

From Learning-by-Doing to Doing-Without-Learning: A Sensemaking Model of Skill Development and Retention in AI-Augmented Work

Alberto Varone¹[0009–0005–6900–0839], Francesco Bolici¹[0000–0002–8532–4981], and
Kevin Crowston²[0000–0003–1996–3600]

¹ University of Cassino and Southern Lazio, Cassino, Italy

alberto.varone@unicas.it f.bolici@unicas.it

² Syracuse University, Syracuse NY, USA crowston@syr.edu

Abstract. This conceptual paper examines how the use of artificial intelligence (AI) systems transform the conditions under which skills develop or are maintained by reshaping the sensemaking processes workers use to interpret, verify and learn from their actions. Using sensemaking and socio-technical systems theories, we theorize that unless verification capabilities are cultivated, AI shifts work from “learning by doing” to “doing without learning”. We develop a four-phase orchestration model (framing, prompting, interpreting, verifying) that identifies where in the human-AI interaction cycle learning opportunities emerge or disappear. Verification emerges as the developmental bottleneck linking AI use to divergent skill trajectories. We present propositions for empirical study and discuss implications for interaction design and organizational practice.

Keywords: AI Orchestration Skills · Verification Intent · Skill Development and Erosion · Human-AI Interaction · Organization Sensemaking

1 Introduction

Industry 5.0 envisions a future where technological advancement serves human flourishing, placing worker well-being and autonomy at the center of production processes [27]. As artificial intelligence (AI) systems³, large language models (LLMs) specifically, become more capable, they risk inducing workers to abdicate cognitive agency to machine intelligence, discounting their own problem-solving and interpretive judgment [12, 14, 35]. Understanding how to design human-machine interaction that genuinely augments rather than replaces human cognition has become central to realizing Industry 5.0’s promise of harmonizing technological capability with human well-being and control [1].

³ We define AI as systems “that build on machine learning, computation and statistical techniques, as well as rely on large data sets to generate responses, classifications or dynamic predictions” [19, p. 62].

This challenge is particularly acute in knowledge-intensive work, where emerging evidence reveals a troubling pattern: AI use raises short-term performance while simultaneously altering the cognitive activities that historically supported learning and expertise development [17, 30, 33, 4]. Workers across domains from programming to professional services to medicine appear to engage in “doing without learning”: tasks are completed successfully with AI assistance, yet the interpretive and diagnostic routines that underpin skill development and autonomous capability are bypassed [41, 36]. This pattern represents not just a technical issue but a fundamental threat to worker agency and long-term well-being, as individuals lose the ability to operate independently of system support.

What remains poorly understood is how individuals construct meaning when working with AI and how this process shapes expertise development. Sensemaking research suggests that learning depends on active engagement with cues and discrepancies, yet AI tools may suppress these triggers by offering fluent outputs that appear authoritative even when flawed [25, 6]. Emerging work indicates that verification—interrogating AI outputs rather than accepting them at face value—is the central differentiator between genuine augmentation and *de facto* automation. We treat *verification intent*, that is, the motivation and capacity to verify, as the core mechanism in our model. Individuals who interrogate and reinterpret system outputs develop richer mental models, retain greater agency over their work and maintain well-being through a sense of competence and control, whereas those who rely uncritically on AI plateau or even regress in expertise development despite performance remaining superficially high [22, 37].

Our contribution is twofold. First, we generalize prior findings beyond individual domains, integrating sensemaking and STS theories into a four-phase orchestration model that specifies where in the human-AI interaction cycle learning opportunities emerge or disappear. Second, we articulate organizational conditions under which AI can support long-term human capability and well-being rather than erode them. These include task portfolio design, workflow transparency, feedback arrangements and managerial practices that preserve opportunities to verify and reflect on AI-generated material [19, 34, 20].

The paper thus provides both a theoretical foundation for understanding divergent skill trajectories under seemingly identical AI usage and practical guidance for organizations seeking to cultivate sustainable human capability. In doing so, it addresses Industry 5.0’s core challenge: how to design human-machine interaction that genuinely augments human cognition while ensuring workers retain control, governance and meaningful engagement with their work.

2 Theoretical Background

We develop our model through theory synthesis [21], integrating constructs and findings from the literature into a unified developmental model that specifies mechanisms, boundary conditions and testable propositions not derivable from any single source literature. This section develops the two theoretical pillars informing our model: sensemaking and socio-technical systems (STS). These two

serve as method theories [21] that we apply to the domain of AI-augmented knowledge work, synthesizing constructs from both to develop the orchestration model and its propositions. Together, these perspectives allow us to analyze how AI alters not only task execution but the learning conditions shaping long-term skill development.

2.1 Sensemaking and the Construction of Meaning in Work

Several traditions converge on the principle that expertise develops through effortful engagement rather than passive exposure. Skill acquisition theory traces progression from rule-following to intuitive judgment through accumulated experience [16], a trajectory we draw on in Section 4 to differentiate verification patterns across expertise levels. Deliberate practice research emphasizes that improvement requires activities at the boundary of current ability with informative feedback [18, 42]. Experiential learning models hold that reflection on concrete experience drives conceptual development [23, 39]. Sensemaking research, our focus, emphasizes that individuals learn by noticing cues, interpreting discrepancies and enacting responses that create feedback for reflection [43, 44]. Learning is thus an active process of meaning construction depending on attention to cues, interpretation of their relevance and actions revealing whether interpretations hold. In complex work, discrepancies and uncertainties trigger deeper engagement and revision of assumptions [26]. Such episodes are central for learning because they expose gaps in understanding and motivate inquiry. When individuals confront ambiguous situations, they experiment, seek information and refine mental models [38].

AI introduces a new dynamic. Generative AI models produce fluent, seemingly authoritative outputs that reduce ambiguity. While these contributions lower cognitive effort, they may mask the triggers that initiate sensemaking. When outputs appear correct, individuals may engage less in questioning and verification, proceeding with a feeling of clarity ungrounded in deeper understanding, potentially allowing errors to remain hidden [41, 22]. Moreover, if AI tools absorb too much problem-framing and exploration work, users receive fewer opportunities to test assumptions or interact with others. Studies show individuals often over-rely on AI suggestions unless they have strong domain knowledge or explicit incentives to verify [37, 24]. Skills thus become contingent on verification intent: without it, individuals bypass interpretive cycles essential for skill development; with it, AI can enable richer sensemaking.

2.2 Socio-Technical Systems and the Reconfiguration of Work

Socio-technical systems theory (STS) views work as shaped by interaction between technology, tasks, structures and human capabilities [40, 11, 3]. Rather than treating technology as neutral, STS emphasizes that technological effects depend on how systems are configured and enacted within organizational practices [31, 5]. The introduction of AI shifts the distribution of cognitive work across people and systems. AI tools generate intermediate representations, draft solutions and propose interpretations, occupying roles once held by humans. These changes

alter workflows, feedback loops and opportunities for individuals to engage with tasks in ways supporting expertise development [19, 34].

STS research underscores the importance of congruence: performance improves when task demands, technology capabilities and human skills align [28, 40]. Misalignment produces coordination difficulties, ambiguity or reduced learning opportunities. In the AI context, congruence depends on whether individuals have interpretive and evaluative skills to complement what technology does well and compensate for what it does poorly. Recent work shows that nominally “augmentative” AI systems can slide into *de facto* automation when users lack expertise to scrutinize suggestions [20, 14]. In contrast, experts can use the same systems to extend capabilities, explore alternatives and learn faster. This difference reinforces the idea that AI alters socio-technical configurations in ways privileging some expertise profiles over others.

From an STS perspective, AI’s impact on learning depends on how organizations design tasks, workflows and practices around technology. Systems that expose reasoning steps, foreground exceptions and invite intervention preserve learning opportunities. Systems that hide intermediate reasoning or produce apparently finished outputs reduce the need for interpretation and weaken the socio-technical foundation for skill development.

Together, sensemaking and STS perspectives explain why the same AI system may support learning for some users but undermine it for others: skill development requires exposure to meaningful cues and opportunities to verify them and AI modifies both processes.

3 Human–AI Task Accomplishment and the Centrality of Verification

In knowledge-intensive work, AI now participates in tasks previously handled solely by humans, from writing and programming to analysis and decision support [8, 14, 33]. What matters is not simply that AI generates content, but how people interact with the system: which tasks they retain, which they delegate and where they intervene. These interaction patterns shape both immediate performance and longer-term expertise development [41, 9].

We conceptualize human-AI interaction as an AI Orchestration Cycle (Figure 1) consisting of four recurring phases: Framing the problem, Delegating through prompting, Interpreting system outputs and Verifying and integrating outputs into work. Although steps may unfold quickly, they represent distinct cognitive demands and learning opportunities. The cycle provides an analytical framework for examining AI-augmented work across different domains and skill levels. The decomposition enables analysts to identify where workers might bypass critical cognitive work, assess skill requirements at each phase, evaluate whether organizational conditions support or suppress engagement and guide interaction design accordingly.

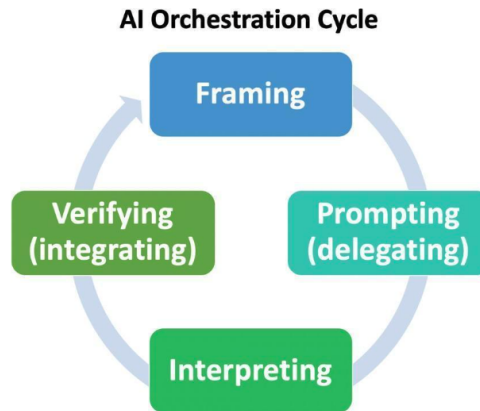


Fig. 1. AI Orchestration Cycle: steps of AI-Human interaction representing distinct cognitive demands and opportunities for learning

3.1 Framing and Prompting: How Users Define the Problem Space

Human-AI interaction begins before any model output is generated. The framing step involves defining goals, constraints, success criteria and assumptions. Research shows framing shapes not only AI suggestion quality but also user understanding of the problem [15, 14, 8]. Poorly-framed tasks lead to plausible but irrelevant outputs, prompting correction cycles that rarely deepen expertise.

Delegation through prompting, the second step, is often discussed, yet empirical findings suggest “prompting skill” has a more modest role than assumed. In several studies, domain expertise rather than prompting proficiency predicted performance [45, 9]. Experienced users gave shorter, simpler prompts because they already possessed mental models to evaluate and adapt outputs. Less experienced users often produced long, detailed prompts without improving results, indicating prompting is a complement to expertise rather than a substitute.

Framing and prompting are not merely technical actions: they express underlying cognitive structures, representational resources and organizational norms regarding how problems should be approached. They also set conditions under which interpretation and verification become possible.

3.2 Interpretation: Locating AI Outputs Within Context

Once outputs are generated, users must interpret them, determining their meaning and assessing relevance and coherence with existing knowledge or project constraints. Here AI’s limitations become particularly salient. Studies show LLMs often produce superficially coherent solutions lacking robust causal or domain grounding [6, 25]. As tasks become more interdependent or open-ended, interpretation demands increase because AI outputs must be situated within organizational procedures, stakeholder expectations or technical dependencies.

Sensemaking theory clarifies that interpretation is not only cognitive but also social. Users draw on shared routines, local norms and interactional cues to judge whether an AI suggestion “makes sense” in their setting [43, 7]. Interpretation operates at the intersection of personal expertise and collective practice. Where interpretive resources are thin—because workers are novices, feedback is sparse or the task is highly ambiguous—the orchestration cycle becomes fragile and users may short circuit the cycle entirely.

3.3 Verification: Checking and Learning

Despite widespread claims that prompting is the new core skill, emerging evidence shows verification intention and ability are far more predictive of meaningful learning and expertise development [9, 46, 47]. When verification weakens, the cycle collapses into passive use. For instance, studies in medicine have documented decreases in individual performance after a period of using AI [29, 10]. In contrast, when verification is active and reflective, users re-engage with the problem space, refine representations, internalize new skills and reinforce existing ones [43, 26].

Users vary widely in whether they verify AI outputs, how deeply and under which conditions. Four clusters of drivers shape verification intent: individual, task, system and organizational.

Individual drivers: Users differ in epistemic stance. Those with stronger prior expertise, richer mental models or greater epistemic humility are more likely to scrutinize outputs, detect inconsistencies and iterate. Novices often lack schemas to recognize errors and exhibit high reliance even when suspicious. Time pressure, fatigue and cognitive load also reduce active verification, pushing users toward superficial acceptance.

Task drivers: Tasks with high uncertainty, ambiguous goals, non-routine dependencies or significant downstream consequences make verification more salient. For instance, in programming, debugging tasks trigger more verification than boilerplate generation. In HR analytics or financial modeling, verification increases when outputs affect legal compliance or risk exposure. Low-stakes or highly-standardized tasks suppress verification because error costs appear lower.

System drivers: AI systems differ in cues they provide to signal uncertainty or invite scrutiny. Models that surface assumptions, reveal intermediate reasoning steps or provide multiple candidate solutions elicit deeper verification. Interfaces presenting outputs as authoritative, final or overly polished suppress them. Research shows transparency and appropriate uncertainty signaling improve oversight, while opacity fosters over-reliance [46, 32].

Organizational drivers: Verification is shaped by social and organizational environment. Cultures emphasizing learning, constructive critique and psychological safety encourage deeper engagement. Environments rewarding speed, efficiency or error-free appearances discourage verification, especially for novices who fear signaling incompetence. Team norms around documentation, peer review or reflective practice can also amplify or suppress verification behaviour.

3.4 Preliminary Observations

These theoretical drivers find early support in a study we are conducting of AI-augmented programming (N=16 interviews with university faculty and students using AI tools to develop research software). When asked about their interaction patterns, nearly all interviewees spontaneously emphasized verification’s importance. One developer explained: *“I feel if I have to be accurate, I need to check the code more in detail if we’re talking about AI. If AI writes the code, I have to be very cautious whether it’s right or wrong”*. However, participants also acknowledged that verification intent varies. One noted that: *“I don’t dive into the code until I find that my results don’t pass a sniff test.”* Another noted context dependency, e.g., that the need to verify depended on the task: *“for data munging of participant surveys, I really don’t care how much I understand the details of what’s going on.”* One admitted to periods of passive use under time pressure, stating: *“I can’t find any other solution or if this problem is an emergency, I can use AI.”* These preliminary observations suggest verification intent is both salient to practitioners and variable across individuals, tasks and organizational contexts, motivating the formal propositions developed next.

4 A Sensemaking Model of Human–AI Skill Development

This section develops a set of propositions for future studies of how interactions within the AI orchestration cycle shape learning trajectories. Building on sensemaking and socio-technical perspectives [43, 31], we conceptualize verification as the moment where users confront discrepancies between expected and observed outcomes and decide whether to reaffirm or revise interpretations. Through this lens, learning and skill retention emerge from how individuals make sense of AI-generated cues while enacting the orchestration cycle.

4.1 Verification as a Trigger for Interpretive Work

When users actively interrogate an AI output by asking why the system produced a particular solution, contrasting it with alternative frames or probing missing constraints they engage in interpretive effort that sensemaking theory identifies as the foundation of learning [44, 26]. By contrast, when verification is reduced to superficial checking, users bypass these interpretive processes, achieving acceptable performance but learning little about the underlying task: the augmentation-deskilling paradox [41]. This leads to our first proposition:

Proposition 1. *When users engage in active verification during the AI orchestration cycle, they are more likely to develop and retain task-specific expertise than when verification is superficial or absent.*

4.2 Verification Across Expertise Levels

Expertise shapes how individuals approach verification and how much they learn from AI use [13, 4].

- Novices often lack interpretive schemas to scrutinize AI outputs. Even when they intend to verify, they may not recognize subtle inconsistencies or omitted constraints, leading to rapid surface performance improvement without deepening underlying competence.
- Intermediates represent the most consequential group. They possess emerging schemas allowing them to sense when an output seems “off”, but their ability to diagnose why remains fragile. For them, AI can powerfully amplify learning if verification is active: discrepancies become resources for interpretation and iterative sensemaking refines developing mental models. However, if intermediates defer too quickly to AI suggestions, their progression plateaus and dependency increases.
- Experts use AI to expand the range of available cues. Their strong mental models allow efficient verification and integration of AI contributions without losing interpretive control.

This differentiated pattern implies learning effects of verification are not uniform, leading to our second proposition:

Proposition 2. *Verification behaviours vary systematically across expertise levels, with intermediates showing the greatest divergence between accelerated learning when verification is active and stagnation when verification is weak.*

4.3 Organizational Conditions That Shape Verification

Whether users engage in verification is not only individual preference but shaped by organizational factors, as discussed above. How AI tools are introduced, whether foregrounding speed and convenience or reasoning and critique, creates different norms for verification. When organizations institutionalize moments for reflection (e.g., audit steps, rationale articulation, comparison between alternatives), verification becomes routine rather than exception. This leads to our final proposition:

Proposition 3. *Organizational conditions that support reflective practice (e.g., incentives for accuracy, feedback mechanisms and workflow structures requiring checking) will increase active verification and strengthen long-term skill development and retention.*

5 Discussion and Implications

Traditional work analysis approaches examine how workers make sense of complex work domains, identify constraints and demands and develop strategies for effective performance. Our orchestration model extends work analysis to AI-augmented contexts by identifying four distinct cognitive steps that structure human-AI interaction, revealing where in the cycle AI alters cognitive demands, introduces new coordination challenges or enhances or removes learning opportunities. From a work-domain-analysis perspective, AI systems function as both

tools and collaborators, creating work arrangements where task allocation becomes dynamic and verification requirements multiply. Learning occurs unevenly across these phases: skills grow when individuals remain engaged in framing and verification, where sensemaking, uncertainty and interpretive judgment are concentrated.

Over time, divergent verification patterns create distinct development paths within the same organization. Users who approach AI interactively accumulate representational resources supporting future learning; those relying on it as an answer engine experience progressive narrowing of interpretive capacity. If organizations structure work so AI performs most framing and verification, orchestration becomes predominantly passive. Performance may rise short-term, but foundational capabilities decay. Early-career employees may become proficient at managing AI outputs while lacking the ability to solve problems independently [2]. Conversely, workflows preserving active orchestration, e.g., by requiring workers to articulate problem boundaries and examine intermediate reasoning, sustain learning. Work design should therefore preserve access to ‘learning-band’ tasks: those containing enough complexity to require framing and verification but remaining tractable without over-reliance on AI [42]. Rotational assignments, manual-first exercises or workflows surfacing intermediate reasoning steps help maintain these conditions.

Universities face similar challenges. Students now complete complex tasks through efficient prompting and selective editing, circumventing precisely the interpretive work curricula are designed to develop. Embedding instruction on verification strategies, decomposition, interpretive judgment and uncertainty handling becomes essential. The orchestration cycle offers a practical scaffold for structuring these interventions, helping learners understand not only how to use AI tools, but how to learn with them.

Finally, skill divergence is not merely individual: it also reshapes organizational capability. Responsible AI adoption is not only about ethics or compliance but organizational sustainability. Teams drifting toward passive orchestration become dependent on AI for everyday reasoning, weakening resilience. Teams cultivating active orchestration develop adaptive capacity: they can question outputs, correct failures and recombine knowledge in novel situations.

6 Conclusions

This paper develops a lens for understanding AI’s impact on human capabilities through a four-phase orchestration model that identifies where in human-AI interaction learning opportunities emerge or disappear. We present propositions for empirical study and note several boundaries for future research.

First, the orchestration model abstracts from variation across domains, industries and task types. While grounded in established theories, its dynamics may unfold differently in highly regulated settings (e.g., healthcare) compared to fast-moving digital sectors. Second, the paper focuses on cognitive tasks and knowledge work. Applicability to hybrid digital-physical environments (e.g., man-

ufacturing, logistics, construction) remains open, particularly where embodied skills and sensorimotor learning play central roles. Third, the model emphasizes individual-level processes. Organizational factors—team structures, workflow design, incentive systems and managerial decisions—are discussed but not formally incorporated into conceptual propositions. Finally, the paper does not specify how AI system design itself (e.g., interface transparency, memory, stepwise reasoning, constraint setting) moderates orchestration patterns. For instance, systems could better express levels of uncertainty, prompting more verification of results. These system-level features likely shape how much framing and verification humans must retain and deserve future attention.

The conceptual nature of this paper presents numerous opportunities for empirical work. First, the research propositions are stated at a conceptual level and so need to be translated into more operational hypotheses, for instance, how specifically to assess “verification intent” or skill development. Second, empirical work is needed to observe how orchestration patterns (verification in particular) evolve over time and shape expertise trajectories. Empirical studies of AI-mediated work can use the orchestration cycle to structure data collection and analysis. Interaction logging can capture timestamped records of prompting, output generation and verification, quantifying how workers distribute attention across steps and how outputs are used. Critical incident interviews can ask workers to recount episodes where AI helped or hindered their work, with specific probes for each orchestration phase. Researchers can thus test whether active orchestrators develop more durable and transferable skills than passive ones.

Third, as noted above, different professions may distribute framing, interpretation and verification activities in distinct ways. Comparative studies of programmers, consultants, data analysts, medical practitioners or students could identify which domain characteristics amplify or mitigate the risk of “doing without learning”. Workplace ethnography can observe how organizational context shapes verification practices. Research may explore how career progression, competence frameworks and incentive structures influence verification intent. Finally, future studies could evaluate the impact on orchestration skill development of the suggested interventions, such as job rotation, manual-first assignments, reflective sessions or structured peer review. In education, researchers can examine how different pedagogical strategies (e.g., requiring justification of prompts, mandating verification protocols) affect student learning and whether early exposure to active orchestration shapes later professional competence.

These research directions will contribute to Industry 5.0’s vision of harmonizing human capability with technological advancement. The orchestration model offers a framework for operationalizing core Industry 5.0 principles: maintaining worker autonomy through active orchestration practices, ensuring human control and governance through workflow designs that preserve interpretive engagement and sustaining worker well-being by protecting the cognitive activities that support both competence development and meaningful work. Rather than asking workers to serve as passive supervisors of increasingly autonomous AI systems, a configuration that risks eroding both capability and professional identity, the

model points toward socio-technical designs where AI genuinely augments human cognition while workers retain agency over problem-framing, interpretation, verification and ultimate decision-making. This approach is not merely a technical design challenge but a fundamental question of human values, one that human-work interaction design and Industry 5.0 frameworks are uniquely positioned to address through integration of human sciences, organizational theory and technological understanding.

Acknowledgments. This work was partly funded by a grant from the Alfred P. Sloan Foundation, G-2025-79202.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Alfes, K., Avgoustaki, A., Beauregard, T.A., Cañibano, A., Muratbekova-Touron, M.: New ways of working and the implications for employees: A systematic framework and suggestions for future research. *The International Journal of Human Resource Management* **33**(22), 4361–4385 (2022). <https://doi.org/10.1080/09585192.2022.2149151>
2. Anthony, C.: When knowledge work and analytical technologies collide: The practices and consequences of black boxing algorithmic technologies. *Administrative Science Quarterly* **66**(4), 1173–1212 (2021). <https://doi.org/10.1177/00018392211016755>
3. Appelbaum, S.H.: Socio-technical systems theory: An intervention strategy for organizational development. *Management Decision* **35**(6), 452–463 (1997). <https://doi.org/10.1108/00251749710173823>
4. Bolici, F., Crowston, K., Varone, A., Fudge, M.: Rethinking programming skills in the age of generative AI. In: Proceedings of the 59th Hawaii International Conference on System Sciences (2026), <https://hdl.handle.net/10125/112269>
5. Boudreau, M.C., Robey, D.: Enacting integrated information technology: A human agency perspective. *Organization Science* **16**(1), 3–18 (2005). <https://doi.org/10.1287/orsc.1040.0103>
6. Bowen, C., Sætre, R., Miyao, Y.: A comprehensive evaluation of inductive reasoning capabilities and problem solving in large language models. In: Graham, Y., Purver, M. (eds.) Findings of the Association for Computational Linguistics: EACL 2024. pp. 323–339 (2024), <https://aclanthology.org/2024.findings-eacl.22>
7. Brown, A.D., Colville, I., Pye, A.: Making sense of sensemaking in organization studies. *Organization Studies* **36**(2), 265—277 (2015). <https://doi.org/10.1177/0170840614559259>
8. Brynjolfsson, E., Li, D., Raymond, L.: Generative AI at work. *The Quarterly Journal of Economics* **140**(2), 889–942 (2025). <https://doi.org/10.1093/qje/qjae044>
9. Buçinca, Z., Malaya, M.B., Gajos, K.Z.: 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), 1–21 (2021). <https://doi.org/10.1145/3449287>

10. Budzyń, K., Romańczyk, M., Kitala, D., Kołodziej, P., Bugajski, M., Adami, H.O., Blom, J., Buszkiewicz, M., Halvorsen, N., Hassan, C., Romańczyk, T., Holme, Ø., Jarus, K., Fielding, S., Kunar, M., Pellise, M., Pilonis, N., Kamiński, M.F., Kalager, M., Bretthauer, M., Mori, Y.: Endoscopist deskilling risk after exposure to artificial intelligence in colonoscopy: A multicentre, observational study. *The Lancet Gastroenterology & Hepatology* **10**(10), 896–903 (Oct 2025). [https://doi.org/10.1016/S2468-1253\(25\)00133-5](https://doi.org/10.1016/S2468-1253(25)00133-5)
11. Cherns, A.: The principles of sociotechnical design. *Human Relations* **29**(8), 783–792 (1976). <https://doi.org/10.1177/001872677602900806>
12. Crowston, K., Bolici, F.: Impacts of machine learning on work. In: *Proceedings of the Hawaii International Conference on System Sciences* (2019). <https://doi.org/10.24251/HICSS.2019.719>
13. Crowston, K., Bolici, F.: Deskilling and upskilling with AI systems. *Information Research an International Electronic Journal* **30**, 1009–1023 (Mar 2025). <https://doi.org/10.47989/ir30iConf47143>
14. Dell’Acqua, F., McFowland, E., Mollick, E., Lifshitz, H., Kellogg, K.C., Rajendran, S., Krayer, L., Candelon, F., Lakhani, K.R.: Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality. *Organization Science* p. orsc.2025.21838 (Mar 2026). <https://doi.org/10.1287/orsc.2025.21838>
15. Denny, P., Becker, B.A., Leinonen, J., Prather, J.: Chat overflow: Artificially intelligent models for computing education: RenAIssance or apocAIypse? In: *Proceedings of the 2023 Conference on Innovation and Technology in Computer Science Education V. 1*. pp. 3–4 (2023). <https://doi.org/10.1145/3587102.3588773>
16. Dreyfus, S.E.: The five-stage model of adult skill acquisition. *Bulletin of Science, Technology & Society* **24**(3), 177–181 (2004). <https://doi.org/10.1177/0270467604264992>
17. Ehsan, U., Passi, S., Saha, K., McNutt, T., Riedl, M.O., Alcorn, S.: From future of work to future of workers: Addressing asymptomatic AI harms for dignified human-AI interaction (Jan 2026). <https://doi.org/10.48550/arXiv.2601.21920>
18. Ericsson, K.A., Krampe, R.T., Tesch-Römer, C.: The role of deliberate practice in the acquisition of expert performance. *Psychological Review* **100**(3), 363 (1993). <https://doi.org/10.1037/0033-295X.100.3.363>
19. Faraj, S., Pachidi, S., Sayegh, K.: Working and organizing in the age of the learning algorithm. *Information and Organization* **28**(1), 62–70 (2018). <https://doi.org/10.1016/j.infoandorg.2018.02.005>
20. Grønsund, T., Aanestad, M.: Augmenting the algorithm: Emerging human-in-the-loop work configurations. *The Journal of Strategic Information Systems* **29**(2), 101614 (2020). <https://doi.org/10.1016/j.jsis.2020.101614>
21. Jaakkola, E.: Designing conceptual articles: Four approaches. *AMS Review* **10**(1-2), 18–26 (Jun 2020). <https://doi.org/10.1007/s13162-020-00161-0>
22. Jakesch, M., Hancock, J.T., Naaman, M.: Human heuristics for AI-generated language are flawed. *Proceedings of the National Academy of Sciences* **120**(11), e2208839120 (2023). <https://doi.org/10.1073/pnas.2208839120>
23. Kolb, D.A.: *Experiential Learning: Experience as the Source of Learning and Development*. Prentice Hall, Englewood Cliffs, NJ (1984)
24. Lai, V., Tan, C.: On human predictions with explanations and predictions of machine learning models: A case study on deception detection. In: *Proceedings of the Conference on Fairness, Accountability, and Transparency*. pp. 29–38 (2019). <https://doi.org/10.1145/3287560.3287590>

25. Liu, E., Neubig, G., Andreas, J.: An incomplete loop: Instruction inference, instruction following, and in-context learning in language models (2024). <https://doi.org/10.48550/arXiv.2404.03028>
26. Maitlis, S., Christianson, M.: Sensemaking in organizations: Taking stock and moving forward. *Academy of Management Annals* **8**(1), 57–125 (2014). <https://doi.org/10.5465/19416520.2014.873177>
27. Mayer, H., Yee, L., Chui, M., Roberts, R.: Superagency in the workplace: Empowering people to unlock AI’s full potential. Tech. rep., McKinsey & Company (2025), <https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/superagency-in-the-workplace-empowering-people-to-unlock-ais-full-potential-at-work>
28. Nadler, D.A., Tushman, M.L.: A model for diagnosing organizational behavior. *Organizational Dynamics* **9**(2), 35–51 (1980). [https://doi.org/10.1016/0090-2616\(80\)90039-X](https://doi.org/10.1016/0090-2616(80)90039-X)
29. Natali, C., Marconi, L., Dias Duran, L.D., Cabitza, F.: AI-induced deskilling in medicine: A mixed-method review and research agenda for health-care and beyond. *Artificial Intelligence Review* **58**(11), 356 (Aug 2025). <https://doi.org/10.1007/s10462-025-11352-1>
30. Noy, S., Zhang, W.: Experimental evidence on the productivity effects of generative artificial intelligence. *Science* **381**(6654), 187–192 (2023). <https://doi.org/10.1126/science.adh2586>
31. Orlikowski, W.J., Iacono, C.S.: Research commentary: Desperately seeking the “IT” in IT research—A call to theorizing the IT artifact. *Information Systems Research* **12**(2), 121–134 (2001). <https://doi.org/10.1287/isre.12.2.121.9700>
32. Parasuraman, R., Manzey, D.H.: Complacency and bias in human use of automation: An attentional integration. *Human Factors* **52**(3), 381–410 (2010). <https://doi.org/10.1177/0018720810376055>
33. Peng, S., Kalliamvakou, E., Cihon, P., Demirer, M.: The impact of AI on developer productivity: Evidence from GitHub Copilot (2023). <https://doi.org/10.48550/arXiv.2302.06590>
34. Raisch, S., Krakowski, S.: Artificial intelligence and management: The automation-augmentation paradox. *Academy of Management Review* **46**(1), 192–210 (2021). <https://doi.org/10.5465/amr.2018.0072>
35. Renkema, M., Tursunbayeva, A.: The future of work of academics in the age of artificial intelligence: State-of-the-art and a research roadmap. *Futures* **163**, 103453 (2024). <https://doi.org/10.1016/j.futures.2024.103453>
36. Rinta-Kahila, T., Penttinen, E., Salovaara, A., Soliman, W., Ruissalo, J.: The vicious circles of skill erosion: A case study of cognitive automation. *Journal of the Association for Information Systems* **24**(5), 1378–1412 (2023). <https://doi.org/10.17705/1jais.00829>
37. Salimzadeh, S., He, G., Gadiraju, U.: Dealing with uncertainty: Understanding the impact of prognostic versus diagnostic tasks on trust and reliance in human-ai decision making. In: *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. pp. 1–17 (2024). <https://doi.org/10.1145/3613904.3641905>
38. Sandberg, J., Tsoukas, H.: Sensemaking reconsidered: Towards a broader understanding through phenomenology. *Organization Theory* **1**(1), 2631787719879937 (2020). <https://doi.org/10.1177/2631787719879937>
39. Schön, D.: *The reflective practitioner*. Basic Books, New York (1983)
40. Trist, E.L., Bamforth, K.W.: Some social and psychological consequences of the longwall method of coal-getting: An examination of the psychological situation and defences of a work group in relation to the social structure and

- technological content of the work system. *Human Relations* **4**(1), 3–38 (1951). <https://doi.org/10.1177/001872675100400101>
41. Varone, A., Crowston, K., Bolici, F.: Generative AI and the evolution of skills: A conceptual model for skill development and retention. In: 41st EGOS Colloquium (2025)
 42. Vygotsky, L.S., Cole, M.: *Mind in Society: Development of Higher Psychological Processes*. Harvard University Press (1978)
 43. Weick, K.E.: *Sensemaking in Organizations*. Sage, Thousand Oaks, CA (1995)
 44. Weick, K.E., Sutcliffe, K.M., Obstfeld, D.: Organizing and the process of sensemaking. *Organization Science* **16**(4), 409–421 (2005). <https://doi.org/10.1287/orsc.1050.0133>
 45. Wu, Z., Qiu, L., Ross, A., Akyürek, E., Chen, B., Wang, B., Kim, N., Andreas, J., Kim, Y.: Reasoning or reciting? Exploring the capabilities and limitations of language models through counterfactual tasks. In: *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. vol. 1: Long Papers. Association for Computational Linguistics (2024). <https://doi.org/10.18653/v1/2024.naacl-long.102>
 46. Yin, M., Wortman Vaughan, J., Wallach, H.: Understanding the effect of accuracy on trust in machine learning models. In: *Proceedings of the CHI Conference on Human Factors in Computing Systems*. pp. 1–12 (2019). <https://doi.org/10.1145/3290605.3300509>
 47. Zamfirescu-Pereira, J.D., Wong, R.Y., Hartmann, B., Yang, Q.: Why Johnny can't prompt: how non-AI experts try (and fail) to design LLM prompts. In: *Proceedings of the CHI Conference on Human Factors in Computing Systems*. pp. 1–21 (2023). <https://doi.org/10.1145/3544548.3581388>