

AI Adoption Across a Multinational Workforce: Sociotechnical Conditions for GenAI Acceptance in Human Resources

Dalia Ali^{1*†}, Maria José Rodríguez Velázquez^{1,*}, Manoel Horta Ribeiro³,
Q. Vera Liao², Orestis Papakyriakopoulos¹

¹Technical University of Munich

²University of Michigan

³Princeton University

Abstract

Generative AI (GenAI) deployment in the workplace is accelerating rapidly. Nevertheless, questions of *who adopts, who benefits, and who is left behind and why* are still understudied. In this paper, we investigate these dynamics in the context of a multinational tech company transitioning from a legacy Human Resources (HR) search system to a GenAI-supported system, analyzing search log data, survey data ($n=25$), and ten semi-structured interviews. Our findings show that adoption depended on the fit between the GenAI system’s design assumptions and employees’ work positionalities (role, spoken language, tenure). Further, we find that employees’ trust in GenAI answers was built through source-checking, comparison among systems, and seeking input from colleagues or HR when in doubt. Our contribution is twofold. First, we provide empirical evidence of workplace GenAI adoption during a live organizational transition, showing that adoption is influenced by individual and contextual factors such as situational fit, search literacy, and trust calibration. It is also shaped by knowledge conditions such as the system’s content quality, employee training, and guidance. Second, we translate these findings into design considerations for inclusive deployment and adoption in high-stakes environments such as HR. We argue that organizations should design systems considering the role- and context-sensitive benefits they yield to different social groups. They also need to treat the organizational knowledge infrastructure as AI infrastructure to improve the accountability and usability of GenAI systems.

Extended version: <https://arxiv.org/pdf/2606.17887>

Introduction

Artificial Intelligence (AI) is becoming a general-purpose layer of organizational infrastructure, transforming processes and decision-making (Enholm et al. 2022). Recent studies suggest that AI capabilities are improving rapidly, enabling systems to handle increasingly complex tasks with limited human intervention (Douglas and Verstyuk 2025; Kwa et al. 2025). Improvements in AI capabilities can therefore affect organizations at multiple levels. At the task level, AI can increase productivity by automating routine cognitive work (Brynjolfsson, Li, and Raymond 2025; Lee et al.

2022; Ali et al. 2025). At the employee level, these systems may offer real-time assistance, feedback, and communication support, facilitating learning and increasing job satisfaction, especially for less-experienced employees (Brynjolfsson, Li, and Raymond 2025; Bhargava, Bester, and Bolton 2021). Last, at the organizational level, AI can make information more widely available across teams and support coordination between departments, reducing some of the silos and expertise imbalances that often limit organizational learning and collaboration (Dell’Acqua et al. 2025).

Yet, adoption lags behind technological progress. Recent industry reports show gaps in AI access, exploration, and exploitation. While workforce access to AI technologies is growing, few companies manage to move from the pilot to the production stage; 65% of senior leaders struggle to tie productivity gains directly to AI adoption (Institute 2026; MIT Project NANDA 2025; Ernst & Young LLP 2025). Previous work shows that adoption is shaped by different factors such as perceived usefulness and ease of use, social influence, data, technology infrastructure, organizational culture, ethics, and regulation (Enholm et al. 2022; Venkatesh 2022; Davis 1989). However, these explanations do not fully capture the complex, interconnected, and dynamic environments in which technology is implemented; further, they offer limited insight into the complex situated user experiences and their fit with available technologies (Lee, Ramasamy, and Subbarao 2025; Shachak, Kuziemy, and Petersen 2019; Mogaji et al. 2024).

To bridge this gap, we need to understand not only how employees use AI but also what drives different employees’ decisions to adopt, reject, or selectively use AI during organizational transitions. Without such understanding, organizations risk failing to realize the value of their investments, missing the benefits of increased efficiency and performance, and reproducing uneven patterns of adoption among employees (Enholm et al. 2022; Mlekus et al. 2020). From the employee perspective, poor GenAI design may mean that a system does not support the task at hand, produces outputs that are difficult to evaluate or trust, or excludes some employees due to a lack of language fit, accessibility, and work-context needs (Gu et al. 2024; Tankelevitch et al. 2024; Liao and Sundar 2022; Weidinger et al. 2021; Budhwar et al. 2023; Sarkar et al. 2022). Since barriers

*Equal contribution.

†Corresponding author: dalia.ali@tum.de

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

ers to GenAI adoption already exacerbate existing inequalities across workers, inclusive design is needed to ensure that efficiency benefits are not concentrated among some employees while others are left behind (Humlum and Vestergaard 2025).

To examine AI adoption from a user-centered perspective and understand why some employees benefit from GenAI systems while others are left behind, we study a real organizational transition within a multinational tech company from an existing knowledge system for HR to one that incorporates GenAI capabilities into its retrieval process. The new GenAI-based search system relies on HR documents from the internal organizational repository to enable employees to access HR-related information faster and more accurately. Existing organizational leadership's willingness to implement an innovative AI solution and the availability of necessary technological and financial resources allow us to investigate actual sociotechnical factors that facilitate or impede adoption, rather than boundary conditions of technological availability. Specifically, we raise the following research question:

Why do users adopt or resist a GenAI-enhanced HR knowledge system during an organizational transition?

By combining one year of the employees' search log data from an existing knowledge search system, two weeks of data from the new GenAI-based search system, an exploratory survey, and ten semi-structured interviews with project team members and end-users, we show that AI adoption was not a simple migration from an old to a new system. Instead, it was characterized by selective use across parallel systems, situated fit across user roles, social contexts, and the technologies, user's existing GenAI search literacy, how users calibrated their trust in the new system, and organizational conditions such as content quality and user guidance.

Our paper makes two main contributions: (1) we identify sociotechnical mechanisms through which workplace GenAI produces unequal benefits, showing that when systems are designed around assumptions of office-based access, digital confidence, organizational familiarity, and language fit, workers outside these assumptions are systematically less well served despite formal access; (2) we translate the findings into design considerations for inclusive and accountable workplace GenAI systems.

Background

AI Adoption in Organizations

As AI capabilities continue to advance rapidly, adoption remains uneven across both organizations and workers (Humlum and Vestergaard 2025; Institute 2026; MIT Project NADA 2025). To understand this variation, it becomes necessary to consider adoption at different levels. At the organizational level, for instance, the Technology-Organization-Environment (TOE) framework helps identify factors related to technology, organization, and the environment which influence adoption of new technology (Tornatzky and Fleischer 1990; Baker 2011). However, TOE does not capture individual-level adoption decisions within

organizations (Awa, Ojiabo, and Orokor 2017; Oliveira and Martins 2011).

At the individual level, the Technology Acceptance Model (TAM) and the Unified Theory of Acceptance and Use of Technology (UTAUT) explain adoption through perceived usefulness, ease of use, performance expectancy, effort expectancy, social influence, and facilitating conditions (Davis 1989; Venkatesh et al. 2003). While these models show high predictive power for traditional software, they have not been designed for AI systems, which, due to their probability, opacity, and contextualization, change the nature of how users interact with the technology (Venkatesh 2022). Extensions of these models to AI domain have taken into account issues related to trust, transparency, and perceived risk, recognizing that users must cope with probabilistic outputs, limited explainability, and unpredictable behavior (Venkatesh 2022; Enholm et al. 2022; Wolfe et al. 2025). Nonetheless, the efficient use of GenAI technologies introduces additional interaction requirements that cannot be entirely accommodated by the previous trust models or adoption models (Tankelevitch et al. 2024; Zamfirescu-Pereira et al. 2023; Sarkar et al. 2022; Weisz et al. 2023).

Research on AI in the workplace consistently shows that these systems do not impact all workers equally. Several studies document how GenAI reshapes knowledge work (Woodruff et al. 2024), and how algorithmic systems create failure loops that systematically undermine certain types of work (Kawakami et al. 2024). In HR research, most scrutiny has focused on algorithmic hiring (Raghavan et al. 2020; Sanchez-Monedero, Dencik, and Edwards 2020), and emotion AI in interviews has been shown to result in lower perceptions of justice among gender minorities (Ingber and Andalibi 2025).

Despite evidence that AI affects workers unequally, existing adoption models still provide limited insight into how adoption unfolds through everyday interactions among users, technology, and organizational context (Shachak, Kuziemy, and Petersen 2019; Mogaji et al. 2024). They do not adequately explain how users learn, adapt, develop workarounds, or incorporate AI systems into practice (Shachak, Kuziemy, and Petersen 2019); nor how adoption is shaped by users' expertise, job role, task suitability, and organizational conditions (Bankins et al. 2024).

User-centered considerations of AI systems

Addressing this gap requires moving beyond current adoption models and examining how users engage with AI systems in practice. User-centered AI adoption depends on more than access. How well employees benefit from the AI system depends on their ability to interpret its outputs, build trust, learn a new interaction logic, and rely on the organizational knowledge infrastructure underlying the system (Tankelevitch et al. 2024; Liao and Sundar 2022; Alavi, Leidner, and Mousavi 2024). Although these factors are largely studied when people use systems, prior research does not take them into consideration when studying adoption dynamics. Rather, the focus lies more on defining isolated social and technical factors that shape adoption (Venkatesh 2022; Wolfe et al. 2025).

Research on Human-AI interactions shows that the level of trust in AI systems is not a function of system accuracy; is also influenced by how users interact with the system (Liao and Sundar 2022). Liao and Sundar (2022) differentiate between trustworthy AI, which is related to system accuracy and fairness, and human trust, which requires perception and communication. They propose the MATCH model, in which trust arises from the interplay among model features, system affordances, trustworthiness cues, and users' cognitive processing. (Bansal et al. 2021) also show that explanations, such as supporting information features that promote trust in AI, do not automatically produce appropriate reliance; they can increase users' acceptance of AI recommendations even when those recommendations are incorrect. This means that reliance on the AI output needs to be studied empirically, rather than assumed based on the system's accuracy.

A second factor is the interaction competencies GenAI demands of its users. Unlike traditional information retrieval systems, GenAI requires users to prompt, evaluate outputs, reformulate queries, and decide whether answers are reliable enough to act on (Tankelevitch et al. 2024; Zamfirescu-Pereira et al. 2023). The development of such competence is not straightforward, as employees must develop proficiency in prompting, output evaluation, and knowing when GenAI is appropriate (Xia et al. 2026). This is made more challenging because workers use GenAI within existing tasks rather than as a standalone activity (Xia et al. 2026).

Another factor to consider is the organizational knowledge that GenAI systems draw from. Information systems are only as effective as the knowledge supporting them (Alavi and Leidner 2001); and GenAI systems are no exception, since they retrieve and synthesize content from underlying content such as articles, metadata, and localized content instead of generating answers independently (Alavi, Leidner, and Mousavi 2024). Recent work argues that GenAI may change how organizations manage knowledge by making retrieval and recombination more conversational and automated (Alavi, Leidner, and Mousavi 2024).

Jointly, this literature identifies user-centered factors that shape how people use and evaluate GenAI systems. Given this, we investigate how the above and additional sociotechnical factors interact to shape adoption during a live HR organizational transition across a structurally diverse workforce, in which a GenAI system is introduced alongside a legacy system.

Methodology

To understand how employees adopt or resist GenAI in organizational knowledge work, we conducted a mixed-method single-case study at a multinational technology organization with more than 300,000 employees worldwide. This case centered on the organization's introduction of PEOPLE TOOL¹, an HR platform with an embedded GenAI On Search feature, alongside the continued availability of its legacy HR search system, STEVE² PEOPLE TOOL worked

¹PEOPLE TOOL is a pseudonym for the newer GenAI-enhanced HR system.

²STEVE is a pseudonym for the legacy HR knowledge system.

as a retrieval-augmented system where the queries of employees resulted in either GenAI output with sources or search recommendations alone, based on whether there was enough material for generating a response.

We conducted a three-phase sequential design. First, we analyzed STEVE and PEOPLE TOOL search logs to identify behavioral patterns across query formulation, result availability, click-through behavior, and GenAI answer exposure. Based on these patterns, we conducted an exploratory survey with 25 employees across six countries to assess usage, usefulness, usability, trust, system preference, and employee reactions when the answer generated by GenAI was insufficient. Afterward, we conducted semi-structured interviews with 10 employees from four countries to explore how employees perceive the two systems and how personal, organizational, and technological factors affect adoption paths. Since each method served a different purpose, the three data sources cover different country contexts and are treated as complementary rather than directly comparable country-level samples. This study received ethical approval from our institution, and all data were handled in accordance with GDPR requirements.

Search Log Analysis

Search logs are system-level logs that capture information about employees' search activities. Each search log recorded details of queries entered by employees, the systems in which queries were searched, the search timing, country context, and system feedback, such as whether the system returned results or displayed a GenAI response. Search log analysis served as the pre-study phase of our sequential mixed-methods design. Its purpose was exploratory: to surface behavioral patterns in how employees queried each system, with findings directly shaping the design of the survey instrument and the interview protocol.

We obtained search logs from both systems covering six countries: Egypt, the Netherlands, the Philippines, Pakistan, Qatar, and Saudi Arabia. STEVE logs spanned approximately one year (September 2024–August 2025). The logs of PEOPLE TOOL collected over a period of two weeks (July 24–August 7, 2025) due to a hard constraint placed by the provider of the system with regard to their data retention policy. For cross-system volume comparisons, we constrained data from both systems to the observation window of (July 24–August 2, 2025), and we kept the extended baseline data for STEVE separately to prevent conflating differences in system usage with variations in observation period.

Non-English searches were translated to English by a native speaker in order to enable cross-country comparisons at the topic level. The study considered volume and distribution of queries within the two systems, the structure and length of the queries, availability of results, the exposure rate to GenAI responses, and overlap of topics between the two systems. The log data gives an account of what the employees searched for and the results they obtained from each system; however, it does not give any information regarding their reasons for using one system over another, nor their perception of the results; those questions were addressed in the survey and interviews.




Figure 1: Sequential mixed-methods study design. Each phase informed the next, with search log patterns shaping the survey and survey responses shaping the interview protocol.

Survey

Following up from the initial search log analysis, we conducted an exploratory survey among employees who can access the PEOPLE TOOL and the search features of the GenAI integrated into it. Employees were sampled from seven different countries, namely Germany, the US, India, Austria, Switzerland, Mexico, and France. The purpose of this survey was two-fold: the first was data gathering about their experience with STEVE and PEOPLE TOOL, and the second was developing an interview guide.

The survey collected contextual information, including country of work, job level, and tenure, and employees were required to report on their usage frequency for each system, their top HR topics searched for, and their ratings on the systems' ease of use, helpfulness, and trustworthiness. Particular emphasis was given to the use of the GenAI search through the PEOPLE TOOL, during which the participants were asked about the frequency with which they received the answer from the GenAI function, the quality of the provided answers, the frequency of clicking sources after getting the answer, and the action they would take if the answer was insufficient. The participants were free to point out any issues they encountered, their preference for one system over another, topics that did not yield useful results, and areas for improvement. Finally, a call was made to those willing to participate in the interviews.

The survey was answered by 25 employees. For the closed-ended questions, we conducted descriptive statistics using counts and percentages. For the open-ended questions, we analyzed responses to identify recurring concerns that might reappear in interviews. The survey results cannot be considered statistically representative due to their limited sample size and unequal representation. Instead, they serve as descriptive evidence for the interviews and help connect behavior identified in the logs with the process discussed in the interviews (Appendix B, Table 1, Limitations Section).

Measure	PEOPLE TOOL	STEVE
Daily or weekly use	19/25 (76%)	14/25 (56%)
Preferred system	14/25 (56%)	3/25 (12%)
Very/extremely helpful	16/25 (64%)	8/25 (32%)
Easy/very easy to use	16/25 (64%)	16/25 (64%)
Mean helpfulness	3.76	3.12
Mean trust	3.92	3.72
<i>GenAI-specific items</i>		
Never/rarely clicked sources in GenAI answers		13/25 (52%)
Would rephrase when GenAI answer was poor		19/25 (76%)

Table 1: Key descriptive survey results used to contextualize the findings.

Note. Survey data are descriptive, not statistically representative.

Preference responses were mixed or unclear for 7/25 (28%).

Helpfulness and trust means are on 5-point scales (see Appendix B).

Interview Study

Semi-structured interviews were the primary source of interpretive data in this research. While search logs and surveys provided information on behavior and contextual descriptions, interviews helped examine how people interpreted STEVE, PEOPLE TOOL, and the GenAI search feature in practice. Interviewees were selected among those who agreed to participate in the interviews in addition to the survey, as well as from employees involved in everyday use of PEOPLE TOOL.

We conducted ten semi-structured interviews using an on-line videoconferencing tool. Each interview was recorded, lasted between 20 and 35 minutes, and was conducted in English or Spanish, at the participant's preference. At the beginning of each interview, participants were reminded of the purpose of the study, informed that participation was voluntary, and asked to consent to the recording for transcription and analysis.

The interview guide was organized around six key themes, which involved participants’ background and familiarity with digital technologies; participant experience and comparison of STEVE and PEOPLE TOOL; perceived usefulness and limitations of GenAI on Search; trust, source-checking, and information behavior; adoption dynamics and organizational context; and reflections on future improvement.

Interview participants represented two broad perspectives. The first group comprised employees involved in strategy, rollout, knowledge management, system ownership, testing, or support for PEOPLE TOOL; however, these participants also used the system in their own work, and could therefore speak both as practitioners with implementation knowledge and as users seeking HR-related information. The second group comprised employees whose engagement with the systems was primarily as users seeking HR-related information. We use this distinction to contextualize participants’ accounts, but we do not conduct a systematic comparative analysis between these groups. Participants varied in tenure, ranging from less than 1 year to more than 15 years, and were distributed across four countries. Table 2 provides an overview of the interview participants.

We generated the interview transcriptions automatically and manually checked them for accuracy. We anonymized them by removing all personal identifiers and assigned new participant IDs ranging from P1 to P10. Minor corrections were made to enhance readability without altering the content bearing the message. To analyze the interview data, we used reflexive thematic analysis following Braun and Clarke (Braun and Clarke 2006). We first read the transcripts repeatedly to familiarize ourselves with the data, then generated descriptive codes capturing participants’ accounts of system comparison, perceived usefulness, ease of use, trust, source checking, fallback practices, AI-search literacy, training and communication, organizational support, content quality, work context, language, and attitudes or fears towards AI. These initial codes were then iteratively grouped into more focused sub-themes.

This analysis was conducted over multiple rounds of discussions among four researchers. Instead of considering coding as a single classification process, the discussions helped us compare understandings, merge redundant codes, push against poorly formed themes, and link sub-codes to broader analytical themes that relate to the research question. The outcome of this discussion led us to the final themes presented in the Findings section: selective use across parallel systems, situated fit across workers and contexts, learning a new GenAI search logic, reliance on checking and fallback, and knowledge conditions.

Findings

By analyzing interviews, survey responses, and the search logs of both systems, we observed that adoption was not a linear process of migration from legacy search to GenAI-enhanced search. As both systems, PEOPLE TOOL and STEVE, were available, employees could switch between them, compare their outcomes, verify the responses, and return to alternative routes when necessary. Adoption was

ID	Generalized role context	Tenure (yrs)	Country
<i>Project team participants</i>			
P1	HR IT / digital transformation; application ownership	1–5	Germany
P2	Transition project member	5–15	Germany
P3	HR IT intern; content migration and AI testing	1–5	Germany
P4	Product ownership for GenAI and knowledge management	15+	Germany
P7	HR support leadership	5–15	Spain
P8	Regional HR IT coordination	5–15	USA
<i>End-user participants</i>			
P5	Digitalization-related role in adjacent ServiceNow environment	5–15	USA
P6	Compensation and HR tooling role	5–15	USA
P9	HR systems project management and cloud migration	1–5	Spain
P10	Early-career P&O systems professional	<1	Norway

Table 2: Overview of interview participants ($n = 10$).

Note. Role descriptions are intentionally generalized to preserve participant anonymity.

shaped by the interaction between employees, the system, and the organization.

In this section, we first show that adoption involved selective use across parallel systems rather than a complete migration to the new system. Then, we examine why this use varied across employees and contexts. Next, we argue that effective use requires employees to learn a new GenAI search logic. This, in turn, raises questions of trust, as employees had to determine when an answer was reliable enough to act upon. Finally, we show that these everyday adoption practices were shaped by organizational knowledge conditions.

Selective Use Across Parallel Systems

Since STEVE continued to function after the transition, employees were not required to rely on a single route for HR information at all times. Instead, they chose among the available routes depending on the situation. For instance, employees would use PEOPLE TOOL if it provided speed, integration, and clearer direction. They would use STEVE if the search process were familiar, the results were predictable and article-based, or if it were a route they already knew how to navigate. Source articles provided precise information when needed, and employees reached out to colleagues, managers, or the HR department for personal or complex answers. Thus, adoption was not a single decision point, but an ongoing practice.

Survey responses reflected this parallel-use pattern. Indeed, 76% of the respondents (19/ 25) used PEOPLE TOOL either daily or weekly. However, STEVE was not discontinued as 56% of respondents (14/25) used STEVE daily or weekly. Further, regarding preferences, 56% of respondents preferred PEOPLE TOOL, whereas 12% preferred STEVE.




Figure 2: Sociotechnical dimensions of workplace GenAI adoption. Adoption was not a simple migration from the legacy STEVE system to the GenAI-enhanced PEOPLE TOOL. Employees selectively combined the two systems with source articles and human support. The usefulness of these routes varied with employees’ work conditions, information needs, organizational familiarity, and language context. For effective use, it was necessary to acquire new search skills, while dependence was adjusted by the means of source checking, comparison with other systems and asking colleagues or HR specialists for help. Across these dimensions, reliable use depended on the quality of the underlying organizational content and the availability of timely guidance and training.

It appears that although PEOPLE TOOL was widely used, the legacy system was still used as part of the routine HR information-seeking workflow.

Interviews also elucidate the underlying logic by showing that switching to a new system was worthwhile. Employees were not asking only whether PEOPLE TOOL was new or AI-powered, they were interested in its advantages relative to familiar information-search methods. One participant explained his concerns as follows:

“Oh man, now there’s another HR system (PEOPLE TOOL) I have to use. You know, I’m comfortable with this one. Why do we have to change again?” (P5)

However, this response cannot be regarded as a sign of general resistance to AI. Learning a new route takes time and effort that employees simply do not have to devote to a task of lower priority than their main work. Since a familiar route could be accessed with a single click, the new route needed to provide sufficient value. If PEOPLE TOOL offered clear answers and useful information, employees could find a valid reason to use it; but when it added more navigation and cognitive burden without a clear benefit, familiar routes remained the easier choice.

Therefore, the simultaneous availability of PEOPLE TOOL and STEVE meant that adoption was not a single migration decision from the legacy to the new GenAI system, but an ongoing practice of choosing among routes. This raises a question: *what made PEOPLE TOOL the preferred route for some employees but not others?*

Situated Fit Across Workers and Contexts

The reasons behind selective adoption were rooted in situational fit. We define situational fit as the degree to which the tool’s offerings match users’ actual work conditions, needs, and contexts. PEOPLE TOOL was designed for laptop and computer access, assumed organizational familiarity, and prioritized some languages over others, resulting in a fit that varied considerably across the workforce.

A recurring concern throughout the interviews was that PEOPLE TOOL was better suited to office-based employees who regularly had access to laptops, were accustomed to the company’s internal knowledge system, and could frame their HR inquiries in organizational terminology. For shift-based, factory workers, and shop-floor employees, this was not the case. Participants suggested that these worker groups relied on their managers, team assistants, local HR staff, or informal channels. For instance, one participant noted that blue-collar workers may not be familiar with organizational HR terminology, making it difficult to even formulate a question *“when you’re a blue-collar, you’re not as exposed to that side of things and so even knowing what to ask”* (P6). Another participant suggested that PEOPLE TOOL was *“not ideal”* for factory workers (P1). These responses suggest that although the system was technically available to all employees, its design for office-based, digitally familiar, and organizationally embedded users could lead to exclusion, with employees with diverse work conditions at a disadvantage in using the tool.

Thus, making PEOPLE TOOL formally available to employees did not mean that all employees could use it equally

in practice. Some employees lacked regular access to a suitable device “workers who go on site to work...they sometimes don’t even have a laptop” (P5); others did not know how to phrase HR questions in organizational terminology (P6), and others continued to rely on managers, local HR staff, or colleagues as their main route to information. These employees could disappear from system usage logs or aggregate adoption figures, even if the tool was technically available to them. As a result, aggregate adoption figures therefore overstate inclusion, as a high query volume from office-based employees masks the absence of shift-based, blue-collar, and multilingual workers who never meaningfully engaged with the system.

Tenure and life stage played an important role in shaping situational fit as well. Because new recruits were still learning the organization’s internal processes and the location of HR documentation, they could gain the most from PEOPLE TOOL. On the other hand, tenured employees used the PEOPLE TOOL less, not because they resisted it, but because their information uncertainty had declined, and they already knew the company’s policies, the personnel responsible for addressing HR concerns, and the routine mechanisms to resolve issues. As one participant explained:

“The older you are, the less changes you have in life. I would say the kids are gone, house is built, car is there. What shall you ask HR? Maybe the younger people have more interactions with HR because they are still building their lives.” (P4)

The GenAI tool’s situated fit with employees was also evident from the logs. First, search behavior was distinct by country. For instance, in the Netherlands, users searched for policy and process-related topics, in Saudi Arabia and Qatar, searches focused on documents and certificates, in Egypt, on benefits and insurance, and in Pakistan, on travel claims. System-level indicators were distinct as well, for instance, in Pakistan, click-through was the lowest across all countries at 18.5%, suggesting employees were less likely to move from search results to further interaction with the system. These patterns show that a global HR search system encountered diverse local task environments, with employees searching for HR artifacts using distinct terminologies and addressing diverse requirements.

Language was another factor that impacted adoption. If a system failed to accommodate the user’s language or did not reflect local HR terminologies, its effectiveness suffered despite its technical capabilities. In multinational rollouts, language support is typically added incrementally, with non-English and non-European languages often deprioritized or absent altogether; one participant emphasized that “if you don’t have the languages, then people sometimes don’t use it” (P4). Together, the interview and log data suggest that a one-size-fits-all design for GenAI is ineffective in multinational HR knowledge work. The same system provided convenience to some users while creating friction or exclusion for others.

Thus, situational fit shaped the selective use of PEOPLE TOOL. When the tool’s offerings aligned with users’ device access, language, work environment, life stage, and HR in-

formation needs, it gained greater traction. Otherwise, employees relied on legacy systems or local sources that they were already familiar with. This raises the next question, *even when the tool aligned with a user’s context, what did it take for employees to effectively use GenAI HR search?*

Effective Use Required Learning a New GenAI Search Logic

Access to PEOPLE TOOL did not automatically translate into effective use. Even when the system fit a user’s context, employees still had to learn a new way of searching that differed fundamentally from STEVE. This involved phrasing questions, adding context, rephrasing failed queries, interpreting summaries, and knowing when to move from a GenAI response to source articles or human support. Thus, adoption was based not only on what the system offered but also on whether employees could translate their HR needs into questions the system could answer; and as the survey data show, these two dimensions, usefulness and ease, did not always align.

This distinction helps explain why usefulness and ease did not necessarily occur in parallel. The survey showed that PEOPLE TOOL was rated as “very” or “extremely” useful by 64% of respondents (16/25), compared with 32% (8/25) for STEVE. However, both systems were rated as “easy” or “very easy” by the same proportion of respondents, 64% (16/25). In other words, employees recognized the added value of PEOPLE TOOL GenAI search, but realized that value required learning a new interaction logic, including how to ask, interpret, and adjust their searches.

The survey further illustrates this interaction effort. When GenAI failed to provide a satisfactory answer, 76% of the respondents (19/25) reported they would rephrase their question. This demonstrates that users viewed poor answers as something that could be improved by asking differently, rather than a system failure. By doing so, employees participated in creating the response, making the answer dependent not only on the system itself, but also on employees’ knowledge of what to ask and how. The remaining 24% did not rephrase, suggesting that some employees disengaged entirely when the system failed to provide a useful answer.

The logs show that this new search logic had not yet fully replaced older keyword-search habits. We measured query length in terms of tokens, defined as units similar to words in a search request. The average query length increased only slightly, from 1.87 tokens in STEVE to 2.11 tokens in PEOPLE TOOL. Similarly, the share of short one- to two-token queries stayed practically unchanged, 83.4% in STEVE and 83.3% in PEOPLE TOOL. This suggests that employees continued to use traditional keyword-based search practices even in the GenAI-enabled environment, although short keyword queries may also have remained effective given the system’s retrieval-augmented architecture. The interviews explain this, as a participant described the change as a shift in search logic:

“Once they [the user] have understood that the logic changes, that the way of searching and finding is different, I think there is greater acceptance” (P7)

Effective use, therefore, required more than familiarity with the system; employees had to develop a new form of search competence, including learning not only where to search, but how to search with AI, how to recover when the answer was poor, and how to judge whether an answer was enough for their HR question. The survey evidence supports this; while 64% of respondents rated PEOPLE TOOL as useful, the same proportion rated both systems as equally easy to use, suggesting that perceived value and interaction competence are distinct dimensions of adoption. This competence was not about learning where to search, but about developing an ongoing practice of prompting, evaluating, and deciding. This leads to the next issue: even when employees learned how to ask, *how employees decided whether a GenAI answer was reliable enough to use?*

Users Built Trust by Checking, Comparing, and Fallback

Receiving a GenAI answer was not sufficient for adoption. Employees also had to decide whether that answer was accurate enough to act on. In HR, acting on incorrect or partial information can have serious implications for salary, leave, and personal circumstances. Therefore, adoption depended on employees' ability to move from obtaining an answer from the GenAI tool to deciding whether to rely on it. This reliance was not based on blind acceptance but was calibrated through checks, comparisons, and fallback.

The survey suggests a gap between usefulness and trust. Mean helpfulness was higher for PEOPLE TOOL than STEVE: 3.76 compared with 3.12. By contrast, the difference in mean trust was more modest, 3.92 compared with 3.72. This suggests that employees may perceive greater value in PEOPLE TOOL while remaining cautious about whether its answers are complete, current, and applicable to their situation. This caution also appeared in the interviews as one participant stated:

"I will not sell my house because of an AI answer. There's always a double net. There's always checks and balances, talking to managers, talking to colleagues, talking to my wife, consulting someone, friends who made this already. So it's not just the answer in [People Tool] leading to the decision. It's always a mix of information and then you come to your own conclusion." (P4)

Source-checking was one way of managing this uncertainty, although it was not used consistently. In the survey, 52% of respondents (13/25) reported that they never or rarely clicked the sources shown in GenAI-generated answers, while 48% (12/25) clicked them sometimes or often. This difference should not be viewed simply as a split between trusting and distrusting users; sometimes, employees did not click sources because the GenAI answer was sufficient for their needs and required no further verification. Others coped with this uncertainty in different ways, such as inspecting sources, relying on plausibility, restructuring their question, comparing the answer with what is known from other systems, and seeking help from others.

At the same time, trust was not always actively calibrated. Some employees appeared to lend credibility to PEOPLE TOOL because it was an internal company system. As one participant explained: *"I think people innately trust things internally in [the Company] that it's going to give them the correct answer, and I think people are trusting AI too much."* (P6). Yet, the organization placed responsibility for interpreting outputs back onto employees as one participant pointed out: *"We always tell the users, 'You can use AI, but the risk is on you.' How can I force them to make use if they know that they have the responsibility?"* (P4). Employees sometimes trust PEOPLE TOOL GenAI responses not because they have reviewed their source or assessed their completeness, but because they were delivered through an approved organizational platform; this, in turn, creates the risk that fluent but incorrect answers would be accepted without sufficient verification.

The interviews also identified the continued need for a fallback. A GenAI summary could be useful for quick orientation, particularly when the question was general or low-risk. But when questions are specific, complex, or high-risk, employees find greater satisfaction in seeking reassurance from source documents, STEVE, peers, supervisors, or HR. Hence, trust was distributed across the answer itself and the wider set of information routes surrounding it.

Trust was also shaped by prior experiences and first impressions. Some employees approached GenAI with curiosity, while others began with skepticism, privacy concerns, or worries about unreliable HR guidance. These starting points did not determine adoption on their own, but they influenced how employees interpreted their initial interactions with the system. Repeated use could revise initial expectations, but only when the system behaved consistently enough to support confidence. When GenAI answers were wrong or irrelevant, employees would disengage quickly. When answers appeared plausible but remained uncertain, employees had to decide whether to verify, rephrase, or ask someone else.

In general, reliance on PEOPLE TOOL's GenAI search was built through repeated use rather than assumed a priori. Employees delegated some information work to PEOPLE TOOL, but also checked, compared, or returned to other routes when uncertain. This raises another question: *what made GenAI's answers reliable or unreliable in the first place?*

Knowledge as a Condition for Reliable GenAI Use

The previous analysis showed that employees needed mechanisms to validate, compare, and trace answers back to different sources when GenAI-produced answers are ambiguous. Thus, there was an even deeper issue about the trustworthiness of PEOPLE TOOL that was not dependent on the GenAI feature alone. Employees accessed HR information through GenAI, which in turn depended on the quality of the underlying HR articles, metadata, tagging, taxonomy, localization, training, communication, and escalation routes. When these were well-structured, this allowed the generation of useful guidance via the GenAI process; however, uncoordinated input data, outdated articles, poor tagging, and vague explanations made it difficult for the GenAI tool to

compensate for these issues.

This is supported by survey data, which shows that while the PEOPLE TOOL proved to be more effective than STEVE (mean helpfulness of the tool has risen from 3.12 to 3.76), trust increased only moderately (up from 3.72 to 3.92). In other words, PEOPLE TOOL made HR information easier to find, but employees remained uncertain about whether the answers were complete, current, and applicable to their specific situation. Employee responses focused on the quality of information architecture more than on the AI technology itself, as one employee noted:

“In PEOPLE TOOL the information is not that user-friendly and as each item might be under a different article, it’s hard to find what you want.” (Survey respondent)

Another employee linked poor search results to weak tagging, *“Not every document or article has the tags that I would search for and therefore they are not shown in the search results”* (Survey respondent). These responses illustrate why GenAI reliability cannot be separated from content quality. Employees sometimes attributed poor answers to the AI when the underlying problem was poorly organized, outdated, or localized HR content.

Training and guidance were also part of these organizational conditions, not just rollout activities. Employees needed to know more than that PEOPLE TOOL existed; they needed to understand what had changed, how to search, when GenAI answers would appear, when to check sources, and when to escalate to human support. Interviewees suggested that support was most useful when provided close to the moment of use, through examples, short guidance, Q&A sessions, office hours, prompts, and clear contact routes. One-off pre-launch communication was less effective because it could be forgotten before employees actually needed the tool.

Therefore, content quality, guidance, and training were not background conditions for GenAI adoption; they determined whether the answers employees received were reliable enough to act on. PEOPLE TOOL could make HR search faster and more convenient, but only when the underlying content was structured, maintained, localized, clearly communicated, and employees were supported at the point of use.

Discussion

Who Benefits from Workplace GenAI?

Recent research on AI fairness and sociotechnical design argues that potential harms arising from AI cannot be understood only by examining model behavior; the social and organizational context in which these systems are deployed must also be considered (Selbst et al. 2019; Costanza-Chock 2020). Similarly, work in workplace HCI has shown how algorithms affect workers’ experiences within organizations and the management settings and infrastructure of workplaces (Cheon and Erickson 2025). Research on digital inclusion has also highlighted the mistake of considering technological access alone is sufficient for understanding inclusion and emphasized the roles of skills, social context,

conditions of use, and the ability to translate technologies into concrete gains (Warschauer 2003; van Dijk 2005; Helsper 2021). Our findings extend this argument to workplace GenAI deployment: making a GenAI system available across an organization does not mean that all employees are equally positioned to benefit from it.

In our study, formal availability of the GenAI HR system did not translate equally into useful HR support. As highlighted in Section 4.2, employees’ ability to take advantage of the system depended on its alignment with work processes, linguistic preferences, HR-related information requirements, and access to support channels. This raises an equity issue that remains challenging to operationalize. *When the system is designed for office-based, digitally confident, and English-speaking users, how will organizations recognize which workers are at a disadvantage?* This question is difficult to answer through aggregate usage metrics alone, because those who face the greatest friction may be characterized by their absence, inactivity, and infrequent use. A system may therefore appear broadly adopted while still providing more reliable and actionable support to some groups and excluding the others (Selbst et al. 2019; Humlum and Vestergaard 2025). The relevant question is not only who can access the system, but who can obtain useful, trustworthy, and actionable support from it in everyday work.

Design recommendations. Organizations should evaluate their GenAI deployment in the workplace by considering role-sensitive and context-sensitive benefit measures. This includes examining which groups are missing from usage logs, which users receive search results but not GenAI answers, which queries fail or require repeated reformulation, who continue to rely on local intermediaries, and which groups do not receive useful or actionable support. Then, the less visible employee segment should be included continuously in testing, feedback, and redesign, rather than being treated as edge cases after deployment.

Who Carries the Risk of Getting It Wrong?

Previous studies on human-AI interaction show that effective use of AI involves calibrated reliance, avoiding both over-reliance on incorrect outputs and under-reliance on useful ones (Lee and See 2004; Bansal et al. 2021). The factors that impact trust are not limited to system efficiency or performance but also include information communicated about the AI system, such as explanations, source visibility, and institutional credibility (Liao and Sundar 2022). With GenAI specifically, users face additional metacognitive demands, including prompting, evaluating outputs, and deciding whether an answer is reliable enough to act on (Tankelevitch et al. 2024). However, in reality, as illustrated by our findings (Section 4.4), reliance extended beyond the user-system interaction; we observed that employees were moving between the GenAI system PEOPLE TOOL, legacy systems STEVE, documents, coworkers, and HR personnel, depending on task priority and uncertainty. This raises a question worth discussing: *When an employee navigates this network for HR information and still acts on incorrect information, who is responsible? The system designer, the organization that deployed it, or the employee who used it?*

This question is consequential in HR contexts, where employees may act on AI-generated information about high-stakes decisions such as pay, leave, benefits, or employment conditions (Budhwar et al. 2023; Dasaklis et al. 2025). When the AI-generated answers appear fluent and delivered through an approved system by the organization, institutional trust can replace the careful verification by the employee (Knowles and Richards 2021; Liao and Sundar 2022). Without clear information on source legitimacy, established escalation procedures, and robust accountability structures within the organization, responsibility for accuracy rests with employees rather than with the organization that deployed the system.

Design recommendations. At a system design level, workplace GenAI systems should make provenance, uncertainty, and fallbacks apparent as default features. Source URLs need to be seen as an accountability architecture, not simply a UI-based feature. Also, the GenAI answers themselves must provide information about their origin, verifiability, and whether human intervention is recommended. At the organizational level, it is important to maintain existing human channels for assistance where possible, rather than replacing them with AI. It is also critical for organizations to establish clear lines of accountability when AI systems are involved in decision-making in high-risk areas, so that the extent to which the system can and cannot be relied upon is clear, and interpretive risk does not fall on individual workers by default.

Knowledge Readiness as AI Infrastructure

As stated in (Section 4.5), participants did not view PEOPLE TOOL as a stand-alone AI model. Users interacted with PEOPLE TOOL through HR artifacts, including articles, metadata, tagging, taxonomy, localization, source references, training guides, communication, and escalations. When these elements were well structured, GenAI could make HR information faster and easier to access. On the contrary, when some of these artifacts were absent, out-of-date, poorly labeled, or difficult to navigate, the GenAI system would generate incomplete answers or answers that did not address employees' needs.

This means that knowledge management becomes a critical part of accountability for workplace GenAI implementations (Alavi, Leidner, and Mousavi 2024). In high-stakes domains, content quality is more than a back-office implementation issue; it determines whether the information summarized by AI systems is complete, accurate, up-to-date, and reliable for action (Tankelevitch et al. 2024). Similarly, training and communications are not limited to mere rollout considerations. Employees need to be able to frame their questions, interpret the results generated, check their sources, and raise uncertainties. Without such considerations, companies may place AI-related uncertainties on individuals' shoulders.

Design recommendations. Knowledge infrastructure should be considered part of the core AI infrastructure. This includes keeping HR content up to date, assigning ownership of articles, improving tagging and taxonomies, localizing terminology, auditing unsuccessful queries, and estab-

lishing feedback loops between poor AI answers and knowledge owners. Training must be continuous and timely with respect to its use. This involves using brief instructions, examples, rewording help, Q&A, and readily available help lines. In corporate GenAI applications, responsible deployment requires preparing the organization around the AI, not just for the AI itself.

Limitations

This study has several limitations. First, it is based on a single organizational case. This limits statistical generalizability, but the case is valuable because it captures GenAI adoption during a live workplace transition, where employees could compare the new system with an existing legacy system. Second, the number of survey respondents ($n = 25$) was relatively small and unevenly distributed across nationality, job roles, and work experience within the company. For that reason, the survey data were treated as descriptive only and triangulated with data from interviews and search logs to enable interpretation. Third, the logs captured different usage periods for both systems. While STEVE's logs capture a longer span of searches, PEOPLE TOOL logs had to be confined due to ServiceNow's retention policies. To address this limitation, we restricted cross-system comparisons to the same observation window and used the longer STEVE logs only as historical context. Finally, some worker groups central to situational fit, such as shop-floor, shift-based, or less digitally embedded employees, were not systematically sampled as direct participants. Some interviewees held roles that exposed them to these work contexts, so their accounts provide informed yet indirect evidence of uneven fit. Future studies would do well to include such groups directly to study whether workplace GenAI systems reproduce or reduce inequalities across different labor conditions.

Conclusion

In this study, we examine how employees across a multinational workforce adopt, resist, and selectively use a GenAI-enhanced HR knowledge system during a live organizational transition from a legacy HR knowledge system to a new GenAI-enhanced system. Through our analysis of both systems' search logs, a survey, and ten semi-structured interviews, we show that several factors shaped the adoption process, including situational fit, search literacy, trust calibration, and deployment conditions, such as content quality, user guidance, and accountability structures. Drawing on these findings, we argue that equitable workplace GenAI deployment cannot be assessed solely by access to the system. Our results raise questions about how organizations should define adoption success, who is responsible for wrong AI-generated HR information, and what conditions are necessary for GenAI to be deployed reliably and equitably in high-stakes domains.

References

Alavi, M.; and Leidner, D. E. 2001. Review: Knowledge Management and Knowledge Management Systems: Con-

- ceptual Foundations and Research Issues. *MIS Quarterly*, 25(1): 107–136.
- Alavi, M.; Leidner, D. E.; and Mousavi, R. 2024. A Knowledge Management Perspective of Generative Artificial Intelligence. *Journal of the Association for Information Systems*, 25(1): 1–12.
- Ali, D.; Ahmed, M.; Wang, H.; Khan, A.; Jordan, N. P. A.; Kim, S. S.; Muchhala, M. D.; Merkle, A. K.; and Papakyrakopoulos, O. 2025. AI Adoption Across Mission-Driven Organizations. *arXiv preprint arXiv:2510.03868*.
- Awa, H. O.; Ojiabo, O. U.; and Orokor, L. E. 2017. Integrated technology-organization-environment (TOE) taxonomies for technology adoption. *Journal of Enterprise Information Management*, 30(6): 893–921.
- Baker, J. 2011. *The technology–organization–environment framework*, 231–245. Springer.
- Bankins, S.; Ocampo, A. C.; Marrone, M.; Restubog, S. L. D.; and Woo, S. E. 2024. A multilevel review of artificial intelligence in organizations: Implications for organizational behavior research and practice. *Journal of Organizational Behavior*, 45(2): 159–182.
- Bansal, G.; Wu, T.; Zhou, J.; Fok, R.; Nushi, B.; Kamar, E.; Ribeiro, M. T.; and Weld, D. 2021. Does the Whole Exceed Its Parts? The Effect of AI Explanations on Complementary Team Performance. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*.
- Bhargava, A.; Bester, M.; and Bolton, L. 2021. Employees’ Perceptions of the Implementation of Robotics, Artificial Intelligence, and Automation on Job Satisfaction, Job Security, and Employability. *Journal of Technology in Behavioral Science*, 6: 106–113.
- Braun, V.; and Clarke, V. 2006. Using Thematic Analysis in Psychology. *Qualitative Research in Psychology*, 3(2): 77–101.
- Brynjolfsson, E.; Li, D.; and Raymond, L. 2025. Generative AI at Work. *The Quarterly Journal of Economics*, 140(2): 889–942.
- Budhwar, P.; Chowdhury, S.; Wood, G.; Aguinis, H.; Bamber, G. J.; Beltran, J. R.; Boselie, P.; Cooke, F. L.; Decker, S.; DeNisi, A.; Dey, P. K.; Guest, D.; Knoblich, A. J.; Malik, A.; Paauwe, J.; Papagiannidis, S.; Patel, C.; Pereira, V.; Ren, S.; Rogelberg, S.; Saunders, M. N. K.; Tung, R. L.; and 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.
- Cheon, E.; and Erickson, I. 2025. Fulfillment of the Work Games: Warehouse Workers’ Experiences with Algorithmic Management. *Proceedings of the ACM on Human-Computer Interaction*, 9.
- Costanza-Chock, S. 2020. *Design Justice: Community-Led Practices to Build the Worlds We Need*. Cambridge, MA: The MIT Press.
- Dasaklis, T. K.; Giannopoulos, P. G.; Koutras, D.; Malamas, V.; and Chountalas, P. T. 2025. Large Language Models in Human Resource Management: A Systematic Literature Review of Applications, Open Issues and Future Research Directions. Preprint submitted to Elsevier.
- Davis, F. D. 1989. Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3): 319–340.
- Dell’Acqua, F.; Ayoubi, C.; Lifshitz, H.; Sadun, R.; Mollick, E.; Mollick, L.; Han, Y.; Goldman, J.; Nair, H.; Taub, S.; and Lakhani, K. R. 2025. The Cybernetic Teammate: A Field Experiment on Generative AI Reshaping Teamwork and Expertise. Technical Report w33641, NBER.
- Douglas, M. R.; and Verstyuk, S. 2025. Progress in Artificial Intelligence and its Determinants. *arXiv preprint arXiv:2501.17894*.
- Enholm, I. M.; Papagiannidis, E.; Mikalef, P.; and Krogstie, J. 2022. Artificial Intelligence and Business Value: a Literature Review. *Information Systems Frontiers*, 24: 1709–1734.
- Ernst & Young LLP. 2025. The Dividend Age: How AI Is Turning Promise into Payoff. Us ai pulse survey: Wave 4, EY. Survey of 500 US senior business leaders.
- Gu, K.; Grunde-McLaughlin, M.; McNutt, A.; Heer, J.; and Althoff, T. 2024. How Do Data Analysts Respond to AI Assistance? A Wizard-of-Oz Study. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, CHI ’24. ACM.
- Helsper, E. J. 2021. *The Digital Disconnect: The Social Causes and Consequences of Digital Inequalities*. London: SAGE Publications.
- Humlum, A.; and Vestergaard, E. 2025. The unequal adoption of ChatGPT exacerbates existing inequalities among workers. *Proceedings of the National Academy of Sciences*, 122(1): e2414972121.
- Ingber, A. S.; and Andalibi, N. 2025. Emotion AI in Job Interviews: Injustice, Emotional Labor, Identity, and Privacy. In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAccT)*.
- Institute, D. A. 2026. State of AI in the Enterprise: The Untapped Edge. Technical report, Deloitte AI Institute.
- Kawakami, A.; Taylor, J.; Fox, S.; Zhu, H.; and Holstein, K. 2024. AI Failure Loops in Feminized Labor: Understanding the Interplay of Workplace AI and Occupational Devaluation. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES)*.
- Knowles, B.; and Richards, J. T. 2021. The Sanction of Authority: Promoting Public Trust in AI. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 262–271.
- Kwa, T.; West, B.; Becker, J.; Deng, A.; Garcia, K.; Hasin, M.; Jawhar, S.; Kinniment, M.; Rush, N.; Von Arx, S.; Bloom, R.; Broadley, T.; Du, H.; Goodrich, B.; Jurkovic, N.; Miles, L. H.; Nix, S.; Lin, T.; Parikh, N.; Rein, D.; Sato, L. J. K.; Wijk, H.; Ziegler, D. M.; Barnes, E.; and Chan, L. 2025. Measuring AI Ability to Complete Long Tasks. *arXiv preprint arXiv:2503.14499*.
- Lee, A. T.; Ramasamy, R. K.; and Subbarao, A. 2025. Understanding Psychosocial Barriers to Healthcare Technology Adoption: A Review of TAM and UTAUT Frameworks. *Healthcare*, 13(3): 250.

- Lee, J. D.; and See, K. A. 2004. Trust in Automation: Designing for Appropriate Reliance. *Human Factors*, 46(1): 50–80.
- Lee, Y. S.; Kim, T.; Choi, S.; and Kim, W. 2022. When does AI pay off? AI-adoption intensity, complementary investments, and R&D strategy. *Technovation*, 118: 102590.
- Liao, Q. V.; and Sundar, S. S. 2022. Designing for Responsible Trust in AI Systems: A Communication Perspective. In *Proceedings of the 2022 ACM conference on fairness, accountability, and transparency*. ACM.
- MIT Project NANDA. 2025. The GenAI Divide: State of AI in Business 2025. Technical report, Massachusetts Institute of Technology. Preliminary findings.
- Mlekus, L.; Bentler, D.; Paruzel, A.; Kato-Beiderwieden, A.-L.; and Maier, G. W. 2020. How to raise technology acceptance: user experience characteristics as technology-inherent determinants. *Gruppe. Interaktion. Organisation. Zeitschrift für Angewandte Organisationspsychologie*, 51: 273–283.
- Mogaji, E.; Viglia, G.; Srivastava, P.; and Dwivedi, Y. K. 2024. Is it the end of the technology acceptance model in the era of generative artificial intelligence? *International Journal of Contemporary Hospitality Management*, 36(10): 3324–3339.
- Oliveira, T.; and Martins, M. d. R. F. d. O. 2011. Literature review of information technology adoption models at firm level. *The Electronic Journal Information Systems Evaluation*, 14(1): 110–121.
- Raghavan, M.; Barocas, S.; Kleinberg, J.; and Levy, K. 2020. Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAccT)*, 469–481.
- Sanchez-Monedero, J.; Dencik, L.; and Edwards, L. 2020. What Does It Mean to “Solve” the Problem of Discrimination in Hiring? Social, Technical and Legal Perspectives from the UK on Automated Hiring Systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAccT)*, 458–468.
- Sarkar, A.; Gordon, A. D.; Negreanu, C.; Poelitz, C.; Raghavan, S. S.; and Zorn, B. 2022. What is it like to program with artificial intelligence? *arXiv preprint arXiv:2208.06213*.
- Selbst, A. D.; boyd, d.; Friedler, S. A.; Venkatasubramanian, S.; and Vertesi, J. 2019. Fairness and Abstraction in Sociotechnical Systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAccT)*, 59–68.
- Shachak, A.; Kuziemsky, C.; and Petersen, C. 2019. Beyond TAM and UTAUT: Future directions for HIT implementation research. *Journal of Biomedical Informatics*, 100: 103315.
- Tankelevitch, L.; Kewenig, V.; Simkute, A.; Scott, A. E.; Sarkar, A.; Sellen, A.; and Rintel, S. 2024. The Metacognitive Demands and Opportunities of Generative AI. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, CHI ’24. ACM.
- Tornatzky, L. G.; and Fleischer, M. 1990. *The Processes of Technological Innovation*. Lexington, MA: Lexington Books.
- van Dijk, J. A. G. M. 2005. *The Deepening Divide: Inequality in the Information Society*. London: SAGE Publications.
- Venkatesh, V. 2022. Adoption and use of AI tools: a research agenda grounded in UTAUT. *Annals of Operations Research*, 308(1): 641–652.
- Venkatesh, V.; Morris, M. G.; Davis, G. B.; and Davis, F. D. 2003. User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3): 425–478.
- Warschauer, M. 2003. *Technology and Social Inclusion: Rethinking the Digital Divide*. Cambridge, MA: MIT Press.
- Weidinger, L.; Mellor, J.; Rauh, M.; Griffin, C.; Uesato, J.; Huang, P.-S.; Cheng, M.; Glaese, M.; Balle, B.; Kasirzadeh, A.; Kenton, Z.; Brown, S.; Hawkins, W.; Stepleton, T.; Biles, C.; Birhane, A.; Haas, J.; Rimell, L.; Hendricks, L. A.; Isaac, W.; Legassick, S.; Irving, G.; and Gabriel, I. 2021. Ethical and Social Risks of Harm from Language Models. *arXiv preprint arXiv:2112.04359*.
- Weisz, J. D.; Muller, M.; He, J.; and Houde, S. 2023. Toward general design principles for generative AI applications. *arXiv preprint arXiv:2301.05578*.
- Wolfe, D.; Price, M.; Choe, A.; Kidd, F.; and Wagner, H. 2025. Revisiting UTAUT for the Age of AI: Understanding Employees AI Adoption and Usage Patterns Through an Extended UTAUT Framework. *arXiv preprint arXiv:2510.15142*.
- Woodruff, A.; Shelby, R.; Kelley, P. G.; Roussou-Schindler, S.; Smith-Loud, J.; and Wilcox, L. 2024. How Knowledge Workers Think Generative AI Will (Not) Transform Their Industries. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*.
- Xia, Q. N.; Constantinides, M.; Sarkar, A.; Brumby, D. P.; and Cox, A. L. 2026. “If You’re Very Clever, No One Knows You’ve Used It”: The Social Dynamics of Developing Generative AI Literacy in the Workplace. *arXiv preprint arXiv:2602.01386*. Accepted at CHIWORK 2026.
- Zamfirescu-Pereira, J. D.; Wong, R. Y.; Hartmann, B.; and Yang, Q. 2023. Why Johnny can’t prompt: how non-AI experts try (and fail) to design LLM prompts. In *Proceedings of the 2023 CHI conference on human factors in computing systems*, 1–21.