# Democratizing Artificial Inteligence via Artificial Intelligence: an Artificial Intelligence-to-Human Knowledge Transfer Framework to Cultivate Deep Organizational Artificial Intelligence Literacy

Veselin Petrov
*Coca-Cola HBC*
Athens, Greece
veselin.petrov@cchellenic.com

Ameer Hamza Siraj
*Coca-Cola HBC*
Amsterdam, Netherlands
ameer.hamza.siraj@cchellenic.com

Amr Sultan
*Coca-Cola HBC*
Cairo, Egypt
amr.sultan@cchellenic.com

Alessandra Virga
*Coca-Cola HBC*
Rome, Italy
alessandra.virga@cchellenic.com

Maximilian Christoph Kohnen
*Coca-Cola HBC*
Zurich, Switzerland
maximilian.kohnen@cchellenic.com

Georgia Klapsa
*Coca-Cola HBC*
Thessaloniki, Greece
georgia.klapsa@cchellenic.com

Haonan Wu
*Coca-Cola HBC*
London, United Kingdom
haonan.wu@hotmail.com

***Abstract*— Organizations worldwide are grappling with a dual challenge: "Hunger": employees are simultaneously feeling an urgent need to understand Artificial Intelligence and how to adopt it; And "Fear": people also increasingly fear being displaced by Artificial Intelligence. The traditional corporate Data and Analytics training programs, when targeting entire organizations, have struggled with low participation rates, poor content relevance, and limited scalability. This paper presents a practitioner-driven case study from a large international manufacturing organization (10,000+ employees) that addresses this challenge by leveraging Artificial Intelligence itself as the medium for delivering Artificial Intelligence literacy. Our solution combines two complementary Artificial Intelligence-powered mechanisms: a suite of Artificial Intelligence-driven learning games and a personalized learning path navigator. Drawing on early pilot results, related work in Artificial Intelligence literacy frameworks, and organizational learning theory, we argue that this self-referential knowledge transferring approach - democratizing Artificial Intelligence via Artificial Intelligence - offers a scalable, contextually relevant, and engagement-sustaining model for organizational Artificial Intelligence literacy. We discuss outcomes, limitations, and a roadmap for evolving the solution toward a lifelong learning ecosystem.**

***Keywords — Artificial Intelligence, AI, AI Literacy, Organizational Learning, Generative AI, Adventure Learning, Gamified Learning, Personalized Learning Paths, Enterprise Training, Human-AI Interaction***

## I. INTRODUCTION

The rapid proliferation of Generative and Agentic Artificial Intelligence (AI) systems has created a paradox inside large organizations: employees from all job roles are simultaneously hungry for and fearful of AI. They recognize that AI is reshaping job roles and decision-making processes, yet many feel ill-equipped to engage with it meaningfully. For learning and development functions, this creates a complex task - how do you deliver AI literacy at scale, in a way that is role-relevant, engaging, and continuously up to date? This challenge is compounded by the fact that while foundational AI literacy competencies have been deeply defined in academic literature [1], [2], recent scholarship emphasizes that the arrival of generative models requires an entirely new framework for responsible, conversational literacy [3]. Furthermore, empirical strategies for scaling these frameworks across diverse, large-scale enterprise workforces remain notably scarce [4], [5].

This paper reports on a practitioner-led initiative within a large international consumer goods manufacturer employing more than 10,000 people globally. Over the past three years, the organization invested substantially in Data and Analytics training - building a modular curriculum, blending digital self-paced modules with in-person classroom sessions and providing practical assignments to the most dedicated learners. Despite initial moderate successes, outcomes plateaued: participation rate increase dwindled after initial surge and content completion

percentage did not reach 50%. Surveys consistently highlighted two barriers: employees did not feel they had sufficient time to dedicate to upskilling, and the content did not feel relevant to their specific roles. To resolve these friction points, organizations must look beyond technical skills to address the functional, critical, and sociocultural dimensions of workplace literacy [6]. This aligns with broader industry data showing that nearly 70% of organizational AI transformation hurdles stem not from technical limitations, but from cultural, human, and capability readiness boundaries [7]

The emergence of generative and agentic AI systems in 2023–2025 and their rapid adoption across industry heightened both the fear and the hunger among the workforce, and elevated AI literacy from a "nice-to-have" to a strategic organizational priority. It also unlocked a new possibility: using AI itself to solve the AI literacy problem. Rather than treating AI as the subject of training, we asked whether AI could become the medium through which training is created, delivered and sustained.

Our proposed solution comprises two interconnected components. The first is an AI-powered learning game ("AI Adventure") that transforms conceptual AI knowledge into scenario-based game-like experiences. The second is an AI-driven personalized learning path navigator that routes each employee to the specific content most relevant to their role and current knowledge level.

A secondary observation of this paper is worth noting here: both components were built using generative AI. This matters beyond the immediate context, as it illustrates how AI can compress content development timelines substantially - in this case, reducing cycles that previously took months to a matter of weeks. The initiative therefore serves as both a practical intervention and a small-scale demonstration of what AI-assisted instructional design might look like in practice.

The remainder of this paper is structured as follows. Section 2 provides organizational background and describes the challenges encountered with prior training approaches. Section 3 reviews related work on AI literacy frameworks, enterprise learning, and AI-mediated instruction. Section 4 describes the proposed solution in detail. Section 5 presents early pilot results and projected outcomes. Section 6 discusses findings, limitations, and the road ahead..

## II. Organizational Background and the Training Challenge

The organization at the centre of this study is a large, geographically dispersed international consumer goods manufacturer with operations across multiple continents and a workforce exceeding 10,000 employees. Like many enterprises of its scale, it operates with a guiding philosophy that links technology investment directly to human capital, roughly summarized as investing proportionally in technology and in the talent required to leverage it.

Beginning in 2022, the organization launched a structured, employee-driven Data and Analytics training initiative targeting all employee levels. The program was deliberately designed with recognized learning design best practices: a modular structure to allow self-paced consumption, a blended model combining digital e-learning with facilitated classroom sessions, and practical assignments intended to bridge conceptual learning and on-the-job application.

The organization opted to develop the curriculum internally rather than relying on external content providers. This decision was driven by the need to maintain high content relevance and to eliminate the latency associated with waiting for third-party vendors to develop training for rapidly evolving internal needs. Furthermore, by placing the employee in the 'driver's seat' of their own upskilling, the organization sought to empower individuals to take capability building into their own hands. The goal was to foster intrinsic motivation by allowing employees to navigate the learning journey based on their immediate professional needs and personal career aspirations.

The initial outcomes served as a critical learning phase that revealed the complexities of scaling across a global workforce. Of the 10,000 eligible employees, the program successfully engaged an initial cohort of 1,000 enrolees over a period of 3 months post-launch. Within this group, a dedicated core of approximately 100 learners achieved full sequence completion, while live classroom sessions maintained a focused attendance of 20-30 participants. These figures demonstrated that while the curriculum was effective for a specific segment of highly motivated 'power users,' a more tailored approach was required to reach the broader organization. Exit and dropout surveys provided valuable insights into this gap: while the quality was recognized, many employees found it challenging to balance the time-intensive modules with daily operations or noted that the material remained too abstract for their specific role-based tasks (see Figure 1). These results highlighted a clear demand for greater personalization and lower friction, setting the stage for the next phase of the digital transformation.

A meaningful step forward came when the organization transitioned to an even more modularized architecture - breaking longer training sequences into shorter, standalone units. Additionally, the live classroom curriculums were redesigned with examples and tasks from the organization's own data, for example mirroring what a sales analyst might encounter in their day-to-day job. This change alone produced approximately a 30% uplift in completion rates for eLearning modules and doubled the classroom demand, suggesting that format and accessibility matter considerably (Figure 1).

However, the underlying relevance and personalization gaps remained unaddressed. The question of how to scale a training experience that felt truly tailored — across 10,000 employees with diverse roles, digital fluency levels and time constraints — had no easy answer within a conventional learning design paradigm.

The external environment was making this more urgent by the day. As generative and agentic AI tools moved rapidly into the mainstream, anxiety was building across the workforce. Frontline workers felt overwhelmed by a steady stream of new AI-enabled tools landing in their day-to-day work, while middle managers were grappling with a more ambiguous challenge — how to meaningfully introduce AI into the way their teams operate, without clear guidance on where to start. The concerns were diverse,

role-specific, and shifting fast, which is precisely what a static, centrally-designed curriculum struggles to handle.

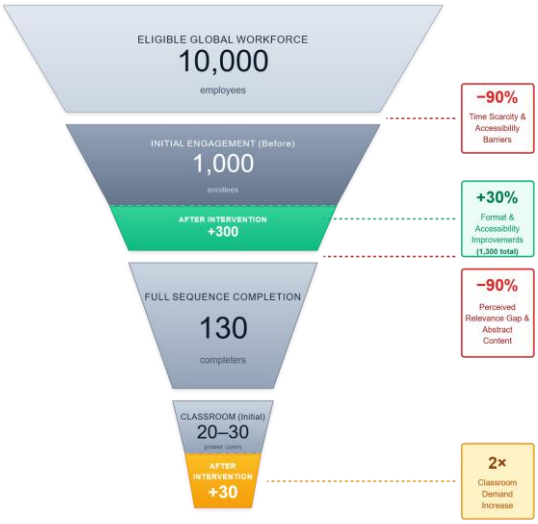


Figure 1. Training participation funnel showing post-intervention impact of format and accessibility improvement. From an eligible workforce of 10,000 employees, two friction points drive 90% attrition at each stage (red annotations). Split layers compare baseline (gray) to post-intervention (colored) outcomes: format and accessibility changes yielded +30% eLearning engagement (green, +300 enrolees) and 2x classroom demand (amber, +30 participants). Data from 3-month post-intervention period.

## III. Related Work

### A. Defining AI Literacy

Scholarly interest in AI literacy has grown considerably since Long and Magerko's foundational contribution [1], which defined it as a set of competencies enabling individuals to critically evaluate AI technologies, communicate and collaborate with AI, and apply AI tools in both professional and everyday context. Building on this, Ng et al. [2] proposed a four-dimensional model, covering know and understand, use and apply, evaluate and create, and ethical issues, that has since become widely cited across the field.

More recent scholarship has pushed these foundations further. Cox [3] argues that the arrival of large language models necessitates a richer conception of what it means to be AI literate. His proposed framework of "responsible generative AI literacy" moves beyond technical competence or prompt engineering to foreground ethical reasoning, awareness of societal and personal consequences, and critically, a form of reflective self-knowledge — the idea that before engaging with generative AI, users should understand their own role, needs, and vulnerabilities in that interaction. For Cox, the shift from using AI as a tool to engaging with it conversationally changes not just the skills required, but the nature of the literacy itself.

Noteworthy for the present study, systematic reviews have consistently found that while AI literacy research has proliferated in K-12 and higher education settings, organisational and workplace contexts remain underserved (Laupichler et al. [4]; Cetindamar et al. [5]). It is precisely this gap that the initiative described here sets out to address.

### B. AI Literacy in Organizational Contexts

Kong and Zhang [6] identify three dimensions of organizational AI literacy - functional, critical, and sociocultural, and argue that the sociocultural dimension is most often neglected. This dimension concerns how AI reshapes power, roles, and collaboration norms within organizations, and is directly relevant to the fear-and-hunger dynamic our employees experience. Recent articles highlight the “Fear” aspect as one of the barriers to the workforce utilizing AI confidently [8], while popular industry commentary by van Geelen and Shah [9] provides a speculative, scenario-based glimpse into 2028 and the negative impact AI prevalence could have on job security. Whether or not those projections prove accurate, their resonance with employees is itself organizationally significant: perceived threat shapes behavior and engagement with learning, regardless of whether the threat materializes.

[5] identify four capability clusters for workplace AI literacy, spanning technology, human, work, and learning dimensions. The learning cluster is particularly relevant here - sustaining engagement with AI upskilling over time was as much an objective of our intervention as building technical fluency itself.

This pattern is echoed in the organizational change management literature. BCG [7] found that approximately 70% of AI implementation hurdles are people- and process-related rather than technical — a finding that shifts the conversation away from tool availability and toward the harder question of human readiness.

### C. AI as a Vehicle for Its Own Instruction

Research into AI-driven AI literacy is still in its early stages, though empirical evidence is beginning to emerge. Xiao et al. [10] designed a Large Language Model (LLM)-based learning module using AI itself to teach prompting literacy through scenario-based deliberate practice which showed measurable gains in terms of competency for K-12 settings. The closest analog in adult professional development is the chatbot-as-tutor literature: Okonkwo and Ade-Ibijola [11] reviewed chatbot applications in education and found consistent evidence of personalization benefits, while Behforouz and Al Ghaithi [12] demonstrated empirically that interactive AI tools improve self-directed learning outcomes.

Beyond standalone modules, the integration of generative and agentic systems into standard enterprise workflows introduces a critical shift in how literacy is acquired. Tools, such as Microsoft Copilot, suggest that ambient AI - assistance embedded directly within the user’s primary workspace, diminishes the cognitive friction typically associated with adopting new technologies. By lowering this initial activation barrier, these systems move AI from an external "tool to be learned" to an invisible component of the work environment. The shift toward agentic systems, which leverage user data to provide proactive, just-in-time learning prompts, challenges traditional models of organizational training in favor of continuous, personalized development.

## IV. The Proposed Solution: AI as Medium

Our solution operates on a central premise: if the barrier to AI literacy is a combination of time scarcity and gaps in

relevance or motivation, then the solution must be fast to consume, highly personalized, and intrinsically engaging. We address each of these through two interconnected AI-powered components.

### *A. AI Adventure: The AI Learning Game*

AI Adventure is an AI-powered learning game designed to transform abstract concepts into lived, scenario-based experiences. The player's focus is to experience working alongside a personal AI Agent compared to the tedious "old way" of executing tasks manually, embedding relevant human-in-the-loop gates to directly address employee fears. To ground this experience in academic theory, these game interactions map directly to foundational literacy dimensions: budget verification scenarios require users to use and apply knowledge, while email cross-referencing forces them to evaluate context. Furthermore, rather than relying solely on high-level engagement metrics, the platform captures a direct, learning-oriented retention and conversion metric by monitoring exact time-on-platform and the rate at which players voluntarily transition from game completion to the Navigator to generate a personalized learning path. This combination addresses both main concerns identified earlier in the paper, helping to alleviate fears while boosting the hunger to learn how AI can support daily operational tasks.

The game was designed with the following principles. Content generation was accelerated using generative AI, compressing structural asset development cycles from months to approximately a week per module. To achieve this velocity, the development lifecycle utilized a specialized, multi-modal AI tool stack: instructional design scripts and dialogue branching logic were engineered using Anthropic's Claude, while localized video production was accelerated via HeyGen to generate synthetic presenter assets. Prototyping was executed in a matter of hours within Figma, enabling a rapid feedback loop with stakeholders. Later deployment of the interactive components was driven by developers utilizing Claude Code for rapid agentic command-line execution and code generation. To prevent errors or hallucinations from impacting the final learning experience, a strict three-tier quality assurance protocol was implemented: first, Claude's script generation was strictly bounded by a pre-approved internal content taxonomy; second, code generated via Claude Code was subjected to local syntax validation; and third, a mandatory human oversight gate required instructional designers to audit and sign off on all Figma flows, scripts, and HeyGen video outputs before deployment, ensuring complete semantic accuracy.

Players are prompted to return to the game through push notifications and reminders calibrated to their behavioural patterns — morning learners receive morning prompts, evening learners receive evening prompts. A pipeline of centrally-devised scenarios is maintained and scheduled for a monthly release to keep the attention of users and ensure relevance (Figure 2).

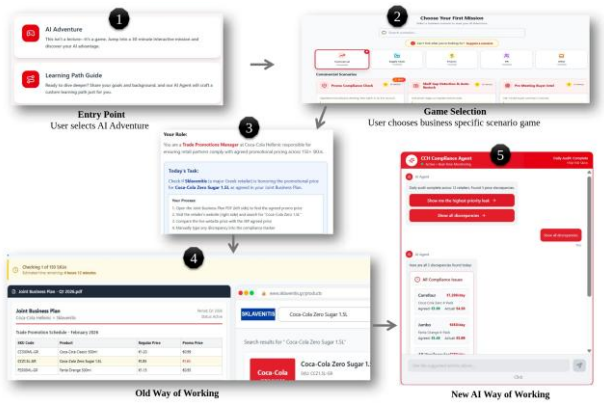


Figure 2. Sequential user engagement flows through the AI Adventure learning platform. The four-stage journey illustrates the path from initial entry point to post-simulation impact demonstration. Red action indicators denote primary user interactions at each stage. Screenshots represent actual platform interface captured during pilot implementation period.

Finally, the game can surface real-world use cases sourced from employee feedback, creating a community learning loop where the workforce contributes to its own curriculum. With the short development lifecycle, a learner can expect to see a scenario they proposed within the next monthly release of new scenarios.

### *B. The Personalized Learning Path Navigator*

The second component addresses the relevance problem directly. Rather than requiring employees to self-navigate a library of training content, an AI-powered navigator conducts a short diagnostic conversation — assessing the employee's role, existing AI familiarity, current challenges, and learning preferences — and then curates a tailored learning path from the available content portfolio. Of course, this recommendation relies on the user's self-described profile and thus could offer unsuitable suggestions. However, the system is designed to be highly iterative; employees can create as many learning paths as they desire and fundamentally shift their desired capabilities until they find the path best suited for their needs.

The navigator is built on an LLM backbone, enabling natural-language interaction that mirrors a conversation with a knowledgeable colleague rather than a form-filling exercise. It draws on an internal content taxonomy that tags each learning unit by role family, data and AI capability domain, proficiency level, and estimated time investment. As the employee progresses, the navigator keeps track of the learning path completion percentage and time investment needed to finish. Each learning journey created can be saved so the user can revisit it at a later stage or fundamentally shift the desired capabilities.

Importantly, the navigator also serves as a rapid content development platform. When new AI capabilities emerge in the market — as they do with increasing frequency — a human content designer can brief the generative AI system on the new topic, and a first-draft learning module can be produced in approximately two weeks, having been reviewed, refined, and deployed. This dramatically reduces the latency between AI developments in the world and the organization's training response (see Figure 3)

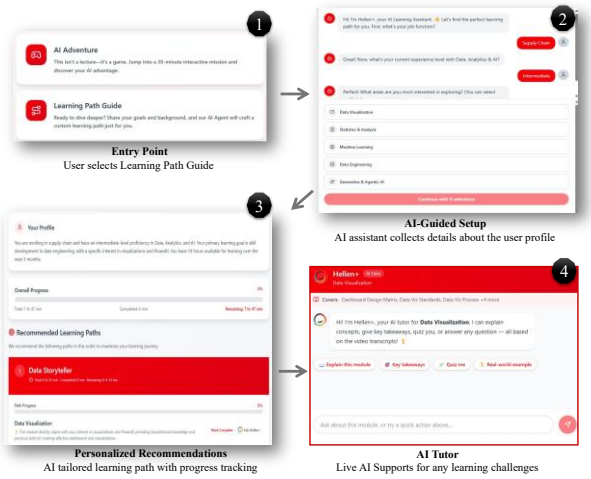


Figure 3. AI-powered learning path creation flow facilitated by AI Assistant. The four-stage process demonstrates personalization from initial entry through AI-guided preference collection to tailored curriculum recommendations with integrated progress tracking and AI Tutoring.

### C. *Integration and Architecture*

The AI Adventure and the Learning Path Navigator are designed to be mutually reinforcing rather than standalone tools. The AI Adventure shows the employees how much more they could do if they took the time to listen to the "Hunger" rather than the "Fear" and directly prompts the players to find out what learning content from the existing library might be useful for them. The Learning Path Navigator then takes their input and provides a personalized suggestion, keeping them engaged on the journey.

Both tools feed a central learning analytics dashboard that provides Learning and Development teams with visibility into engagement patterns, completion rates, knowledge gaps, and content performance — enabling continuous improvement of the curriculum.

## V. Early Pilot Results

The solution described in this paper is in active development and early pilot deployment at the time of writing. Consequently, this work is presented as an early deployment and design-focused contribution; final longitudinal results are not yet available, and the data reported here reflects an early pilot cohort performance snapshot.

The initial pilot cohort comprised a representative sample of 100 employees across various departments, actively engaged and interested in the area of data, analytics and AI. Within this tracking group, the AI Adventure game demonstrated an early 85% completion rate for initial scenarios, with users averaging 12 minutes of active platform interaction per session. Rather than relying on long-term projections, our early observed evidence focused on a key behavioral indicator of learning intent: the rate at which users voluntarily transition from game completion to the Personalized Learning Path Navigator to generate their own structured training pathways. Initial logs show that approximately half of the pilot participants completed this exact transition sequence during their initial sessions. Based on this preliminary trajectory and incorporating early pilot engagement data from the AI Adventure - which shows higher click-through and completion rates than prior digital learning formats - we estimate the potential for expanded effective reach, corresponding to an encouraging early trend toward voluntary training engagement relative to the pre-modularization baseline.

Qualitative feedback gathered from the pilot cohort, including input from senior stakeholders, reflects three consistent themes: the game format significantly reduces the psychological barrier to engaging with AI topics; the personalized path navigator reduces the time burden of finding relevant content; and both tools create a sense of momentum — employees describe returning to the platform voluntarily rather than experiencing it as a compliance obligation.

Detailed quantitative results from the full pilot deployment — including pre/post competency assessments, net promoter scores, and longitudinal engagement metrics — will be reported in subsequent work as the program matures.

## VI. Discussion

### A. *The Self-Referential Insight*

The central insight of this work is both practical and conceptually significant: the most effective tool for teaching AI literacy could be AI itself. This "self-referential loop" where AI acts as both the subject and the instructor is not merely a rhetorical device but an architectural necessity that fundamentally changes the learning experience. By using AI-mediated instruction, the organization moves away from explaining capabilities in the abstract toward instantiating them in the learner's direct experience.

In this model, the medium becomes the message; when an employee interacts with an AI tutor or game agent, they are practicing human-AI collaboration as a byproduct of the learning process itself. This immersion eliminates the "knowledge latency" typical of static training, as the system remains inherently updatable and available on demand, demonstrating its own evolving capabilities in real-time. Ultimately, this approach dismantles the psychological barrier of "Fear" not through theoretical reassurance, but through the low-stakes, practical mastery of the technology as a constant learning companion.

### B. *Limitations and Caveats*

Several important limitations of this work warrant acknowledgment. First, the results reported are from an early pilot, and the projected uplift numbers, while grounded in prior program data, are projections rather than confirmed outcomes. Readers should interpret the quantitative claims with appropriate caution pending full deployment results.

Second, this paper reports on a single-organization case study. While the context — a large, geographically dispersed international manufacturer — is representative of many global enterprises, generalizability to other organizational types, sectors, or workforce demographics should be tested rather than assumed.

Third, the use of AI to generate training content at speed raises quality assurance questions that deserve ongoing attention. Rapid content production through generative AI carries risks of hallucinations, inaccuracies

or oversimplification, especially relevant in the learning context. Addressing those proactively requires continuous, robust human review processes. To guarantee implementation trustworthiness, the platform enforces a strict data governance and privacy framework where all conversational interactions and diagnostic logs are processed locally, ensuring compliance with corporate data protection standards. Furthermore, a strict human oversight protocol governs the system: no AI-generated instructional scenario or curriculum recommendation is pushed live to employees without mandatory verification and sign-off by an expert human content designer, effectively mitigating the risk of unvetted AI explanations.

### C. *Toward a Lifelong Learning Ecosystem*

The roadmap for this initiative extends beyond the current pilot. The vision is for the platform to evolve into a lifelong AI learning ecosystem — one that grows with each employee, surfaces new content as AI capabilities evolve, and creates a community-driven feedback loop in which employees contribute real use cases and challenges that become the basis for future content. Planned enhancements include new game scenarios targeting emerging AI capabilities, expanded learning path personalization leveraging richer behavioural data, customized learning reminders based on individual engagement patterns, and a "community voice" feature enabling employees to flag AI topics they want the platform to address. The aspiration is for a continuous learning companion that keeps pace with the fastest-moving technology domain the modern workforce has ever encountered.

## VII. Conclusion

This paper has argued that the growing demand for organizational AI literacy is best served not by traditional training about AI, but by an immersive pedagogical approach of training through AI. By positioning AI as the medium rather than merely the subject of organizational learning, the framework described here achieves a level of personalization, scalability, and content velocity that conventional instructional design cannot match. The transition from a static, centrally-managed curriculum to an AI-mediated "self-referential" model—where the technology explains itself—allows organizations to move past the binary of "Hunger" and "Fear". In this environment, employees do not just acquire theoretical knowledge; they develop practical fluency through direct, low-stakes interaction with the very systems they are learning to master.

The implications of this shift extend beyond the immediate pilot results. As AI capabilities continue to evolve at a pace that outstrips traditional content development cycles, the "Self-Referential" insight provides a roadmap for creating learning ecosystems that are inherently updatable and contextually grounded in real-world employee challenges. We invite the research community to engage with these findings, critique the methodology, and contribute to the sparsely studied intersection of organizational learning and AI-mediated instruction. While the challenges of quality assurance and generalizability remain, the stakes for a rapidly transforming global workforce are high. The field is at an early and exciting stage, and the move toward treating AI as a continuous learning companion marks a fundamental step in democratizing these capabilities across the modern enterprise.

## Multimedia Appendix

The AI Adventure Game and Personalized Learning Path Navigator were developed using an agentic-native framework, leveraging AI-assisted coding to synthesize front-end and back-end architectures within a single development sprint. This accelerated methodology facilitated a daily feedback loop, enabling the team to implement user-driven modifications rapidly throughout the study period. The complete source code, along with implementation documentation, is available at [13].

## References


[1] 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, New York, NY, USA: ACM, May 2020, pp. 1–16.

[2] D. T. K. Ng, J. K. L. Leung, S. K. W. Chu, and M. S. Qiao, "Conceptualizing AI literacy: An exploratory review," Computers and Education: Artificial Intelligence, vol. 2, id. 100041, Dec. 2021, pp. 1–12.

[3] A. Cox, "Algorithmic literacy, AI literacy and responsible generative AI literacy," Journal of Web Librarianship, vol. 18, no. 3, Sep. 2024, pp. 93–110.

[4] M. C. Laupichler, A. Aster, J. Schirch, and T. Raupach, "Artificial intelligence literacy in higher and adult education: A scoping literature review," Computers and Education: Artificial Intelligence, vol. 3, id. 100072, Aug. 2022, pp. 1–15

[5] D. Cetindamar, K. Kitto, M. Wu, Y. Zhang, B. Abedin, and S. Knight, "Explicating AI literacy of employees at digital workplaces," IEEE Transactions on Engineering Management, vol. 71, Jan. 2024, pp. 810–823.

[6] S.-C. Kong and G. Zhang, "A Conceptual Framework for Designing Artificial Intelligence Literacy Programmes for Educated Citizens," in Global Chinese Conference on Computers in Education Main Conference Proceedings, Jun. 2021, pp. 11-15.

[7] N. de Bellefonds, T. Charanya, J. Walker, S. D'Silva, and M. Reeves, "Where's the Value in AI?" Boston Consulting Group, Research Report, Oct. 2024. [Online]. Available: https://www.bcg.com/publications/2024/wheres-value-in-ai, [retrieved: May, 2026].

[8] J. Bellamy, J. Douglas, J. Wickett-Whyte, and N. Bransden, "AI Skills for Life and Work: Summary Report," Ipsos, Department for Science, Innovation and Technology (DSIT) and Department for Culture, Media and Sport (DCMS), London, UK, Tech. Rep., Jan. 2026. [Online]. Available: https://www.gov.uk/government/publications/ai-skills-for-life-and-work-summary-report/ai-skills-for-life-and-work-summary-report--2. [retrieved: May, 2026].

[9] J. van Geelen and A. Shah, "The 2028 Global Intelligence Crisis," *Citrini Research Substack*, Feb. 22, 2026. [Online]. Available: https://www.citriniresearch.com/p/2028gic, [retrieved: May, 2026].

[10] R. Xiao et al., "Learning to Use AI for Learning: Teaching Responsible Use of AI Chatbot to K-12 Students Through an AI Literacy Module," arXiv preprint arXiv:2508.13962, Aug. 2025. [Online]. Available: https://arxiv.org/abs/2508.13962, [retrieved: May, 2026].

[11] C. W. Okonkwo and A. Ade-Ibijola, "Chatbots applications in education: A systematic review," Computers and Education: Artificial Intelligence, vol. 2, id. 100033, Nov. 2021, pp. 1–11.

[12] B. Behforouz and A. Al Ghaithi, "The impact of using interactive chatbots on self-directed learning," Studies in Self-Access Learning Journal, vol. 15, no. 3, Sep. 2024, pp. 317–344.

[13] https://github.com/hamzakpt/CCH-AI-Academy