SECURITY JOURNEY

AI Literacy and Workplace Performance:

The Case for Structured Training



AI Literacy and Workplace Performance: The Case for Structured Training

Michael Burch
Security Journey

A narrative evidence review

Abstract. Generative artificial intelligence can improve the speed and quality of work, creating a substantial opportunity for organizations that develop their employees' ability to use it effectively. To realize that opportunity, employees must choose suitable tasks, give clear instructions, check outputs, and decide how to use the results. This paper examines the case for investing in workplace AI training. It brings together workplace experiments, studies of AI instruction and decisions about trusting AI, and research on training and assessment. The evidence identifies valuable uses of AI and recurring difficulties in judging its outputs. Studies that test instruction show that targeted training can improve AI use and skills for specific tasks. Broader training research supports connecting instruction with practice, feedback, and support at work. Together, these findings support quality instruction and relevant practice to develop the understanding and judgment needed for effective AI use. Employers should plan training as part of introducing AI, connect it to employees' responsibilities, and assess what employees can do and how their work changes.

Keywords: AI literacy; workplace training; generative AI; human–AI collaboration; training transfer; performance assessment.

1 Introduction

The business case for workplace AI creates a corresponding case for investing in the people who use it. Organizations adopting generative AI are asking employees to choose suitable tasks, communicate requirements, assess unfamiliar outputs, and decide what to implement. Structured training gives employers a way to develop these capabilities deliberately and connect them to expectations for quality, productivity, and responsible use.

The scale of adoption makes this an immediate workforce issue. In Microsoft's and LinkedIn's 2024 survey of 31,000 knowledge workers, 75% of respondents reported using AI at work, while 39% of AI users reported company-provided training [1]. BCG's 2025 survey found regular use among 51% of frontline employees, while about one third of its broader respondent group reported being properly trained [2]. In the New York Fed's November 2025 survey, 39% of employed respondents reported using AI in their job during the previous year and 15.9% reported that their employer offered AI training [3]. These surveys use different populations and measures; together, they indicate substantial use alongside uneven training support.

The opportunity is already visible in performance research. An AI assistant increased issues resolved per hour by about 15% in a customer-support deployment [4]. In an experiment with 453 professionals, ChatGPT reduced the time required for writing assignments by about 40% and improved rated quality by 18% [5]. Results such as these establish the value of developing effective AI use. They also make it important to understand why performance varies across tasks and how employees can learn to recognize those differences.

This paper argues that organizations should invest in quality instruction, supported by hands-on practice, as a core component of workplace AI deployment. It brings together research identifying the capabilities required for effective AI use, studies testing training approaches, and broader evidence on how people apply their learning at work. The review explains why training matters, what judgments employees need to learn, and how practice and assessment can support that learning.

1.1 AI literacy and the scope of training

AI literacy concerns the ability to understand, critically evaluate, and work effectively with AI. Long and Magerko's framework includes skills that go beyond operating a tool [6]. For generative AI, Tankelevitch et al. describe the demands involved in setting goals, writing prompts, checking outputs, and deciding how to use the system in a work process [7]. This perspective connects AI literacy to the judgments employees make while doing their work.

Workplace AI training is understood here as structured instruction in the knowledge, skills, and judgment required for effective AI use, supported by hands-on practice. Instruction explains relevant concepts and approaches; labs and applied exercises give learners opportunities to use them, examine results, and improve their work.

2 Review Approach and Scope

This narrative review examines what existing evidence tells employers about the value and design of workplace AI training. It covers workplace AI performance, instruction, judgments about AI outputs, the application of learned skills at work, and assessment. Sources were selected for their relevance to these topics. An October 2026 supplemental web search broadened coverage of instructional studies, including those reporting benefits, no clear effects, or harm. Workplace studies were prioritized; student and specialist studies were included where they informed relevant skills or teaching approaches, with their populations identified. Broader training reviews and authoritative security guidance supplied additional context. The review provides a selective synthesis rather than an exhaustive search or a statistical pooling of study results.

Studies were assessed in light of their design, participants, comparison groups, outcomes, and reported limitations.

Claims were checked against available primary texts and official publication records. Published articles were preferred, and preprints and author manuscripts are identified where they supply the reported evidence. Survey findings describe organizational conditions; studies of AI access and reliance identify performance opportunities and risks. Instructional studies inform the effects of the training tested, while broader training research supports recommendations about practice, feedback, and workplace application. Security studies and guidance identify relevant responsibilities and review practices. These distinctions guide the synthesis: improvements from AI access are not treated as estimates of training effectiveness, and recommendations informed by general learning research are distinguished from findings of AI-specific instructional studies.

3 Why Training Matters for Workplace AI

3.1 From available capability to effective use

The organizational value of AI depends on how its capabilities are applied. Dell'Acqua et al.'s consultant experiment illustrates this relationship [8]. In the published article's main statistical analysis, human ratings of quality on tasks within GPT-4's capabilities were 29.9% higher with access and 33.9% higher with access plus a brief overview, compared with the average in the group without AI. On a separate task beyond those capabilities, correctness fell from 84.5% without AI to 70.6% with access and 60.0% with access plus the overview. The overview group therefore scored 10.6 percentage points below the AI-only group. The direct comparison met a 10% threshold for statistical significance but not the stricter 5% threshold ($p = .082$–$.095$ across statistical models). This is suggestive evidence that the particular overview could worsen decisions on that task, alongside its positive contribution to quality on tasks within the system's capabilities.

The instructional lesson is that orientation must be evaluated for the decisions it changes. The study establishes neither that more instruction would have prevented the error nor that training generally harms performance. It motivates teaching and testing task selection and output evaluation as explicit capabilities. Tool orientation can be part of this preparation, but a program also needs to address the decisions surrounding the tool. An employee who can generate a persuasive answer still needs a way to determine whether that answer meets the task's requirements.

Other workplace studies reinforce the importance of interpreting performance in context. In the Procter & Gamble experiment, 776 professionals worked on product-innovation tasks, and AI-enabled participants received an hour of instruction [9]. The study tests AI use together with initial instruction, so it cannot show how much of the result came from instruction alone. Cui et al.'s developer study also requires care in interpretation [10]. Its June 2025 manuscript reports a 26.08% gain using an instrumental-variable estimate. This method estimates the effect of using AI for developers whose use was changed by the experimental offer of access. The figure is not the average gain from offering AI to every developer. Training programs should use such findings to identify relevant opportunities and then define the performance expected in their own work setting.

3.2 Judgment is an instructional objective

Research on reliance explains why judgment belongs in the curriculum. Bansal et al. found that explanations could increase acceptance of AI advice whether that advice was correct or incorrect [11]. Vasconcelos et al. show that explanations can reduce excessive trust in AI when they make its advice easier to check [12]. Fok and Weld emphasize the broader difficulty: many AI-assisted decisions lack an easy independent check [13]. Employees therefore need to learn how to evaluate an output, what evidence would justify using it, and when uncertainty requires review.

Appropriate reliance has two sides. Employees should make effective use of sound advice and challenge flawed advice. In model-literacy research, interactive instruction about a particular AI model benefited some participants in a house-price prediction task [14]. Another study found that training worsened decisions about relying on AI among participants who initially underestimated their ability. This result was statistically significant. The improvement among those who initially overestimated their ability was not statistically significant [15]. These findings support instruction that develops sound judgment and assesses both kinds of error: accepting a flawed answer and rejecting a useful one.

This is a substantive reason to invest in training. A curriculum can make the reasoning behind these decisions explicit, give learners examples to examine, and provide opportunities to demonstrate what they would do. The intended outcome is usable judgment: selecting a suitable application, recognizing a questionable result, and taking an appropriate next step.

3.3 Training and organizational support

AI literacy also has an organizational dimension. Tambe documents the growing importance of algorithmic skills—working with computational methods—in occupations requiring specialist subject knowledge [16]. This places capability development across business functions, where employees combine subject knowledge with decisions about AI use. A Danish study similarly finds higher reported adoption in workplaces combining tools, training, and encouragement [17]. These comparisons describe an association; they do not establish that training caused higher use. They support considering training alongside the other conditions that shape AI use at work.

Interviews with knowledge workers identify learning needs concerning output evaluation, data protection, and the practical use of AI [18]. Workplace research also describes informal learning, uncertainty about expectations, and concealment of AI use [19]. Structured training offers employers a means of communicating shared expectations and making learning opportunities accessible across roles. Managers can reinforce that work by clarifying acceptable uses and providing opportunities to apply newly developed skills.

Table 1. Selected instructional studies and their implications for workplace training.

Study and setting	Instruction and observed result	Meaning and limits
Chen and Bao [20]; 164 law students	Brief video and quiz increased adoption and improved grades relative to untrained AI access.	Supports task-specific instruction, including brief formats. One mock exam; one-sided tests; application at work remains to be studied.
Ma et al. [21]; 30 participants	Requirement-focused practice and feedback improved prompting and output measures relative to tutorial and free practice.	Supports teaching learners to specify requirements and assess results. Small sample and a narrow task setting.
Xiao et al. [22]; 431 participants with complete data	Four instructional formats produced different prompting-skill gains, with some retention at seven weeks.	Supports checking skills and whether learning lasts. Many participants did not finish; all groups received instruction; training time differed. The best duration remains unclear.
Puppart and Aru [23]; 36 students	A brief explanatory text increased rejection of correct answers without significantly reducing acceptance of incorrect answers.	Highlights the need to teach and assess the ability to distinguish useful from flawed outputs. A small, short student task.

4 What the Instructional Evidence Shows

4.1 Targeted instruction can change behavior and skill

Direct studies provide evidence that specific forms of AI instruction can improve specific outcomes. In a working paper involving 164 law students, a nine-and-a-half-minute video and quiz increased AI adoption from 26.32% to 41.38% and improved mock-examination grades relative to untrained access [20]. The adoption and grade comparisons used one-sided statistical tests, which test for an effect in a specified direction. The trained group did not perform statistically significantly better than the group without AI. Within that setting, the finding demonstrates the potential for brief, relevant instruction to change how learners use an available tool.

Ma et al.'s study of 30 participants compared requirement-focused prompting practice and feedback with a tutorial followed by free practice [21]. The results favor making task requirements explicit in instruction and connect training to observable changes in prompts and outputs. Both conditions included practice and used equal total session time in the consulted author manuscript. The comparison supports the value of the instructional structure and feedback, rather than isolating the effect of hands-on activity alone.

A larger student study compared four prompting-instruction formats and found differences in skill gains, with some retention at seven weeks [22]. Its analysis included 431 of the 979 people who started. Training time also differed across formats. These features limit conclusions about the best delivery method or duration. Taken together, the studies support a practical proposition: relevant instruction can influence AI use and task-specific capability, and programs should connect their learning activities to clearly defined outcomes.

A recent preprint, a paper not yet peer reviewed, adds evidence from working professionals. Kamali et al. studied 32 intelligence analysts receiving 30 minutes of image-literacy instruction and reported a 9-percentage-point improvement in image-classification accuracy (95% confidence interval: 2.7–15.4 percentage points) [24]. Participants saw different image sets before and after training, with the order reversed for the other group. This counterbalanced design helps account for differences in image difficulty. All participants received training at the midpoint, so practice or the passage of time could also have contributed. The immediate improvement is encouraging; whether it lasts and carries over to work remains to be assessed.

4.2 Instructional quality matters

Variation in outcomes gives program designers useful information about what to improve. A preprint involving 36 high-school students found that a brief explanatory text increased rejection of correct ChatGPT answers without significantly reducing acceptance of incorrect ones [23]. A corporate workshop study found lower completion and quality when required to follow a set process for working together with AI [25]. The groups attended at different times of day, which could have affected their results. That difference and the workshop setting limit what the study tells us about everyday work. Both findings make the same design issue salient: training should help learners perform the intended decisions, and the resulting behavior should be assessed.

Table 1 summarizes selected instructional evidence. Its practical implication is to align content, activity, and assessment. If the objective is to evaluate an output, learners need to encounter outputs of differing quality and explain their judgments. If the objective is to formulate useful instructions, they need to express requirements and inspect whether those requirements were met. This alignment gives an organization a basis for improving its program when a particular activity falls short.

4.3 A foundation in training research

The broader training literature provides a substantial foundation for these choices. Arthur et al. combined results from organizational training studies in a meta-analysis. They reported average effects of $d = 0.63$ for learning and $d = 0.62$ for behavior, weighted by sample size [26]. These standardized effect sizes express differences on a common statistical

Table 2. Instructional priorities and examples of observable behavior and assessment evidence.

Instructional priority	Observable behavior	Assessment evidence
Select appropriate uses	Identifies where the available AI system fits the task and recognizes when another method or specialist is needed.	Explains the choice of tool and approach for representative cases.
Specify and refine the work	Provides relevant context and requirements, inspects the result, and revises the approach when needed.	Produces instructions and outputs that meet stated assessment criteria.
Evaluate outputs	Checks material claims or results using suitable sources, tests, assessment criteria, or qualified review.	Identifies errors, supports corrections, and explains unresolved uncertainty.
Judge when to rely on AI	Uses sound advice, challenges flawed advice, and adjusts confidence to the available evidence.	Responds appropriately to a mixture of correct and incorrect suggestions.
Apply organizational boundaries	Uses approved tools and data and recognizes disclosure, review, and authorization requirements.	Makes appropriate decisions in scenarios involving data and actions with significant consequences.

scale. They describe general training outcomes, not percentages of improvement in AI productivity or financial returns. Keith and Frese's meta-analysis finds benefits from training that teaches people to learn from errors, including better performance on new tasks, largely in software and laboratory settings [27]. Leadership-training research links assessing learning needs, feedback, practice, and spreading learning over time with stronger outcomes [28]. Research on applying learned skills at work links the work environment, including support and opportunities to apply learning, with the use of trained skills [29].

The recommendation to support instruction with labs and applied practice draws on this broader training literature and the narrower AI-specific evidence on structured practice and feedback [21, 27, 28]. Explanations and demonstrations can make the reasoning behind effective use visible. Supporting labs can ask learners to apply that reasoning, encounter mistakes, and revise their approach with feedback. This connection gives practice an instructional purpose and gives learners a way to use what they have been taught. The scope and pace of training should reflect the required capability and evidence of learning.

5 Implications for Workplace AI Instruction

The reviewed evidence points to recurring instructional priorities in the decisions employees make before, during, and after AI use. Table 2 connects these priorities to observable behaviors and examples of assessment evidence. Their application should reflect the tasks and responsibilities of each role, so that employees can connect what they learn to the work they are expected to perform.

5.1 Connect instruction with applied practice

Training should begin with representative tasks and the capability gaps that affect them. A shared foundation can introduce AI concepts and organizational expectations, while role-relevant examples show how those principles apply to actual decisions. Prioritization should consider the value of the work, the consequences of error, and demonstrated learning needs. Experience with a job or with an AI interface can inform this assessment, but neither substitutes for examining the relevant capability.

Instruction should explain why an approach is suitable, what a satisfactory result requires, and how to recognize an error [7, 21]. A supporting lab can then ask learners to apply those principles to a representative task and revise their work against stated requirements. For example, instruction in source checking can be followed by an exercise requiring learners to identify and correct unsupported statements in an AI-generated summary. Feedback should address the result and the reasoning behind it. This alignment gives employees a clear connection between the lesson and its application.

5.2 Teach evaluation and security-relevant application

The depth of checking should match the task and the consequences of error. A factual summary may require comparison with the source; a calculation may require an independent check; a recommendation with significant consequences may require qualified review. Some tasks have a clear answer key, while others require a defensible judgment under uncertainty. Training should make that difference explicit and teach the appropriate response.

Organizational boundaries belong in this same applied context. Employees need to know which tools and data are approved, what they may disclose, who reviews an output, and when an action requires authorization. Practice should use approved materials and environments. These learning objectives connect productive use with the controls and responsibilities already governing the work.

Security makes these decisions concrete. NIST's generative AI profile identifies privacy and information-security risks, including indirect prompt injection: instructions hidden in content an AI system retrieves that try to redirect its behavior [30]. Instruction can connect these risks to employee decisions: selecting information appropriate for the approved tool, distinguishing source content from authorized instructions, and checking the destination and scope of a requested action. An exercise might ask a learner to identify an unnecessary disclosure in a draft response or refer an unexpected

request for access to the appropriate reviewer. These are examples of applying the learning objectives, rather than claims of measured incident reduction.

For employees producing code or automations, review must include security as well as whether the output works as intended. Perry et al. found that participants using a particular AI coding assistant produced less secure code overall and were more likely to believe their code was secure [31]. The result concerns the studied assistant and tasks; its instructional relevance is the gap between believing code is secure and showing that it is secure. Role-specific practice can require inspection, testing, and review of an output before use.

Training should operate alongside technical controls: restricted permissions, approved data access, and review of actions with significant consequences. Employees need to understand when to stop and seek help, while system owners remain responsible for controls that constrain what the tool can do. This connects secure use to the same judgment and evaluation skills needed for productive work.

5.3 Support application after instruction

The work environment helps determine whether a learned capability is used. Research on applying learned skills supports providing opportunities to use them and support from supervisors [29]. Managers can clarify where AI is appropriate, make time for relevant learning, and discuss examples of successful use and avoidable error. These activities reinforce the expectations taught in training and help employees apply them to local workflows.

Learning should also develop with employees' responsibilities. Foundational instruction can establish shared understanding, with more specialized teaching and practice as employees take on complex tasks or support colleagues. When a tool, policy, or workflow changes, the relevant examples and assessments should be reviewed. Continuing instruction keeps capability development connected to the work employees are doing.

6 Evaluating Training and Workplace Value

Evaluation should connect participation to learning and learning to work performance. Completion records establish who received instruction. Knowledge questions, applied tasks, and workplace outcomes establish different aspects of what changed. Satisfaction and confidence are useful for understanding the learner's experience, while broader training research shows that liking a course is a weak indicator of how much someone learned [32].

6.1 Measure what employees can do

Research on assessment shows why we should distinguish what people believe they can do from what they demonstrate. In a study validating the AICOS test with 514 German-speaking adults, test scores had almost no relationship to self-ratings on the short MAILS questionnaire ($r = .04$) [33]. A study of teachers also found weak relationships between self-ratings and objective assessments [34]. These findings support using both kinds of measure for their intended purposes, with performance assessment central when the question concerns what someone can do.

A 2024 review identified 16 AI literacy scales, of which 13 were self-reports and three had evidence that they could detect changes in AI literacy [35]. When choosing an assessment, organizations should consider its language, who it was tested with, its relevance to the task, and whether it can detect the intended learning. They should also check for newer versions published since the review.

An assessment program should cover three areas. First, objective questions can examine relevant knowledge and task understanding. Second, cases containing both sound and flawed AI advice can test evaluation and reliance decisions. Third, representative work products can be assessed for quality, errors, completion, and the total time required for production, checking, and correction. Assessments using new cases after a suitable delay can examine retention and application beyond the immediate lesson.

6.2 Distinguish assisted performance from independent knowledge

Both tool-assisted performance and independent knowledge can matter, depending on the role. In the Wiles et al. author manuscript, consultants with AI access and onboarding performed better on unfamiliar technical tasks but showed no detectable advantage on five knowledge questions answered immediately afterward without AI [36]. The customer-support study, by contrast, interprets performance during outages as evidence of learning through prior AI use [4]. These findings show why the outcome must be specified: completing work with assistance, learning a concept, and retaining the ability to perform independently are related but distinct objectives.

Training evaluation should reflect that distinction. A role that requires oversight or continuity without the tool needs an independent-performance measure. A role primarily using AI to produce a reviewed output needs an assessment of that assisted workflow. Matching the measure to the responsibility makes the evidence useful to both learners and the organization.

The assessment and training-transfer research reviewed here supports two complementary priorities: measuring demonstrated capability and supporting its application at work [29, 33]. For workplace AI training, the practical implication is to align assessment with learning objectives and use the results to identify where instruction, practice, or workplace support needs improvement.

7 Discussion and Limitations

Several findings together support the case for training. Workplace experiments establish valuable opportunities for AI use and show that outcomes depend on the task and the way the system is used. Research on trust in AI identifies decisions that employees repeatedly face. Instructional studies show that targeted teaching can change adoption and task-specific skills. Broader training research supplies a foundation for

connecting instruction with practice and workplace application. Together, these findings support deliberate investment in teaching employees how to use AI effectively.

The practical implication is to fund instruction against defined capability needs, connect it to relevant practice, and evaluate application at work. Table 2 and Section 6 translate that recommendation into observable expectations.

The evidence also identifies questions for further research. Studies of AI training differ in their methods and measures. Many involve students or short tasks, and some are working papers that have not yet been peer reviewed. In the law-student study, one participant's group classification changed after assignment. Many participants in the prompting-course study did not complete it. Broader training meta-analyses cover different technologies and occupations. Links between training design and outcomes do not establish how much each design choice caused those outcomes. Future studies should examine sustained workplace performance, retention, and differences across roles, including how instruction and supporting practice contribute to those outcomes.

Training operates alongside tool and work design. Some performance difficulties call for a better interface, clearer requirements, stronger review, or a different use case. An effective deployment program should address these conditions together. Training helps an organization develop effective AI use. Evaluation helps determine where instruction and other changes can contribute most.

8 Conclusion

AI literacy training deserves a central place in workplace AI strategy. The opportunity created by generative AI is substantial, and employees need the knowledge and judgment to apply it to the work for which they are responsible. The reviewed evidence supports quality instruction focused on those decisions, with hands-on labs and relevant practice helping employees apply and strengthen what they learn.

Employers should invest in structured, role-relevant training and assess the resulting capability and work performance. Training providers should design and maintain instruction and practice around those outcomes. Continued learning and workplace support make this investment part of how the organization develops its people as AI use evolves. The purpose of training is to help employees turn available AI capabilities into effective, responsible work.

Declaration of Interest

Security Journey provides AI and cybersecurity training and has a commercial interest in workforce education.

References

[1] Microsoft and LinkedIn. AI at Work Is Here. Now Comes the Hard Part. 2024 Work Trend Index Annual Report, 2024. URL https://www.microsoft.com/en-us/worklab/work-trend-index/ai-at-work-is-here-now-comes-the-hard-part. Published 8 May 2024.

[2] V. Beauchene, S. Duranton, N. Kalra, and D. Martin. AI at Work 2025: Momentum Builds, but Gaps Remain. Boston Consulting Group, 2025. URL https://www.bcg.com/publications/2025/ai-at-work-momentum-builds-but-gaps-remain. Published 26 June 2025.

[3] A. Hashim, G. Kosar, and W. van der Klaauw. Use of Gen AI in the Workplace and the Value of Access to Training. Federal Reserve Bank of New York, Liberty Street Economics, 2026. URL https://libertystreeteconomics.newyorkfed.org/2026/04/use-of-gen-ai-in-the-workplace-and-the-value-of-access-to-training/. 14 April 2026; survey conducted November 2025.

[4] E. Brynjolfsson, D. Li, and L. Raymond. Generative AI at Work. *The Quarterly Journal of Economics*, 140(2):889–942, 2025. doi: 10.1093/qje/qjae044.

[5] S. Noy and W. Zhang. Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654):187–192, 2023. doi: 10.1126/science.adh2586.

[6] D. Long and B. Magerko. What is AI Literacy? Competencies and Design Considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–16, 2020. doi: 10.1145/3313831.3376727.

[7] L. Tankelevitch, V. Kewenig, A. Simkute, A. E. Scott, A. Sarkar, A. Sellen, and S. Rintel. The Metacognitive Demands and Opportunities of Generative AI. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 1–24, 2024. doi: 10.1145/3613904.3642902.

[8] F. Dell'Acqua, E. McFowland, III, E. Mollick, H. Lifshitz, K. C. Kellogg, S. Rajendran, L. Krayer, F. Candelon, and K. R. Lakhani. Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality. *Organization Science*, 37(2):403–423, 2026. doi: 10.1287/orsc.2025.21838.

[9] F. Dell'Acqua, C. Ayoubi, H. Lifshitz, R. Sadun, E. Mollick, L. Mollick, Y. Han, J. Goldman, H. Nair, S. Taub, and K. R. Lakhani. The Cybernetic Teammate: A Field Experiment on Generative AI and Teamwork. *Organization Science*, 37(4):1217–1242, 2026. doi: 10.1287/orsc.2025.20702. Detailed intervention design checked in NBER Working Paper 33641 (2025).

[10] K. Z. Cui, M. Demirer, S. Jaffe, L. Musolff, S. Peng, and T. Salz. The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers. *Management Science*, 2026. doi: 10.1287/mnsc.2025.00535. Articles in Advance, 27 February 2026. Detailed estimate checked in the June 2025 author manuscript.

[11] G. Bansal, T. Wu, J. Zhou, R. Fok, B. Nushi, E. Kamar, M. T. Ribeiro, and D. Weld. 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*, pages 1–16, 2021. doi: 10.1145/3411764.3445717.

[12] H. Vasconcelos, M. Jörke, M. Grunde-McLaughlin, T. Gerstenberg, M. S. Bernstein, and R. Krishna. Explanations Can Reduce Overreliance on AI Systems During Decision-Making. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW1):1–38, 2023. doi: 10.1145/3579605.

[13] R. Fok and D. S. Weld. In search of verifiability: Explanations rarely enable complementary performance in AI-advised decision making. *AI Magazine*, 45(3):317–332, 2024. doi: 10.1002/aaai.12182.

[14] C. W. Chiang and M. Yin. Exploring the Effects of Machine Learning Literacy Interventions on Laypeople's Reliance on Machine Learning Models. In *27th International Conference on Intelligent User Interfaces*, pages 148–161, 2022. doi: 10.1145/3490099.3511121.

[15] G. He, L. Kuiper, and U. Gadiraju. Knowing About Knowing: An Illusion of Human Competence Can Hinder Appropriate Reliance on AI Systems. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–18, 2023. doi: 10.1145/3544548.3581025.

[16] P. B. Tambe. Reskilling the Workforce for AI: Domain Expertise and Algorithmic Literacy. *Management Science*, 72(1):515–537, 2026. doi: 10.1287/mnsc.2022.03968.

[17] A. Humlum and E. Vestergaard. Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI. Working Paper 33777, National Bureau of Economic Research, 2025. Revised March 2026.

[18] A. Zhang and M. K. Lee. Knowledge Workers' Perspectives on AI Training for Responsible AI Use. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–18, 2025. doi: 10.1145/3706598.3714100.

[19] Q. Xia, M. Constantinides, A. Sarkar, D. P. Brumby, and A. L. Cox. "If You're Very Clever, No One Knows You've Used It": The Social Dynamics of Developing Generative AI Literacy in the Workplace. In *Proceedings of the 5th Annual Symposium on Human-Computer Interaction for Work*, pages 1–18, 2026. doi: 10.1145/3808045.3808050.

[20] B. M. Chen and H. Bao. Training for Technology: Adoption and Productive Use of Generative AI in Legal Analysis. arXiv:2603.04982, 2026. URL https://arxiv.org/abs/2603.04982v3. Working paper, version 3, 5 June 2026.

[21] Q. Ma, W. Peng, C. Yang, H. Shen, K. Koedinger, and T. Wu. What Should We Engineer in Prompts? Training Humans in Requirement-Driven LLM Use. *ACM Transactions on Computer-Human Interaction*, 32(4):1–27, 2025. doi: 10.1145/3731756. Detailed intervention design checked in arXiv:2409.08775v3 (April 2025).

[22] R. Xiao, R. Ye, X. Hou, J. Wen, H. Kumar, M. Liut, and J. Stamper. Transforming GenAI Policy to Prompting Instruction: An RCT of Scalable Prompting Interventions in a CS1 Course. arXiv:2602.16033, 2026. URL https://arxiv.org/abs/2602.16033v1. Working paper, version 1, 17 February 2026.

[23] B. Puppart and J. Aru. A Brief AI Literacy Intervention Does Not Significantly Reduce Over-Reliance and Increases Under-Reliance on ChatGPT: A Randomized Study. arXiv:2503.10556, 2026. URL https://arxiv.org/abs/2503.10556v2. Preprint, version 2, 18 September 2026.

[24] Negar Kamali, Candice Rockell Gerstner, Jessica Hullman, and Matthew Groh. Generative AI Literacy Training Improves Intelligence Analysts' Discrimination of Real and AI-Generated Images. arXiv:2606.28510, 2026. URL https://arxiv.org/abs/2606.28510v1. Preprint, version 1, June 2026.

[25] A. Farach, A. Cambon, L. Tankelevitch, C. Hsueh, and R. Janssen. Scaffolding Human-AI Collaboration: A Field Experiment on Behavioral Protocols and Cognitive Reframing. arXiv:2604.08678, 2026. URL https://arxiv.org/abs/2604.08678v2. Working paper, version 2, 19 April 2026.

[26] W. Arthur, W. Bennett, P. S. Edens, and S. T. Bell. Effectiveness of Training in Organizations: A Meta-Analysis of Design and Evaluation Features. *Journal of Applied Psychology*, 88(2):234–245, 2003. doi: 10.1037/0021-9010.88.2.234.

[27] N. Keith and M. Frese. Effectiveness of Error Management Training: A Meta-Analysis. *Journal of Applied Psychology*, 93(1):59–69, 2008. doi: 10.1037/0021-9010.93.1.59.

[28] C. N. Lacerenza, D. L. Reyes, S. L. Marlow, D. L. Joseph, and E. Salas. Leadership Training Design, Delivery, and Implementation: A Meta-Analysis. *Journal of Applied Psychology*, 102(12):1686–1718, 2017. doi: 10.1037/apl0000241.

[29] B. D. Blume, J. K. Ford, T. T. Baldwin, and J. L. Huang. Transfer of Training: A Meta-Analytic Review. *Journal of Management*, 36 (4):1065–1105, 2010. doi: 10.1177/0149206309352880.

[30] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. Technical Report NIST AI 600-1, National Institute of Standards and Technology, 2024.

[31] Neil Perry, Megha Srivastava, Deepak Kumar, and Dan Boneh. Do Users Write More Insecure Code with AI Assistants? In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, pages 2785–2799, 2023. doi: 10.1145/3576915.3623157.

[32] G. M. Alliger, S. I. Tannenbaum, W. Bennett, JR, H. Traver, and A. Shotland. A Meta-Analysis of the Relations Among Training Criteria. *Personnel Psychology*, 50(2):341–358, 1997. doi: 10.1111/j.1744-6570.1997.tb00911.x.

[33] A. Markus, A. Carolus, and C. Wienrich. Objective measurement of AI literacy: Development and validation of the AI competency objective scale (AICOS). *Computers and Education: Artificial Intelligence*, 9:100485, 2025. doi: 10.1016/j.caeai.2025.100485.

[34] S. Zhang, R. Xiao, A. F. Botelho, G. Liao, T. K. F. Chiu, J. Stamper, and K. R. Koedinger. How to Assess AI Literacy: Misalignment Between Self-Reported and Objective-Based Measures. In *Proceedings of the LAK26: 16th International Learning Analytics and Knowledge Conference*, pages 405–414, 2026. doi: 10.1145/3785022.3785088.

[35] T. Lintner. A systematic review of AI literacy scales. *npj Science of Learning*, 9(1):50, 2024. doi: 10.1038/s41539-024-00264-4.

[36] E. Wiles, L. Krayer, M. Abbadi, U. Awasthi, R. Kennedy, P. Mishkin, D. Sack, and F. Candelon. Generative AI and the Temporary Upskilling of Knowledge Workers. Author manuscript, 2026. URL https://emmawiles.github.io/storage/reskill.pdf. Current manuscript linked from the author's research page; accessed October 2026.

About the Author

Michael Burch is Vice President of AI Enablement and Acceleration at Security Journey. His work focuses on helping employees develop practical AI and security skills through instruction and hands-on learning. His background includes software development, application security leadership, cyber defense with the North Carolina National Guard, and service as a U.S. Army Special Forces medical sergeant. He earned a B.S. in Cybersecurity, magna cum laude, from Colorado Technical University and eight GIAC certifications spanning security leadership, auditing, incident response, cyber defense, and Python programming.