

AI Adoption in the Workplace: A Scoping Review on Benefits, Risks, Mitigations, and Readiness Frameworks

Abstract

With the increasing use of AI in organizations to enhance or automate occupational tasks, there is a poorly structured problem: how to trade off genuine productivity gains against the documented risks AI poses to workers and organizational performance. This scoping review synthesizes 48 sources (20 empirical studies, 13 practice and policy documents, and 15 theoretical or review papers) to identify factors shaping expected benefits, risks, mitigation strategies, and organizational readiness when AI is adopted for an occupational task. We identify four interdependent themes: Expected Benefits (efficiency, skill augmentation, consistency, accelerated innovation, employee wellbeing); Anticipated Risks (skill erosion, IP and copyright gaps, occupational stress, algorithmic opacity, accountability diffusion, professional identity shifts); Mitigation Strategies (human oversight, IP governance, lifecycle governance, transparency, participatory design, co-skilling); and Organizational Readiness (infrastructure, AI literacy, leadership alignment and cultural readiness, reskilling strategy, regulatory readiness). We apply the framework to AI-assisted concept generation in packaging design, drawing on design-specific evidence.

Keywords

AI adoption, organizational readiness, AI risks, creative work, packaging design, EU AI Act, workplace AI, reskilling, scoping review

1 INTRODUCTION

Artificial intelligence is no longer a boundary technology. From logistics scheduling to medical imaging, organizations are actively integrating AI tools into occupational workflows. Yet the decision to adopt AI for a specific task remains poorly supported by generalizable, evidence-based guidance. Most research either evaluates a single tool in a single sector, or addresses adoption at a level of abstraction too high to inform practice. Policy instruments, including the EU AI Act (2024) [13], address post-decision compliance obligations but not the pre-decision assessment of whether adoption is advisable and under what conditions. This scoping review addresses the question: What factors shape the expected benefits, risks, mitigations, and organizational readiness when an organization adopts AI to augment or automate an occupational task? The question is deliberately general. A framework whose factors travel across occupations, tools, and sectors is our goal. Profession-specific content appears exclusively in the worked use case (Section 4).

2 Method

This scoping review was conducted in accordance with the PRISMA-ScR (Preferred Reporting Items for Systematic Reviews and Meta-Analyses extension for Scoping Reviews) guidelines.

2.1 Information Sources and Search Strategy

Searches were conducted between March 24th and April 22nd, 2026, across Scopus and Google Scholar, supplemented by direct retrieval of policy documents from institutional repositories (NIST, OECD, EU-OSHA, EU legislative database). To ensure our literature mapping systematically addressed our research question, we structured our Boolean search queries around the four dimensions of our Contextual Integrity framework, while specifically targeting knowledge and creative work contexts (e.g., packaging design) to address known literature gaps. The core search clusters included can be found in Table 1.

2.2 Eligibility Criteria

To ensure relevance and rigor, strict eligibility criteria were applied.

Inclusion Criteria: Papers published between 2015 and 2026 (Figure 2); peer-reviewed academic literature, industry reports, or official policy documents; English-language publications; studies addressing human-AI interaction, organizational policy or workflow changes in knowledge, creative or service work contexts.

Exclusion Criteria: Purely technical computer science papers lacking organizational context; clinical medical/diagnostic AI studies; papers focusing on physical robotics (e.g., manufacturing/welding robots) rather than cognitive/generative AI; studies exclusively situated in university/student environments rather than occupational settings; and papers with zero citations (unless published within the last 18 months).

2.3 Selection Process

From an initial pool of 2,384 records identified via database searches, 2,017 records were removed prior to screening (1237 duplicates and 780 papers excluded based on a zero-citation count). Title screening was conducted on the remaining 372 records to remove clearly irrelevant technical or medical studies, resulting in the exclusion of 238 papers. Subsequently,

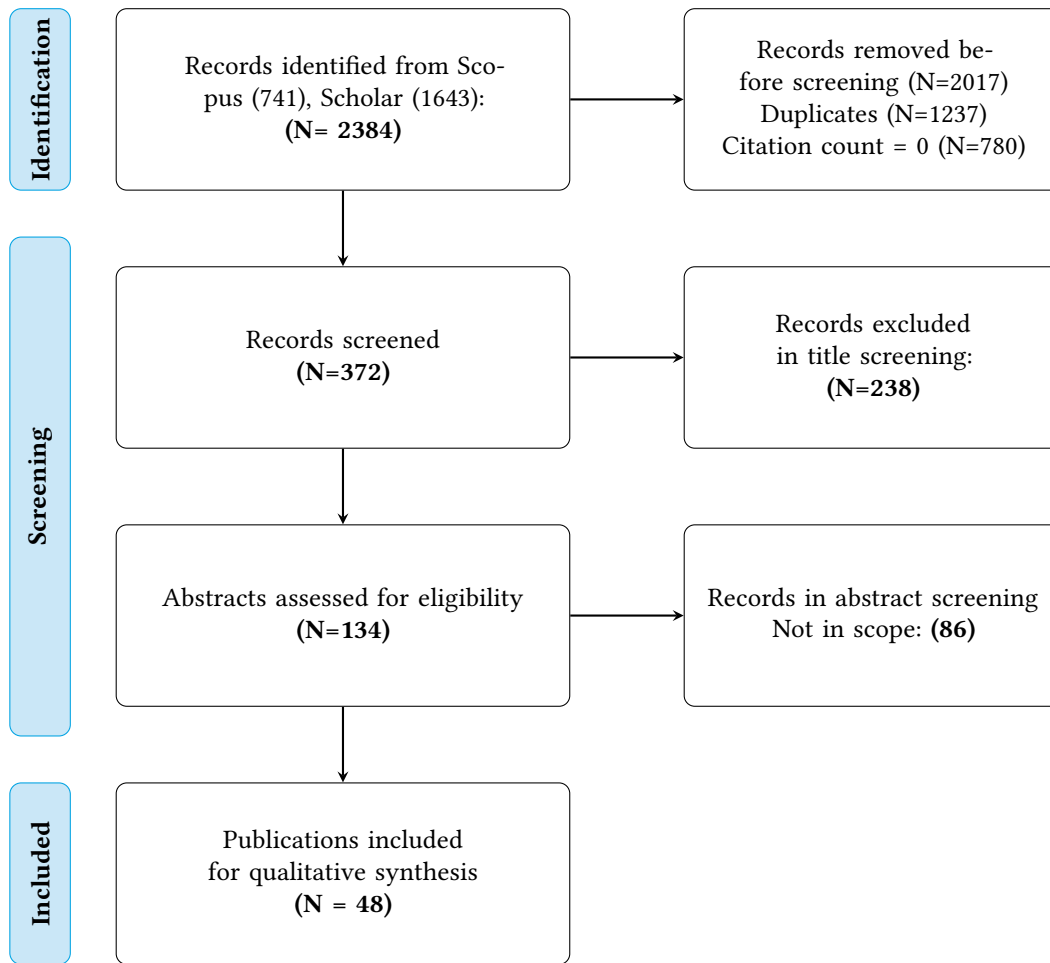


Figure 1: PRISMA-ScR Flow Chart of our Literature Selection Procedure.

134 abstracts were assessed for eligibility. To reduce individual bias, screening was not conducted alone. Abstracts and subsequent full texts were independently evaluated, resulting in the exclusion of 86 records that were not in scope. Uncertain or edge cases were discussed by the research team until consensus was reached. This process yielded a final corpus of 48 included sources (Figure 1).

2.4 Data Charting and Synthesis

We employed Reflexive Thematic Analysis (RTA) to synthesize the literature. Data extraction was documented in a centralized coding matrix (available in the supplementary ZIP archive). For each paper, we extracted key variables including the Contextual Integrity dimension (Roles, Activities, Norms, Values), the specific industry context (e.g., Creative Industry, Customer Service, General), the specific factor identified (e.g., skill erosion, algorithmic opacity), and the evidence type (Empirical, Theoretical, or Practice/Policy). These codes

were iteratively grouped into the four overarching themes presented in Section 3.

3 RESULTS

Coding produced four interdependent themes (Table 2). Organizational readiness determines which risks occur and which benefits are achievable. Reducing these risks requires prior preparation, and both risks and benefits change together as adoption deepens. When evaluating an organization, these factors must be analyzed simultaneously rather than as a step-by-step process.

3.1 Expected Benefits

The most robustly documented benefit is throughput: AI processes data and tasks at a volume, speed, and marginal cost inaccessible to human labor alone. Brynjolfsson, Li, and Raymond [1] studied 5,172 customer-support agents exposed to a generative AI assistant and observed a 15% mean productivity gain, with the most pronounced effect (30%) accruing

to novice and lower-skilled workers. This finding is replicated in the taxi-driver context by Kanazawa et al. [2], who document that low-skilled workers gain the most from AI navigation assistance, as the system effectively disseminates the knowledge of expert peers. Both studies confirm the skill augmentation mechanism: AI encodes and distributes expertise that previously required years of on-the-job learning. Frank et al. [3] provide the theoretical basis, documenting that most occupations consist of task bundles, only some of which are suitable for machine learning, meaning augmentation, not replacement, characterizes the modal adoption scenario for knowledge and creative work. Brynjolfsson and Mitchell [15] extend this with a practical rubric identifying eight properties that make tasks ML-suitable, arguing that most occupations will require deliberate redesign rather than substitution. Beyond throughput, AI adoption in creative and knowledge-intensive domains is associated with accelerated innovation cycles: generative AI tools expand the explorable design space per unit time, compressing ideation-to-prototype cycles [1, 30]. From a business-model perspective, Kanbach et al. [30] document how generative AI is restructuring the innovation value chain across different sectors, enabling smaller organizations to access ideation capacity previously available only to larger companies. A less anticipated benefit, documented by Hopkins and Gallagher [47], is the potential of AI assistance to reduce the cognitive load off employees and support their mental well-being, when deployment is well-designed and clearly framed as augmentative rather than evaluative.

3.2 Anticipated Risks

The productivity gains documented above carry structural costs. Brynjolfsson et al. [1] observe that their AI tool had marginal impact on experienced workers, raising the concern that expert employees who no longer exercise independent judgment to match AI performance may gradually lose the capability that made them experts: AI reduces not the need for human judgment, but the frequency with which workers are needed to exercise it, eroding the deliberate practice that sustains expert performance. This deskilling dynamic has been further articulated in design contexts by Hou et al. [32], whose experimental study on generative AI in the creative process identifies an *Einstellung* effect: established routines of experienced designers are disrupted, while novices benefit most.

A second cluster concerns IP, copyright, and authorship. Generative AI systems are trained on vast datasets that may include copyrighted material; their outputs can inadvertently reproduce protected visual, textual, or structural elements. Chesterman [25] analyses the foundational policy questions (training data rights, fair use, and the inadequacy of current

IP frameworks for GenAI) drawing on major litigation cases (Getty Images vs Stability AI; NY Times vs OpenAI). Novelli et al. [28] confirm that EU law, despite the AI Act, retains significant gaps in addressing liability, privacy, and IP in generative systems. Lucchi [29] demonstrates through the ChatGPT case study that copyright laws regarding fair use and adapted works must be updated. Al-kfairy et al. [23] and Huang et al. [24] further document how authorship ambiguity and attribution failures constitute systemic ethical risks that voluntary frameworks have failed to resolve.

Algorithmic opacity constitutes the third major risk cluster. Burrell [4] demonstrates empirically how machine learning systems generate outputs through processes that are fundamentally illegible, not merely complex, but structurally opaque even to their designers. When workers cannot understand why a system produced a particular output, they cannot exercise meaningful oversight, and accountability diffuses across the worker, the manager, and the vendor. Munn [5] argues that dominant ethical-principles responses to this risk are insufficient: principles remain meaningless without enforcement mechanisms, a gap partially addressed by the EU AI Act [13] through binding obligations, and by the NIST AI RMF [7] through structured governance functions.

Occupational stress and psychosocial harm represent a fourth cluster with strong empirical support. Battal [6] identifies multiple AI-induced stressors: role ambiguity when AI competes for identity-defining tasks, work intensification when output expectations exceed human capacities, and surveillance anxiety when AI outputs are used to evaluate workers. The PEROSH synthesis [17] confirms these patterns across 39 studies (2022–2024), documenting consistent depletion of job autonomy and social support. The EU-OSHA report [11] notes that AI systems intended as productivity tools frequently function as monitoring mechanisms. Salama et al. [43] demonstrate through hospitality sector data how AI awareness drives job burnout and psychological contract breach, raising turnover intentions even among workers who are not directly replaced. Wu et al. [48] add an experimental dimension: collaboration with generative AI consistently enhances task performance but simultaneously reduces perceived sense of control and intrinsic motivation, with the largest autonomy drop occurring at the moment of transitioning from solo to AI-assisted work.

Finally, Hsu and Bechky [38] theorize how generative AI reshapes social categorizations in creative work, erodes traditional markers of creative expertise, shifts power dynamics between junior and senior workers and threatens the professional identity of experienced practitioners, a risk particularly salient in professions where creative distinctiveness is central to market positioning. Xu et al. [26] additionally

document AI-enabled monitoring as a driver of unethical employee behavior, highlighting how surveillance applications of AI can increase the misconduct risks (knowledge hiding and time theft) they are designed to prevent.

3.3 Mitigation Strategies

The NIST AI Risk Management Framework [7] provides the most comprehensive mitigation structure, organizing risk management into four lifecycle functions: govern (establish accountability structures, policies, and risk tolerances), map (contextualize AI risks for the specific deployment), measure (assess risk probability and impact) and manage (respond through control measures and contingency plans). The EU AI Act [13] operationalizes several NIST-aligned mitigations as binding legal obligations for high-risk AI deployments: mandatory human oversight, technical documentation, bias monitoring, worker notification prior to deployment, and conformity assessments. Leon [27] extends this with a lifecycle-based governance model that proposes structured accountability mechanisms across the entire AI system lifecycle, from design through decommissioning, specifically targeting the reliability and ethical consistency gaps.

On IP specifically, Chesterman [25] and Novelli et al. [28] converge on the recommendation that organizations deploying generative AI must conduct pre-deployment IP audits covering training-data origin, licensing arrangements, and opt-out documentation, recommendations the EU AI Act begins to formalize through transparency requirements for general-purpose AI models. Mahajan [22] frames this within a broader argument that AI governance must address what he calls the “soul” of AI systems — their embedded values and ethical orientations — rather than treating governance as purely procedural compliance.

The CHI 2024 study by Lee [12] emphasizes that participatory design is not a soft addition to the mitigation toolkit but a structural requirement: when workers are excluded from adoption decisions, they develop adversarial AI relationships that undermine efficiency gains. This finding is reinforced by Nkanta [44], whose qualitative study of organizational communication strategies during AI-driven change identifies transparent, two-way internal communication as a critical factor in preventing resistance and managing AI-related anxiety. Zhang et al. [45] propose co-skilling, a model where AI adoption is coupled with organizational support for workers to develop AI-collaboration skills, as an effective approach to mitigate job insecurity. Dobre [31] specifically addresses the design of AI tools that prevent cognitive shortcuts, arguing that interface and workflow design choices can prevent the passive AI-reliance that accelerates deskilling.

Impact monitoring (post-deployment review for skill degradation, bias, and wellbeing) emerges consistently across the NIST RMF [7], EU-OSHA [11], and OECD Employment Outlook [14] as necessary for sustained trustworthy adoption. The McKinsey State of AI report [8] documents that organizational resistance, and not technical failure, is the primary cause of AI adoption project failure, underestimating the complementary importance of cultural readiness alongside technical governance.

3.4 Organizational Readiness

Bunjak et al. [18] document how transformational leadership and management innovation jointly predict technology adoption, establishing leadership commitment as a foundational readiness condition. The McKinsey report [8] confirms this at scale, documenting a pervasive adoption-strategy gap: most organizations approve AI budgets without establishing proper measurable governance, a misalignment that predicts adoption failure more reliably than any technical deficiency. Zhang and colleagues [16] theorize this through an algorithmic management lens, showing that leadership’s explicit stance on human oversight and accountability shapes the organization cultural receptiveness to AI integration.

The second critical readiness factor is workforce AI literacy, the capacity not merely to use AI tools, but to critically evaluate their outputs, recognize their limitations, and maintain independent judgment. Tambe [21] provides the most rigorous econometric evidence: AI and algorithms act as complements to domain expertise, creating greatest value when algorithmic literacy is broadly diffused among workers. Critically, reskilling by combining domain expertise with algorithmic literacy, is more effective than narrow AI-specific role specialization. This finding is supported by Morandini et al. [41], whose organizational review documents that effective AI workforce transitions require proactive upskilling investment rather than reactive adaptation, and by Bennett and Bennett [42], who theorize a shift from technical skills to strategic and adaptive competencies as the horizon of human resource development in AI-integrated organizations.

Appiah [46] evaluates the effectiveness of corporate upskilling and reskilling strategies, finding that skill-gap analysis before adoption, not after, is the fundamental prerequisite for effective workforce transition. The OECD [9] and McKinsey [8] both independently identify data governance deficiencies as the most persistently underestimated readiness factor: organizations assume their existing data is of sufficient quality for AI deployment and discover otherwise only after adoption failures. Regulatory and legal preparedness rounds out the readiness picture [7, 13, 14].

3.5 Contextual Integrity and the four Themes

The contextual integrity framework maps onto this review's four themes. Roles (who) determine who exercises oversight and bears accountability, themes that McKinsey [8], Lee [12], and Hsu and Bechky [38] illuminate from organizational, worker, and professional-identity perspectives. Activities (what) define which tasks are augmented, shaping which benefits are achievable and which risks are most severe. Norms (how) prescribe the governance and regulatory standards constraining acceptable AI behavior — documented by the NIST RMF [7], EU AI Act [13], and IP literature [25, 27–29]. Values (why) orient the entire analysis: whose interests drive adoption and whose must be protected. Contextual integrity is preserved when AI adoption respects these norms, violated when algorithmic systems breach occupational expectations: for instance when AI outputs erode the creative autonomy a designer's role is defined by, or when surveillance mechanisms reframe a productivity tool as an evaluation instrument.

4 WORKED USE CASE: COMMERCIAL & INDUSTRIAL DESIGNER - AI-ASSISTED PACKAGING CONCEPT GENERATION

Profession and task. Commercial and Industrial Designers in packaging design conceptualize the structure, material, and visual identity of product packaging. The AI use examined is AI-assisted concept generation: the deployment of text-to-image and generative 3D tools to produce early-stage visual and structural concepts from natural-language briefs, which designers then select, refine, and develop into finished proposals.

Expected benefits (T1). The main benefit is accelerated ideation, which directly expands the designer's explorable design space for each project. This finding is consistent with the experimental evidence of Hou et al. [32], who demonstrate that generative AI co-creation produces the greatest speed gains in the early phases of creative work, precisely the phase packaging concept generation occupies. For junior designers specifically, Li et al. [19] document that AI tools can function as tacit-knowledge transfer mechanisms, making the implicit visual vocabulary of senior practitioners accessible to less experienced team members and accelerating their skill development trajectory.

Anticipated risks (T2). Three interconnected risks are acute in the packaging design context. First, IP and copyright: generative AI tools are trained on datasets that may include copyrighted design work; AI-generated packaging concepts may reproduce protected visual or structural elements without the generating designer's awareness [25, 29]. The EU AI Act [13] does not yet mandate full training-data disclosure for

all generative AI systems. Second, creative skill erosion: packaging design requires the integration of aesthetic judgment, material knowledge, sustainability expertise, and manufacturing feasibility, a complex tacit skill bundle [3]. Shukla et al. [20] specifically warn against the paradox of AI-assisted design: cognitive offloading and misplaced responsibilities emerge as designers delegate generative decisions to AI systems, gradually losing the independent generative capacity that defines senior expertise [1, 2]. Hou et al. [32] add that the Einstellung effect is particularly pronounced for experienced designers: existing creative heuristics, consolidated through years of practice, are disrupted by generative AI co-creation, while novices, lacking such consolidated routines, benefit most. Third, professional identity and occupational stress: design agencies are experiencing pressure to increase concept volume while maintaining quality. As Hsu and Bechky [38] document in the context of the "digital undertow," AI erodes traditional markers of creative expertise in ways that affect senior designers' professional positioning and job satisfaction. Li et al. [19] confirm through interviews with 20 UX designers that fears about job replacement and creative exhaustion are real concerns among professionals.

Mitigations (T3). Applying Theme 3, the organization should: (1) require a qualified designer to critically evaluate and substantially transform all AI-generated outputs before client delivery, implementing genuine human creative oversight [7, 13]; (2) conduct pre-deployment IP audits of AI tools, establish contractual clauses clarifying AI-assisted content ownership with clients, and monitor vendor training-data disclosure as EU AI Act transparency requirements take effect [25, 28]; (3) apply Dobre's [31] design-for-expertise principles in selecting and configuring AI tools, specifically preferring systems that surface options for human selection over systems that deliver single outputs, thereby preserving the designer's evaluative judgment; (4) involve designers in tool selection and protocol definition, as Li et al. [19] find that professionals who participate in adoption decisions show substantially higher engagement and lower resistance; (5) monitor designer creative output quality and independent concept generation capability on a semi-annual basis to detect early deskilling signals [7, 20]; and (6) implement structured internal communication about AI adoption rationale and role boundaries, as Nkanta [44] documents this as a critical factor in preventing anxiety-driven resistance.

Organizational readiness (T4). Applying Theme 4 to a packaging design agency: does the agency's IT infrastructure support the AI tools without data-privacy violations? Do designers understand AI output limitations well enough to critically evaluate outputs (AI literacy)? Has design leadership established a governance position on AI use aligned with the agency's positioning and client commitments (strategic

alignment)? Is the studio culture open to experimentation and tolerant of failure to support AI-assisted process change (cultural readiness)? Tambe [21] provides the most direct readiness prescription: agencies should invest in broad AI literacy programs that combine algorithmic understanding with deepened packaging-domain expertise, rather than narrow prompt-engineering skills, since the former is more resilient to tool changes and better preserves the balance between human judgment and AI outputs. Appiah [46] adds that skill-gap analysis should precede tool procurement, not follow it. Finally, the legal counsel must review EU AI Act obligations applicable to AI-generated commercial packaging content, particularly for regulated categories such as food, pharmaceutical, and child-product packaging (regulatory readiness).

Framework judgment. AI adoption for packaging concept generation is conditionally viable: efficiency and innovation-cycle benefits are real and replicable. However, viability requires active IP governance, preservation of a human creative lead over AI concept selection and development, proactive reskilling investment in domain expertise, and ongoing monitoring of designer skill trajectories. Without these conditions, legal liability from IP infringement, reputational risk from homogenized or derivative design output, and medium-term professional capability degradation collectively outweigh throughput gains.

5 DISCUSSION

The application of our four-theme framework to the packaging design use case illustrates a central finding of this scoping review: the benefits of AI adoption are inseparable from the organizational, regulatory, and cultural structures in which the technology is embedded. By synthesizing 48 sources across the dimensions of roles, activities, norms, and values, we have constructed a tool for assessing both risks and organizational readiness. However, the robustness of any such framework is bound by the scope and biases of its underlying literature. In this section, we critically evaluate the blind spots within the current evidence base and summarize the final implications for organizational practice.

5.1 GAPS AND FUTURE WORK

Four significant gaps limit the framework's current evidence base. First, empirical evidence remains concentrated in service work, logistics, and customer support [1, 2]; creative and knowledge-intensive occupations, where AI adoption is advancing most rapidly, are underrepresented with design-specific evidence [19, 20, 32, 38] predominantly focused on UX rather than industrial or packaging design. Second, the geographic distribution is heavily set toward the United States and Western Europe; how the framework's factors

manifest under different labor-relations regimes and cultural norms remains open. Third, longitudinal evidence on skill erosion is essentially absent: the deskilling hypothesis, while increasingly supported by cross-sectional and experimental data [1, 2, 20, 32], requires further study designs to establish causal pathways and timescales. Fourth, the interaction between specific organizational readiness profiles and specific risk-materialization probabilities has not been quantitatively modeled, limiting the framework's use.

5.2 CONCLUSION

This scoping review has synthesized 48 sources to produce a general, transferable framework for assessing AI adoption in organizational and occupational contexts across four interdependent themes: expected benefits, anticipated risks, mitigation strategies, and organizational readiness. The framework's central finding is that AI adoption outcomes are not determined by the technology alone, but by the interaction between technical capability, organizational governance, workforce capacity, and the regulatory environment in which deployment occurs. Efficiency gains are real and empirically robust, but they are conditional on governance structures that preserve meaningful human oversight, on reskilling investments that deepen rather than replace domain expertise, and on IP frameworks adequate to the novel challenges of generative AI.

The packaging design use case demonstrates the framework's practical utility: it produces a structured, profession-specific adoption assessment that is derived from the evidence rather than assumed. The same analytical moves are available to any organization assessing AI adoption for any occupational task.

Appendix

Research Team

All contributing members of this research team are affiliated with **Politecnico di Torino** (Turin, Italy).

- Edoardo Lamantea, s339548@studenti.polito.it
- Elia Mandrelli, s322108@studenti.polito.it
- Ali Adham, s309742@studenti.polito.it
- Sofia Pesare, s342909@studenti.polito.it
- Cyril de Wandeleer, s343255@studenti.polito.it
- Pablo Panariti, s323832@studenti.polito.it
- Andrea Vozza, s339885@studenti.polito.it
- Alberto Benedicenti, s326539@studenti.polito.it
- Emanuele Dulla, s337030@studenti.polito.it
- Alessandro Gabriele Solomon, s341327@studenti.polito.it
- Gemma Stumpo, s340329@studenti.polito.it

Tables and Figures

Table 1: Example Boolean Search Strings by Concept Cluster

Concept Cluster	Boolean Search String
Roles (Power & Identity)	("artificial intelligence" OR "generative AI" OR "augmentation vs automation") AND ("workplace" OR "organization*" OR "design*") AND ("power dynamics" OR "professional identity" OR "autonomy" OR "designer autonomy" OR "human-in-the-loop" OR "HITL")
Activities (Task & Skill Transformation)	("artificial intelligence" OR "generative AI") AND ("workplace" OR "organization*" OR "design*") AND ("deskilling" OR "creative skill erosion" OR "skill erosion" OR "task restructuring" OR "design ideation speed" OR "concept generation variants" OR "algorithmic aesthetics")
Norms (Governance & Accountability)	("artificial intelligence" OR "generative AI") AND ("workplace" OR "organization*" OR "design*") AND ("algorithmic opacity" OR "accountability boundaries" OR "copyright" OR "legal liability" OR "IP infringement" OR "copyright governance")
Values (Wellbeing & Readiness)	("artificial intelligence" OR "generative AI") AND ("workplace" OR "organization*" OR "design*") AND ("reskilling strategies" OR "AI literacy" OR "worker psychological safety" OR "cross-cultural bias" OR "cross-cultural AI bias" OR "non-Western AI adoption" OR "psychological safety")




Figure 2: Distribution of included papers by publication year.

The notable idea of the graph is that most of the data comes from recent papers (2024-2025, highlighted in blue). Considering this literature production rate, we believe that the gaps in our framework can be quickly filled by future literature.

Table 2: Summary Framework: Themes, Factors, Evidence Status, and Key Sources

Theme	Factor	Description	Evidence	Key Sources
T1 - Expected Benefits	Operational efficiency & throughput	AI processes high task volumes faster and at lower marginal cost	Strong (empirical)	[1–3]
	Skill augmentation	AI surfaces information and best practices, extending worker capability without replacing judgment	Strong (empirical + theoretical)	[1–3, 10, 19]
	Decision consistency	Uniform application of criteria reduces inconsistency; depends on data quality	Moderate	[2, 15]
	Accelerated innovation cycles	AI shortens ideation-to-prototype cycles; expands explorable design space	Emerging (empirical)	[1, 30, 32, 35]
	Employee wellbeing	Well-designed AI assistance reduces cognitive load and supports mental wellbeing when framed as augmentative	Emerging (empirical)	[47, 48]
T2 - Anticipated Risks	Skill erosion & deskilling	Over-reliance on AI output atrophies independent professional judgment; experienced employees especially are at risk of complacency	Moderate-Strong	[1, 2, 6, 20, 32, 34]
	IP, copyright & authorship gaps	Unclear ownership of AI-assisted outputs; potential for unintentional IP infringement because of training data	Moderate (emerging)	[5, 13, 23, 25, 28, 29]
	Occupational stress & psychological harm	AI-induced work intensification, job insecurity, and constant evaluation amplify stress and reduce wellbeing	Strong (empirical synthesis)	[6, 11, 17, 43, 48]
	Algorithmic opacity	AI decision logic is illegible to workers and managers; black-box outputs impede meaningful oversight	Strong (theoretical + empirical)	[4, 5, 7]
	Accountability diffusion	Responsibility for AI-assisted decisions becomes unclear	Strong (regulatory + qualitative)	[5, 12, 13, 28]
	Professional identity shifts & power dynamics	Digital undertow erodes traditional expertise markers; reshapes seniority hierarchies in creative work	Moderate (theoretical + qualitative)	[19, 26, 38, 39]
T3 - Mitigation Strategies	Meaningful human oversight	Genuine capacity to review, query, and override AI outputs; requires time, competence, and authority	Strong (regulatory + empirical)	[7, 13, 14, 31]
	IP and ethics governance	Pre-deployment IP audit; contractual clarity on AI-generated content ownership; alignment with EU AI Act obligations	Strong (practice)	[13, 25, 27–29]
	Lifecycle governance	Structured accountability mechanisms across the full AI system lifecycle from design to decommissioning	Moderate (practice)	[7, 22, 27]
	Transparency & explainability	AI tools communicate decision logic to workers, who can then interrogate and understand the rationale behind the given output	Strong (empirical + practice)	[4, 7, 14]
	Participatory design & worker involvement	Affected workers involved in deployment decisions; reduces resistance	Moderate (empirical)	[12, 17, 19, 44]
	Co-skilling & AI-expertise integration & impact on monitoring	Reskilling programs coupling domain expertise deepening with algorithmic literacy; co-skilling to reduce job insecurity	Moderate-Strong (empirical)	[21, 41, 45, 46]
T4 - Organizational Readiness	Technical infrastructure & data quality	Compatible systems, quality data, and computational resources; data readiness is the most consistently underestimated factor	Strong (empirical)	[8, 9, 18]
	Workforce AI literacy	The ability to critically evaluate AI outputs, maintain independent judgment and recognize system limitations	Strong (empirical + practice)	[9, 18, 21, 41]
	Leadership & strategic alignment	Leaders establish governance expectations, set measurable objectives and model accountability	Strong (empirical + review)	[8, 16, 18, 37]
	Reskilling strategy & HRD alignment	Pre-adoption skill-gap analysis; investment in broad expertise+literacy programs; HRD adaptation to AI	Strong (empirical + practice)	[21, 41, 42, 46]
	Regulatory & legal readiness	Legal expertise, vendor validation, compliance infrastructure for EU AI Act and sector-specific obligations	Strong (practice)	[7, 13, 14, 27]

Evidence strength of table 2 was evaluated as following: - strong evidences are supported by large studies (with large evidence pool), show consistent and rigorous results, are supported by many papers - moderate evidences present a strong theoretical explanation but are supported by smaller scale studies (with smaller sample sizes) and do not have longitudinal support - emerging evidences are backed by recent studies and refer to newer areas of study They may still be object to debates.

References

- [1] Brynjolfsson, E.; Li, D.; Raymond, L.R. (2025). Generative AI at Work. *The quarterly journal of economics*, Oxford academy. DOI: 10.3386/w31161.
- [2] Kanazawa, K.; Kawaguchi, D.; Shigeoka, H.; Watanabe, Y. (2026). AI, Skill, and Productivity: The Case of Taxi Drivers. *Management Science*. DOI: 10.1287/mnsc.2023.01631.
- [3] Frank, M.R.; Autor, D.; Bessen, J.E.; Brynjolfsson, E. et al. (2019). Toward understanding the impact of artificial intelligence on labor. DOI: 10.1073/pnas.1900949116.
- [4] Burrell, J. (2016). How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big Data & Society*, 3(1). DOI: 10.1177/2053951715622512.
- [5] Munn, L. (2023). The uselessness of AI ethics. *AI and Ethics*. DOI: 10.1007/s43681-022-00209-w.
- [6] Battal, S. (2025). Impact of artificial intelligence usage on occupational stress: an exploratory study of risk factors and prevention strategies. *AI & Society*. DOI: 10.1007/s00146-025-02385-9.
- [7] NIST (2023). *AI Risk Management Framework (AI RMF 1.0)*. DOI: 10.6028/NIST.AI.100-1.
- [8] McKinsey & Company (2023). *The State of AI in 2023: Generative AI's Breakout Year*. McKinsey Global Institute.
- [9] Lane, M.; Saint-Martin, A. (2021). *The Impact of Artificial Intelligence on the Labour Market: What do we know so far?* OECD Social, Employment and Migration Working Papers, No. 256. DOI: 10.1787/7c895724-en.
- [10] OECD (2019). *Artificial Intelligence in Society*. OECD Publishing. DOI: <https://doi.org/10.1787/eedfee77-en>.
- [11] EU-OSHA (2024). *Impact of Artificial Intelligence on Occupational Safety and Health*. European Agency for Safety and Health at Work.
- [12] Lee, H.R. et al. (2024). Contrasting perspectives of workers: exploring labor relations in workplace automation and potential intervention. *CHI '24*, ACM. DOI: 10.1145/3613904.3642907.
- [13] EU (2024). *Regulation (EU) 2024/1689 – Artificial Intelligence Act*. Official Journal of the EU, 13 June 2024.
- [14] OECD (2023). *Ensuring Trustworthy AI in the Workplace*. In *OECD Employment Outlook 2023*. DOI: <https://doi.org/10.1787/08785bba-en>.
- [15] Brynjolfsson, E.; Mitchell, T. (2017). What can machine learning do? Workforce implications. *Science*, 358(6370). DOI: 10.1126/science.aap8062.
- [16] Zhang, M.M.; Cooke, F.L.; Ahlstrom, D.; McNeil, N. (2025). The rise of algorithmic management and implications for work and organizations. *New Technology, Work and Employment*. DOI: 10.1111/ntwe.12343.
- [17] PEROSH Network (2025). *Algorithmic management and psychosocial risks at work: an emerging occupational safety and health challenge*. *Scandinavian Journal of Work, Environment & Health*. DOI: 10.5271/sjweh.4270.
- [18] Bunjak, A.; Bruch, H.; Cerne, M. (2022). Context is key: The joint roles of transformational and shared leadership and management innovation in predicting employee IT innovation adaptation. *International Journal of Information Management*. DOI: 10.1016/j.ijinfomgt.2022.102516.
- [19] Li, J.; Cao, H.; Lin, L.; Hou, Y.; Zhu, R.; El Ali, A. (2024). *User Experience Design Professionals' Perceptions of Generative Artificial Intelligence*. *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. DOI: 10.1145/3613904.3642114.
- [20] Shukla, P.; Bui, P.; Levy, S.S.; Kowalski, M.; Baigelenov, A.; Parsons, P. (2025). De-skilling, cognitive offloading, and misplaced responsibilities: potential ironies of AI-assisted design. *CHI EA '25*, ACM. DOI: 10.1145/3706599.3719931.
- [21] Tambe, P.B. (2026). *Reskilling the Workforce for AI: Domain Expertise and Algorithmic Literacy*. *Management Science*. DOI: 10.1287/mnsc.2022.03968
- [22] Mahajan, P. (2025). *The Soul of the AI: Governance, Ethics, and the Future of Human-AI Integration*. DOI: 10.5281/zenodo.15789678.
- [23] Al-Khairy, M.; Mustafa, D.; Kshetri, N.; Insiew, M.; Alfandi, O. (2024). *Ethical Challenges and Solutions of Generative AI: An Interdisciplinary Perspective*. *Informatics*. DOI: 10.3390/informatics11030058.
- [24] Huang, Y.; Arora, C.; Huang, W.C.; Kanij, T. et al. (2026). *Ethical concerns of generative AI and mitigation strategies: A systematic mapping study*. *Applied Soft Computing*. DOI: 10.1016/j.asoc.2026.114789.
- [25] Chesterman, S. (2025). *Good models borrow, great models steal: intellectual property rights and generative AI*. *Policy and Society*. DOI: 10.1093/polsoc/puae006.
- [26] Xu, C.; Yao, Z.; Huo, W. (2025). *Every coin has two sides: The application of artificial intelligence on employees' unethical behaviours*. *Knowledge Management Research & Practice*. DOI: 10.1080/14778238.2024.2399021.
- [27] Leon, M. (2026). *Lifecycle-Based Governance to Build Reliable Ethical AI Systems*. *Systems Research and Behavioral Science*. DOI: 10.1002/sres.70014.
- [28] Novelli, C.; Casolari, F.; Hacker, P.; Spedicato, G.; Floridi, L. (2024). *Generative AI in EU law: Liability, privacy, intellectual property, and cybersecurity*. *Computer Law & Security Review*. DOI: 10.1016/j.clsr.2024.106066.
- [29] Lucchi, N. (2023). *ChatGPT: A Case Study on Copyright Challenges for Generative AI Systems*. *European Journal of Risk Regulation*. DOI: 10.1017/err.2023.59.
- [30] Kanbach, D.K.; Heiduk, L.; Blueher, G.; Schreiter, M. (2024). *Generative AI from a business model innovation perspective*. *Review of Managerial Science*. DOI: 10.1007/s11846-023-00696-z.
- [31] Dobre, J. (2025). *Designing AI for human expertise: Preventing cognitive shortcuts*. *UX Matters*.
- [32] Hou, J.; Wang, L.; Wang, G.; Wang, H.J.; Yang, S. (2025). *The Double-Edged Roles of Generative AI in the Creative Process: Experiments on Design Work*. *Information Systems Research*. DOI: 10.1287/isre.2024.0937.
- [33] Dell'Acqua, F.; McFowland III, E.; Mollick, E.; Lifshitz, H. et al. (2026). *Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality*. *Organization Science*. DOI: 10.1287/orsc.2025.21838.
- [34] Izak, M.; Barros, A.; Prasad, A.; Šliwa, M. (2025). *Generative artificial intelligence and learning: At the dawn of Idiocracy?* *Management Learning*. DOI: 10.1177/13505076251348575.
- [35] Bilgram, V.; Laarmann, F. (2023). *Accelerating innovation with generative AI: AI-augmented digital prototyping*. *IEEE Engineering Management Review*. DOI: 10.1109/EMR.2023.3272799.
- [36] Doshi, A.R.; Hauser, O.P. (2024). *Generative AI enhances individual creativity but reduces the collective diversity of novel content*. *Science Advances*. DOI: 10.1126/sciadv.adn5290.
- [37] Mehta, S.; Sharma, N. (2025). *Theory Z+1: Redefining Managerial Trust and Human-AI Synergy in the Hybrid Work Era*. SSRN. DOI: 10.2139/ssrn.5826062.
- [38] Hsu, G.; Bechky, B.A. (2024). *Exploring the digital undertow: how generative AI impacts social categorizations in creative work*. *Organization Theory*, SAGE. DOI: 10.1177/26317877241275118.
- [39] Süner-Pla-Cerdà, S.; Şen, G.; Kumbasar, E.; Şahin, B.; Ünal, C.E. (2025). *Designer experiences and perspectives on the role of generative AI in industrial design*. *AI & Society*. DOI: 10.1007/s00146-025-02504-6.
- [40] Sharma, S. (2024). *Benefits or concerns of AI: A multistakeholder responsibility*. *Futures*. DOI: 10.1016/j.futures.2024.103328.
- [41] Morandini, S.; Fraboni, F.; De Angelis, M.; Puzzo, G.; Giusino, D. (2023). *The Impact of AI on Workers' Skills: Upskilling and Reskilling in Organisations*. *Interdisciplinary Description of Complex Systems*. DOI: 10.28945/5078.
- [42] Bennett, C.; Bennett, J. (2026). *Human resource development in the age of artificial intelligence: a theoretical synthesis*. *AI & Ethics*. DOI: 10.1007/s43681-025-00944-w.
- [43] Salama, W.M.E.; Khairy, H.A.; Mansour, A.A.; Alrefae, W.M.; Afify, H. (2025). *How AI awareness drives employee intention to leave in hotels: the mediating roles of job burnout and psychological contract breach*. *Geo Journal of tourism and geosites*. DOI: 10.30892/gtg.602spl18-1494.
- [44] Nkanta, E. (2025). *"Don't Panic, it's Just AI": A Qualitative Study on Internal Communication Strategies for Managing AI-Driven Change*. DOI: 10.30707/ETD.1763755358.227483.
- [45] Zhang, Y.; Liu, X.; Yan, Q.; Meng, N. (2025). *Empowering workforces in AI-driven environments: co-skilling, organizational support, and mitigating job insecurity*. DOI: 10.3389/fpsyg.2025.1700129.
- [46] Appiah, R.I. (2025). *Corporate Strategies for Successful Workforce Upskilling and Reskilling in Response to AI Adoption-What works, what does not and why*. DOI: <https://doi.org/10.63282/3050-922X.IJERET-V6H1P112>
- [47] Hopkins, J.; Gallagher, S. (2024). *Generative AI at work: Empowering employee mental wellbeing*. DOI: 10.25916/sut.27252948.v1.
- [48] Wu, S.; Liu, Y.; Ruan, M.; Chen, S.; Xie, X.Y. (2025). *Human-generative AI collaboration enhances task performance but undermines intrinsic motivation*. *Scientific Reports (Nature)*. DOI: 10.1038/s41598-025-98385-2.

AI USE DISCLOSURE

Tools used: Claude (Anthropic) was used to assist with search query formulation, initial thematic organization of coded factors, and draft editing of prose sections.

How used: The tool was used to suggest search terms aligned with each analytical dimension and to propose initial thematic groupings during coding. All substantive analytical judgments — inclusion/exclusion decisions, code assignments, theme labels, framework structure, and factual claims — were made and verified by team members independently of the tool. **Source verification:** Every source cited in this paper was located by a team member, accessed in full text where available, and read before citation. No source was cited based solely on an AI-generated description. **Confirmation:** All 48 cited sources have been verified by team members. The team confirms no AI tool was used to fabricate, invent, or hallucinate any citation.

REFERENCE MATERIAL

All the reference files (readme.txt for clarifications, LaTeX source code and included and excluded papers table) can be found at the following Google Drive link: <https://drive.google.com/drive/folders/1iHQJPX5t1ks-zGB-DTf1ftOBpmKLe8SZ>