

Practical Criteria for AI Task Delegation: A Scoping Review of Task Complexity, Performance Reliability, and Regulatory Compliance

Sila Gulec
s334743@studenti.polito.it
Politecnico di Torino
Italy

Nicolò Fiori
s325013@studenti.polito.it
Politecnico di Torino
Italy

Behrad Galedari
s324382@studenti.polito.it
Politecnico di Torino
Italy

Ahmet Tolga Birand
s335068@studenti.polito.it
Politecnico di Torino
Italy

Benedetta Stellato
s341071@studenti.polito.it
Politecnico di Torino
Italy

Kutay Gunel
s307467@studenti.polito.it
Politecnico di Torino
Italy

Giacomo Simbula
s341576@studenti.polito.it
Politecnico di Torino
Italy

Pietro de Francesco
s340079@studenti.polito.it
Politecnico di Torino
Italy

Jacopo Caravaggio
s339988@studenti.polito.it
Politecnico di Torino
Italy

Gabriele Gallo
s339929@studenti.polito.it
Politecnico di Torino
Italy

Abstract

This scoping review identifies the practical criteria that determine whether workers choose to delegate tasks to AI. Drawing on 39 empirical, theoretical, and policy sources retrieved between 25 March and 25 April 2026, we synthesize the literature into five interrelated themes: Characteristics and Properties of Task, AI Performance, Cost and Efficiency, Accountability, and Human–AI Collaboration. We apply the resulting framework to the worked use case of a computer programmer using LLM-powered coding assistants such as GitHub Copilot, and conclude that full delegation is unwarranted, while highly conditional partial delegation is feasible. We close by identifying four gaps for future research: limited job focus, lack of long-term studies, weak empirical evidence on the worker-level impact of regulation, and systematic neglect of combined effects between factors.

1 Introduction

As tasks are increasingly delegated to AI systems, the question is *when and how far to* rather than *whether to* delegate. Delegation is governed by a complex interplay of factors that determine when AI assistance is beneficial and when it introduces unacceptable risk. This scoping review addresses the research question: *Which practical criteria, including task complexity, performance reliability, and regulatory compliance, determine whether workers choose to delegate tasks to AI?*

The three criteria stated here do not operate in isolation because task complexity shapes verification cost, and performance reliability shapes the worker–AI dynamic. Our framework expands them into

five interrelated operational themes that capture how the criteria interact in practice:

Characteristics & Properties of Task: complexity, structure, and verifiability;

Cost & Efficiency: whether time saved outweighs verification costs;

AI Performance: accuracy, transparency, and error visibility;

Accountability: legal and ethical liability for AI-generated errors; and

Human–AI Collaboration: trust, AI literacy, and the capacity to modify outputs.

The resulting framework provides a domain-agnostic lens for evaluating the appropriate boundaries of AI delegation.

2 Methodology

This study uses a scoping review approach that adheres to the Preferred Reporting Items for Systematic Reviews and Meta-Analyses Extension for Scoping Reviews (PRISMA-ScR) framework. A scoping review was chosen over a systematic review because the goal is to map the range of practical factors influencing workers' decisions to delegate tasks to AI, rather than to assess the effectiveness of a specific intervention.

2.1 Search Strategy

The initial searches were conducted between 25 March and 9 April 2026, yielding an initial corpus of 24 sources. A secondary search was conducted between 16 April and 25 April 2026, expanding the corpus to the final 39 sources. The searches were conducted across two electronic databases: Google Scholar and the ACM Digital

Library. Google Scholar was selected for its broad multidisciplinary coverage, while the ACM Digital Library was chosen for its depth in human–computer interaction research.

Boolean search strings were constructed around the following six concept areas: AI Delegation, Characteristics & Properties of Task, AI Performance, Cost & Efficiency, Accountability, and Human–AI Collaboration. For the exact and complete list of all search strings, refer to Appendix A.

2.2 Inclusion and Exclusion Criteria

Sources were evaluated using predefined eligibility criteria. To be considered, a source had to be published between 2015 and April 2026, written in English, address at least one practical factor influencing AI task delegation, and be a peer-reviewed article, conference paper, technical standard, policy document, or major industry report. Sources that focused solely on AI model development without any mention of delegation, as well as blog posts and opinion editorials, were eliminated. An exception to the date criterion was made for source [16], which was retained as a foundational source because it established a theoretical basis for human–AI task allocation that remains widely cited in subsequent work. Refer to Appendix B for the full criteria table.

2.3 Screening Process

Screening followed three stages that are summarized in the complete PRISMA-ScR flow diagram (refer to Appendix C).

Stage 1 – Identification. The combined search strings across both databases yielded 782 sources.

Stage 2 – Title and duplicate screening. 142 duplicates, 68 outdated sources, and 437 clearly off-topic results were removed, yielding 135 sources.

Stage 3 – Abstract and full-text screening. Abstracts of the 135 remaining papers were reviewed for substantive engagement, eliminating 86. The remaining 49 papers then underwent full-text assessment for depth relevance to the research question. A further 10 sources were removed, resulting in a final corpus of 39 publications: empirical studies ($n = 23$), theoretical papers ($n = 7$), policy or standards documents ($n = 8$), and foundational works ($n = 1$). At each stage, team members independently assessed eligibility.

2.4 Data Extraction and Synthesis

Each included source was coded using a standardized extraction form that captured the author(s), year, study type, and the specific delegation factors covered. Thematic analysis was then applied to identify recurring patterns across the extracted data. Through constant comparison across sources, codes were organized into five main themes that together shape the workers’ AI task delegation decisions. The extraction form and full coded dataset are available in the accompanying CSV file (see Appendix D for the ZIP file link).

3 Results

The review included 39 sources that explored practical criteria for AI task delegation in the workplace. The corpus consists of 23 empirical studies, 7 theoretical and conceptual works, 8 policy, standards, or industry documents, and 1 foundational source. Empirical

methodologies included controlled experiments, user studies, large-scale surveys, and simulation-based analyses. The fields covered include software engineering, general knowledge work, international AI governance, and multidisciplinary workforce analysis. The majority of sources were published between 2022 and 2026, demonstrating the rapid growth of this research area following the broad adoption of LLM-based tools.

From the analysis, we developed a framework of five interrelated themes that collectively determine whether, and under what conditions, workers choose to delegate tasks to AI. The following subsections describe each theme and summarize the supporting evidence. Table 1 provides a unified summary of all themes and sources. We next apply the whole framework to a specific use case to demonstrate how the themes interact in practice.

4 Framework

4.1 Characteristics & Properties of Task

The “Characteristics and Properties of Task” theme underscores that AI integration is not a “one-size-fits-all” endeavor; the specific nature of the work dictates the success of the collaboration. A primary consideration is *task delegability*, which in knowledge work environments is often determined by criteria such as the clarity of data inputs and the inherent structuredness of the workflow [10]. When these properties align with AI capabilities, users are more likely to offload work, yet this act of delegation modulates both the objective quality of the final performance and the human’s subjective task satisfaction [17]. Task familiarity further shapes delegation mode: workers switch between an *acceleration mode* for routine tasks and an *exploration mode* for unfamiliar ones, adjusting their verification effort accordingly [3].

Task complexity acts as a critical gatekeeper for technology adoption. In service-oriented roles, the technical and social intricacies of a problem-solving scenario can either drive or deter a user’s intention to engage with AI [39]. Interestingly, research suggests that AI-related task complexity does not just impact efficiency; under manageable cognitive loads it can stimulate innovative work behavior, while excessive difficulty causes psychological paralysis that discourages delegation [7].

Certain task properties introduce significant psychological risks. Under high difficulty or extreme time pressure, workers are more susceptible to automation bias, accepting AI outputs uncritically [16]. Recent evidence further shows that time pressure and complexity interact non-linearly, producing a multiplier effect on uncritical delegation rather than an additive one [32]. This vulnerability can be mitigated at the interface level: cognitive forcing functions that require workers to make their own judgement before viewing AI output meaningfully reduce overreliance under cognitive load [6]. In specialized technical domains, the property of *verifiability* is paramount: when verification effort is high, blind trust quickly degrades into costly errors. These diverse characteristics suggest that AI must be calibrated to the unique demands and pressures of the specific task, rather than deployed uniformly across roles.

Table 1: Summary table of the themes.

Theme Name	Theme Definition	All Used Sources
T1 – Characteristics & Properties of Task	Examines how task factors such as complexity, familiarity, verifiability, and time constraints affect a worker’s desire and ability to properly transfer work to AI.	[3, 6, 7, 10, 16, 17, 32, 39]
T2 – AI Performance	Investigates how actual AI reliability, user expectations, calibration gaps, and explanation restrictions influence the dynamic evolution of human trust and algorithm appreciation.	[2, 11, 15, 16, 23, 37]
T3 – Cost & Efficiency	Examines the economic and time trade-offs of AI delegation, emphasizing the conflict between quick output generation, verification costs, and the potential of long-term skill degradation.	[6, 12, 25, 27, 32, 33, 35]
T4 – Accountability	Highlights the legal, regulatory (e.g., EU AI Act, OECD), and practical issues of determining culpability and maintaining effective human control for AI-driven results.	[4, 9, 13, 22, 26, 28, 29]
T5 – Human–AI Collaboration	Investigates the interdependent dynamics of machine-in-the-loop workflows, as influenced by user literacy levels, multidimensional trust, and the worker’s ability to change and regulate AI outputs.	[8, 11–13, 15, 18, 20, 24, 31]

4.2 AI Performance

This theme focuses on how the actual and perceived performance of AI outputs jointly shape delegation decisions, since workers act not on AI’s true reliability but on what they believe it to be.

Trust as a dynamic, experience-based phenomenon. AI trust is not a fixed, unchangeable quality but rather a dynamically evolving experience: mistakes weaken trust, whereas consistent performance strengthens it, and the trend changes depending on the kind of AI involved [15]. Employees’ expectations influence this process; people tend to expect perfection from machines, and thus AI errors disappoint more than comparable human errors. Initial automation bias turns into algorithm aversion, and employees become much less willing to hand over tasks to algorithms [16]. On the other hand, large-scale empirical research indicates a different side of the story: lay users often prefer algorithms to human judgement even before any failure occurs, with the direction of bias moderated by user expertise and task type [23]. Consequently, the same person might trust an AI system on Monday but refuse to use it on Friday, depending on their current experiences.

The calibration gap. A central finding across the corpus is that the perception of AI reliability does not always match reality. In laboratory experiments, employees working with AI support generated less accurate outcomes than their peers working without such assistance, even though they felt confident about their AI-supported solutions [11]. Critically, this miscalibration occurs even when AI provides explanations, as human–AI collaborations hardly ever surpass AI’s performance on its own—implying that explanation interfaces are no guarantee against the trust–reliability mismatch [2].

Implication for delegation. Together these findings indicate that AI performance operates through a trust-mediated feedback loop in which actual reliability, perceived reliability, and the worker’s

ability to detect errors jointly determine whether delegation succeeds. The bias–aversion cycle and the calibration gap are well documented in safety-critical and high-precision domains, but coverage in service, creative, and managerial work remains thin and is a key target for future research.

4.3 Cost & Efficiency

This theme examines the practical economic and temporal drivers that govern the delegation of tasks to AI, operating simultaneously at the micro level (individual worker choices) and the macro level (organizational productivity).

Performance–cost considerations. At the individual level, delegation represents a constant trade-off between cognitive savings and verification effort, with performance–cost considerations as the primary driver of delegation decisions [14]. Workers are naturally inclined to delegate complex tasks to AI to conserve mental energy and manage their cognitive load [27, 35]. However, this behavioral shift is economically sustainable only if the cost of verifying the AI’s output remains substantially lower than the effort required to perform the task manually [35]. If reviewing and correcting algorithmic errors requires excessive time or expertise, the expected efficiency gains are negated, and workers revert to manual execution [37]. Industry-side observation of AI usage at scale confirms this pattern: the bulk of occupational AI use is augmentation rather than full automation, suggesting workers continually negotiate where the verification trade-off becomes economical [1].

Temporal context. The temporal context in which delegation occurs significantly affects its effectiveness. When workers operate under severe time constraints, the rigorous checks required for safe delegation are often the first to be overlooked, leading to an uncritical acceptance of flawed outputs [32]. Although interventions such as “cognitive forcing functions”—which introduce artificial friction

to prompt the user to reflect—can effectively reduce this overreliance [6], they inherently conflict with the primary efficiency goals of delegation by intentionally slowing down the decision-making process [6, 32].

Organizational level. At the organizational level, generative AI has the potential to act as a massive multiplier of human capital, particularly in high-paying cognitive labor sectors [12, 25, 38]. Real-world field evidence corroborates this potential: deployment of a generative AI assistant raised customer-support productivity by 14% on average, with the largest gains for less-experienced workers [5]. Global labour analyses further suggest that generative AI is more likely to augment than replace most occupations [19]. However, maximizing this potential requires navigating a hidden long-term cost: the vicious cycle of skill erosion [33]. As organizations aim for maximum short-term efficiency through widespread delegation, workers may gradually lose the fundamental skills needed to effectively verify AI outputs [33]. As a result, initial productivity gains can lead to systemic inefficiencies, leaving the workforce fundamentally ill-equipped to manage and oversee the algorithmic systems on which it relies [33]. Similar synergistic effects have been observed across other technical domains [34].

4.4 Accountability

The “Accountability” theme addresses the regulatory, legal, ethical, and behavioral structures that shape whether, and how far, tasks can be delegated to AI. It focuses on who is held accountable when AI-delegated outputs are inaccurate, as well as how this responsibility distribution restricts delegation in practice.

A major challenge is the development of *accountability gaps*, in which no single actor—developer, deployer, or end-user—is clearly liable for an AI-driven error [22]. For example, when a worker accepts an AI-generated recommendation that later turns out to be incorrect, it is sometimes unclear who is responsible. This ambiguity does more than just increase legal risk; it actually discourages delegation in high-stakes situations. Individual liability perceptions are reconfigured by regulatory frameworks such as the EU AI Act, which directly alters workers’ daily delegation choices. Loss-aversion theory predicts that when personal liability for AI errors becomes prominent, workers grow reluctant to delegate even to capable systems [4, 9]; however, direct workplace evidence remains insufficient (see Section 6).

Accountability is thus both a structural feature of governance systems and a psychological factor that influences individual delegation decisions.

International regulators are increasingly defining the scope of permitted delegation. The EU AI Act classifies AI systems used in employment, worker management, and access to self-employment as high-risk, requiring human oversight, transparency, and systematic audits [13]. The OECD AI Principles similarly mandate accountability, traceability, and responsible management throughout the AI lifecycle [28]. At the organizational level, the NIST AI Risk Management Framework operationalizes these obligations through four governance functions—Govern, Map, Measure, and Manage—which specify how trustworthiness attributes such as validity, dependability, fairness, and accountability should be maintained for delegated outputs [26].

High-level principles are insufficient. The OECD’s workplace-specific research on AI deployment argues that successful oversight necessitates not just legal compliance, but also employer-provided AI training, dedicated verification time, and worker consultation channels [29]. This bridges the gap between abstract policy and the practical conditions under which delegation decisions are made—conditions such as whether workers understand what they are responsible for, whether they have been trained to evaluate AI outputs, and whether organizational processes allow enough time for meaningful reviews.

These converging regulatory, organizational, and behavioral forces point to a fundamental constraint. Computer systems can execute complicated tasks but cannot be held morally or legally accountable. Responsibility stays with human actors, and delegation is thus limited not just by technological skill, but also by the clarity of liability structures.

4.5 Human–AI Collaboration

The Human–AI relationship is fundamentally rooted in collaboration, transitioning from a binary choice between human and machine to an interdependent “automation–augmentation paradox” [12]. In this context, a key requirement for delegation is the maintenance of a “machine-in-the-loop” model, in which the human retains primary control and intentional accountability, while the AI functions as a supportive tool [12, 24]. Trust—specifically in the areas of performance, process, and purpose—sustains this collaborative dynamic. Trust is a psychological aspect for the worker’s rational assessment of uncertainty during task execution [24]. Thus, delegation works best when people do not see it as replacing tasks but as a coevolutionary process in which people provide the intentionality and objective-orientedness that machines lack [12].

Beyond structural paradigms, the practical decision to delegate is heavily influenced by the worker’s cognitive background and the system’s capacity for user intervention. Research indicates that the relationship between AI experience and reliance follows a non-linear “automation bias curve,” where individuals with moderate, surface-level knowledge are more likely to over-rely on AI agents, whereas those with deeper technical backgrounds show more calibrated, appropriate use [18]. This finding suggests that AI literacy is itself a delegation criterion: workers in the middle of the literacy spectrum represent the highest-risk group, because they have enough familiarity to trust the system but not enough expertise to detect its errors.

To bridge the gap between algorithm aversion and functional delegation, a primary practical criterion is “modifiability”: the technical affordance that allows users to adjust AI outputs [20]. Modifying results satisfies the human need for control and agency, significantly increasing the likelihood that workers will choose to collaborate with imperfect algorithmic systems instead of rejecting them entirely [20]. Theoretical work further confirms that current AI systems remain bounded in scope, particularly for tasks requiring contextual judgment or creative problem-solving, reinforcing the complementary rather than substitutive role of AI in collaborative workflows [8].

5 Worked Use Case: Computer Programmer

5.1 Scenario

Consider a junior developer who is assigned two tasks in the same week: first, writing a simple function that sorts a list of customer names alphabetically, and second, writing the code for a login page that stores and checks user data. She uses GitHub Copilot, a coding assistant powered by an LLM that offers code as she types. The framework predicts different delegation decisions for these two tasks.

5.2 Applying the Framework

Characteristics & Properties of Task. The sorting function is common and structured; the developer has written similar code many times and can verify instantly whether the output list is in correct order. She accepts Copilot’s suggestion in “acceleration mode” with a quick glance. The login data storage code, by contrast, involves encryption rules that she has never used before, forcing her into “exploration mode” where she needs to carefully study each generated line, which can take longer.

AI Performance. A mistake in the sorting function is visible—the names are either in order or they are not. A flaw in password storage is not. Copilot may generate a login page that works perfectly on the surface, so users can login, but stores passwords in a format that hackers can read or access.

Cost & Efficiency. Copilot saves time on the sorting function and it takes little time to examine the output, thus defining a benefit. For the login page the initial generation is equally fast, but it can take hours to thoroughly examine the password logic. The developer is less likely to find security problems under a strict schedule. So, the saved time could end up costing far more if a vulnerability reaches real users.

Accountability. A sorting bug is harmless—someone notices names are in the wrong order. On the other hand, every user’s credentials can be leaked with a bug in the login page. Both the EU AI Act and the NIST framework mandate that code handling personal data undergo documented human review prior to deployment. In the event of a breach, the developer and her company have complete legal responsibility, not Copilot. Therefore, whereas the sorting task does not require a formal review gate, the password task must.

Human–AI Collaboration. As a junior developer, she sits in one of the riskiest zones of the automation bias curve—familiar enough with Copilot to trust it, but lacking the security expertise to catch its mistakes in password handling. She can update risky areas since Copilot generates editable output. Yet only if her team views Copilot output as a first draft that needs human review, not a final result.

5.3 Judgement

The framework gives a split decision. For routine, easily checked tasks like sorting a list, delegation brings a clear efficiency gain with minimal risk. For the user data handling code, three themes converge against delegation. The developer’s experience level prevents security problems from being discovered (Human–AI Collaboration), the review time exceeds the generation benefit (Cost &

Efficiency), and the human bears full responsibility for any data leaks (Accountability).

Therefore, full delegation is unjustified; the acceptable limit is extremely conditional partial delegation, which is restricted to straightforward, easily verifiable subtasks and dependent on the programmer’s capacity to audit outputs. More specifically, the framework results in a sub-task-tiered recommendation:

- (i) full delegation is acceptable for routine tasks (sorting, formatting, autocompletion);
- (ii) partial delegation with mandatory review is appropriate for ordinary business logic; and
- (iii) delegation is actively discouraged for security-critical or unfamiliar logic such as authentication, cryptography, or concurrency.

6 Future Gaps & Works

This review identified five main themes guiding AI delegation, but significant gaps remain for future investigation.

Limited Occupational and Geographical Scope. Evidence heavily relies on technical domains like code generation [3, 30, 36], underrepresenting service or managerial contexts [39] and lacking cross-cultural diversity. *Future Work:* Researchers should conduct cross-sectoral and regional studies to determine if task complexity and algorithmic trust manifest differently across diverse professional and cultural environments.

Scarcity of Longitudinal Research. Current literature predominantly uses short-term experiments [11, 21, 32], leaving the long-term risks of cognitive automation and skill erosion [33] largely unexplored. *Future Work:* Studies must employ longitudinal designs to measure how sustained, daily AI reliance impacts actual skill retention, trust calibration, and evolving AI literacy over time.

Disconnect Between Policy and Worker Impact. While major frameworks [13, 26, 28] are well-documented, empirical data on how these top-down regulations affect an individual’s daily choices is scarce. *Future Work:* Research should utilize observational methods to investigate whether explicit awareness of legal liability practically deters workers from delegating high-stakes tasks in real-world scenarios.

Systematic Neglect of Interacting Factors. Existing studies typically isolate variables like complexity [7], trust [15], or cost [14]. However, a worker simultaneously navigating time pressure, complexity, and regulations will not behave according to the simple sum of isolated experiments. *Future Work:* Research must adopt multi-variable methodologies to test how these practical criteria interact and collectively influence delegation boundaries.

7 Conclusion

This scoping review identified the practical criteria that determine how and when workers delegate tasks to AI. Through the synthesis of 39 sources, we developed a framework comprising five interrelated themes: Characteristics and Properties of Task, AI Performance, Cost and Efficiency, Accountability, and Human–AI Collaboration. Applying this framework to a computer programmer utilizing an LLM-powered coding assistant demonstrates that the

practical limits of AI delegation remain strictly context-dependent. Currently, the threshold for safe delegation is restricted to partial, supervised collaboration, which is fundamentally shaped by the user's ability to accurately review generated results and manage the associated verification costs.

Verification cost emerges as a limiting threshold: when the cognitive effort of reviewing AI output approaches the effort of performing the task directly, delegation loses its practical justification. Ultimately, while AI systems bring significant efficiency benefits, human actors remain the ultimate carriers of legal and ethical responsibility. To fully understand the evolving dynamics of augmentation in the workspace, future work should examine how delegation thresholds shift as AI literacy develops over time, and how sector-specific regulations reshape the boundaries of acceptable delegation.

References

- [1] Anthropic. 2025. *Anthropic Economic Index Report — Measuring AI's Effects on the Economy*. Technical Report. Anthropic, San Francisco, CA. <https://www.anthropic.com/research/anthropic-economic-index>
- [2] Gagan Bansal, Tongshuang Wu, Joyce Zhou, Raymond Fok, Besmira Nushi, Ece Kamar, Marco Tulio Ribeiro, and Daniel Weld. 2021. Does the Whole Exceed Its Parts? The Effect of AI Explanations on Complementary Team Performance. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Article 81, 16 pages. <https://doi.org/10.1145/3411764.3445717>
- [3] Shraddha Barke, Michael B. James, and Nadia Polikarpova. 2023. Grounded Copilot: How Programmers Interact with Code-Generating Models. *Proceedings of the ACM on Programming Languages* 7, OOPSLA1 (April 2023), 1–27. <https://doi.org/10.1145/3586030>
- [4] Jesse Bockstedt and Tyler Buckman. 2025. Loss Aversion and AI Delegation. *Information Systems Research* 36, 2 (June 2025), 512–534. <https://doi.org/10.1287/mnsc.2024.05585>
- [5] Erik Brynjolfsson, Danielle Li, and Lindsey Raymond. 2025. Generative AI at Work. *Quarterly Journal of Economics* 140, 2 (May 2025), 889–942. <https://doi.org/10.1093/qje/qjae044>
- [6] Zana Bućinca, Maja Barbara Malaya, and Krzysztof Z. Gajos. 2021. To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI in AI-Assisted Decision-Making. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1, Article 188 (April 2021), 21 pages. <https://doi.org/10.1145/3449287>
- [7] Hongyi Cai, Yuhui Ge, and Heng Zhao. 2025. How AI-Related Task Complexity Shapes Innovative Work Behavior: A Coping Theory Perspective. *Behavioral Sciences* 15, 11 (Oct. 2025), 1467. <https://doi.org/10.3390/bs15111467>
- [8] Jean-François Caron. 2025. *Artificial Intelligence and Life in the Fourth Industrial Revolution: Summoning the Demon?* Routledge. <https://doi.org/10.4324/9781003546283>
- [9] Paulius Čerka, Jurgita Grigienė, and Gintarė Sirbikytė. 2015. Liability for Damages Caused by Artificial Intelligence. *Computer Law & Security Review* 31, 3 (2015), 376–389. <https://doi.org/10.1016/j.clsr.2015.03.008>
- [10] Izabel Cvetkovic and Eva Bittner. 2022. Task Delegability to AI: Evaluation of a Framework in a Knowledge Work Context. In *Proceedings of the 55th Hawaii International Conference on System Sciences*. <https://doi.org/10.24251/hicss.2022.021>
- [11] Berkeley J. Dietvorst, Joseph P. Simmons, and Cade Massey. 2018. Overcoming Algorithm Aversion: People Will Use Imperfect Algorithms If They Can (Even Slightly) Modify Them. *Management Science* 64, 3 (March 2018), 1155–1170. <https://doi.org/10.1287/mnsc.2016.2643>
- [12] Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. 2023. GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models. *arXiv preprint arXiv:2303.10130* (March 2023). <https://doi.org/10.48550/arXiv.2303.10130>
- [13] European Parliament and Council of the European Union. 2024. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act), Annex III, Section 4(b). Official Journal of the European Union. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
- [14] Andrew Fuchs, Andrea Passarella, and Marco Conti. 2024. Optimizing Delegation in Collaborative Human-AI Hybrid Teams. *ACM Transactions on Autonomous and Adaptive Systems* 19, 4 (Nov. 2024), 1–33. <https://doi.org/10.1145/3687130>
- [15] Ella Glikson and Anita Woolley. 2020. Human Trust in Artificial Intelligence: Review of Empirical Research. *Academy of Management Annals* 14, 2 (July 2020), 627–660. <https://www.researchgate.net/publication/340605601>
- [16] Kate Goddard, Abdul Roudsari, and Jeremy C. Wyatt. 2012. Automation Bias: A Systematic Review of Frequency, Effect Mediators, and Mitigators. *Journal of the American Medical Informatics Association* 19, 1 (Jan. 2012), 121–127. <https://doi.org/10.1136/amiajnl-2011-000089>
- [17] Patrick Hemmer, Monika Westphal, Max Schemmer, Sebastian Vetter, Michael Vössing, and Gerhard Satzger. 2023. Human-AI Collaboration: The Effect of AI Delegation on Human Task Performance and Task Satisfaction. In *Proceedings of the 28th International Conference on Intelligent User Interfaces*. 453–463. <https://doi.org/10.1145/3581641.3584052>
- [18] Michael C. Horowitz and Lauren Kahn. 2024. Bending the Automation Bias Curve: A Study of Human and AI-Based Decision Making in National Security Contexts. *International Studies Quarterly* 68, 2 (June 2024), sqae020. <https://doi.org/10.1093/isq/sqae020>
- [19] International Labour Organization. 2023. *Generative AI and Jobs: A Global Analysis of Potential Effects on Job Quantity and Quality*. Technical Report ILO Working Paper 96. International Labour Office, Geneva. <https://doi.org/10.54394/FHEM8239>
- [20] International Organization for Standardization/International Electrotechnical Commission. 2023. ISO/IEC 42001:2023 — Information Technology — Artificial Intelligence — Management System. Geneva: ISO/IEC. <https://www.iso.org/standard/81230.html>
- [21] S. Mo Jones-Jang and Yoonsang Park. 2023. How Do People React to AI Failure? Automation Bias, Algorithmic Aversion, and Perceived Controllability. *Journal of Computer-Mediated Communication* 28, 1 (Jan. 2023), zmac029. <https://doi.org/10.1093/jcmc/zmac029>
- [22] Min Hun Lee and Martyn Zhe Tok. 2025. Towards Uncertainty Aware Task Delegation and Human-AI Collaborative Decision-Making. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*. 2274–2289. <https://doi.org/10.1145/3715275.3732155>
- [23] Jennifer M. Logg, Julia A. Minson, and Don A. Moore. 2019. Algorithm Appreciation: People Prefer Algorithmic to Human Judgment. *Organizational Behavior and Human Decision Processes* 151 (March 2019), 90–103. <https://doi.org/10.1016/j.obhdp.2018.12.005>
- [24] Brian Lubars and Chenhao Tan. 2019. Ask Not What AI Can Do, but What AI Should Do: Towards a Framework of Task Delegability. In *Advances in Neural Information Processing Systems*, Vol. 32. 57–67. <https://doi.org/10.48550/arXiv.1902.03245>
- [25] Alexandra-Ioana Murariu, Silviu Mihai Dragomir, and Bogdan Mihai Frătăilă. 2024. Artificial Intelligence as a Human Capital Multiplier: A New Economic Paradigm. *Bucharest University of Economic Studies / Valahia University of Târgoviște Working Paper* (2024).
- [26] National Institute of Standards and Technology. 2023. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. Technical Report NIST AI 100-1. NIST. <https://doi.org/10.6028/NIST.AI.100-1>
- [27] Suley Fischer Omura and Eduardo de Rezende Francisco. 2025. Hybrid Business Decisions: Does Task Complexity Uplift Delegation to Artificial Intelligence? *Revista de Administração Contemporânea* 29, 6 (Jan. 2025), e250032. <https://doi.org/10.1590/1982-7849rac2025250032.en>
- [28] Organisation for Economic Co-operation and Development. 2024. Recommendation of the Council on Artificial Intelligence. OECD Legal Instruments. <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>
- [29] Organisation for Economic Co-operation and Development. 2024. *Using AI in the Workplace: Opportunities, Risks and Policy Responses*. Technical Report OECD Artificial Intelligence Papers No. 14. OECD Publishing, Paris. <https://doi.org/10.1787/73a0e3b4-en>
- [30] Neil Perry, Megha Srivastava, Deepak Kumar, and Dan Boneh. 2023. Do Users Write More Insecure Code with AI Assistants?. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*. 2785–2799. <https://doi.org/10.1145/3576915.3623157>
- [31] Sebastian Raisch and Sebastian Krakowski. 2021. Artificial Intelligence and Management: The Automation–Augmentation Paradox. *Academy of Management Review* 46, 1 (Jan. 2021), 192–210. <https://doi.org/10.5465/amr.2018.0072>
- [32] Charvi Rastogi, Yunfeng Zhang, Dakuo Wang, Kush R. Varshney, Amit Sharma, and Richard Tomsett. 2022. Deciding Fast and Slow: The Role of Cognitive Biases in AI-Assisted Decision-Making. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW1 (March 2022), 1–22. <https://doi.org/10.1145/3512930>
- [33] Tapani Rinta-Kahila, Esko Penttinen, Antti Salovaara, Wael Soliman, and Joonas Ruusalo. 2023. The Vicious Circles of Skill Erosion: A Case Study of Cognitive Automation. *Journal of the Association for Information Systems* 24, 5 (Sept. 2023). <https://doi.org/10.17705/1jais.00829>
- [34] Fareed Salih and Islam H. El-adaway. 2025. Investigating the Correlation and Synergy of Artificial Intelligence Techniques, Construction Technologies, and Their Key Benefits. *Journal of Architectural Engineering* 31, 4 (2025). <https://doi.org/10.1061/JAEIED.AEENG-1979>
- [35] Nenad Tomašev, Ulrich Paquet, Demis Hassabis, and Scott Reed. 2026. Intelligent AI Delegation. *Nature Machine Intelligence* 8, 4 (April 2026), 301–314. <https://arxiv.org/pdf/2602.11865>
- [36] Priyan Vaithilingam, Tianyi Zhang, and Elena L. Glassman. 2022. Expectation vs. Experience: Evaluating the Usability of Code Generation Tools Powered by LLMs.

In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, 1–7. <https://doi.org/10.1145/3491101.3519665>

- [37] Helena Vasconcelos, Matthew Jörke, Madeleine Grunde-McLaughlin, Tobias Gerstenberg, Michael Bernstein, and Ranjay Krishna. 2023. Explanations Can Reduce Overreliance on AI Systems During Decision-Making. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1, Article 119 (April 2023), 36 pages. <https://doi.org/10.1145/3579605>
- [38] World Economic Forum. 2023. *The Future of Jobs Report 2023*. Technical Report. World Economic Forum. <https://www.weforum.org/reports/the-future-of-jobs-report-2023>
- [39] Yide Xu, Li-Ren Shieh, Patrick van Esch, and I-Ling Ling. 2020. AI Customer Service: Task Complexity, Problem-Solving Ability, and Usage Intention. *Australasian Marketing Journal* 28, 3 (Aug. 2020), 189–199. <https://doi.org/10.1016/j.ausmj.2020.03.005>

A Search Strings

Table 2: Search strings used.

Concept	Search Strings
AI Delegation	“AI delegation” OR “task delegation to AI” OR “AI task allocation”
Characteristics & Properties of Task	“task complexity” OR “task properties” OR “task difficulty”
AI Performance	“AI accuracy” OR “AI trust” OR “algorithm performance”
Cost and Efficiency	“performance cost trade-off” OR “time pressure” OR “cost efficiency”
Accountability	“AI governance” OR “accountability” OR “AI regulation” OR “transparency”
Human-AI Collaboration	“human-AI collaboration” OR “human-in-the-loop” OR “partial delegation”

B Inclusion/Exclusion Criteria

Table 3: Inclusion/Exclusion criteria used.

Criterion	Inclusion	Exclusion
Date	2015–April 2026	Before 2015 (exception: [16])
Language	English	Non-English
Relevance	Factors influencing AI task delegation	Pure technical AI model development with no delegation discussion
Source Type	Peer-reviewed articles, standards, policy documents, major industry reports	Blog posts, opinion pieces, editorials

C PRISMA-ScR Diagram

PRISMA 2020 flow diagram for new systematic reviews which included searches of databases and registers only

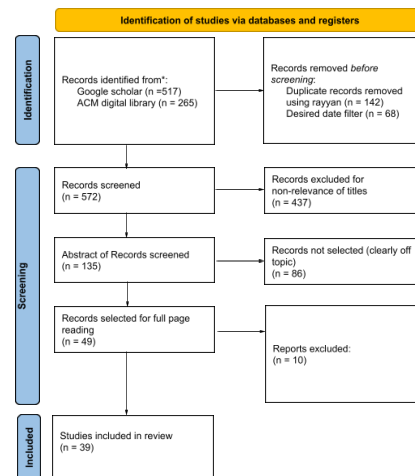


Figure 1: PRISMA-ScR flow diagram.

D Coding Scheme

The link to the ZIP archive including our coding table in CSV format and a txt file explaining the production process is given below.

https://drive.google.com/file/d/1aBK5hI08itO5qTVD6X1oHbGBr_IWdSrW/view?usp=sharing

E Framework Diagram

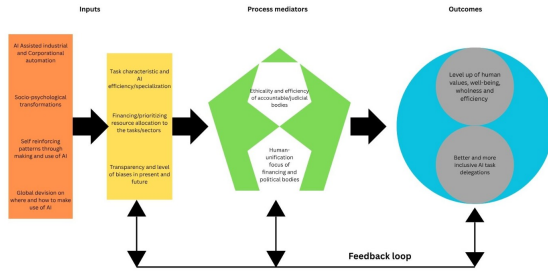


Figure 2: Framework Diagram.

F AI Use Disclosure

Table 4: AI Use disclosure.

Question	Answer
Which tools were used?	Claude and Gemini
How were they used?	They were used as aids for coming up with search strings, the overall structure of the paper, and creating the citations in ACM format with the chosen sources.
What verifications were applied?	All AI suggestions were checked by team members (strings were tried individually, organization changes were proofread).
All cited sources read by a team member?	Yes, we confirm that all the cited sources were read by team members.

G Trustworthy 3 Members

- (1) **Sila Gulec** – s334743@studenti.polito.it – Politecnico di Torino, Italy
- (2) **Ahmet Tolga Birand** – s335068@studenti.polito.it – Politecnico di Torino, Italy
- (3) **Jacopo Caravaggio** – s339988@studenti.polito.it – Politecnico di Torino, Italy
- (4) **Giacomo Simbula** – s341576@studenti.polito.it – Politecnico di Torino, Italy
- (5) **Nicolò Fiori** – s325013@studenti.polito.it – Politecnico di Torino, Italy
- (6) **Benedetta Stellato** – s341071@studenti.polito.it – Politecnico di Torino, Italy
- (7) **Pietro de Francesco** – s340079@studenti.polito.it – Politecnico di Torino, Italy
- (8) **Behrad Galedari** – s324382@studenti.polito.it – Politecnico di Torino, Italy
- (9) **Kutay Gunel** – s307467@studenti.polito.it – Politecnico di Torino, Italy
- (10) **Gabriele Gallo** – s339929@studenti.polito.it – Politecnico di Torino, Italy