

# Designing and Implementing Agentic Generative AI Professional Learning for the U.S. Federal Workforce: An Education-focused Framework for AI Literacy, Ethical Practice, and Transfer to Work

Satyadhar Joshi<sup>1</sup>, Noor Zulfiqar<sup>2,\*</sup>, Muhammad Usman Asif<sup>3</sup>, Sana Kazi<sup>4</sup>

<sup>1</sup>Department of Information Technology, MSIT Alumnus, Touro College, New York, NY 10010, USA.

<sup>2</sup>Department of Chemistry, Faculty of Sciences, University of Agriculture, Faisalabad 38000, Pakistan.

<sup>3</sup>Berlin School of Business and Innovation, Berlin 10437, Germany.

<sup>4</sup>Global Banking School, Greenford, London, UB6 0HE, UK.

**How to cite this paper:** Satyadhar Joshi, Noor Zulfiqar, Muhammad Usman Asif, Sana Kazi. (2026). Designing and Implementing Agentic Generative AI Professional Learning for the U.S. Federal Workforce: An Education-focused Framework for AI Literacy, Ethical Practice, and Transfer to Work. *The Educational Review, USA*, 10(2), 103-110.  
DOI: 10.26855/er.2026.02.007

**Received:** December 28, 2025

**Accepted:** January 30, 2026

**Published:** February 28, 2026

**Corresponding author:** Noor Zulfiqar, Department of Chemistry, Faculty of Sciences, University of Agriculture, Faisalabad 38000, Pakistan.

© 2026 by the author(s).

This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution-NonCommercial-NoDerivatives (CC BY-NC-ND) license, which permits non-commercial use, distribution, and reproduction in any medium, provided the original work is properly cited and is not modified or adapted.

<https://creativecommons.org/licenses/by-nc-nd/4.0/>

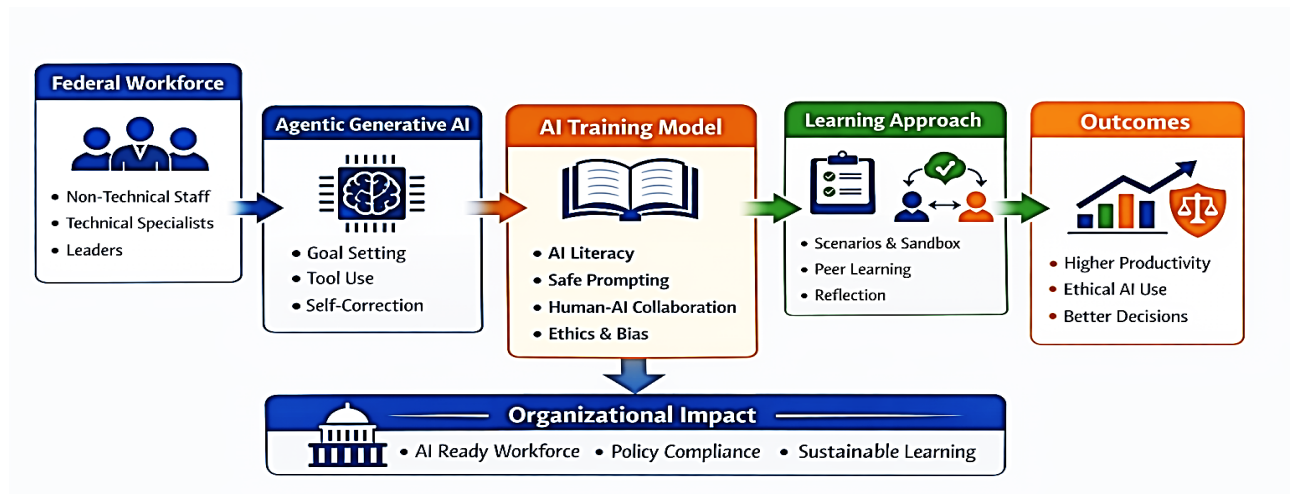
## Abstract

Agentic generative AI systems extend large language models by enabling goal-setting, multi-step planning, tool use, and iterative self-correction—capabilities that can significantly enhance workplace productivity but also require new forms of professional learning. In workplace settings, these capabilities can accelerate drafting, synthesis, and analysis tasks; however, the educational challenge is to prepare adult learners to collaborate with AI responsibly while maintaining professional judgment, data stewardship, and accountability. For the U.S. federal workforce, AI capability is not only a technology issue but also a workforce education issue: agencies require scalable professional learning programs that build AI literacy, safe-use habits, and measurable transfer of learning to job performance. This paper reframes agentic generative AI adoption as an educational design problem and proposes an education-focused framework for training federal employees. Drawing on AI literacy research, adult learning theory, experiential and workplace learning, and evaluation science, we outline (a) a competency model for agentic AI collaboration, (b) a modular curriculum blueprint aligned with adult learners' needs and federal operational constraints, (c) pedagogical strategies for secure, scenario-based practice in sandboxed environments, and (d) an assessment and program evaluation plan that emphasizes performance evidence, learning transfer, and ethical compliance. The framework supports differentiated pathways for non-technical staff, technical specialists, and leaders. By positioning agentic AI training as lifelong learning and professional development, the proposed approach helps educational leaders and public-sector training units design programs that improve learner confidence, reduce automation bias, strengthen information governance, and build organizational capacity for trustworthy AI use.

## Keywords

Agentic generative AI; AI literacy; adult education; professional learning; federal workforce development; instructional design; training evaluation

## Graphical Abstract



## 1. Introduction

Generative AI has become a mainstream productivity tool across knowledge-work sectors, including government. Recent advances in agentic generative AI extend conventional chat interfaces by enabling systems to break down goals, plan sequences of actions, call external tools (e.g., search, databases, analytics), maintain short-term memory across steps, and iterate toward a target output (Wang et al., 2024). These capabilities can support writing, analysis, and decision support tasks that are common in public administration, such as drafting policy memos, summarizing public comments, and synthesizing evidence.

Yet a critical constraint on responsible adoption is educational: without structured learning opportunities, employees may lack the competencies needed to use these tools safely and effectively. Without structured learning opportunities, employees may over-trust AI outputs, mishandle sensitive information, or apply tools in ways that conflict with legal or ethical obligations. Research on automation shows that people may misuse or over-rely on automated outputs under time pressure or uncertainty (Parasuraman & Riley, 1997). In addition, large language models may generate fluent but incorrect content (“hallucinations”), requiring learners to develop verification routines and epistemic vigilance (Ji et al., 2023).

*The Educational Review, USA* publishes scholarship spanning educational technology, adult education, educational training, and assessment and evaluation. Within that scope, federal AI upskilling can be studied as a large-scale adult education initiative: a system-wide professional learning effort that must balance instructional quality, equity, security constraints, and measurable outcomes. Accordingly, this paper reframes agentic AI integration as a curriculum, pedagogy, and assessment challenge. The purpose is to provide a practical, education-grounded framework that training units, higher education partners, and educational researchers can use to design and evaluate agentic generative AI professional learning for the U.S. federal workforce.

## 2. Education-Focused Literature Review

### 2.1 AI literacy and digital competence as learning outcomes

AI literacy research describes the knowledge and skills needed to understand AI systems, interpret their outputs, and make informed decisions about use. Long and Magerko (2020) propose core AI literacy competencies that include recognizing AI capabilities and limitations, interpreting outputs, and considering ethical and social impacts. International organizations similarly emphasize competency frameworks that foreground responsible use, privacy, and human agency. For example, UNESCO has issued guidance and competency frameworks intended to help education systems build capacity for safe and meaningful AI use (UNESCO, 2023; UNESCO, 2024). Although those frameworks target formal education, their competency language is transferable to adult professional learning when adapted to workplace roles and risk contexts.

## 2.2 Adult learning and workplace learning

Federal employees are adult learners with diverse prior experience and a strong orientation to job relevance. Andragogy highlights adults' need for self-direction, problem-centered learning, and immediate application (Knowles, 1980). Experiential learning theory further suggests that adults learn through cycles of concrete experience, reflection, conceptualization, and active experimentation (Kolb, 1984). In practice, agentic AI training should privilege authentic tasks, reflective debriefing, and iterative practice rather than lecture-heavy instruction.

Workplace learning also depends on social participation. Communities of practice provide a mechanism for sustaining learning beyond formal courses, supporting peer coaching, shared artifacts (prompt libraries, checklists), and norm-setting around responsible AI use (Wenger, 1998). These social structures are especially important for rapidly evolving tools where "one-and-done" training quickly becomes obsolete.

## 2.3 Instructional design principles for complex digital tools

Because agentic systems combine natural language interaction with multi-step tool use, learners can experience cognitive overload if instruction is not scaffolded. Cognitive load theory shows that instructional design can reduce extraneous load and improve schema acquisition (Sweller, 1988). Merrill's first principles further argue that learning is promoted when learners solve real problems and receive opportunities to activate prior knowledge, observe demonstrations, apply skills, and integrate learning into their contexts (Merrill, 2002). For agentic AI training, this implies a progression from guided exemplars and worked examples to increasingly independent, job-aligned projects.

## 2.4 Training transfer and evaluation

Public-sector training programs are often judged by completion rates rather than performance change. Transfer of training research, however, shows that learning outcomes depend on the alignment of training design, learner characteristics, and workplace supports (Baldwin & Ford, 1988). Evaluation models such as Kirkpatrick's levels remain widely used, but scholars have also called for approaches that examine training benefits in terms of organizational outcomes, equity of access, and long-term sustainability (Bates, 2004). For agentic AI, evaluation should include performance evidence (e.g., accuracy, time-to-completion, documentation quality), safety outcomes (e.g., reduced policy violations), and learners' calibrated trust in automation (Kohn *et al.*, 2021).

## 3. Conceptual Framework and Method

This article is a design-oriented synthesis rather than an empirical intervention study. We integrate research from educational technology, adult learning, workplace training, and AI governance to propose a practical framework for curriculum and program design. The framework is informed by a sociotechnical view that technology adoption and organizational practice co-evolve (Clegg, 2000; Trist & Bamforth, 1951), and by organizational learning perspectives that emphasize iterative adaptation in complex systems (Tosey *et al.*, 2012).

From an educational research perspective, the proposed work aligns with design-based research sensibilities: it identifies a real-world problem (rapid AI diffusion with uneven workforce preparation), synthesizes theory to propose design principles, and offers artifacts (competency model, curriculum blueprint, evaluation plan) that can be tested and refined in future studies.

## 4. Competency Model for Agentic Generative AI Collaboration

To fit the journal's focus on curriculum and teaching, we specify competencies as assessable learning outcomes that can be mapped to instructional activities and evidence. The model below is designed for adult learners and can be adapted across agencies. It distinguishes foundational competencies (expected for all users) from advanced competencies (for technical specialists and leaders). Table 1 shows the Competency Framework for AI Literacy and Agentic Systems.

**Table 1. A Competency-Based Framework for AI Literacy, Governance, and Human–AI Collaboration**

Competency domain	What learners should be able to do (examples)	Sample learning activities	Assessment evidence
AI literacy for agentic systems	Explain how agentic AI differs from standard chat systems; describe limitations such as non-determinism and hallucinations; select appropriate use cases (Wang <i>et al.</i> , 2024; Ji <i>et al.</i> , 2023).	Concept checks; case discussions of successful and failed AI use.	Short-answer quiz; annotated use-case justification.
Information and data stewardship	Apply privacy and records-management rules; classify data sensitivity; use safe prompting and redaction routines (NIST, 2023).	Policy-based scenarios; “what can/can’t go into the prompt” exercises.	Performance task in a sandbox; compliance checklist.
Human-AI collaboration and verification	Design prompts and workflows; verify outputs using authoritative sources; document assumptions; avoid automation bias (Parasuraman & Riley, 1997).	Prompt studios; verification drills; peer review.	Rubric-scored memo with citations; error-detection rate.
Ethical reasoning and fairness	Identify bias risks; conduct basic impact reasoning; communicate limitations transparently (Mehrabian <i>et al.</i> , 2021; UNESCO, 2023).	Structured ethical dilemmas; role-play on citizen-facing communication.	Reflective brief; fairness checklist completion.
Implementation and leadership (advanced)	Specify governance roles; evaluate vendors; set learning supports (communities of practice, coaching); monitor outcomes (Baldwin & Ford, 1988).	Leadership workshops; program planning labs.	Implementation plan; dashboard of evaluation metrics.

The competency model can be differentiated by role. For example, most employees need foundational AI literacy and safe collaboration routines, while technical specialists require deeper capabilities in evaluation, security testing, and system integration. Leaders require skills in learning governance, resource allocation, and sustaining learning transfer.

## 5. Curriculum Design Blueprint

Curriculum design for agentic AI should begin with a needs analysis that treats workplaces as learning environments. Agencies can segment learners into personas (e.g., policy analysts, program managers, customer-service staff, technical teams) and map high-frequency tasks that are candidates for AI augmentation. Instruction can then be organized as a modular pathway that supports just-in-time learning while enabling deeper specialization.

### 5.1 Design principles

Five instructional design principles are recommended:

- 1) Job authenticity: prioritize real tasks and artifacts (memos, briefs, summaries) over abstract exercises (Merrill, 2002).
- 2) Scaffolded complexity: move from worked examples to independent, multi-step agent workflows to manage cognitive load (Sweller, 1988).
- 3) Reflective practice: include structured reflection to calibrate trust, document uncertainty, and surface ethical concerns (Kohn *et al.*, 2021).
- 4) Social learning: embed peer review and communities of practice to sustain learning and share prompt patterns (Wenger, 1998).
- 5) Equity and accessibility: design materials that support diverse learners and reduce barriers to participation (UNESCO, 2023).

### 5.2 Suggested module sequence

This modular training framework is designed to develop comprehensive competencies in generative and agentic AI through a structured, practice-oriented approach. It progresses from foundational knowledge to advanced leadership skills, ensuring that learners not only understand AI concepts but can also apply them effectively and responsibly in real-world contexts.

The framework begins with Module 1, which introduces the core principles of generative and agentic AI, including how such systems operate, their probabilistic nature, and their inherent limitations. This foundational understanding enables learners to critically evaluate when and how AI tools should be used.

Module 2 focuses on safe and responsible AI use by emphasizing information governance. Learners are trained to recognize data sensitivity, apply privacy and records-management rules, and practice secure prompting techniques, ensuring compliance with established standards and policies.

In Module 3, the emphasis shifts to human–AI collaboration. Learners develop practical skills in designing prompts, refining outputs iteratively, and verifying results using reliable sources. This module promotes critical thinking and reduces risks such as automation bias by reinforcing validation and documentation practices.

Module 4 addresses ethical considerations, including bias, fairness, and transparency. Learners engage in applied ethical reasoning and are trained to communicate AI limitations clearly, fostering responsible and trustworthy use of AI systems in diverse contexts.

Module 5 provides hands-on, role-specific learning through labs and a capstone project. Learners work on real or realistic scenarios relevant to their roles, producing AI-assisted outputs along with audit trails and reflective analyses. This ensures that learning outcomes translate into tangible organizational value.

Finally, Module 6 (advanced track) targets leadership and governance. It equips learners with the skills to design AI implementation strategies, evaluate vendors, monitor performance, and support organizational change through structured learning and governance mechanisms. Table 2 shows the AI Training Program: Modules, Practice Tasks, and Learning Duration.

**Table 2. Modular Training Framework for Generative and Agentic AI: Topics, Practice, and Duration**

Module	Core topics	Practice tasks	Typical duration
M1. Foundations of generative and agentic AI	How agentic AI works; probabilistic outputs; limitations; appropriate use cases.	Classify tasks as suitable/unsuitable; critique example outputs.	2-3 hours
M2. Safe prompting and information governance	Data sensitivity; privacy; records; prompt hygiene; secure tool use (NIST, 2023).	Rewrite prompts to remove sensitive data; apply redaction rules.	2-3 hours
M3. Human-AI collaboration workflows	Prompting patterns; iterative refinement; tool use; documentation; verification routines.	Draft and verify a short policy brief with citations; peer review.	4-6 hours
M4. Ethics, bias, and transparency in practice	Bias and fairness; explainability; communicating limitations (Mehrabi <i>et al.</i> , 2021).	Complete a mini impact assessment for a proposed AI use case.	3-4 hours
M5. Role-specific labs and capstone project	Agency-relevant scenarios; agent workflows; evaluation and reporting.	Capstone: deliver an AI-assisted work product plus an audit trail and reflection.	6-10 hours
M6. Leadership and governance (advanced track)	Program governance; vendor evaluation; monitoring; change management; learning supports.	Create a rollout plan with metrics and risk controls.	4-6 hours

A blended model is recommended: short asynchronous lessons (concepts and policies) paired with synchronous workshops where learners practice in secure sandboxes and receive coaching. Capstone projects should be selected from real work queues so that learning produces immediate organizational value.

## 6. Pedagogical Strategies and Learning Environment

### 6.1 Secure sandboxes and simulation-based learning

Because federal work often involves sensitive information, training must provide secure environments that allow experimentation without risking data exposure. Simulations and de-identified datasets (consistent with federal data classification standards, e.g., Controlled Unclassified Information) can support practice on realistic cases without risking exposure of sensitive information. Simulation-based approaches also make it easier to assess competence through performance tasks rather than self-report.

## 6.2 Feedback, coaching, and communities of practice

Learning should not end with course completion. Agencies can establish communities of practice, office hours, and peer review structures to support ongoing improvement (Wenger, 1998). Coaching helps learners build confidence and reduces the risk of unsafe improvisation. Self-efficacy theory suggests that mastery experiences and credible feedback can improve learners' willingness to engage with complex tools while maintaining appropriate caution (Bandura, 1997).

## 6.3 Managing trust and avoiding automation bias

A persistent educational goal is calibrated trust. Learners must avoid both over-trust (accepting outputs uncritically) and under-trust (rejecting tools even when they add value). Training should include deliberate practice in identifying errors, cross-checking claims, and documenting uncertainty (Parasuraman & Riley, 1997; Kohn *et al.*, 2021). Structured checklists and peer review can reduce automation bias and normalize verification as a professional habit.

## 7. Institutional Implementation in Large Organizations

Implementing agentic AI education at scale requires attention to instructional capacity, governance, and organizational culture. Prior public-sector AI research notes that adoption is shaped by institutional readiness, data capacity, and risk management (Wirtz *et al.*, 2019; Zuiderwijk *et al.*, 2021). From an educational administration perspective, three implementation features are essential.

First, agencies should adopt a phased rollout: pilot cohorts to test the curriculum and assessment, followed by broader deployment through a centralized learning management system and local facilitators. Second, a train-the-trainer model can develop internal instructional capacity, ensuring that facilitators can teach both tool use and the underlying reasoning habits (verification, documentation, ethical reflection). Third, leadership should align training with performance expectations by defining when and how AI use is appropriate and by rewarding quality, not just speed.

## 8. Assessment and Program Evaluation

Evaluation should prioritize evidence of competence and transfer to work. We recommend a multi-level evaluation plan that combines learner outcomes, workplace performance indicators, and governance metrics. This evaluation framework is designed to measure the effectiveness of the AI training program across multiple levels, ensuring that learning translates into meaningful, safe, and sustainable outcomes. It adopts a multi-level approach, moving from individual learning gains to organizational impact and long-term system sustainability.

At the learning outcomes (course level), the focus is on assessing whether participants have acquired the intended knowledge and skills. Key indicators include understanding AI limitations, the ability to craft high-quality prompts, accuracy in verifying AI-generated outputs, and sound ethical reasoning. These outcomes are typically measured using pre- and post-assessments, along with rubric-based evaluation of performance tasks, providing clear evidence of learning progression.

The transfer to work (job level) evaluates how effectively learners apply their skills in real workplace settings. Indicators such as reduced rework, improved quality of documents, faster task completion without compromising accuracy, and proper documentation (including citations and audit trails) reflect practical impact. Data is gathered through supervisor evaluations, audits of work samples, and workflow analytics, ensuring that training leads to tangible productivity and quality improvements.

At the safety and compliance (governance level), the framework assesses whether AI use aligns with organizational policies and regulatory standards. Indicators include a reduction in policy violations, improved data-handling practices, and clear evidence of human oversight in AI-assisted processes. These are monitored באמצעות incident reports, compliance reviews, and system audit logs, reinforcing accountability and responsible AI use.

Finally, the sustainability (system-level) dimension examines the long-term viability of the training initiative. It looks at continued engagement in communities of practice, participation in ongoing microlearning activities, and regular updates to the curriculum to keep pace with evolving AI technologies. Data sources such as learning platform analytics and community engagement metrics help determine whether the program remains relevant and continuously improves over time. Table 3 shows the evaluation Framework for AI Competency Development and Organizational Impact.

**Table 3. Multi-Level Evaluation Framework for AI Training Effectiveness**

Evaluation focus	Example indicators	Data sources
Learning outcomes (course level)	Knowledge of limitations; quality of prompts; verification accuracy; ethical reasoning.	Pre/post assessments; rubric-scored performance tasks.
Transfer to work (job level)	Reduced rework; improved document quality; faster completion with maintained accuracy; documented citations and audit trails.	Supervisor ratings; work-sample audits; workflow analytics.
Safety and compliance (governance level)	Fewer policy violations; improved data-handling compliance; evidence of human oversight.	Incident reports; compliance checks; audit logs.
Sustainability (system level)	Participation in communities of practice; ongoing microlearning uptake; curriculum updates.	Learning platform analytics; community activity metrics.

In addition to outcome metrics, evaluation should examine equity of access and participation. Agencies should track whether learners across grades, locations, and job families can access practice environments and receive coaching. Findings can inform iterative improvement cycles, consistent with organizational learning approaches (Tosey et al., 2012).

## 9. Discussion: Implications for Educational Research and Practice

This framework positions public-sector agentic AI training as a significant adult education and educational technology initiative. It contributes to education research in three ways. First, it translates AI literacy concepts into role-based, assessable outcomes for workplace learners (Long & Magerko, 2020). Second, it operationalizes adult learning theory through scenario-based practice, reflection, and community supports (Knowles, 1980; Kolb, 1984; Wenger, 1998). Third, it emphasizes evaluation designs that go beyond completion rates by focusing on transfer, safety, and quality.

Future research should test the framework empirically through quasi-experimental or design-based studies across agencies, comparing instructional approaches (e.g., workshop-heavy vs. fully online) and examining how learning supports (coaching, communities of practice) influence transfer. Research is also needed on assessment validity for AI-assisted work products: what constitutes “good work” when AI contributes to drafting, and how can educators distinguish productive augmentation from unsafe delegation?

## 10. Conclusion

Agentic generative AI adoption in government is fundamentally a learning challenge. To realize benefits while preserving accountability, agencies need education-centered professional learning programs that build AI literacy, safe collaboration habits, and measurable transfer to work. The framework proposed here provides a competency model, curriculum blueprint, pedagogical strategies, and evaluation plan aligned with adult learning and educational technology research. By treating agentic AI capability building as lifelong learning, the federal workforce can develop trustworthy practices that support effective public service in an AI-rich environment.

## Funding

The authors declare no external funding for this manuscript.

## References

- Baldwin, T. T., & Ford, J. K. (1988). Transfer of training: A review and directions for future research. *Personnel Psychology*, *41*(1), 63-105. <https://doi.org/10.1111/j.1744-6570.1988.tb00632.x>
- Bandura, A. (1997). Self-efficacy: The exercise of control. W. H. Freeman.
- Bates, R. (2004). A critical analysis of evaluation practice: The Kirkpatrick model and the principle of beneficence. *Evaluation and Program Planning*, *27*(3), 341-347. <https://doi.org/10.1016/j.evalprogplan.2004.04.011>
- Clegg, C. W. (2000). Sociotechnical principles for system design. *Applied Ergonomics*, *31*(5), 463-477.

[https://doi.org/10.1016/S0003-6870\(00\)00009-0](https://doi.org/10.1016/S0003-6870(00)00009-0)

- Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12), Article 248. <https://doi.org/10.1145/3571730>
- Knowles, M. S. (1980). *The modern practice of adult education: From pedagogy to andragogy* (2nd ed.). Association Press.
- Kohn, S. C., de Visser, E. J., Wiese, E., Lee, Y.-C., & Shaw, T. H. (2021). Measurement of trust in automation: A narrative review and reference guide. *Frontiers in Psychology*, 12, 604977. <https://doi.org/10.3389/fpsyg.2021.604977>
- Kolb, D. A. (1984). *Experiential learning: Experience as the source of learning and development*. Prentice-Hall.
- Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1-16). Association for Computing Machinery. <https://doi.org/10.1145/3313831.3376727>
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 115. <https://doi.org/10.1145/3457607>
- Merrill, M. D. (2002). First principles of instruction. *Educational Technology Research and Development*, 50(3), 43-59. <https://doi.org/10.1007/BF02505024>
- NIST. (2023). Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-1>
- Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 39(2), 230-253. <https://doi.org/10.1518/001872097778543886>
- Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science*, 12(2), 257-285. [https://doi.org/10.1016/0364-0213\(88\)90023-7](https://doi.org/10.1016/0364-0213(88)90023-7)
- Tosey, P., Visser, M., & Saunders, M. N. K. (2012). The origins and conceptualisations of 'triple-loop' learning: A critical review. *Management Learning*, 43(3), 289-305. <https://doi.org/10.1177/1350507611426239>
- Trist, E. L., & Bamforth, K. W. (1951). Some social and psychological consequences of the longwall method of coal-getting. *Human Relations*, 4(1), 3-38. <https://doi.org/10.1177/001872675100400101>
- UNESCO. (2023). Guidance for generative AI in education and research. United Nations Educational, Scientific and Cultural Organization. <https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research>
- UNESCO. (2024). AI competency framework for students. United Nations Educational, Scientific and Cultural Organization. <https://www.unesco.org/en/articles/ai-competency-framework-students>
- Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., Zhao, W. X., Wei, Z., & Wen, J.-R. (2024). A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6), 186345. <https://doi.org/10.1007/s11704-024-40231-1>
- Wenger, E. (1998). *Communities of practice: Learning, meaning, and identity*. Cambridge University Press.
- Wirtz, B. W., Weyerer, J. C., & Geyer, C. (2019). Artificial intelligence and the public sector: Applications and challenges. *International Journal of Public Administration*, 42(7), 596-615. <https://doi.org/10.1080/01900692.2018.1498103>
- Zuiderwijk, A., Chen, Y.-C., Salem, F., & Harahap, N. C. (2021). Artificial intelligence and the public sector: Results from a systematic literature review. *Government Information Quarterly*, 38(3), 101577. <https://doi.org/10.1016/j.giq.2021.101577>