

Unaccountable Delegation, Fading Skills: Mapping the Risks of Workplace AI Agents

Gabriele La Malfa¹, Lakmal Meegahapola², Edyta Bogucka², Jie M. Zhang¹, Michael Luck³,
Elizabeth Black¹, Daniele Quercia^{2,4}

¹King’s College London, UK

²Nokia Bell Labs, Cambridge, UK

³University of Sussex, UK

⁴Politecnico di Torino, Italy

{gabriele.la_malfa, jie.zhang, elizabeth.black}@kcl.ac.uk, michael.luck@sussex.ac.uk,
{lakmal.meegahapola, edyta.bogucka, daniele.quercia}@nokia-bell-labs.com

Abstract

To anticipate the socio-technical risks posed by AI agents, organizations first need taxonomies to classify them. However, existing AI risk taxonomies focus on broad risks and do not capture the job-specific risks introduced by agents. To address this gap, we make three main contributions. First, we developed a multi-layer framework based on a review of the literature on AI agents. The framework models three core components and their interactions: the agents, their goals, and their environment. Second, we embedded this framework in a structured prompt and applied it to descriptions of 2,078 job tasks from the O*NET occupational database, producing 8,356 risk scenarios labeled by severity and deployment mode (automation or augmentation). We validated these scenarios with 45 workers across 10 job roles and an independent LLM judge, confirming their high plausibility and alignment with the corresponding job tasks. Finally, we extended an existing taxonomy to create a 15-category taxonomy of workplace AI agent risks that covers all our risk scenarios. Our analysis highlights four findings. First, augmentation is not inherently safe because overreliance on agents can gradually erode workers’ skills and oversight. Second, erroneous actions of agents account for the largest share of risk scenarios and has the highest number of severe risks. Many of these risks arise at the human–agent boundary. Third, automation is associated mainly with organizational risks (operational failures, financial losses), whereas augmentation is associated mainly with risks to workers (capability erosion, psychological and social risks). Fourth, workers found our taxonomy easier to use for a risk classification task than two other risk taxonomies and preferred it in 64% of non-tied comparisons with a recent generative AI risk taxonomy. Taken together, these findings show that workplace AI agent risks do not arise from agents alone; they also depend on how people work with agents and how agents are deployed. Safer workplaces therefore require not only safer agents but also carefully designed human–AI agent collaboration.

Project website —

<https://social-dynamics.net/ai-risks/workplace-agents/>

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

1 Introduction

AI agents are autonomous entities that perceive their environment, interact with humans or other agents, and act to achieve goals with varying degrees of independence (Sapkota, Roumeliotis, and Karkee 2026). They are increasingly deployed to support everyday workplace tasks (Eloundou et al. 2024; Xi et al. 2025; Mohny 2025). AI agents do not operate in isolation: together with the environment they perceive, the goals they pursue, and the humans they work alongside, they form what we call an *agentic AI system* (AAIS) (IBM 2024). As these systems are deployed across jobs and industries, organizations need ways to identify and classify the risks they introduce (Weidinger et al. 2023). This need arises when teams red-team agents before launch (Ganguli et al. 2022), conduct impact assessments (Moss et al. 2021), or conduct risk and incident audits (Raji et al. 2020; McGregor 2021). In all of these cases, organizations rely on risk taxonomies: classification systems that provide a shared language for anticipating risks that have not yet occurred and classifying those that have (Shelby et al. 2023). Without a taxonomy built for the specific context of workplace AI agents, these efforts rely on risk classifications that can leave risks unnamed and untracked.

Existing AI risk taxonomies were not designed for workplace agentic AI systems, leaving three gaps. First, many taxonomies are too generic or focus primarily on technical failures. They can therefore overlook *socio-technical AI risks* that emerge from the dynamic interaction between an AI agent and its social context, including the humans, organizations, and institutions in which it is embedded. Second, most taxonomies were developed for generic AI (Abercrombie et al. 2024) or generative AI (Li et al. 2025b), rather than for AI agents, which can act more autonomously. Even recent AI agent risk taxonomies tend to model agents as purely technical systems (Hammond et al. 2025; Bryan et al. 2025; Khoo, Foo, and Lee 2025; Ghosh et al. 2025), without explicitly capturing how agents, goals, environments, and human workers interact to produce risks. Third, taxonomies based on incident databases and media reports (Abercrombie et al. 2024; Turri and Dzombak 2023; Paeth et al. 2025) capture risks that have already occurred and were visible enough to be recorded. They are therefore less suited to

identifying risks that develop gradually or have not yet produced visible harms, such as the erosion of professional skills caused by prolonged reliance on AI (Richards, Benn, and Zilka 2025; Shao et al. 2025a). Together, these gaps point to the need for a taxonomy that operates at the level of specific job tasks, reflects how agentic AI systems are structured, and anticipates risks before they become harms. To address these gaps, we make three main contributions:

1. We introduce a multi-layer framework for workplace AI agents. Building on prior work on socio-technical safety and compositional risk assessment, the framework organizes risks across three layers (§3.1): Technical Capability (AAIS components and how they work together), Human Interaction (how people engage with AAIS), and Systemic Impact (broader organizational and societal consequences). The framework also captures the relationships among agents, goals, environments, and human workers. We refined the framework through structured interviews with three experts in agentic AI systems and responsible AI.

2. We generate and validate job-specific risk scenarios. We embedded the framework in a structured prompt and applied it to descriptions of 2,078 job tasks from the O*NET occupational database (§3.2). This process generated 8,356 risk scenarios labeled by severity (minimal, limited, high, critical) and deployment mode (§3.3): augmentation, in which an agent supports a worker, or automation, in which the agent performs the job task without worker involvement. We then validated the scenarios in two ways (§3.4). First, 45 workers across 10 job roles evaluated 450 scenarios from their own roles. Second, an LLM-as-a-judge from a different model family assessed the full corpus. Both evaluations found high plausibility (workers: 4.03/5; LLM: 4.81/5) and strong task alignment (workers: 3.94/5; LLM: 4.60/5).

3. We construct and validate a workplace AI agent risk taxonomy. To ground the taxonomy in both generated and observed risks, we supplemented the 8,356 generated scenarios with 222 workplace AI agent incidents from three public databases and 11 failure cases from a live multi-agent red-teaming study. We converted these cases into a common risk scenario format and used the combined corpus to extend Li et al.’s generative AI risk taxonomy into a workplace AI agent risk taxonomy (§3.5). The resulting taxonomy contains 15 categories and 44 sub-categories (Figure 2). We evaluate its structural qualities, compare its coverage with 10 established AI risk frameworks, and test its usability in a comparative study with 26 workers (§3.6). Finally, we analyze the risk corpus along the themes of taxonomy categories, severity, deployment mode, and framework layers.

Across these contributions, our results point to a clear pattern (§4): most workplace AI agent risks arise from how agents and people work together. The two largest taxonomy categories are *Erroneous Actions* of agents (30.6%), including misinterpretations and incorrect recommendations, and *Capability Erosion* (21.3%), including the loss of workers’ skills and oversight capabilities. Together, these categories account for 51.9% of all risk scenarios and occur mostly under augmentation. Under automation, risks shift to the or-

ganizational level, particularly *Operational Failures* and *Financial Losses*. Taken together, these findings suggest that safer workplaces require not only safer agents but also carefully designed human–AI agent collaboration.

2 Related Work

Prior work provides a good foundation for understanding AI risks in general. However, three gaps remain. Existing taxonomies tend to operate at the wrong level of detail for specific job tasks (§2.1), are rarely built around how agents behave (§2.2), and focus on past incidents rather than emerging risks (§2.3).

2.1 From Broad Risk Classifications to Job-Level Risk

Research on AI risk tends to fall into two areas: technical and socio-technical. Both offer valuable perspectives, but neither maps easily onto what happens when a specific AI agent enters a specific job workflow. Technical approaches explain risks primarily at model-level, such as bias in training data or failures of robustness (Mehrabi et al. 2021; Raji et al. 2022; Fabris et al. 2022). These are important concerns, but they tend to abstract away the organizational and human context in which AI operates. Socio-technical approaches take a broader view, framing AI as part of a wider social system and classifying large-scale risks such as unfair resource distribution or the erosion of accountability (Selbst et al. 2019; Macrae 2022; Shelby et al. 2023; Slattery et al. 2025). These approaches show why context matters, but their broad categories can be difficult to apply to a particular task or workflow. Neither perspective is designed to answer a practical question facing organizations: “What risks could emerge when this AI agent is introduced into this specific job task?” What is missing is a framework that connects multi-layer risk concepts to concrete job tasks and workflows.

2.2 Modeling Risks from Agent Behavior

Risks of AAISs can emerge from how agents pursue goals, take multi-step actions, and interact with their environments (Sapkota, Roumeliotis, and Karkee 2026; Chan et al. 2023; Gridach et al. 2025; Liu et al. 2025). When several agents work together within an AAIS, coordination failures, conflicting goals, and emergent behavior create additional risks (Hammond et al. 2025; Altmann et al. 2024; Kong et al. 2025). Autonomy can also make agent behavior harder to monitor and responsibility harder to assign (Anwar et al. 2024; Ranjan, Gupta, and Singh 2025). Moreover, recent work categorizes agent-specific risks such as misalignment with human intent, prompt-injection vulnerabilities, and cascading failures in multi-agent systems (Bryan et al. 2025; Khoo, Foo, and Lee 2025; Ghosh et al. 2025). However, these risk classifications often study agents in isolation or in simulated environments. They rarely connect the three components at the core of an AAIS: the agents, their goals, and their environment, to concrete job tasks performed by workers. Thus, existing agent-safety research identifies important risks and failure mechanisms but does not yet explain how those mechanisms produce risks in specific job tasks.

2.3 From Existing Evidence to Anticipating Risks

AI risk taxonomies draw on several sources of evidence. Documentation-based approaches capture reported incidents and other public records (Li et al. 2025b; Lee et al. 2024), academic and policy literature (Shelby et al. 2023; Slatery et al. 2025), or user-generated data such as social media discussions (Zhang et al. 2025). Human-centered approaches develop or refine risk categories through interviews, expert consultation, workshops, and annotation studies (Steenstra and Bickmore 2025; Abercrombie et al. 2024). Some studies combine these sources. For example, Abercrombie et al. (Abercrombie et al. 2024) build on existing taxonomies and documented incidents, then refine their categories through expert feedback and crowdsourced annotation. These sources offer complementary benefits. Documented evidence grounds taxonomies in observed harms, while human elicitation can reveal contextual knowledge that public records overlook. However, both also have limitations. Incident databases contain only harms that were observed, reported, and made public (Paeth et al. 2025; Richards, Benn, and Zilka 2025). Literature-based taxonomies inherit the scope of prior research, while human elicitation depends on the experience and composition of the participants. As a result, emerging, gradual, and job-specific risks may remain underrepresented, particularly as AI capabilities and workplace uses evolve. Prospective methods offer a complementary way to identify risks before they produce documented harms. Research in labor economics and HCI has used the O*NET occupational database to ground analyses of AI exposure in real job tasks (Eloundou et al. 2024; Felten, Raj, and Seamans 2023). Other work shows that LLMs can support structured scenario generation and red-teaming exercises that anticipate risks before deployment (Perez et al. 2022; Ganguli et al. 2022; Fröhling et al. 2026; Herdel et al. 2025). However, these research strands have largely developed separately. What is missing is a systematic method that combines real job-task data with prospective risk generation to identify workplace AI agent risks before they become documented harms.

Taken together, prior work offers broad risk classifications, detailed accounts of agent failure modes, and strong methods for learning from existing evidence. However, these three strands have not yet been fully integrated. A workplace AI agent risk taxonomy must therefore (1) be developed by considering specific job tasks, (2) consider risks across *Technical Capability*: AAISs including agents, goals, and environments, and how they interact; *Human Interaction*: how workers engage with these systems; and *Systemic Impact*: broader organizational and societal consequences, and (3) anticipate risks that may not yet appear in incident records. Together, these gaps motivate our job-task-grounded and forward-looking approach, which combines anticipated risk scenarios with documented evidence to construct a taxonomy of workplace AI agent risks.

3 Methodology

Our methodology consists of six steps (Figure 1). First, we define a multi-layer framework for analyzing risks in AAISs

(§3.1). Second, we select job tasks in the O*NET database that AI agents are likely to affect (§3.2). Third, we embed the multi-layer framework in a structured prompt to generate risk scenarios for job tasks (§3.3). Fourth, we evaluate the scenarios with workers and an LLM-as-a-judge (§3.4). Fifth, we use the risk scenarios to construct a workplace AI agent risk taxonomy (§3.5). Finally, we evaluate the taxonomy for structural integrity, coverage relative to 10 established frameworks, and practical usability (§3.6).

3.1 Defining the Multi-Layer AAIS Framework

We build the multi-layer framework by extending the sociotechnical framework of Weidinger et al. (Weidinger et al. 2023), which argues that evaluating technical artifacts alone cannot establish whether an AI system is safe. A capability may indicate a potential hazard, but whether it produces harm depends on the people using the system, the context of use, and the broader systems in which it is deployed. The framework therefore distinguishes three targets of evaluation: *capability*, which examines the AI system and its technical artifacts; *human interaction*, which examines how people use and respond to the system; and *systemic impact*, which examines broader organizational, social, economic, and environmental effects. These layers add progressively more context, but they are not chronological stages and can influence one another. To apply this layered approach to workplace AI agents, we make the internal structure of an agentic AI system explicit. Workplace agents pursue goals, act in changing environments, and interact with other agents and human workers. Risks can arise from these components or from their interactions, even when each component appears to function correctly. We therefore retain the three evaluation layers while expanding them to represent the components and interaction pathways of an AAIS.

Layer 1: Capability. This layer represents risks arising from three components and their interactions (Russell and Norvig 2009; Wooldridge 2009): the *AI agents*, their *environment*, and their *goals*.

Component risks. Agent properties such as autonomy, reactivity, and proactiveness can create different risks. An agent may respond poorly to an unfamiliar situation, take an unrequested action, or modify a workflow without sufficient oversight. The environment, including APIs, databases, and physical workspaces (IBM 2024), can create risks when it provides incomplete, outdated, or rapidly changing information. Goals can create risks when they are represented by poor proxy measures. For example, a diagnostic agent optimized for case volume may deprioritize difficult cases to increase throughput.

Interaction risks. Risks can also emerge from interactions among components. In *agent-agent* interactions, missing coordination rules may lead agents to take conflicting actions. In *agent-environment* interactions, an agent may optimize a local objective while violating environmental or organizational constraints. In *agent-goal* interactions, an agent may satisfy an objective in an unintended way. For example, a quality-control agent could classify every item as defective to maximize its detection score, thereby halting production.




Figure 1: **Overview of the six-step methodology, from framework definition through risk scenario generation and validation to taxonomy construction and evaluation.** **Step 1** defines the multi-layer framework, which organizes sociotechnical risk into three layers: *technical capability* (AAIS including the AI agents, goals they pursue, and the environment they act in), *human interaction* (how workers engage with the AAIS and components), and *systemic impact* (consequences that propagate beyond the individual job task). **Step 2** selects 2,078 computer-based O*NET job tasks (209 roles, 20 industries). **Step 3** applies the framework to each task through a structured prompt, anticipating 8,356 risk scenarios labeled with severity and deployment mode. **Step 4** validates the scenarios with 45 workers from the selected roles and with an LLM-as-a-judge, on validity, utility, and discovery metrics. **Step 5** builds the taxonomy from 8,589 scenarios: the 8,356 anticipated ones plus 233 describing risks that have already occurred, drawn from workplace AI incidents and documented agent failure cases; scenarios that do not map onto the GenAI Risk Taxonomy produce new categories and sub-categories, extending it into the 15-category Workplace AI Agent Risk Taxonomy. **Step 6** evaluates the taxonomy on structural integrity, contextual coverage, and practical usability.

Layer 2: Human Interaction. This layer represents how workers interact with agents and the broader system (Hammond et al. 2025). We distinguish five interaction pathways. In *agent–human* interactions, agents may undermine trust, violate workplace norms, or encourage over-reliance. *Human–human* interactions may change when AI-generated targets influence how managers evaluate and direct workers. *Human–environment* interactions capture how agent-oriented workplaces may constrain or endanger workers, such as when a warehouse is organized primarily for robotic efficiency. *Human–goal* interactions capture conflicts between system objectives and professional judgment, such as when clinical staff feel pressured to follow automated triage scores. Finally, *human–AAIS* interactions capture risks arising from the system as a whole, including cases in which responsibility becomes too distributed for workers to understand or contest a decision.

Layer 3: Systemic Impact. This layer represents consequences that extend beyond an individual task or interaction. Such risks may affect organizations, industries, society, the economy, or the natural environment. They often emerge through large-scale adoption, repeated interactions, or feedback loops rather than a single failure. For exam-

ple, interacting AI trading agents may collectively amplify incorrect market signals and contribute to a market disruption. Widespread use of AI management systems may shift bargaining power toward employers or normalize harmful workplace practices. Similarly, repeated use of agents that favor established solutions may gradually reduce the diversity of professional knowledge across a field.

Framework Review and Refinement. We reviewed the framework through three structured interviews with experts in agentic AI systems and responsible AI. Each interview lasted approximately 30 minutes and included a framework briefing, technical questions, and detailed feedback (Supplementary Material B). The experts assessed the coherence of the components, interaction pathways, and three layers and identified areas requiring clarification. We used their feedback to refine definitions, clarify the boundaries between interactions, and improve the presentation of the framework.

3.2 Selecting Job Tasks Impacted by AI Agents

To ground our analysis in real job tasks, we used O*NET (National Center for O*NET Development 2025), a U.S. government database describing the tasks, activities, and work contexts of hundreds of jobs. We used a filtered

subset developed by Shao et al. (Shao et al. 2025b), which focuses on computer-based tasks on which AI agents are likely to be deployed. It contains 2,078 job tasks from 209 job roles across 20 industries, including healthcare, finance, and transportation. Each task is described concretely. Examples include “Diagnose acute conditions that could result in rapid physiological deterioration” for an acute care nurse and “Examine whether the organization’s objectives are reflected in its management activities” for a manager. These descriptions provide the job context needed to generate job task-specific risk scenarios.

3.3 Generating Risk Scenarios for Job Tasks with the Multi-Layer Framework

We use gpt-4o-mini (OpenAI 2024) to generate prospective risk scenarios for each of the 2,078 tasks. This approach is consistent with prior work showing that LLMs can serve as effective tools for structured risk foresight (Bogucka, Šćepanović, and Quercia 2024; Bućinca et al. 2023; Herdel et al. 2025; Wang et al. 2024b; Perez et al. 2022; Fröhling et al. 2026). We selected gpt-4o-mini to balance reasoning quality with the computational demands of processing thousands of prompts at scale. To prevent the model from producing generic or implausible outputs, we used a structured prompt that included a description of our multi-layer framework defining agentic AI systems’ risk surfaces, enabling the generation of risks specific to agents and their interactions (Supplementary Material E; Figure 1). The prompt instructed the model to consider risks at each framework layer (capability, human interaction, systemic impact) and to label each scenario with its corresponding component or interaction pathway (e.g., agent–human, agent–goal). This structure kept the generated risks tied to specific pathways rather than generic failures. For each task, the model was asked to generate up to five risk scenarios, enough to capture a diverse set of risks while avoiding highly speculative edge cases. Each scenario was annotated with two additional labels: (1) *severity*, classified according to four risk tiers (minimal, limited, high, critical) (Yin et al. 2025), and (2) *deployment mode*, specifying whether the risk arises from *automation* (complete replacement of the human task as defined in (Raisch and Krakowski 2021)) or *augmentation* (AI supporting the human in performing the tasks as defined in (Autor 2015)). This process produced 8,356 distinct risk scenarios.

3.4 Validating Risk Scenarios

We assessed the quality of the generated scenarios using both workers and an LLM judge, in line with previous work (Zheng et al. 2023; Li et al. 2024; Chiang and Yi Lee 2023; Li et al. 2025a; Fröhling et al. 2026). In both cases, each scenario was evaluated on 10 metrics across three dimensions: *validity*, covering plausibility and connection to the task; *utility*, covering usefulness, actionability, detail, complexity, and specificity; and *discovery*, covering originality, rarity, and novelty (Si, Yang, and Hashimoto 2024; Hu and Zhou 2024).

Domain Expert Evaluation. We selected 10 job roles and recruited 45 workers whose reported professions matched

those roles (Supplementary Material C). We collected demographic and role-related information from each expert. For each job role, we randomly sampled 10 of the approximately 70 role-specific scenarios. Each expert independently evaluated the 10 scenarios for their role using the 10 metrics described above on a 5-point Likert scale. In total, the workers evaluated 450 risk scenarios.

LLM-as-a-Judge Evaluation. The domain expert evaluation provided useful ratings but covered only the subset of scenarios that the experts could feasibly review. To rate all 8,356 scenarios using the same criteria, we used gemini-3-flash (Google 2025) as an LLM judge. The model rated each scenario using the 10 metrics described above, allowing us to apply a consistent standard across the full corpus. We took two steps to assess the reliability of these ratings. First, we chose a judge from a different model family than the generator (Google Gemini rather than the OpenAI model used to generate the scenarios), since models tend to score their own outputs more favorably (Panickssery, Bowman, and Feng 2024). Second, we tested the judge’s consistency by re-rating a random sample of 370 scenarios with a second model, gpt-4o, and comparing the two. The models rarely gave the exact same 5-point score. However, when we collapsed the 5-point scale into positive, neutral, and negative bands (*sentiment agreement*), the two models landed in the same band almost every time: 98.9% for plausibility and 98.1% for connection to the task. Therefore, the judges agreed on the underlying judgment for the validation sample (Supplementary Material C).

Results. Both the workers and the LLM judge rated the risk scenarios as highly plausible (workers: 4.03/5; LLM: 4.81/5) and closely tied to the corresponding job tasks (workers: 3.94/5; LLM: 4.60/5). They also rated the scenarios as useful (workers: 3.70/5; LLM: 3.38/5) and actionable (workers: 3.62/5; LLM: 3.37/5). Their ratings differed most on novelty. Workers considered the scenarios moderately novel (originality: 2.96/5; rarity: 3.00/5), while the LLM judged them as more familiar (originality: 1.80/5; rarity: 1.28/5). Overall, we interpret these results as a sign that the risk scenarios are well grounded: workers recognized them as plausible and practically relevant, while the LLM, which has seen vast amounts of text, recognized them as patterns it had encountered before (Supplementary Material C).

3.5 Building a Taxonomy from Risk Scenarios

We began with the 8,356 risk scenarios generated in the previous step. To ground the taxonomy in real-world evidence, we added 11 AI agent failure cases from a study of agent behavior (Shapira et al. 2026) and 222 workplace AI incidents from three public incident databases. We converted both the failure cases and incidents into the same risk scenario format as the generated scenarios. Next, our goal was to map as many scenarios as possible to an existing taxonomy. We used the Generative AI (GenAI) risk taxonomy of Li et al. (Li et al. 2025b) as the starting point and followed three steps. First, we used gpt-4o-mini (for the generated risk corpus) or manual inspection (for failure cases and incidents) to map each scenario to the closest category or sub-




Figure 2: **Workplace AI Agent Risk Taxonomy.** 15 categories and 44 sub-categories organized across agent execution risks, human risks, organizational risks, and societal risks.

category in Li et al.’s taxonomy. Scenarios without a suitable match were marked as unmatched. Second, we analyzed the unmatched scenarios and identified recurring themes. Third, we manually reviewed these themes to determine how they should extend the taxonomy. When a theme fit within an existing category, we added it as a new sub-category. When it did not fit any existing category, we created a new top-level category. We then refined the category and sub-category definitions to make their boundaries clear. The resulting Workplace AI Agent Risk Taxonomy contains 15 top-level categories and 44 sub-categories (Figure 2). See Supplementary Material G for the full taxonomy with descriptions and example risk scenarios.

3.6 Evaluating the Taxonomy

We evaluated the taxonomy in three ways. First, we checked whether the categories were distinct (structural integrity). Second, we checked whether the taxonomy captured workplace risk categories that existing frameworks overlook (contextual coverage). Third, we examined whether workers could use it to classify workplace risks (practical usability).

Structural Integrity. We first examined whether the 15 categories and 44 sub-categories were distinct in both their vocabulary and meaning. *Vocabulary check.* We represented the taxonomy with names and definitions for categories and sub-categories. Then, we converted this text into a TF-IDF vector, which gives more weight to words that are distinctive to a category and less weight to commonly used words.

We then calculated the cosine similarity between every pair of categories on a scale from 0 to 1, where higher values indicate greater overlap. The mean inter-category similarity was 0.01, indicating minimal overlap in vocabulary. We also calculated Jaccard similarity, which measures the proportion of words shared by two descriptions. Using 0.30 as the overlap threshold, 99.9% of sub-category pairs fell below this value. *Meaning check.* Categories may use different words while describing similar concepts. For example, “skill erosion” and “deskilling” have similar meanings despite sharing few words. To capture this type of overlap, we encoded the category definitions using the pre-trained sentence-transformer model: all-MiniLM-L6-v2. The model represents texts with similar meanings as nearby vectors, even when they use different words. The mean pairwise semantic similarity was 0.24, indicating limited average overlap in meaning. Together, these checks suggest that the taxonomy definitions are distinct in both their vocabulary and their underlying concepts (Supplementary Material H).

Contextual coverage. To check whether our taxonomy captures workplace risk categories that prior work overlooks, we compared it with ten established AI risk taxonomies and classification frameworks, which we collectively call *risk classifications*. We followed the comparative mapping approach of Abercrombie et al. (Abercrombie et al. 2024). Our comparison included the seven sources examined in their study and we also included three more recent sources (Table 1). *Mapping process.* For each of our 15 categories, we

Workplace AI Agent Risk Category	EPIC	CSET	Turing	Azure	Sony	TASRA	Shelby	AIR	MA Adv AI	AI Risk Repo
Agent Execution Risks										
Manipulative Behaviour	■	□	■	■	■	■	■	■	■	■
Erroneous Actions	□	□	■	□	■	■	■	■	■	■
Defective Outputs	□	□	□	□	□	■	■	□	■	■
Discriminatory Outcomes	■	■	■	■	■	■	■	■	■	■
Human Risks										
Capability Erosion	■	□	■	■	■	■	■	□	■	■
Employment Displacement	■	□	□	■	■	■	■	□	■	■
Psychological and Social Risks	■	■	■	■	■	■	■	□	■	□
Physical Risks	■	■	■	■	■	■	■	■	■	■
Organizational Risks										
Strategic Failures	□	□	□	■	□	□	■	□	■	□
Operational Failures	□	■	□	□	□	□	□	□	■	□
Financial Losses	■	■	□	■	□	■	□	■	■	□
Reputational Risks	■	■	■	□	■	■	■	□	■	□
Societal Risks										
Geopolitical and Security Risks	□	□	□	□	□	■	□	■	■	□
Environmental and Ecosystem Risks	■	■	■	■	□	■	■	□	■	■
Societal Functioning Risks	■	■	■	■	■	■	■	■	■	■

Table 1: **Coverage of our 15 taxonomy categories across ten established AI risk classifications.** Categories are grouped by dimension. We use the color coding: ■ = full match, ■ = partial match, □ = no match. Compared risk classifications are: EPIC (Fergusson et al. 2023), CSET (Hoffman and Frase 2023), the Alan Turing Institute (Leslie 2019), Microsoft Azure (Microsoft 2023), Sony (Hutiri, Papakyriakopoulos, and Xiang 2024), TASRA (Critch and Russell 2023), and Shelby et al. (Shelby et al. 2023), AIR 2024 (Zeng et al. 2024), Hammond et al. (Hammond et al. 2025), and the AI Risk Repository (Slattery et al. 2025). We find that organizational risks (*Strategic Failures*, *Operational Failures*) and agent-execution risks (*Erroneous Actions*, *Defective Outputs*) are consistently absent from or only partially covered by prior work.

examined its coverage in each of the ten risk classifications. We assigned one of three labels: *Full match* (■) when a comparison source contained a category that directly represented ours; *Partial match* (■) when it covered only part of our category; and *No match* (□) when it did not cover the risk (Supplementary Material L). Table 1 shows that prior risk classifications cover human-oriented risks thoroughly: *Discriminatory Outcomes* of agents is fully represented in nine of the ten classifications, *Physical Risks* in seven, and *Psychological and Social Risks* in six. Beyond these, three gaps emerge. First, organizational risks are the least covered in the comparison: *Operational Failures* and *Strategic Failures* are the only two categories with no full match anywhere, and are absent from eight and seven of the ten classifications respectively. Second, the agent-execution risks that dominate our corpus are recognized only sporadically: under agent execution risks, *Defective Outputs* is absent from six classifications, and *Erroneous Actions* is fully represented in only two, appearing elsewhere as generic model inaccuracy rather than as distinct failures of what an agent interprets, recommends, and produces. Third, and most tellingly, *Manipulative Behaviour* of agents is fully represented only

in the two classifications built specifically for agentic systems (Hammond et al. 2025; Slattery et al. 2025); the remaining eight treat deception either as human misuse of a model or not at all. By contrast, our *Societal Functioning Risks* category is well covered by prior work, with seven full matches, which reinforces that its small share in our corpus (0.4%) reflects our task-level elicitation rather than a gap in the literature. Together, these patterns suggest that existing risk classifications were not designed to represent the organizational, operational, and agent-execution risks that emerge when AI agents are introduced into workplaces. Our workplace-specific taxonomy addresses this limitation.

Practical usability. A taxonomy should be easy to use in practice. We therefore designed a study around a common risk-management activity: classifying risks with a taxonomy. This task is routine in impact assessments, model cards, or risk registers (Mitchell et al. 2019; Gebru et al. 2021; Moss et al. 2021; Raji et al. 2020). Using the same professional-background matching criteria as in the scenario-validation study, we recruited 26 workers from five professional backgrounds: software engineering, data science, design, healthcare, and finance. All participants con-

firmed daily AI use and medium-to-high awareness of AI agents and passed three comprehension checks on the definition of AI agents before proceeding. Each participant classified five AI agent risk scenarios drawn at random from a held-out pool not used during taxonomy construction. The pool combined LLM-generated scenarios with documented incidents from public AI incident databases, so that participants saw both anticipated risks and real failures. Each scenario included its workplace context: the risk, industry, job title, and corresponding job task. Participants classified every scenario using three taxonomies: (A) our *Workplace AI Agent Risk Taxonomy* with 15 categories and 44 sub-categories; (B) the *Generative AI Risk Taxonomy* (Li et al. 2025b) with 12 categories and 48 sub-categories; and (C) the *AI Risk Repository Domain Taxonomy* (Slattery et al. 2025) with 7 categories and 23 sub-categories. We presented the taxonomies in a randomized order under the neutral labels Framework 1, Framework 2, and Framework 3 to reduce name recognition bias. For each scenario and taxonomy, participants selected the best fitting category or chose “no good fit” when no category represented the risk. This option allowed us to distinguish between difficulty selecting among relevant categories and the absence of a suitable category. Participants then rated the ease of using each taxonomy on a seven-point scale and provided a brief written reflection on their experience using different taxonomies for risk classification. We report the findings in §4.

4 Main Findings

We classified our 8,356 risk scenarios into the taxonomy and examined how they distribute across taxonomy categories, severity levels, deployment modes, and framework layers (Figure 3). Our main findings are provided below.

Finding 1: Augmentation is not inherently safe; it introduces a slow and largely invisible risk of deskilling. AI augmentation, where the agent assists rather than replaces the human, is often treated as the more responsible deployment choice (Wang et al. 2024a; Shao et al. 2025b). Our findings complicate this assumption. *Capability Erosion* is the second most common risk category in the entire dataset, representing 21.3% (1,782) of all scenarios, and 97.0% of these risks arise from augmentation rather than automation, the most one-sided augmentation profile in the taxonomy. Within it, the dominant sub-category is not the loss of a manual skill but the loss of oversight itself: *Decision and oversight* in *Capability Erosion* alone accounts for 1,420 scenarios (17.0% of the corpus), describing workers who stop questioning, checking, or independently judging what the agent produces. The pattern is consistent with earlier research on human factors, which has long warned that the more reliable a system is, the more people stop practicing the skills they would need if it ever failed (Bainbridge 1983). What makes modern AI agents different is that they are proactive rather than passive: they suggest next steps, flag priorities, and take initiative, which accelerates the rate at which workers hand over their professional judgment to the system. The severity profile confirms that this is not a dramatic or sudden failure: 79% of capability erosion risks

are classified as “limited” severity, and only $\leq 0.1\%$ as “critical”. The risk is real precisely because it is gradual, cumulative, and easy to miss until it is too late to reverse.

Finding 2: Erroneous agent action is both the most common and the most severe risk, and it is not purely a technical problem. *Erroneous Actions* under *Agent Execution Risks*, which covers situations where an agent mishandles data, misreads a situation, advises incorrectly, or executes a task wrongly, is the largest category in the taxonomy at 30.6% (2,558) of all scenarios, and it also carries the highest number of severe risks: 88% are rated “high” or “critical” (2,183 “high”). This might seem like a straightforward technical issue: the model hallucinates, or its training data is wrong. But our analysis shows that the failure happens between the agent and the human, not inside the model itself. Fabrication is rare, accounting for only 57 scenarios (0.7% of the corpus, filed under *Defective Outputs* in *Agent Execution Risks*). The dominant sub-types are instead *erroneous interpretation*, where the agent misreads data, context, or intent (1,158 scenarios, 13.9%), and *erroneous recommendations*, where it advises incorrectly on that basis (946 scenarios, 11.3%). Moreover, 69.5% of these risks arise under augmentation, meaning a human is present and acting on the flawed output. The agent produces an answer that is technically consistent with its inputs but wrong for the situation; the human acts on it because the agent communicated confidence. Improving model accuracy alone will therefore not solve the problem.

Finding 3: Automation concentrates risk in the organization; augmentation concentrates it in the worker, and each requires different governance. The capability erosion and agent error risks above are concentrated in different deployment modes, and a third pattern only becomes visible when we consider automation. When agents fully automate a task, removing humans from the loop, the dominant risks are organizational and institutional: *Financial Losses* (485 risks, 90.9% from automation), *Operational Failures* (744 risks, 83.5%), and *Reputational Risks* (518 risks, 71.8%). The starkest case is *Employment Displacement*, which is 98.8% automation-bound and 92.7% high or critical severity, the most one-sided and most severe category in the taxonomy. These failures tend to be visible, fast, and expensive: legal liability, operational downtime, or public relations damage from an agent acting without oversight. When agents augment rather than automate, the risks land on the individual worker rather than the organization, manifesting as the slow erosion of professional agency described in Finding 1. This structural divide has a direct governance implication: a single approach cannot cover both deployment modes; we translate this into concrete safeguards in the Discussion.

Finding 4: The Workplace AI Agent Risk Taxonomy outperforms existing frameworks in practical usability. In our usability study, the Workplace AI Agent Risk Taxonomy was the clearest of the three taxonomies for classifying workplace AI agent risks. Across 26 workers classifying 130 risks, the Workplace AI Agent Risk Taxonomy achieved a higher ease-of-use rating than the GenAI




Figure 3: Distribution of 8,356 workplace AI agent risk scenarios across the taxonomy categories. (A) Number of scenarios in each category, with its share of the full corpus; categories are ordered by frequency, under 4 themes. (B) Severity composition of each category. (C) Deployment mode: whether the agent assists a worker (augmentation) or replaces one (automation). (D) Framework layer at which the risk arises: technical capability, human interaction, or systemic impact. In panels (B)–(D), cells are row percentages and shading is proportional to the value. *Geopolitical and Security Risks* is omitted because no generated scenario was assigned to it; this category and its sub-categories were formed from incident-based risk scenarios rather than generated ones.

Risk Taxonomy (Li et al. 2025b) (mean 5.75 vs. 5.46 on a 1–7 scale; parametric: $t(129) = 2.23$, $p = .027$, Cohen’s $d = 0.20$; non-parametric: Wilcoxon $W = 1300.5$, $p = .026$). A Friedman test across all three frameworks was also significant ($\chi^2(2) = 6.33$, $p = .042$). Our taxonomy was preferred in 64% of decided pairwise comparisons against the GenAI Risk Taxonomy (54 wins vs. 30, 46 ties) and achieved the highest classification success rate (97.7% vs. 96.2% and 96.9%). The MIT AI Risk Repository’s domain taxonomy (Slattery et al. 2025) was excluded from the pairwise comparison because categories that should have been distinct collapsed into one: 52% of all classifications used a single catch-all category (“*AI system safety, failures, and limitations*”), making the taxonomy look easy to use even though the categories did not match the risks well.

Qualitative reflections attributed the advantage to workplace specificity: our taxonomy names workplace consequences rather than AI mechanisms. A software engineer noted: “*Workplace taxonomy went above and beyond and provided the specific result of the misinformation, which was lowered efficiency*” (18–29, Male, Software Engineer). A family medicine doctor found that it “*clearly captured the patient safety consequences of prescribing the wrong medication; the GenAI Risk Taxonomy focused more on the data/privacy aspect, while the AI Risk Repository’s Domain Taxonomy emphasized the technical AI failure rather than the medical risk to patients*” (30–39, Female, Family Medicine Doctor). A financial analyst highlighted that “*the risk is about workflow disruption and productivity loss; the GenAI Risk Taxonomy was harder because ‘Financial & Business’ only indirectly captures productivity loss*” (18–29, Female, Financial Analyst). Full results and participant quotes are reported in Supplementary Material M.

5 Discussion

We first discuss how our findings extend theory on risk of AI agents (§5.1), then translate them into practical guidance for workplace governance (§5.2), and finally examine the study’s limitations and directions for future research (§5.3).

5.1 Theoretical Implications

Make agentic decomposition explicit. Rather than treating an AI system as a black box, our framework models risk as emerging from the properties of three core components (agents, goals, environments) and from the interactions between them and humans. This distinction matters because interaction-driven risks can arise even when no single component is malfunctioning, a pattern our empirical findings confirm. This extends work focused primarily on agent–agent communication (Hammond et al. 2025) by showing that goal and environment mediation are equally important risk pathways. Further, the framework decomposes human interaction into five distinct relationship types (agent–human, human–human, human–environment, etc.) rather than treating it as a single category. This enables more precise theorizing about how accountability drifts, how reliance develops, and how distributed agency makes it difficult for workers to contest outcomes that affect them.

Connect framework-level concepts to empirical workplace structure. By grounding the framework in a 15-category taxonomy built from real job tasks, this paper provides something that prior work largely lacks: a taxonomy that operates at the level of specific workflows rather than broad societal patterns. The severity profile of *Erroneous Actions* by agents, and the dominance of *Capability erosion* risks, together suggest that workplace risks are driven less

by model failures than by failures in how agents and humans communicate intent, interpret outputs, and share responsibility. That both leading categories are majority-augmentation reinforces the point: the risk concentrates where a human remains in the loop, not where the agent runs alone. This implies that “capability” in agentic systems is not a property of the model alone, but formed by the agent, the human, and the organizational context in which they interact.

5.2 Practical Implications

Beyond showing that our taxonomy is statistically distinct and practically usable (§4, Finding 4), our findings point to three concrete shifts in how organizations should approach AI agent governance.

Focus risk assessment on the human-agent boundary.

The highest-frequency and highest-severity risks in our dataset consistently cluster at the point where humans interpret and act on agent outputs. Governance efforts should therefore not stop at auditing model performance. Assurance processes should also evaluate how uncertainty is communicated, whether workers have clear pathways to question or override agent recommendations, and whether organizational goals and proxy metrics create incentives that the agent system could exploit in harmful ways.

Apply layered controls that match the framework layers.

At technical capability layer, organizations should test how components fail under stress: partial observability simulations, proxy metric sensitivity checks, and multi-agent coordination tests. At human interaction layer, safeguards should address human behavior: role clarity training, contestability channels that allow workers to challenge agent decisions, and structured interventions against over-reliance. At systemic impact layer, monitoring should track slower systemic effects such as deskilling trajectories, shifts in bargaining power, and the adoption dynamics that can normalize unsafe practices across an industry.

Govern automation and augmentation separately. Automated deployments concentrate risk in institutional outcomes: legal exposure, operational failure, reputational damage. These require hard fail-safes, clear audit trails, and well-defined incident response procedures. Augmented deployments concentrate risk in the individual worker: skill erosion, professional judgment displacement, and workplace stress. These require socio-technical safeguards: protected human override, regular skill assessments, and workflow redesigns that ensure workers retain meaningful agency rather than simply approving agent recommendations. The taxonomy itself can serve as a practical tool here, used as a living risk register that is updated iteratively as agent capabilities and organizational practices evolve.

5.3 Limitations and Future Work

Although the scenarios were generated by an LLM, we evaluated them with workers and LLM judges and grounded them in real incident data and a live multi-agent red-teaming study. This methodology is well grounded in established risk foresight and validation practice (Perez et al. 2022; Ganguli et al. 2022; Fröhling et al. 2026; Buçinca et al. 2023; Rao

et al. 2025; Herdel et al. 2025; Wang et al. 2024b), but future work could test whether the same risk patterns emerge under different generation conditions, prompting techniques, and agentic approaches. The divergence in novelty scores between workers and the LLM judge also deserves attention. Workers found the scenarios moderately novel, offering genuine discovery value; the LLM was far stricter, likely recognizing training data patterns. This raises a question on whether there categories of workplace AI agent risk that are genuinely novel and therefore harder to surface through LLM-driven foresight?

The automation vs. augmentation distinction simplifies a continuum. Many real deployments are hybrid: tasks may be partially automated, intermittently supervised, or redistributed across roles in ways that do not fit cleanly into either category. Similarly, O*NET reflects a U.S.-centric occupational framing and may not capture informal work, global labor conditions, or sector-specific practices. Finally, the taxonomy represents a snapshot. Because the scenarios are LLM-generated, the category frequencies and distributional analyses describe only the composition of our risk corpus; they should not be interpreted as estimates of how often these risks occur in real workplaces. These patterns may also reflect the selected job tasks, prompt design, and model used. As agent capabilities evolve, both the distribution of risks and the taxonomy structure will need updating. Future work should triangulate these LLM-generated scenarios with field studies and incident reports, test robustness across models and prompt designs, enrich the occupational grounding with organizational context such as staffing constraints and safety culture, and conduct participatory validation with workers and domain stakeholders to ensure the taxonomy supports fairness-oriented governance in practice.

6 Conclusion

As AI agents become embedded in everyday work, organizations need practical ways to anticipate the risks they may introduce. This paper contributes: a multi-layer framework that models how risks emerge from interactions between agents, goals, environments, and workers; a corpus of 8,356 risk scenarios grounded in 2,078 real job tasks across 209 roles; and a 15-category taxonomy that organizes those risks under 4 themes. Across our corpus, the main pattern is that workplace risks often emerge from breakdowns at the human-agent boundary rather than from isolated model failures. More scenarios were classified as *Erroneous Actions* of agents than under any other category (30.6%), and 88.2% of these scenarios were rated High or Critical. *Capability Erosion* was the second most frequently assigned category (21.3%). Scenarios in both categories were predominantly associated with augmentation, showing that keeping a worker in the loop does not eliminate risk. These patterns point to the same conclusion: the central challenge is not only building more capable agents, but also designing safer ways for humans to work alongside them.

Acknowledgements

This work was supported by UK Research and Innovation (grant number EP/S023356/1), in the UKRI Centre for Doctoral Training in Safe and Trusted Artificial Intelligence (www.safeandtrustedai.org).

References

- Abercrombie, G.; Benbouzid, D.; Giudici, P.; Golpayegani, D.; Hernandez, J.; Noro, P.; Pandit, H.; Paraschou, E.; Pownall, C.; Prajapati, J.; Sayre, M. A.; Sengupta, U.; Suriya-wongkul, A.; Thelot, R.; Vei, S.; and Waltersdorfer, L. 2024. A Collaborative, Human-Centred Taxonomy of AI, Algorithmic, and Automation Harms. *arXiv:2407.01294*.
- Altmann, P.; Schönberger, J.; Illium, S.; Zorn, M.; Ritz, F.; Haider, T.; Burton, S.; and Gabor, T. 2024. Emergence in Multi-Agent Systems: A Safety Perspective. *arXiv:2408.04514*.
- Anwar, U.; Saparov, A.; Rando, J.; Paleka, D.; Turpin, M.; Hase, P.; Lubana, E. S.; Jenner, E.; Casper, S.; Sourbut, O.; Edelman, B. L.; Zhang, Z.; Günther, M.; Korinek, A.; Hernandez-Orallo, J.; Hammond, L.; Bigelow, E.; Pan, A.; Langosco, L.; Korbak, T.; Zhang, H.; Zhong, R.; hÉigeartaigh, S. O.; Recchia, G.; Corsi, G.; Chan, A.; Anderljung, M.; Edwards, L.; Petrov, A.; de Witt, C. S.; Motwan, S. R.; Bengio, Y.; Chen, D.; Torr, P. H. S.; Albanie, S.; Maharaj, T.; Foerster, J.; Tramer, F.; He, H.; Kasirzadeh, A.; Choi, Y.; and Krueger, D. 2024. Foundational Challenges in Assuring Alignment and Safety of Large Language Models. *arXiv:2404.09932*.
- Autor, D. H. 2015. Why are there still so many jobs? The history and future of workplace automation. *Journal of economic perspectives*, 29(3): 3–30.
- Bainbridge, L. 1983. Ironies of automation. *Automatica*, 19(6): 775–779.
- Bogucka, E.; Šćepanović, S.; and Quercia, D. 2024. Atlas of AI Risks: Enhancing Public Understanding of AI Risks. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, 12(1): 33–43.
- Bryan, P.; Severi, G.; de Gruyter, J.; Jones, D.; Bullwinkel, B.; Minnich, A.; Chawla, S.; Lopez, G.; Pouliot, M.; Fournay, A.; Maxwell, W.; Pratt, K.; Qi, S.; Chikanov, N.; Lutz, R.; Dheekonda, R. S. R.; Jagdagdorj, B.-E.; Kim, E.; Song, J.; Hines, K.; Jones, D.; Lundeen, R.; Vaughan, S.; Westerhoff, V.; Zunger, Y.; Kawaguchi, C.; Russinovich, M.; and Kumar, R. S. S. 2025. Taxonomy of Failure Mode in Agentic AI Systems. Technical report, Microsoft. Microsoft AI Red Team (AIRT) Whitepaper.
- Buçinca, Z.; Pham, C. M.; Jakesch, M.; Ribeiro, M. T.; Olteanu, A.; and Amershi, S. 2023. AHA!: Facilitating AI Impact Assessment by Generating Examples of Harms. *arXiv:2306.03280*.
- Chan, A.; Salganik, R.; Markelius, A.; Pang, C.; Rajkumar, N.; Krashennikov, D.; Langosco, L.; He, Z.; Duan, Y.; Carroll, M.; Lin, M.; Mayhew, A.; Collins, K.; Molamohammadi, M.; Burden, J.; Zhao, W.; Rismani, S.; Voudouris, K.; Bhatt, U.; Weller, A.; Krueger, D.; and Maharaj, T. 2023. Harms from Increasingly Agentic Algorithmic Systems. In *2023 ACM Conference on Fairness Accountability and Transparency, FAccT '23*, 651–666. ACM.
- Chiang, C.-H.; and yi Lee, H. 2023. Can Large Language Models Be an Alternative to Human Evaluations? *arXiv:2305.01937*.
- Critch, A.; and Russell, S. 2023. TASRA: a Taxonomy and Analysis of Societal-Scale Risks from AI. *arXiv:2306.06924*.
- Eloundou, T.; Manning, S.; Mishkin, P.; and Rock, D. 2024. GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702): 1306–1308.
- Fabris, A.; Messina, S.; Silvello, G.; and Susto, G. A. 2022. Tackling Documentation Debt: A Survey on Algorithmic Fairness Datasets. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization, EAAMO '22*. New York, NY, USA: Association for Computing Machinery. ISBN 9781450394772.
- Felten, E.; Raj, M.; and Seamans, R. 2023. How will language modelers like ChatGPT affect occupations and industries? *arXiv preprint arXiv:2303.01157*.
- Fergusson, G.; Fitzgerald, C.; Frascella, C.; Iorio, M.; McBrien, T.; Schroeder, C.; Winters, B.; and Zhou, E. 2023. Generating Harms: Generative AI's Impact & Path Forward. Technical report, Electronic Privacy Information Center (EPIC).
- Fröhling, L.; Giaconia, A.; Bogucka, E. P.; and Quercia, D. 2026. Agent-Supported Foresight for AI Systemic Risks: AI Agents for Breadth, Experts for Judgment. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*, 1–48.
- Ganguli, D.; Lovitt, L.; Kernion, J.; Askell, A.; Bai, Y.; Kadavath, S.; Mann, B.; Perez, E.; Schiefer, N.; Ndousse, K.; et al. 2022. Red teaming language models to reduce harms: Methods, scaling behaviors, and lessons learned. *arXiv preprint arXiv:2209.07858*.
- Gebru, T.; Morgenstern, J.; Vecchione, B.; Vaughan, J. W.; Wallach, H.; Daumé III, H.; and Crawford, K. 2021. Datasheets for Datasets. *Communications of the ACM*, 64(12): 86–92.
- Ghosh, S.; Simkin, B.; Shiarlis, K.; Nandi, S.; Zhao, D.; Fiedler, M.; Bazinska, J.; Pope, N.; Prabhu, R.; Rohrer, D.; Demoret, M.; and Richardson, B. 2025. A Safety and Security Framework for Real-World Agentic Systems. *arXiv:2511.21990*.
- Google. 2025. Gemini models. <https://ai.google.dev/gemini-api/docs/models>.
- Gridach, M.; Nanavati, J.; Abidine, K. Z. E.; Mendes, L.; and Mack, C. 2025. Agentic AI for Scientific Discovery: A Survey of Progress, Challenges, and Future Directions. *arXiv:2503.08979*.
- Hammond, L.; Chan, A.; Clifton, J.; Hoelscher-Obermaier, J.; Khan, A.; McLean, E.; Smith, C.; Barfuss, W.; Foerster, J.; Gavenčiak, T.; Han, T. A.; Hughes, E.; Kovařík, V.; Kulveit, J.; Leibo, J. Z.; Oesterheld, C.; de Witt, C. S.; Shah, N.; Wellman, M.; Bova, P.; Cimpanu, T.; Ezell, C.;

- Feuillade-Montixi, Q.; Franklin, M.; Kran, E.; Krawczuk, I.; Lamparth, M.; Lauffer, N.; Meinke, A.; Motwani, S.; Reuel, A.; Conitzer, V.; Dennis, M.; Gabriel, I.; Gleave, A.; Hadfield, G.; Haghtalab, N.; Kasirzadeh, A.; Krier, S.; Larson, K.; Lehman, J.; Parkes, D. C.; Piliouras, G.; and Rahwan, I. 2025. Multi-Agent Risks from Advanced AI. arXiv:2502.14143.
- Herdel, V.; Sćepanović, S.; Bogucka, E.; and Quercia, D. 2025. *ExploreGen: Large Language Models for Envisioning the Uses and Risks of AI Technologies*, 584–596. AAAI Press.
- Hoffman, M.; and Frase, H. 2023. Adding Structure to AI Harm: An Introduction to CSET’s AI Harm Framework. Technical report, Center for Security and Emerging Technology (CSET).
- Hu, T.; and Zhou, X.-H. 2024. Unveiling LLM Evaluation Focused on Metrics: Challenges and Solutions. arXiv:2404.09135.
- Hutiri, W.; Papakyriakopoulos, O.; and Xiang, A. 2024. Not My Voice! A Taxonomy of Ethical and Safety Harms of Speech Generators. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, FAccT ’24, 359–376. ACM.
- IBM. 2024. What is Agentic AI? <https://www.ibm.com/think/topics/agentic-ai>.
- Khoos, S.; Foo, J.; and Lee, R. K.-W. 2025. With Great Capabilities Come Great Responsibilities: Introducing the Agentic Risk & Capability Framework for Governing Agentic AI Systems. arXiv:2512.22211.
- Kong, D.; Lin, S.; Xu, Z.; Wang, Z.; Li, M.; Li, Y.; Zhang, Y.; Peng, H.; Sha, Z.; Li, Y.; Lin, C.; Wang, X.; Liu, X.; Zhang, N.; Chen, C.; Khan, M. K.; and Han, M. 2025. A Survey of LLM-Driven AI Agent Communication: Protocols, Security Risks, and Defense Countermeasures. arXiv:2506.19676.
- Lee, H.-P.; Yang, Y.-J.; von Davier, T. S.; Forlizzi, J.; and Das, S. 2024. Deepfakes, Phrenology, Surveillance, and More! A Taxonomy of AI Privacy Risks. arXiv:2310.07879.
- Leslie, D. 2019. Understanding artificial intelligence ethics and safety: A guide for the responsible design and implementation of AI systems in the public sector.
- Li, D.; Jiang, B.; Huang, L.; Beigi, A.; Zhao, C.; Tan, Z.; Bhattacharjee, A.; Jiang, Y.; Chen, C.; Wu, T.; Shu, K.; Cheng, L.; and Liu, H. 2025a. From Generation to Judgment: Opportunities and Challenges of LLM-as-a-judge. arXiv:2411.16594.
- Li, H.; Dong, Q.; Chen, J.; Su, H.; Zhou, Y.; Ai, Q.; Ye, Z.; and Liu, Y. 2024. LLMs-as-Judges: A Comprehensive Survey on LLM-based Evaluation Methods. arXiv:2412.05579.
- Li, M.; Bickersteth, W.; Tang, N.; Hong, J.; Cranor, L.; Shen, H.; and Heidari, H. 2025b. A Closer Look at the Existing Risks of Generative AI: Mapping the Who, What, and How of Real-World Incidents. arXiv:2505.22073.
- Liu, B.; Li, X.; Zhang, J.; Wang, J.; He, T.; Hong, S.; Liu, H.; Zhang, S.; Song, K.; Zhu, K.; Cheng, Y.; Wang, S.; Wang, X.; Luo, Y.; Jin, H.; Zhang, P.; Liu, O.; Chen, J.; Zhang, H.; Yu, Z.; Shi, H.; Li, B.; Wu, D.; Teng, F.; Jia, X.; Xu, J.; Xiang, J.; Lin, Y.; Liu, T.; Liu, T.; Su, Y.; Sun, H.; Berseth, G.; Nie, J.; Foster, I.; Ward, L.; Wu, Q.; Gu, Y.; Zhuge, M.; Liang, X.; Tang, X.; Wang, H.; You, J.; Wang, C.; Pei, J.; Yang, Q.; Qi, X.; and Wu, C. 2025. Advances and Challenges in Foundation Agents: From Brain-Inspired Intelligence to Evolutionary, Collaborative, and Safe Systems. arXiv:2504.01990.
- Macrae, C. 2022. Learning from the Failure of Autonomous and Intelligent Systems: Accidents, Safety, and Sociotechnical Sources of Risk. *Risk Analysis*, 42(9): 1999–2025.
- McGregor, S. 2021. Preventing repeated real world AI failures by cataloging incidents: The AI incident database. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 15458–15463.
- Mehrabi, N.; Morstatter, F.; Saxena, N.; Lerman, K.; and Galstyan, A. 2021. A Survey on Bias and Fairness in Machine Learning. *ACM Comput. Surv.*, 54(6).
- Microsoft. 2023. Types of harm - Azure Application Architecture Guide.
- Mitchell, M.; Wu, S.; Zaldivar, A.; Barnes, P.; Vasserman, L.; Hutchinson, B.; Spitzer, E.; Raji, I. D.; and Gebru, T. 2019. Model Cards for Model Reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, FAT* ’19, 220–229. New York, NY, USA: Association for Computing Machinery. ISBN 9781450361255.
- Mohney, M. 2025. What AI Means for the Future of Work.
- Moss, E.; Watkins, E. A.; Singh, R.; Elish, M. C.; and Metcalf, J. 2021. Assembling accountability: algorithmic impact assessment for the public interest. *Available at SSRN* 3877437.
- National Center for O*NET Development. 2025. O*NET Database. O*NET Resource Center. <https://www.onetcenter.org/database.html>.
- OpenAI. 2024. GPT-4o mini: advancing cost-efficient intelligence. <https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/>.
- Paeth, K.; Atherton, D.; Pittaras, N.; Frase, H.; and McGregor, S. 2025. Lessons for editors of AI incidents from the AI incident database. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 28946–28953.
- Panickssery, A.; Bowman, S. R.; and Feng, S. 2024. Llm evaluators recognize and favor their own generations. *Advances in Neural Information Processing Systems*, 37: 68772–68802.
- Perez, E.; Huang, S.; Song, F.; Cai, T.; Ring, R.; Aslanides, J.; Glaese, A.; McAleese, N.; and Irving, G. 2022. Red teaming language models with language models. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 3419–3448.
- Raisch, S.; and Krakowski, S. 2021. Artificial intelligence and management: The automation–augmentation paradox. *Academy of management review*, 46(1): 192–210.
- Raji, I. D.; Kumar, I. E.; Horowitz, A.; and Selbst, A. 2022. The Fallacy of AI Functionality. In *2022 ACM Conference on Fairness Accountability and Transparency*, FAccT ’22, 959–972. ACM.

- Raji, I. D.; Smart, A.; White, R. N.; Mitchell, M.; Gebru, T.; Hutchinson, B.; Smith-Loud, J.; Theron, D.; and Barnes, P. 2020. Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, 33–44.
- Ranjan, R.; Gupta, S.; and Singh, S. N. 2025. Fairness in Agentic AI: A Unified Framework for Ethical and Equitable Multi-Agent Systems. arXiv:2502.07254.
- Rao, P. S. B.; Šćepanović, S.; Zhou, K.; Bogucka, E. P.; and Quercia, D. 2025. RiskRAG: A Data-Driven Solution for Improved AI Model Risk Reporting. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI '25. New York, NY, USA: Association for Computing Machinery. ISBN 9798400713941.
- Richards, I.; Benn, C.; and Zilka, M. 2025. From Incidents to Insights: Patterns of Responsibility following AI Harms. In *Proceedings of the 5th ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, 151–169.
- Russell, S.; and Norvig, P. 2009. *Artificial Intelligence: A Modern Approach*. USA: Prentice Hall Press, 3rd edition. ISBN 0136042597.
- Sapkota, R.; Roumeliotis, K. I.; and Karkee, M. 2026. AI Agents vs. Agentic AI: A Conceptual taxonomy, applications and challenges. *Information Fusion*, 126: 103599.
- 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*, FAT* '19, 59–68. New York, NY, USA: Association for Computing Machinery. ISBN 9781450361255.
- Shao, Y.; Huang, C.; Song, Y.; Wang, M.; Song, Y. H.; and Shao, R. 2025a. Using Augmentation-Based AI Tool at Work: A Daily Investigation of Learning-Based Benefit and Challenge. *Journal of Management*, 51(8): 3352–3390.
- Shao, Y.; Zope, H.; Jiang, Y.; Pei, J.; Nguyen, D.; Brynjolfsson, E.; and Yang, D. 2025b. Future of Work with AI Agents: Auditing Automation and Augmentation Potential across the U.S. Workforce. arXiv:2506.06576.
- Shapira, N.; Wendler, C.; Yen, A.; Sarti, G.; Pal, K.; Floody, O.; Belfki, A.; Loftus, A.; Jannali, A. R.; Prakash, N.; Cui, J.; Rogers, G.; Brinkmann, J.; Rager, C.; Zur, A.; Ripa, M.; Sankaranarayanan, A.; Atkinson, D.; Gandikota, R.; Fiotto-Kaufman, J.; Hwang, E.; Orgad, H.; Sahil, P. S.; Taglicht, N.; Shabtay, T.; Ambus, A.; Alon, N.; Oron, S.; Gordon-Tapiero, A.; Kaplan, Y.; Schwartz, V.; Shaham, T. R.; Riedl, C.; Mirsky, R.; Sap, M.; Manheim, D.; Ullman, T.; and Bau, D. 2026. Agents of Chaos. arXiv:2602.20021.
- Shelby, R.; Rismani, S.; Henne, K.; Moon, A.; Roshtamzadeh, N.; Nicholas, P.; Yilla, N.; Gallegos, J.; Smart, A.; Garcia, E.; and Virk, G. 2023. Sociotechnical Harms of Algorithmic Systems: Scoping a Taxonomy for Harm Reduction. arXiv:2210.05791.
- Si, C.; Yang, D.; and Hashimoto, T. 2024. Can LLMs Generate Novel Research Ideas? A Large-Scale Human Study with 100+ NLP Researchers. arXiv:2409.04109.
- Slattery, P.; Saeri, A. K.; Grundy, E. A. C.; Graham, J.; Noetel, M.; Uuk, R.; Dao, J.; Pour, S.; Casper, S.; and Thompson, N. 2025. The AI Risk Repository: A Comprehensive Meta-Review, Database, and Taxonomy of Risks From Artificial Intelligence. arXiv:2408.12622.
- Steenstra, I.; and Bickmore, T. W. 2025. A Risk Ontology for Evaluating AI-Powered Psychotherapy Virtual Agents. In *Proceedings of the 25th ACM International Conference on Intelligent Virtual Agents*, IVA '25. New York, NY, USA: Association for Computing Machinery.
- Turri, V.; and Dzombak, R. 2023. Why we need to know more: Exploring the state of AI incident documentation practices. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 576–583.
- Wang, R.; Cheng, R.; Ford, D.; and Zimmermann, T. 2024a. Investigating and designing for trust in ai-powered code generation tools. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, 1475–1493.
- Wang, Z. J.; Kulkarni, C.; Wilcox, L.; Terry, M.; and Madaio, M. 2024b. Farsight: Fostering Responsible AI Awareness During AI Application Prototyping. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, CHI '24, 1–40. ACM.
- Weidinger, L.; Rauh, M.; Marchal, N.; Manzini, A.; Hendricks, L. A.; Mateos-Garcia, J.; Bergman, S.; Kay, J.; Griffin, C.; Bariach, B.; Gabriel, I.; Rieser, V.; and Isaac, W. 2023. Sociotechnical Safety Evaluation of Generative AI Systems. arXiv:2310.11986.
- Wooldridge, M. 2009. *An Introduction to MultiAgent Systems*. Wiley Publishing, 2nd edition. ISBN 0470519460.
- Xi, Z.; Chen, W.; Guo, X.; He, W.; Ding, Y.; Hong, B.; Zhang, M.; Wang, J.; Jin, S.; Zhou, E.; et al. 2025. The rise and potential of large language model based agents: A survey. *Science China Information Sciences*, 68(2): 121101.
- Yin, F.; Laban, P.; Peng, X.; Zhou, Y.; Mao, Y.; Vats, V.; Ross, L.; Agarwal, D.; Xiong, C.; and Wu, C.-S. 2025. Bingoguard: Llm content moderation tools with risk levels. arXiv preprint arXiv:2503.06550.
- Zeng, Y.; Klyman, K.; Zhou, A.; Yang, Y.; Pan, M.; Jia, R.; Song, D.; Liang, P.; and Li, B. 2024. AI Risk Categorization Decoded (AIR 2024): From Government Regulations to Corporate Policies. arXiv:2406.17864.
- Zhang, R.; Li, H.; Meng, H.; Zhan, J.; Gan, H.; and Lee, Y.-C. 2025. The Dark Side of AI Companionship: A Taxonomy of Harmful Algorithmic Behaviors in Human–AI Relationships. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI '25. New York, NY, USA: Association for Computing Machinery.
- Zheng, L.; Chiang, W.-L.; Sheng, Y.; Zhuang, S.; Wu, Z.; Zhuang, Y.; Lin, Z.; Li, Z.; Li, D.; Xing, E. P.; Zhang, H.; Gonzalez, J. E.; and Stoica, I. 2023. Judging LLM-as-a-judge with MT-bench and Chatbot Arena. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS '23. Red Hook, NY, USA: Curran Associates Inc.