tools, with half of this group stating they would continue doing so even if their employer explicitly banned such use (Software AG, 2024). Hence, deploying proprietary GenAI models in-house even “with enhanced security measures” (Choudhary & Kar, 2025, p. 2) does not resolve the governance challenge, since employees willingly circumvent organizationally governed GenAI systems. Recent research further indicates that employees might actively conceal circumvention (Xhigoli & Schaarschmidt, 2025), which is difficult to detect because external GenAI tools are easily accessible (e.g., via browser or personal devices) and leave few organizational traces (Glorin, 2026; Silic et al., 2025). Unsolicited GenAI use has detrimental consequences for organizations regarding potential cybersecurity (e.g., prompt injections) or compliance incidents (e.g., data exposure) (Gartner, 2025). If organizational governance alone cannot prevent concealed GenAI use, individual competence becomes a critical complement. These developments have intensified information systems (IS) researchers’ attention to the concept of ‘AI literacy,’ which describes the individual continuum of competencies necessary for responsible GenAI use (Long & Magerko, 2020; Passlack et al., 2026; Pinski & Benlian, 2024). AI literacy encompasses engaging with opaque, probabilistic, and often autonomous GenAI systems, and thus requires different competencies for reasoning under uncertainty (Benlian & Pinski, 2025; Söllner et al., 2025). Accordingly, this implies a distinct training challenge (Passlack et al., 2026). Nevertheless, existing research on AI literacy focuses primarily on the individual level of analysis, treating AI literacy as an individual outcome or property of the employee. However, AI literacy is fundamentally stakeholder-specific and context-dependent: different professional roles require different configurations of literacy components (Benlian & Pinski, 2025), and the development of these components is shaped by influencing factors not only at the individual, but also at the organizational and societal level (Passlack et al., 2026). Yet, *how organizations build and activate AI literacy for responsible GenAI use in their workforce*, and thereby link the organizational and individual levels, remains largely unaddressed. We specifically investigate *GenAI sensitization* – the organizational process of developing, activating, and sustaining employees’ AI literacy for responsible GenAI use. In doing so, we make explicit “the dynamics by which collective and individual responsible AI competencies are developed” (Papagiannidis et al., 2025, p. 12). We conceptualize sensitization as distinct from, but complementary to, AI literacy: while literacy describes the individual continuum of competencies (Long & Magerko, 2020; Pinski & Benlian, 2024) and therefore properties at the individual level, sensitization describes the organizational process through which these competencies are contextualized and made actionable. We situate our investigation in Human Resources (HR), a domain where the sensitization challenge is particularly acute: HR is a natural early adopter of GenAI for creative and text-intensive tasks, such as drafting job advertisements, summarizing candidate profiles, or generating onboarding materials (Banh et al., 2026; Budhwar et al., 2023). At the same time, HR operates under critical sensitivity. HR is bound by the General data protection regulation (GDPR) when handling highly personal employee data (Hacker et al., 2023), makes high-stakes decisions in which hallucinations and algorithmic bias carry legal consequences (e.g., Kulal et al., 2025), and operates in a tight regulatory environment in which the European Union (EU) AI Act classifies employment-related AI systems as high-risk and mandates adequate AI literacy among personnel (Art. 4 EU AI Act; Benlian & Pinski, 2025). Furthermore, HR professionals face a distinctive dual agency: they must be sensitized themselves as GenAI users in consequential decision contexts while simultaneously bearing organizational responsibility for sensitizing employees through talent development initiatives. This convergence creates a domain in which sensitizing employees for responsible GenAI use is critical, as “much of what organizations do regarding responsible AI practices hinges on employees’ knowledge” (Papagiannidis et al., 2025, p. 19). Hence, we ask the following research questions (RQs):

RQ1: *Which sensitization practices are deployed in organizations, and which organizational and individual factors influence them?*

RQ2: *Which structural tensions emerge in sensitizing HR employees for responsible GenAI use?*

To answer the RQs, we employ semi-structured interviews with twenty HR professionals to empirically derive a grounded account of the organizational factors, individual factors, sensitization practices, and tensions and barriers that emerge in sensitizing HR professionals for responsible GenAI use. In doing so, we contribute to the discourse on responsible GenAI use by shifting the analytical focus from *which competencies individuals need* (i.e., AI literacy) to *how organizations activate AI literacy through sensitization practices that translate into responsible GenAI use*. Building on these findings, we develop

the sensitization profile matrix, a two-dimensional diagnostic framework that identifies four empirically grounded sensitization archetypes and propose targeted intervention pathways for organizational sensitization. Therefore, we respond to calls for research on underexplored AI user groups (Pinski & Benlian, 2024), multi-level AI literacy research to examine “the interdependencies between various influencing factors at individual [and] organizational levels” (Passlack et al., 2026, p. 29), and the operationalization of responsible AI governance in practice (Papagiannidis et al., 2025).

Conceptual Background

GenAI Governance in Human Resource Management

GenAI has translated into horizontal adoption across virtually all organizational functions – including HR management (Budhwar et al., 2023). Within HR management, GenAI has diffused across the entire employee lifecycle: GenAI supports recruitment (e.g., job advertisements, pre-screening), talent development (e.g., personalized learning content, feedback summarization), and HR operations (e.g., employee communication, compliance documentation) (e.g., Banh et al., 2026; Budhwar et al., 2023). Nonetheless, GenAI adoption is dampened by the consequential and legally sensitive character of HR decisions. HR involves highly sensitive personal data whose disclosure to externally hosted GenAI systems may breach data protection and confidentiality obligations (Hacker et al., 2023), while GenAI’s probabilistic character can introduce hallucinations and bias into decisions with substantial employment consequences protected by anti-discrimination law (e.g., Kulal et al., 2025). In response, policymakers and regulators have begun to codify AI-related competencies into citizenship frameworks (e.g., Vuorikari et al., 2022), educational frameworks (e.g., Holmes & Miao, 2023), and binding obligations, such as through the EU AI Act, which requires ensuring a sufficient level of AI literacy among employees to sensitize for its risks and advocate for responsible use. Similarly, research and practice have responded by developing AI governance frameworks, i.e., structures, policies, and processes through which organizations operationalize ethical, legal, and accountability principles for responsible GenAI use (e.g., Papagiannidis et al., 2025). In HR contexts, these typically materialize as proprietary in-house GenAI deployments (Choudhary & Kar, 2025), use-case whitelists, and mandatory training interventions. Yet, IS research demonstrates that governance frameworks can be systematically undermined through ‘*shadow AI*,’ which refers to the unauthorized or ungoverned use of publicly available (Gen)AI tools to circumvent organizational policies even when proprietary systems are in place (Glorin, 2026; Silic et al., 2025). However, monitoring employees’ conscious circumvention is difficult, as recent research has demonstrated that employees may deliberately conceal their (unauthorized) GenAI use (Xhigoli & Schaarschmidt, 2025). Governance frameworks alone therefore appear insufficient to ensure employees’ responsible GenAI use. The term ‘*responsible GenAI use*’ denotes employees’ use of GenAI for tasks in which it creates value, without breaching external regulation, organizational policies, or job obligations, and is accompanied by a critical evaluation of GenAI outputs and their consequences. Hence, responsible use is situated at the level of employee behavior rather than organizational governance (Papagiannidis et al., 2025). It relates to the notion of effective use (Burton-Jones & Grange, 2013), where the relevant goals, however, include legal and organizational conformance and not only employee performance. Responsible use is accordingly multidimensional. Prior work suggests dimensions such as protecting confidential data, validating outputs, recognizing bias, disclosing GenAI use, and preserving human judgment in consequential decisions (Söllner et al., 2025). Compliance alone is insufficient, as employees can follow formal rules yet still fail to validate outputs (Söllner et al., 2025).

AI Literacy and Sensitization

The concept of literacy, originally the ability to “express [oneself] and communicate using written language” (Long & Magerko, 2020, p. 2), has continually evolved alongside technological advancements (Heimbach et al., 2026). A growing body of research argues that AI literacy is poised to become a key competency for work, everyday life, and lifelong learning (Knoth et al., 2024). IS research has focused dominantly on defining AI literacy components (e.g., Long & Magerko, 2020), developing measurement instruments (e.g., Pinski & Benlian, 2023; Weber et al., 2023), and foregrounding learning mechanisms (e.g., Pinski & Benlian, 2024), converging on the idea that AI literacy is a multidimensional continuum that enables employees to understand, use, critically evaluate, and responsibly interact with AI (Long & Magerko, 2020; Pinski & Benlian, 2024). Building on this view, Passlack et al. (2026) define the building blocks of AI literacy as *skills* (i.e., learnable, task-specific abilities), *knowledge* (i.e., understanding of GenAI, its contexts, and

appropriate use), *personal dispositions* (i.e., orientations such as curiosity), *capabilities* (i.e., the situated mobilization of these resources), and *competencies* (i.e., their integrated application to GenAI-related tasks and responsibilities). AI literacy is thus increasingly framed not only as a desirable individual learning outcome but as a strategic organizational concern that employers need to actively build, maintain, and govern (Papagiannidis et al., 2025) to enable responsible GenAI use in the workplace (Söllner et al., 2025). This organizational framing has, however, outpaced its theorizing, because AI literacy, as currently conceptualized, describes properties of individual employees (e.g., Long & Magerko, 2020; Pinski & Benlian, 2024), whereas responsible GenAI use in organizations constitutes a collective accomplishment. Individual AI literacy is a necessary but not a sufficient condition because collective outcomes are not reducible to the sum of individual competencies but emerge through interactional mechanisms that translate – or fail to translate – individual cognition into collective, responsible practice (Felin et al., 2015). Structured training bridges this divide only partially. Skills and declarative knowledge are amenable to curriculum-based transfer, whereas personal dispositions, capabilities, and their integration into competencies are context-specific and reflexive (Passlack et al., 2026). They develop through engagement with concrete tasks, organizational culture, interaction, and communication rather than structured training alone. This mismatch between individual AI literacy and responsible GenAI use in organizations is intensified by the properties of GenAI itself. Whereas previous AI systems operated within deterministic boundaries (Berente et al., 2021), GenAI reconstructs employees' intent from natural language and produces open-ended, probabilistic outputs (Krakowski, 2025; Storey et al., 2025). Responsible GenAI use therefore cannot rest on predefined procedures but requires iterative judgment, interpretation, and reflexive evaluation in each task (Shapiro et al., 2026). Structured training can convey which criteria matter, but not the situated judgment a concrete task demands (Passlack et al., 2026). Despite growing recognition that AI literacy must be organizationally cultivated (Papagiannidis et al., 2025; Söllner et al., 2025), the organizational practices through which individual competencies are activated, contextualized, and sustained remain underexplored (Papagiannidis et al., 2025). Existing research has predominantly asked which competencies individuals require (e.g., Long & Magerko, 2020). Where organizational practices have been considered, attention has shifted towards the design of structured learning formats and their modalities (e.g., Benlian & Pinski, 2025; Pinski & Benlian, 2024). This narrows AI literacy to a skills agenda (Shapiro et al., 2026) and leaves the provisioning, normative, and orchestrating practices surrounding these formats unexamined. Against this backdrop, we conceptualize *GenAI sensitization* as the organizational process of developing, activating, and sustaining employees' AI literacy for responsible GenAI use. Sensitization not only imparts the trainable AI literacy components directly but also configures the conditions under which all components become actionable in concrete tasks, since awareness-related interventions unfold as interrelated changes across the individual, technological, and organizational levels (Tsohou et al., 2015). Its central mechanism is organizational sensemaking: GenAI confronts employees with issues that are “novel, ambiguous, confusing, or in some other way violate expectations” (Maitlis & Christianson, 2014, p. 57), and sensitization shapes the sensemaking through which employees come to notice GenAI as an object that warrants reflection in the first place (Engström et al., 2024; Wu et al., 2025). Sensitization is therefore an ongoing organizational sensemaking process through which employees co-construct what responsible GenAI use means in their specific role and context, cultivating the dispositional and reflexive layers of AI literacy that structured training cannot reach (Passlack et al., 2026). Thus, GenAI sensitization can be distinguished from four adjacent concepts: *first*, it differs from structured training in scope and temporality: it is continuous rather than episodic, situated rather than codified, and participatory rather than delivered (Scholl, 2023). Structured training is therefore one sensitization practice, but not its equivalent (Naß et al., 2026). *Second*, it differs from AI literacy in the level of analysis: whereas AI literacy describes individual competencies on a continuum, sensitization describes the organizational process through which these competencies are made actionable and translated into collective, responsible practice (Felin et al., 2015). *Third*, it differs from governance in its specificity: whereas governance specifies what is permitted (Papagiannidis et al., 2025), sensitization addresses the discretionary situations that governance can neither specify nor observe ex ante (Silic et al., 2025). *Fourth*, it extends the IS discipline's security awareness tradition (Tsohou et al., 2015) to a technology whose outputs are open-ended and probabilistic. While security awareness ultimately targets compliance with specifiable policies, GenAI sensitization targets judgment in situations for which no policy can prescribe appropriate use. Sensitization is accordingly consequential where responsible use cannot be specified ex ante, where use cases emerge bottom-up, and where (unauthorized) GenAI use leaves few organizational traces (e.g., Glorin, 2026; Silic et al., 2025).

Research Approach

Given the nascent and contextually embedded nature of GenAI sensitization in HR, we chose a qualitative research design centered on expert interviews as the primary data collection method (Myers & Newman, 2007). This approach enables rich, context-sensitive insights into lived experiences from practitioners who navigate the challenges of developing AI literacy and using GenAI in the workplace. Expert interviews are thus particularly well-suited to exploratory research in new technological domains, where the research goal is to elicit the experiences and sensemaking of practitioners (Paré, 2004). **Data collection** took place from August 2025 to April 2026, comprising 20 semi-structured interviews with HR professionals from 13 different organizations in Europe. We utilized purposive sampling via LinkedIn to target HR professionals with direct involvement in recruitment, HR management, or talent development, placing them at the intersection of GenAI use and high-stakes employment decisions. Moreover, interview participants included talent development experts, consultants, and AI-upskilling experts, who provided additional perspectives and reflections on sensitization processes. The sampling strategy sought variation in organizational size, industry sector, and functional focus within the HR department to maximize theoretical richness while maintaining domain coherence required by our RQs. All participants were recruited from German-speaking organizations. Interviews ranged from 24:01 to 78:55 minutes (mean: 39:50 minutes) and were conducted via video call, recorded with participant consent, and subsequently transcribed verbatim. All transcripts were anonymized. Table 1 provides an overview of all 20 participants. The interview guide was developed iteratively and comprised four thematic blocks: (1) organizational context and work background; (2) understanding of GenAI and sensitization; (3) GenAI use scenarios, observed challenges, and existing organizational responses; and (4) strategic recommendations and future perspectives. The guide was pilot-tested before the first interviews and refined based on emerging insights in line with Gioia et al.'s (2013) emphasis on theoretical sampling and iterative data collection. After the first 13 interviews, the guide was updated to probe more deeply into indirect (shadow) GenAI dynamics, workshop strategies, and leadership competencies, which had emerged as particularly significant.

ID	Role	Organization Type	ID	Role	Organization Type
P01	HR People Tech & Data Lead	IT Service Provider	P11	Senior Sourcing Specialist	Energy Corporation
P02	Senior Talent Acquisition Expert	IT Service Provider	P12	Self-Employed Owner	HR Consulting
P03	Managing Director	Executive Search Consultancy	P13	HR Business Partner	Industrial Manufacturing
P04	Employer Branding Expert	Insurance	P14	Lead HR Payroll Accountant	Automotive
P05	Senior Recruiter	Recruitment Agency	P15	CEO Upskilling Advisor	Executive Consultancy
P06	Recruiter	Industrial Manufacturing	P16	Learning Experience Designer	IT Consultancy
P07	People Analytics Expert	Energy Corporation	P17	Head of Learning & Innovation	IT Consultancy
P08	SAP Talent Acquisition Expert	HR Consulting	P18	Chief Executive Officer	IT Consultancy
P09	Senior HR Business Partner	Telecommunications	P19	Product Owner HCM	Energy & Utility
P10	Head of Employer Branding	Telecommunications	P20	HR Manager (Generalist)	Chemical Corporation

Table 1. Interview Study Participants

For **data analysis**, we employed the Gioia methodology (Gioia et al., 2013), a structured approach to grounded theory that enables inductive theory development through three sequential coding stages: first-order concept development through open coding, second-order theme aggregation through axial coding, and aggregate dimension formation through selective coding. The Gioia methodology has been widely adopted in IS research for creating mid-range theories in emerging technological contexts and provides the

rigor-relevance balance appropriate for a study at the intersection of technology adoption and organizational behavior (Birks et al., 2013). All 20 transcripts were coded by the first two authors within the qualitative software MAXQDA. Both authors first open-coded a common subset of three transcripts independently and then discussed every divergence in a joint calibration session. This session produced a common understanding, a shared codebook, and four coding rules for subsequent coding. In the *first-order stage*, all transcripts were then openly coded, staying as close as possible to the language of informants. This process generated 811 first-order codes, capturing specific instances of GenAI use, risk perceptions, organizational sensitization practices, hints on literacy states, and employees' perceptions. In the *second-order stage*, first-order codes were grouped into fifteen themes through abductive reasoning and accompanying memos, identifying common topics and patterns (i.e., axial coding). This stage involved continuous iteration between the raw data and emerging themes to ensure that second-order themes were both internally coherent and adequately grounded in the first-order evidence. Finally, the first two authors revisited every theme together with its constituent codes and applied the four coding rules, which separate conditions the organization is confronted with, interventions it performs, employee-level attributes, and persistent contradictions between sensitization practices and their intended effects. This condensed the fifteen themes into four aggregate dimensions. Throughout all three stages, disagreements were resolved collectively in team discussions until a shared understanding was reached. Additional information on interview participants, coding rules, and the data structure can be retrieved from an online appendix: https://osf.io/2upeb/overview?view_only=c21013664a9e4e5f9213e62e77624e34.

Results

The analysis surfaced four aggregate dimensions that jointly characterize the organizational sensitization process for responsible GenAI use in HR: 1) Organizational Factors, 2) Individual Factors, 3) Sensitization Practices, and 4) Tensions and Barriers. In Figure 1, the arrows represent interpreted associations based on interviewees' accounts. Factors at both levels are associated with organizational sensitization practices, which in turn might be associated with structural tensions and barriers that arise in translating sensitization practices into responsible GenAI use (see solid arrows). Dashed arrows show the recursive association between sensitization practices and responsible GenAI use.




Figure 1. Model of Organizational GenAI Sensitization in HR

Organizational Factors

The **regulatory environment** emerged as the most consistently mentioned organizational factor. Practitioners across industries described GDPR and regulatory compliance as the dominant orienting principle of HR-related GenAI engagement, often invoked before any other consideration. Practitioners emphasized that *“in HR there is no room for compromise – the strict standards of the GDPR must be maintained at all times”* (P13) and that *“data protection is the very top priority, the most important issue, and is governed centrally”* (P01). Beyond data protection, several participants pointed to the broader regulatory framework, encompassing non-discrimination requirements, transparency obligations, and increasingly the EU AI Act, as a moving target that organizations must continuously adapt to and adhere to. In addition to formal legal requirements, participants also cited negative examples of data leakage among competitors as important drivers of governance efforts. Relatedly, P05 stressed that *“companies that cannot or do not want to fully focus on the security and ethical use of GenAI risk losing the trust of both employees and customers”* (P05). The regulatory environment, in this sense, is not merely an externally imposed constraint but a structural condition that defines the legitimate space within which organizational sensitization to GenAI can occur.

Technical environment refers to the infrastructural and toolchain preconditions that determine the feasibility of GenAI deployment and sensitization. The interviews revealed substantial heterogeneity, ranging from organizations with custom-developed GenAI systems and dedicated technical support to organizations whose legacy HR architectures impede meaningful GenAI integration. Several practitioners pointed to the proliferation of available tools (i.e., tool-overload) as a structural challenge: *“[...] there are 80,000 tools from vendors for screening – the selection alone is a challenge”* (P04). After GenAI selection, integration into existing systems emerges as a foundational concern, as P01 framed it: *“the second concern is integration into existing systems: the question is what a frictionless integration into your work looks like”* (P01). P05 complemented this perspective by emphasizing the resource implications of sustained operation: *“GenAI systems are complex and require regular maintenance and updates and that means we need experts who know how to adapt these systems, so they are optimally suited to our specific work environment”* (P05). Where organizations invest in dedicated infrastructure, the technical environment can enable sensitization. Yet its role remains ambivalent: it supports sensitization in some contexts while constraining it in others, especially as organizational complexity increases.

The organizational **culture** surfaced with marked heterogeneity as well. Several organizations focused on a human-centric employer brand, particularly in candidate-facing contexts. P01 captured this orientation in a foundational principle: *“with us, in the end the human makes the decision; that is a foundational company principle”* (P01). Similarly, P04 stressed the relevance of organizational authenticity, since HR is still people’s business: *“we are a very personal company, closely connected with applicants [...] so AI would almost contradict our organizational philosophy”* (P04). A second orientation, observed primarily in technology-near and consultancy contexts, frames GenAI use as culturally normative and openly endorsed: *“we have a very open culture for GenAI in our company [...] basically everyone has a ChatGPT account and uses it”* (P18). Such cultures typically encourage their employees to actively use GenAI. A third orientation, observed in large incumbent firms, manifests as cultural inertia, as outlined by an HR manager in a large multinational corporation: *“Such a corporation is always somewhat like a large oil tanker at sea: it keeps moving forward and forward, without really swaying or changing direction. It simply continues on its course until, at some point, it may sink like the Titanic”* (P20). Cultural orientation could thus be understood as a contingent configuration that influences both the urgency and the modality of sensitization activity.

Leadership encompasses managers’ disposition, AI literacy, and ability to anticipate GenAI’s potential and application scenarios, as well as their capacity to act as credible role models. These leadership competencies constitute preconditions for the organizational reach of sensitization efforts, particularly given top-management support requirements and recurring leadership literacy deficits. Across almost all interviews, leadership was described less as technical expertise than as the ability to legitimize, frame, and accompany GenAI engagement, that is, to set strategic and ethical boundaries and to encourage critical experimentation. Leaders’ own AI literacy frequently fell short of this expectation, as P16 reported that *“around 85% of the managers at the client organizations had not even tried anything involving generative AI in their private lives”* (P16). The asymmetry between leadership expectation and leadership capability was also recognized by practitioners: *“managers sometimes expect something from their own employees*

that they themselves are not able to do" (P07). Beyond literacy levels, leadership emerged as a critical signaling function. As P16 metaphorically said, *"the fish rots from the head down, so if it isn't role-modeled from above, no GenAI mindset will emerge"* (P16). Hence, leadership constitutes an important factor whose configurations shape whether sensitization practices can develop substantive reach in the workforce.

Individual Factors

Motivation surfaced as a recurring concern, addressing employees' willingness to engage with GenAI and their acceptance of it, as well as motivational deficits regarding training participation. Across practitioners, concerns were particularly salient in relation to voluntary participation in structured training, where substantial variation could be observed. P02 reported: *"we employ 2,200 employees in Germany alone and only 70 people were in the training call"* (P02). Several participants emphasized that resistance to GenAI is not necessarily rooted in age or a lack of technical skills, but often in a lack of willingness to engage: *"it is simply an 'I don't want to'"* (P15). Other practitioners identified the perceived technological immaturity of GenAI as a factor that can undermine motivation, particularly when recurring output errors or hallucinations reduce employees' confidence in the usefulness of GenAI. This skepticism was not limited to individual reluctance but sometimes appeared at the team level: *"there were sometimes entire teams that expressed only resistance and had no desire whatsoever to engage with GenAI"* (P15). Motivation therefore emerges as an individual precondition for sensitization: training formats can support engagement or trigger awareness, but they may not fully compensate for a lack of willingness to participate.

Demographic factors condition sensitization paths asymmetrically, particularly age and prior exposure to technologies. Older or long-tenured employees were portrayed as harder to reach, particularly when they had already struggled with earlier digital tools or did not perceive GenAI as relevant for their work. An HR manager illustrated this with reference to colleagues in their mid-fifties who, in their view, *"still cannot really use Google and say that they will never be able to do so, and that they do not care either"* (P20). In contrast, younger employees, especially those with recent educational exposure to GenAI, were described as more familiar with GenAI, quicker to motivate, and more likely to use it productively. However, participants also emphasized considerable heterogeneity within age groups, suggesting that the central issue is not age alone but the intersection of experience with technologies, perceived usefulness of GenAI, and willingness to engage with new ways of working. This is especially relevant because GenAI use requires more than operating a deterministic tool through familiar interfaces: *"AI is something entirely different, because for AI you cannot write documentation in the sense of a manual that says, if you want A, then follow these steps, click here, right-click there, or use this shortcut to reach the desired outcome. Rather, with AI, one has to understand how to interact with it, what prompting should look like, and so on"* (P07). Hence, demographic factors matter as indicators of heterogeneous learning conditions that sensitization practices must actively address.

Fears and pressures describe employees' affective concerns about GenAI and about the pressure to adopt it, particularly job displacement anxiety and the diffuse pressure of accelerated adaptation. Especially older employees feared job displacement by GenAI. Other participants framed these concerns as part of a broader anxiety about competence, redundancy, and change itself. P15 explained that *"in the end, it is often fear that lies behind it: fear of not being able to manage, fear of becoming redundant oneself, but sometimes also simply fear of any kind of change"* (P15). Employees' fears also stemmed from limited understanding of GenAI itself, including concerns about misuse, surveillance, and potentially harmful outputs. This uncertainty extended to reputational risks, as employees worried that sending hallucinated, inappropriate, or unusually written GenAI-generated emails could damage their credibility among colleagues. Almost all practitioners observed a more general pressure to adapt quickly, captured by one participant as *"engage with it now, rather than waiting until it eventually overwhelms you or your workplace"* (P01). These fears and pressures influence whether employees engage with organizational sensitization practices at all.

Sensitization Practices

While organizational factors encompass the conditions, structures, and constraints provided by the organizational environment that are associated with sensitization practices and the individual factors encompass employees' dispositions as preconditions for organizational sensitization reach, the

sensitization practices constitute active organizational interventions through which organizations attempt to develop, activate, and sustain employees' AI literacy for responsible GenAI use.

Compliance sensitization refers to active organizational interventions through which firms translate regulatory requirements into actionable behavioral orientations at the individual level. Frequently mentioned are the integration of AI competency development into employees' target agreements, codes of conduct as a guiding philosophy, ethics committees and ethical guidelines with regular review cycles, and the integration of the data protection function into sensitization processes. Among these, the integration of compliance into binding organizational structures emerges as a particularly relevant pattern. Several organizations couple GenAI use with structural enforcement mechanisms. For instance, at an automotive company, *"the functional Copilot training was a precondition for being authorized to use it at the workplace"* (P14). Yet, practitioners also articulate the foundational purpose of compliance sensitization as protecting employees from unintended risks rather than disciplining them: *"Employees need to be brought along and prepared about data protection and hallucinations or output validity"* (P01). Hence, organizations deploy *"a code of conduct for handling AI [...] more as a foundational philosophy"* (P05) – often via an intranet.

GenAI provisioning refers to the deliberate provision of proprietary, organizationally governed GenAI systems or access points that organizations deploy as alternatives to ungoverned external GenAI systems. Several organizations provided employees with licenses for GenAI systems (e.g., Claude, Copilot, ChatGPT), hosted on European servers or which the organizations developed, trained, and hosted in-house. By providing governed in-house GenAI systems, organizations attempt to substitute shadow AI use with sanctioned use. P18 illustrated the accompanying effort: *"we operate our own instance in which uploading internal data is permitted. At the same time, employees need to be sensitized not to use a private version and upload company-related information there"* (P18). Others combine proprietary deployment with strict access governance and restrict access to external websites and GenAI systems. In this vein, proprietary GenAI systems together with technical support and helpdesks may also engage employees in experimenting with GenAI without the fear of doing something wrong: *"As a result, the threshold for use was comparatively low for employees who wanted to try it out. During this period, the level of support for technical questions and related explanations was very high"* (P01).

Structured training and experiential learning refer to the planned, organization-orchestrated learning formats through which firms deliver sensitization content systematically. It combines didactic knowledge-transfer formats (e.g., functional sessions, blended-learning programs, certification training, mandatory onboarding modules) with experiential learning formats that activate engagement through challenge-based, gamified, or simulation-based designs (e.g., prompting challenges, security test cases, exploratory contests). All formats share the defining property that they are designed learning interventions with explicit pedagogical intent. Interviewees reported a broad repertoire of measures, including introductory courses on GenAI, prompting and data protection training, tool-specific workshops for Copilot or internal GenAI systems, modules related to certifications, information sessions, office hours, and practice-oriented workshops with concrete use cases. These formats were intended not only to explain what GenAI is, but also to enable employees to assess GenAI's potential, limitations, risks, and implications for their own work. A key pattern was that structured training alone was often insufficient. Practitioners emphasized the importance of hands-on learning, smaller discussion settings, and actionable organizational examples. As P16 explained: *"we conduct hands-on sessions so that participants really experience the effect and see how it is done correctly"* (P16). This was particularly relevant because many organizations did not yet have clearly predefined GenAI use cases. Contrary to initial assumptions, use cases often emerged bottom-up through employees' experimentation in specific departments and everyday work contexts, rather than being fully specified top-down in advance (e.g., by their managers). Hence, experiential and gamified formats reinforced this exploratory logic. For instance, P01 described that *"through playful activities, such as generating the most attractive Easter image, raffles with prizes, prompting challenges, and a test case in which employees had to write prompts to extract a password, the organization tried to actively push engagement"* (P01). Overall, structured training and experiential learning form the pedagogical core of sensitization, turning abstract GenAI discourse into practical understanding and situated experimentation.

Strategic steering and orchestration capture the organizational structures through which sensitization is planned, governed, and embedded into the organizational strategy. Across the interviews,

practitioners emphasized that sensitization cannot be reduced to isolated training formats; it requires decisions about responsibilities, resources, tool access, governance boundaries, and the extent to which organizations assume an explicit educational mandate for building GenAI competencies. This included cross-functional coordination among Information Technology departments, HR, learning teams, legal experts, and, in some cases, in-house advisory units or external consultants. Some organizations also created dedicated GenAI working groups or GenAI innovation departments to bundle expertise, evaluate existing GenAI systems, and support HR in identifying meaningful applications and use cases. A central strategic concern was purposeful deployment. Practitioners repeatedly stressed that GenAI should not be introduced because of its current ‘hype,’ but because it addresses concrete organizational problems. As P07 argued: *“I would advise organizations to really examine their recruiting processes first [...] Only then should they consider whether GenAI can address these pain points in some form and whether this is economically worthwhile, including license costs and the training of employees”* (P07). This problem-orientation also shaped tool selection: organizations had to assess whether a tool fits the specific use case, whether it creates new risks, where data are stored, which vendors are trustworthy, and whether the expected value justifies the costs of implementation, licenses, training, and maintenance. Top-management support was described as a precondition for rendering this orchestration legitimate and sustainable. Finally, continuous communication about GenAI projects was seen as crucial because it *“creates trust and acceptance among employees and helps reduce fear of the unknown”* (P05).

Tensions and Barriers

Although organizations invest substantially in sensitization practices, these practices repeatedly fail to translate into responsible GenAI use. The analysis revealed four tensions and barriers that are (recursively) produced by the same organizational forces that motivate sensitization in the first place.

First, the **adoption-engagement gap** captures the structural conditions under which sensitization practices exist at the organizational level but cannot be meaningfully consumed at the individual level. Structural conditions block employees’ engagement with sensitization practices. However, this gap is not driven by employees’ lack of motivation, ignorance, or resistance. Instead, it is driven by the relationship between sensitization design and work conditions. Through the interviews, *temporal misalignment* of sensitization practices, *recursive pressures*, and *leadership deficits* were identified as important mechanisms. Structured training programs or workshops are frequently scheduled in periods of peak workload or vacation, voluntary participation declines steeply over time, and the time required for substantive engagement competes with operational demands the organization continues to enforce (top-down). P14, whose organization reduced her weekly working hours from 40 to 35 in response to economic pressures, described the misalignment in personal terms: *“Training sessions were offered but during time slots when many people were on vacation [...] I would have preferred being able to choose an available date [...] although this is also difficult in terms of time because of work intensification”* (P14). These access barriers are not failures of sensitization practices alone but an expression of a deeper organizational logic in which the same competitive and efficiency pressures that motivate organizations to engage in sensitization simultaneously generate the increased workload that prevents employees from consuming it. P02 confirmed this paradox: *“the time in daily business is simply missing to approach these tools, to apply them, to playfully try to create things with them. And as long as this time is not available, it will simply be difficult to engage”* (P02). Several practitioners echoed this dynamic with explicit reference to the dependency of engagement on managers’ top-down enforcement: *“[...] we often lack the time for [structured trainings]. If it isn’t, so to speak, ordered from above that one should do it, it is very, very difficult to take the time for it”* (P19). The gap is further amplified where leaders communicate the importance of GenAI sensitization without conveying engagement themselves, producing an interpretive incoherence that disincentivizes employees who recognize the asymmetry between leadership demand and leadership engagement. *“We recently had a workshop with executives, and of these executives, I believe 85%, with our clients, had not even done anything with generative AI in their private lives”* (P16). Arguably, this creates tensions, as P07 mentioned: *“I believe that executives partly expect from their own employees something that they themselves cannot do. Because I don’t believe that executives are good role models in this regard, to say ‘Hey, I am the AI expert. I know exactly how to use AI tools in everyday work’”* (P07). Sensitization, in this configuration, becomes positioned as additional to rather than constitutive of daily work, and the organizational logic of sensitization-as-overhead reproduces the conditions that make engagement structurally more costly.

Second, the **knowledge-practice disconnection** captures the structural conditions under which sensitization practices succeed cognitively, but fail to translate into compliant, governance-conformant practice. Although organizations invest heavily in sensitization practices with the aim of channeling GenAI use toward responsible and governed pathways, employees systematically circumvent those through shadow AI, particularly when organizational GenAI systems are inferior in capability, access is restricted by license barriers, or surveillance mechanisms are missing. Proprietary GenAI systems were frequently reported to lag behind the rapidly evolving external ones, so that the very productivity knowledge sensitization practices deliver intensifies the incentive to circumvent governance. P19 described this as follows: *“I knew what is available on the market, what I can use in my free time. And then I saw what we built internally and thought, well, all right. We could have spared ourselves that”* (P19). P15 framed the universality of the resulting circumvention dynamic: *“I really believe there is no company where there is no shadow AI. [...] Sometimes more, sometimes less”* (P15). Hence, sensitization may produce governance-aware circumvention. The disconnection persists because the practices that organizations would need to detect remain invisible to them. Unlike the adoption-engagement gap, which leaves visible traces such as declining attendance, knowledge-practice disconnection produces few organizational signals. Employees use ungoverned tools through personal browsers, mobile devices, or private accounts, and governance frameworks are often communicated as advisories rather than enforced through tool blocking or behavioral monitoring. P16 captured the diagnostic moment when structured training surfaced the disconnection: *“we do indeed have shadow AI usage at some of our clients, where it actually says AI use is not allowed – and then you go into a training thinking you’ll explain it to them, and they say: ‘But we all do it that way because the browser isn’t blocked’”* (P16). This also appears when employees identify meaningful GenAI use cases in training but cannot apply them because the required tools are not yet licensed or approved: *“They were able to transfer the topics directly to their work [...] The problem, however, is that at the client organization the GenAI tools have not yet been approved or are still in the approval process”* (P17). Beyond capability and visibility, the disconnection is amplified by operational pressure: the competitive and efficiency demands that motivate sensitization persist after sensitization has been delivered, and under acute time pressure sensitization-derived knowledge becomes a constraint that employees experience as costly. *“When in doubt, opportunity makes love [...] I already have an 80-hour week and I’ll just dump an unredacted client document in there”* (P16). Operational pressure not only hinders sensitization but converts its knowledge into governance-aware shortcuts.

Third, the **heterogeneity misalignment** captures the structural inadequacy of uniform sensitization formats for a diverse workforce. Although organizations often rely on scalable formats such as standardized e-learning modules, common curricula, or workforce-wide workshops, employees differ substantially in prior GenAI exposure, skills, role-specific use cases, hierarchical position, and domain expertise. This makes sensitization difficult to standardize, as a *“watering-can principle where [a sensitization practice] applies to everyone equally is not the way”* (P20). In HR, GenAI needs to be examined *“from the perspective of employees, but also from the perspective of executives and applicants”* (P01), and a one-size-fits-all format can hardly address practical opportunities, ethical risks, and leadership responsibilities in equal depth. Uniform formats fail because the same intervention may seem too basic or too abstract depending on employees’ prior GenAI exposure. At the higher-literacy end, P18 noted that an internal workshop was *“for 70% [...] already very familiar”* (P18). At the lower end, P14 described a 45-minute online module as *“not for hands-on practice or anything”* (P14). Thus, the problem is not merely poor calibration, but that no single level can serve all groups. Making structured training more advanced loses the lower-literacy end (e.g., *“there was a range of lectures [...], but at some point, I cognitively dropped out of these highly technical topics”* (P01)), while making it more accessible loses the higher-literacy end. This misalignment intensifies with the size of an organization: larger and more diverse workforces create stronger economic incentives for uniform formats. In multinational corporations with multiple subsidiaries, multi-tenant complexity further amplifies this misalignment: the organization with the greatest sensitization needs faces the strongest cost pressures toward uniformity and the strongest empirical reasons against it.

Lastly, the **identity-efficiency tension** is conceptually different from the previous three tensions because it does not operate at the level of how sensitization is organizationally delivered or individually received, but at the level of what is being sensitized to. Sensitization practices simultaneously prepare employees to use GenAI productively and seek to preserve the values that GenAI threatens to erode (e.g., human-centric brand commitments, authenticity of digital communication, candidate-centric processes, and human-in-

the-loop principles in HR contexts). Hence, this tension must be navigated and cannot be resolved through sensitization alone. It becomes visible when GenAI-created outputs conflict with employees' normative expectations of the organization or their organizational self-understanding. A participant illustrated this through the example of employer branding in his organization: *"[Company Name] essentially has its branding globally uniform. With GenAI they generated contextually different images for each country. [...] But at some point, they realized the images appear so artificial. [...] [Company Name] stands for being an authentic company, and they then halted the whole thing. Because the generated images do not match our brand"* (P01). The same tension surfaces at the level of decisional authority in recruiting, where employees reported that GenAI use for curriculum vitae screening conflicts with human contextual judgment: *"AI ultimately performs a very black-and-white analysis in recruiting. [...] He was on a sabbatical for three years, but you find it interesting that he did this – that you wouldn't have noticed with AI because you carry a much more standardized evaluation"* (P01). Beyond brand and decisional authority, the tension extends into employees' subjective perception of authorship and authenticity. P16 mentioned that *"one quickly has this train of thought: 'Ah, this is somehow cheating when I use AI. That is the shortcut. That is then not my work'"* (P16). P06 extended the same dynamic to the candidate-document side, where authorship ambiguity destabilizes professional evaluation itself: *"[...] to no longer see what I receive as the intellectual property of the sender. If you receive an application letter from a candidate that is great and excites you, but you always have to keep in mind: you don't know if he wrote it himself"* (P06). The same tension also appeared in everyday workplace communication, where efficient GenAI-generated text could violate local communication norms. P14 described a GenAI-assisted email as *"really exaggerated"* and filled with *"too much embellishment,"* adding that she would *"never speak in such an inflated way"* and that a business letter should *"get to the point"* (P14). Because she feared colleagues might think differently of her, she added that the email was *"partially AI-generated"* as a form of self-protection in case they noticed any unconventional paragraphs (P14). Thus, the outcome of sensitization does not necessarily increase GenAI use. Even highly sensitized and AI-literate employees may deliberately refrain from using GenAI when identity-efficiency tensions emerge.

Discussion

Contributions and Implications for Research

Insights from semi-structured interviews with HR experts (i.e., an underexplored GenAI user group; Pinski & Benlian, 2024) suggest that GenAI sensitization may be understood as a distinct process occupying the theoretically underdeveloped middle ground between structured training and AI literacy. While IS research has conceptualized AI literacy as a multidimensional continuum of competencies (e.g., Long & Magerko, 2020; Pinski & Benlian, 2024) and structured training as a mechanism to impart knowledge about GenAI systems, the organizational process through which employees' AI literacy is activated for responsible GenAI use has remained undertheorized. HR practitioners themselves recognize this distinction: whereas training *"can be completed, checked off, and then we forget the content,"* sensitization creates *"a genuinely deeper understanding"* that is *"more effective for the organization in the long term"* (P19). Sensitization may thus be understood as the operational counterpart to the multi-literacy "toolbox" (Passlack et al., 2026, p. 30), meaning that sensitization is neither training alone nor AI literacy, but the recursive bridge between both through which competencies are made actionable. Organizational sensitization practices do not, however, automatically translate into responsible GenAI use. Four structural tensions (i.e., the *adoption-engagement gap*, the *knowledge-practice disconnection*, the *heterogeneity misalignment*, and the *identity-efficiency tension*) emerged that characterize the conditions under which this translation fails. These tensions are not correctable implementation gaps but properties of the relationship between sensitization practices and their associated "factors at individual [and] organizational" (Passlack et al., 2026, p. 29) levels. Participants' accounts indicated that the same organizational forces that motivate organizational sensitization practices in the first place (e.g., competitive pressure, productivity expectations, regulatory policies) simultaneously produce the structural conditions under which sensitization might be undermined. By specifying how organizational and individual factors as well as sensitization practices interact and which structural tensions emerge at their interface, we extend recent AI literacy conceptualizations in two directions. First, where prior research asks whether personal dispositions are present as a condition for responsible GenAI use, our findings show that responsibility gaps emerge even when those dispositions are present. The *knowledge-practice disconnection* demonstrates that

governance-aware, AI-literate employees are systematically routed around governed pathways by operational pressure and capability gaps between proprietary and external GenAI systems. The *identity-efficiency tension* shows, conversely, that high AI literacy can also *constrain* GenAI use when efficiency gains conflict with professional identity or organizational values. Second, where culture has been theorized as a moderator of literacy development (Passlack et al., 2026), our results show it can simultaneously produce structural tensions. Human-centric cultures motivate sensitization while intensifying the *identity-efficiency tension* (e.g., P01, P04), whereas tech-affine cultures facilitate experimentation while exposing literate employees to generic sensitization formats that appear redundant (e.g., P17, P18). Our four tensions thus articulate structural conditions under which the operationalization of employees' responsible GenAI use needs to be addressed (Papagiannidis et al., 2025). Interestingly, our findings challenge existing findings on shadow AI. IS research has framed shadow AI either as a sociotechnical governance failure that can be addressed through controlled enablement (Silic et al., 2025) or as a behavioral risk reducible through privacy awareness training (Glorin, 2026). Both presuppose that circumvention results from inadequate governance design or insufficient employee awareness. Our findings challenge this premise. Across several interviews, highly AI-literate employees who fully understood governance rules nonetheless bypassed them because governed alternatives were perceived as inferior in capability or convenience. In this configuration, sensitization may not reduce circumvention but instead amplify it. By increasing awareness of GenAI's productivity potential, organizations may inadvertently increase the incentive to use ungoverned tools when governed alternatives do not suffice. Theoretically, this pattern sharpens the understanding of why "governance drift zones": a condition in which policies exist, but fail to align with the informal, emergent nature of AI tool use" (Silic et al., 2025, pp. 11–12), emerge even in organizations with mature governance policies and proprietary GenAI systems. The *knowledge-practice disconnection* identifies the structural mechanism producing this drift: a productivity-knowledge asymmetry between governed and ungoverned tools that communicative intervention may not be able to resolve. Intervention therefore requires structural changes, such as governed tools competitive with ungoverned alternatives, dedicated time for engagement, and continuous integration of sensitization into everyday work, not just better training design.

Implications for Practice: The Sensitization Profile Matrix

To synthesize individual sensitization profiles that can guide practitioners in assessing their pathway to responsible GenAI use, the first two authors independently analyzed each participant's GenAI use (frequency, routinization, variety, and work integration) and AI literacy (understanding of capabilities and limitations, governance awareness, critical output evaluation, and reflexive engagement). Through iterative discussions, four empirically grounded GenAI sensitization profiles emerged inductively: the *Uninitiated*, *Naïve Adopters*, *Cautious Gatekeepers*, and *Responsible Innovators* (see Figure 2; observed participants in brackets). For practitioners, the sensitization profile matrix may serve as both a diagnostic tool to assess individual sensitization profiles and a typology of sensitization pathways.




Figure 2. Sensitization Profile Matrix

Uninitiated employees combine low GenAI use with low AI literacy. They have limited practical exposure to GenAI and lack a sufficiently developed understanding of its functionality, risks, and use cases. This profile often reflects weak or inaccessible sensitization architectures, where training exists but is not meaningfully consumed due to time constraints, poor scheduling, insufficient leadership modeling, or lack of perceived relevance. The *adoption-engagement gap* is particularly consequential here: sensitization opportunities may formally exist, but structural conditions prevent employees from engaging with them. To move toward responsible GenAI use, organizations should therefore prioritize structural enablement over content delivery as a precondition for any intervention, e.g., allocating manager-sanctioned time, anchoring programs to role-specific use cases, and deploying relatable peer champions. This can ensure that sensitization reaches employees at the point of use rather than requiring initiative they structurally lack. Without such preconditions, the absence of organizational sensitization does not prevent use but redirects it: employees may begin experimenting through private or ungoverned tools, moving toward the naïve adopter quadrant instead of responsible use.

Naïve adopters demonstrate comparatively high GenAI use but low AI literacy. They use GenAI routinely, often because it offers immediate productivity gains, but without sufficiently reflecting on data protection, hallucinations, bias, output validity, authorship, or task appropriateness. This sensitization profile may be the most critical from a responsible GenAI use perspective, because the risks of GenAI are amplified when frequent use is not accompanied by evaluative and governance-related competence. Naïve adopters may not have participated in organizational sensitization practices due to low awareness, engagement, or other factors. However, if employees participated in or received sensitization practices, the *knowledge-practice disconnection* may explain why naïve adopters potentially understand formal rules or risks associated with GenAI use, but still deviate from them under operational pressure, especially when ungoverned tools appear faster, more capable, or more convenient than organizationally approved alternatives. Thus, movement from naïve adopter to responsible innovator requires sensitization that prioritizes meta-competencies (e.g., output validation, bias recognition, and governance-oriented judgment) over functional tool training. Complementarily, governed in-house tools that credibly compete with external alternatives can reduce the productivity gap that drives circumvention, allowing sensitization-derived knowledge to be applied within governed pathways. Depending on the intensity and risk exposure of naïve adopters' use, it might be suitable to first reduce their GenAI use and to focus on strengthening their AI literacy, before once again moving them towards expanding GenAI use.

Cautious gatekeepers combine high AI literacy with low GenAI use. Unlike the uninitiated, their limited use does not stem primarily from ignorance or lack of exposure. Rather, they understand GenAI's potential but deliberately restrict their own use because of governance concerns, regulatory sensitivity, cultural commitments, or unresolved identity-efficiency tensions. This profile is especially relevant in HR, where GenAI use intersects with high-stakes decisions, candidate relationships, personal data, and human-centric organizational values. The *identity-efficiency tension* helps explain why higher AI literacy does not automatically lead to higher GenAI use. Employees may be sufficiently sensitized to recognize that certain efficiency gains would undermine professional authorship, authenticity, candidate-centric processes, or human-in-the-loop decision principles. In such cases, non-use or selective use can represent a responsible outcome rather than a sensitization failure. To support movement toward the responsible innovator, organizations should provide explicit normative roadmaps explaining where GenAI use is endorsed, where human judgment is non-negotiable, and where context determines the balance. Where such clarity is absent, high AI literacy paradoxically reinforces restraint, leaving the most governance-aware and AI-literate employees as the least productive GenAI users.

Responsible innovators use GenAI regularly and often across multiple tools or tasks, but do so with critical awareness of its limitations, governance requirements, and organizational fit. They validate outputs, avoid inappropriate data disclosure, distinguish suitable from unsuitable use cases, and preserve human judgment in consequential HR contexts. This profile therefore does not represent unrestricted adoption but calibrated and context-sensitive use. Responsible innovators emerge when sensitization practices are supported by enabling organizational conditions: governed infrastructure, practical training, experiential learning, leadership support, strategic steering, and sufficient space for reflection. They can navigate rather than eliminate the structural tensions that arise around GenAI use. They may also act as relatable peer champions for the uninitiated or for cautious gatekeepers.

Conclusion, Limitations, and Avenues for Future Research

This study demonstrates that formal training and governance alone are insufficient to ensure responsible GenAI use, as structural tensions often compel even highly AI-literate employees to bypass corporate systems under operational pressure. To translate individual AI literacy into responsible organizational practice, firms must move beyond standardized interventions and implement tailored, context-specific sensitization frameworks. Yet, this study has three limitations that simultaneously open avenues for future research. **First**, the research design is cross-sectional: we document sensitization practices and structural tensions through semi-structured interviews but cannot assess which interventions actually mitigate which tensions over time. Longitudinal research tracking sensitization practices and responsible use outcomes across organizational cycles would allow causal claims that our grounded account cannot support. **Second**, our sample is bounded by German-speaking, HR-specific organizational contexts, which are characterized by strong regulatory conformance pressures (e.g., GDPR, EU AI Act). The structural tensions identified here may manifest differently in less regulated national contexts or in organizational functions with lower data protection sensitivity. Comparative studies across departments, industries, and regulatory environments would test the transferability of the sensitization framework and reveal whether the tensions are features of the HR domain specifically or of GenAI sensitization more broadly. **Third**, the profile matrix may not cover all sensitization archetypes. For instance, P19 described a colleague who organizes his entire working routine through a public GenAI system, entering the notes of every meeting to generate summaries and task lists. P19 characterized him as more knowledgeable about GenAI than P19 himself and confirmed that the two had discussed hallucination and data protection risks explicitly. The colleague nonetheless continued, withholding only the company name as a self-imposed safeguard. This case combines high AI literacy with high GenAI use, the constellation our matrix assigns to the responsible innovator, yet the underlying conduct is its normative opposite. The two dimensions therefore do not discriminate between calibrated use and informed, deliberate circumvention, which calls for further research on additional profiles and the tensions they produce. This aligns with Kim and Ryoo's (2026, p. 1) finding of the paradox that "AI literacy, while intended to foster critical thinking, can instead heighten heuristic judgments, amplifying polarized perceptions of AI-generated [outputs]." Future research should examine whether intentional, literacy-informed circumvention constitutes additional profiles or a third dimension (e.g., task-specific criticality) and explore the conditions under which AI literacy ceases to function as a safeguard for responsible use.

References

- Banh, L., Rex, A., Strobel, G., & Urbach, N. (2026). Hiring Tomorrow's Talents: How Generative Artificial Intelligence Transforms Human Resources Recruitment. *Proceedings of the 59th Hawaii International Conference on System Sciences*, 5482–5492. <https://doi.org/10.24251/HICSS.2026.653>
- Banh, L., & Strobel, G. (2023). Generative artificial intelligence. *Electronic Markets*, 33(1), 63. <https://doi.org/10.1007/s12525-023-00680-1>
- Benlian, A., & Pinski, M. (2025). The AI literacy development canvas: Assessing and building AI literacy in organizations. *Business Horizons*. Advance online publication. <https://doi.org/10.1016/j.bushor.2025.10.001>
- Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Managing Artificial Intelligence. *MIS Quarterly*, 45(3), 1433–1450. <https://doi.org/10.25300/MISQ/2021/16274>
- Birks, D. F., Fernandez, W., Levina, N., & Nasirin, S. (2013). Grounded theory method in information systems research: Its nature, diversity and opportunities. *European Journal of Information Systems*, 22(1), 1–8. <https://doi.org/10.1057/ejis.2012.48>
- Budhwar, P., Chowdhury, S., Wood, G., Aguinis, H., Bamber, G. J., Beltran, J. R., Boselie, P., Lee Cooke, F., Decker, S., DeNisi, A., Dey, P. K., Guest, D., Knoblich, A. J., Malik, A., Paauwe, J., Papagiannidis, S., Patel, C., Pereira, V., Ren, S., ... Varma, A. (2023). Human resource management in the age of generative artificial intelligence: Perspectives and research directions on ChatGPT. *Human Resource Management Journal*, 33(3), 606–659. <https://doi.org/10.1111/1748-8583.12524>
- Burton-Jones, A., & Grange, C. (2013). From Use to Effective Use: A Representation Theory Perspective. *Information Systems Research*, 24(3), 632–658. <https://doi.org/10.1287/isre.1120.0444>
- Choudhary, S. K., & Kar, A. K. (2025). Modeling the role of generative AI in organizational privacy and security. *Decision Support Systems*, 196, 114500. <https://doi.org/10.1016/j.dss.2025.114500>

- Dell'Acqua, F., McFowland III, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayner, L., Candelon, F., & Lakhani, K. R. (2026). Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality. *Organization Science*, 37(2), 403–423. <https://doi.org/10.1287/orsc.2025.21838>
- Engström, A., Pittino, D., Mohlin, A., Johansson, A., & Edh Mirzaei, N. (2024). Artificial intelligence and work transformations: Integrating sensemaking and workplace learning perspectives. *Information Technology & People*, 37(7), 2441–2461. <https://doi.org/10.1108/ITP-01-2023-0048>
- Felin, T., Foss, N. J., & Ployhart, R. E. (2015). The Microfoundations Movement in Strategy and Organization Theory. *Academy of Management Annals*, 9(1), 575–632. <https://doi.org/10.5465/19416520.2015.1007651>
- Feuerriegel, S., Hartmann, J., Janiesch, C., & Zschech, P. (2024). Generative AI. *Business & Information Systems Engineering*, 66(1), 111–126. <https://doi.org/10.1007/s12599-023-00834-7>
- Gartner. (2025, November 19). *Gartner Identifies Critical GenAI Blind Spots That CIOs Must Urgently Address*. <https://www.gartner.com/en/newsroom/press-releases/2025-11-19-gartner-identifies-critical-genai-blind-spots-that-cios-must-urgently-address>
- Gioia, D. A., Corley, K. G., & Hamilton, A. L. (2013). Seeking Qualitative Rigor in Inductive Research: Notes on the Gioia Methodology. *Organizational Research Methods*, 16(1), 15–31. <https://doi.org/10.1177/1094428112452151>
- Glorin, S. (2026). Digital shadow AI risk theory (DART): A framework for managing data disclosure and privacy risks of AI tools at work. *Technological Forecasting and Social Change*, 229, 124697. <https://doi.org/10.1016/j.techfore.2026.124697>
- Grange, C., Demazure, T., Ringeval, M., Bourdeau, S., & Martineau, C. (2026). The Human-GenAI Value Loop in Human-Centered Innovation: Beyond the Magical Narrative. *Information Systems Journal*, 36(1), 29–51. <https://doi.org/10.1111/isj.12602>
- Hacker, P., Engel, A., & Mauer, M. (2023). Regulating ChatGPT and other Large Generative AI Models. *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency*, Chicago, USA, 1112–1123. <https://doi.org/10.1145/3593013.3594067>
- Heimbach, I., Abramova, O., Baumann, A., Krasnova, H., Jussupow, E., Legner, C., Meythaler, A., Müller, O., & Pinski, M. (2026). Algorithmic Literacy: A Compass to Successfully Navigate the Algorithm-Driven World? *Business & Information Systems Engineering*, 68(2), 465–471. <https://doi.org/10.1007/s12599-025-00984-w>
- Holmes, W., & Miao, F. (2023). *Guidance for generative AI in education and research* (Vol. 1). Unesco Publishing.
- Kim, W., & Ryoo, Y. (2026). The literacy paradox: How AI literacy amplifies biases in evaluating AI-generated news articles. *International Journal of Information Management*, 86, 102992. <https://doi.org/10.1016/j.ijinfomgt.2025.102992>
- Knoth, N., Decker, M., Laupichler, M. C., Pinski, M., Buchholtz, N., Bata, K., & Schultz, B. (2024). Developing a holistic AI literacy assessment matrix – Bridging generic, domain-specific, and ethical competencies. *Computers and Education Open*, 6, 100177. <https://doi.org/10.1016/j.caeo.2024.100177>
- Krakowski, S. (2025). Human-AI agency in the age of generative AI. *Information and Organization*, 35(1), 100560. <https://doi.org/10.1016/j.infoandorg.2025.100560>
- Kulal, A., Dinesh, S., & Abhishek, N. (2025). Organizational Impact of AI-Driven Recruitment Practices: A Mixed-Methods Study. *Journal of Computer Information Systems*, 1–16. <https://doi.org/10.1080/08874417.2025.2508860>
- Long, D., & Magerko, B. (2020). What is AI Literacy? Competencies and Design Considerations. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, Honolulu, 1–16. <https://doi.org/10.1145/3313831.3376727>
- Maitlis, S., & Christianson, M. (2014). Sensemaking in Organizations: Taking Stock and Moving Forward. *Academy of Management Annals*, 8(1), 57–125. <https://doi.org/10.5465/19416520.2014.873177>
- Myers, M. D., & Newman, M. (2007). The qualitative interview in IS research: Examining the craft. *Information and Organization*, 17(1), 2–26. <https://doi.org/10.1016/j.infoandorg.2006.11.001>
- Naß, M., Pohlmann, N., Scheja, S., & Harter, A. (2026). Security Awareness als institutionelles Anliegen: Voraussetzungen, Herausforderungen und Gestaltungsräume für NRW-Hochschulen. *HMD Praxis der Wirtschaftsinformatik*. <https://doi.org/10.1365/s40702-026-01261-4>

- Papagiannidis, E., Mikalef, P., & Conboy, K. (2025). Responsible artificial intelligence governance: A review and research framework. *The Journal of Strategic Information Systems*, 34(2), 101885. <https://doi.org/10.1016/j.jsis.2024.101885>
- Paré, G. (2004). Investigating Information Systems with Positivist Case Research. *Communications of the Association for Information Systems*, 13(1), 233–264. <https://doi.org/10.17705/1CAIS.01318>
- Passlack, N., Hammerschmidt, T., & Posegga, O. (2026). With Great Power Comes Great Responsibility: What Shapes AI Literacy for Responsible Interactions of Knowledge Workers With AI? *Information Systems Frontiers*, 28(1), 11–46. <https://doi.org/10.1007/s10796-025-10648-5>
- Pinski, M., & Benlian, A. (2023). AI Literacy—Towards Measuring Human Competency in Artificial Intelligence. *Proceedings of the 56th Hawaii International Conference on System Sciences*, Maui, Hawaii, 165–175. <https://doi.org/10.24251/HICSS.2023.021>
- Pinski, M., & Benlian, A. (2024). AI literacy for users – A comprehensive review and future research directions of learning methods, components, and effects. *Computers in Human Behavior: Artificial Humans*, 2(1), 100062. <https://doi.org/10.1016/j.chbah.2024.100062>
- Retkowsky, J., Hafermalz, E., & Huysman, M. (2024). Managing a ChatGPT-empowered workforce: Understanding its affordances and side effects. *Business Horizons*, 67(5), 511–523. <https://doi.org/10.1016/j.bushor.2024.04.009>
- Scholl, M. (2023). Sustainable Information Security Sensitization in SMEs: Designing Measures with Long-Term Effect. *Proceedings of the 56th Hawaii International Conference on System Sciences*, Maui, Hawaii, 6058–6068. <https://doi.org/10.24251/HICSS.2023.735>
- Shapiro, H., Souto-Otero, M., & Watermeyer, R. (2026). Metacognitive AI literacy: Going beyond the AI skills gap agenda. *Learning, Media and Technology*, 1–15. <https://doi.org/10.1080/17439884.2026.2652638>
- Sharma, A. J., Rathore, B., Biswas, B., Xiao, X., & Gupta, R. (2026). Generative AI Implementation in Enterprises: Lessons From a Case Study of Enhanced IT Service Management. *Information Systems Journal*, 36(6), 661–690. <https://doi.org/10.1111/isj.70029>
- Silic, M., Silic, D., & Kind-Trüller, K. (2025). From Shadow IT to Shadow AI—Threats, Risks and Opportunities for Organizations. *Strategic Change*, 1–16. <https://doi.org/10.1002/jsc.2682>
- Software AG. (2024). *Chasing Shadows: Understanding and Managing Shadow AI*. Software AG. <https://www.it-finanzmagazin.de/wp-content/uploads/2024/12/Shadow-AI-White-Paper.pdf>
- Söllner, M., Arnold, T., Benlian, A., Bretschneider, U., Knight, C., Ohly, S., Rudkowski, L., Schreiber, G., & Wendt, D. (2025). ChatGPT and Beyond: Exploring the Responsible Use of Generative AI in the Workplace: An Interdisciplinary Perspective. *Business & Information Systems Engineering*, 67(2), 289–303. <https://doi.org/10.1007/s12599-025-00932-8>
- Storey, V. C., Yue, W. T., Zhao, J. L., & Lukyanenko, R. (2025). Generative Artificial Intelligence: Evolving Technology, Growing Societal Impact, and Opportunities for Information Systems Research. *Information Systems Frontiers*, 27(5), 2081–2102. <https://doi.org/10.1007/s10796-025-10581-7>
- Tsohou, A., Karyda, M., Kokolakis, S., & Kiountouzis, E. (2015). Managing the introduction of information security awareness programmes in organisations. *European Journal of Information Systems*, 24(1), 38–58. <https://doi.org/10.1057/ejis.2013.27>
- Vuorikari, R., Kluzer, S., Punie, Y., & Dahl, M. R. (2022). *DigComp 2.2, The Digital Competence framework for citizens: With new examples of knowledge, skills and attitudes*. Publications Office of the European Union. <https://doi.org/10.2760/115376>
- Weber, P., Pinski, M., & Baum, L. (2023). Toward an Objective Measurement of AI Literacy. *Proceedings of the Pacific-Asia Conference on Information Systems*, Nanchang, China. <https://aisel.aisnet.org/pacis2023/60>
- Wu, H., Li, D., & Mo, X. (2025). Understanding GAI risk awareness among higher vocational education students: An AI literacy perspective. *Education and Information Technologies*, 30(10), 14273–14304. <https://doi.org/10.1007/s10639-024-13312-8>
- Xhigoli, D., & Schaarschmidt, M. (2025). The Dark Side of AI: Unveiling the Impact of Dark Triad Traits on GenAI Concealment Behavior. *Proceedings of the Academy of Management Conference*. Academy of Management Annual Meeting, Copenhagen, Denmark. <https://doi.org/10.5465/AMPROC.2025.13675abstract>