


WHITEPAPER

Practice, Not Content

The Evidence for AI Coaching in Learning
and Development



AiReady
Powered by AiReady™ Technology

AUTHOR

Garima Gupta

Founder & CEO, Artha Learning Inc.
Published 2026 | Artha Learning Inc.

WITH CONTRIBUTIONS FROM

Iowa State University • Cyclone Support: Creating a Culture of Care

The France Foundation • AI Case Coach for Clinical Education

GardaWorld • Frontline and Operational Training at Scale



Table of Contents

Executive Summary	3
01. The Broken Promise of Online Learning	4
02. What AI Coaching Actually Is	6
03. The Industry Is Moving	8
04. The Evidence: What Learners Told Us	9
05. Limitations	14
06. Design Principles for Responsible AI Coaching	16
07. Conclusion: The Shift That's Already Happening	19
References	20
About	21



Executive Summary

The professional learning industry invests billions annually in online training. The returns are startlingly low. A landmark Cochrane review of 215 studies found that educational meetings produce a median improvement of just **6% in professional practice**. Over 80% of continuing education still relies on passive methods-read, watch, click next-that research has shown to have little or no effect on real-world behavior.

The problem is not a lack of content. It is a lack of practice and feedback.

AI coaching - personalized, rubric-driven feedback on learner-generated responses, embedded directly within the learning experience-addresses this gap. It gives every learner what only a live coach could previously offer: real-time, individualized feedback on their own words, at the scale and cost of eLearning.

This paper presents converging evidence from three sources: peer-reviewed academic research on why passive learning fails and active practice succeeds; industry survey data from over 1,000 L&D practitioners showing that AI coaching is the profession's fastest growing frontier; and primary learner data from three real-world deployments where AI coaching was embedded in production eLearning programs: university staff at **Iowa State University**, clinicians educated by **The France Foundation**, and frontline workers and supervisors at **GardaWorld**.

The findings are consistent across all three evidence layers. When AI coaching replaces passive content delivery:

87%

of learners report satisfaction with AI-powered practice.

94.6%

of learners reported confidence in the target skill after completing the course

60%

of learners spontaneously identified the AI practice as the most valuable part of the course-without being asked.

Emerging evidence suggests AI coaching can be an effective way to incorporate practice and feedback into learning experiences. The question is how organizations can apply these approaches responsibly and effectively.



01. The Broken Promise of Online Learning

A Billion-Dollar Industry with Single-Digit Results

Enterprise learning is a massive global category, with corporate training alone projected to expand from \$361.5 billion in 2023 to \$805.6 billion by 2035.¹² The continuing medical education market is smaller but still substantial, forecast to rise from \$8.66 billion in 2022 to \$13.52 billion by 2028.¹³ Organizations in every sector invest heavily in learning technology, course development, and delivery platforms.

The outcomes do not match the investment.

A Cochrane systematic¹⁴ review examined 215 studies involving over 28,000 healthcare professionals. It found that educational meetings produce a median improvement of just 6% in practice compliance and 3% in patient outcomes. These are not outlier findings. They represent the accumulated evidence across decades of research.

The majority of continuing education activities rely on passive learning methods - lectures, readings, video presentations, static eLearning - despite evidence showing that didactic presentations have little or no beneficial effect on professional practice. The industry has optimized delivery speed, visual design, and platform technology while leaving the fundamental learning model unchanged.

Why “Teach and Hope” Doesn’t Work

Behavioral science has comprehensively debunked the assumption that providing information changes behavior. This “information deficit model” - the idea that if people simply knew the right thing to do, they would do it - persists as the operating assumption of most eLearning despite decades of evidence to the contrary.



The forgetting curve is steep. Research on the Ebbinghaus forgetting curve shows that people forget a large share of newly learned information quickly without reinforcement, often with the steepest drop in the first day. Passive consumption - watching a video, reading a module, clicking through slides - does not create the neural pathways required for durable memory.



The knowing-doing gap is wide. Research on implementation intentions shows that intentions often fail to translate into action; many people who intend to change behavior do not follow through unless they form specific if-then plans. (most people can perhaps correlate this with their own lived experiences).¹⁵



Self-assessment is unreliable. Systematic reviews of physician self-assessment have found that self-ratings often do not correlate with observed competence, with the majority of comparisons showing little, no, or even inverse relationships.¹⁶ Learners cannot be trusted to accurately identify their own skill gaps.

As Victoria Ekstrom argued in JMIR Medical Education (2026)², there is a fundamental paradox at work: we design clinical systems around behavioral science because we know professionals are human-susceptible to cognitive biases, decision fatigue, and status quo inertia-then we design learning systems as if those same professionals are perfectly rational actors who will systematically engage with content, self-assess accurately, and independently translate knowledge into practice change.

The Checkbox Trap

Most eLearning optimizes for completion, not competence. Multiple-choice questions test recognition (the ability to identify a correct answer when presented with options) rather than production (the ability to generate an appropriate response from memory). Learners can click through modules, select answers from pre-written options, and receive a completion certificate without ever engaging in the cognitive effort that produces real learning.

Research confirms the scope of this problem. The majority of accredited continuing professional development activities do not even target behavior change - they focus on knowledge transfer and satisfaction measures. The industry measures what is easy (completion rates and satisfaction scores) rather than what matters (behavior change and performance improvement).

✦ KEY TAKEAWAY

The eLearning industry has spent two decades optimizing content delivery. The research says the problem was never content. It was the absence of practice and feedback.



02. What AI Coaching Actually Is

Defining AI Coaching in Learning


AI coaching, as used in this paper, refers to a specific learning interaction: AI-powered, personalized feedback on learner-generated responses, embedded within the learning experience and evaluated against expert-designed rubrics.

The mechanism is straightforward. A learner encounters a scenario – a student in distress, a patient with a complex medical situation, a customer with an objection. Rather than selecting from pre-written response options, the learner types what they would actually say, in their own words. An AI system evaluates that response against a rubric designed by subject matter experts and returns specific, personalized feedback: what the learner did well, what they missed, and what they might say differently. The learner can then try again.

This isn't a chatbot, and it's not a stand-in for a human coach. It takes the best part of live coaching – real-time feedback while someone is actually practicing – and delivers it at eLearning cost.

Why It's Different

vs. Multiple choice	AI coaching requires production, not recognition. Learners must think and write, not scan and select. This is the most important distinction. Cognitive science consistently shows that producing an answer from memory creates stronger, more durable learning than recognizing a correct answer from a list.
vs. Branching scenarios	Branching scenarios offer pre-written paths. No matter how many branches exist, the learner is choosing between options someone else wrote. AI coaching responds to what the learner actually wrote/said – there are no pre-determined paths.
vs. Open AI Models (like ChatGPT)	AI coaching operates in a closed sandbox with expert-designed rubrics. It references only approved content. Learners stay in the flow of the course. There is no prompt engineering, no risk of hallucinated medical advice, and no need for learners to understand how AI works.
vs. Human coaching	AI coaching offers unlimited practice at any hour, without scheduling, cost, or quality variation. A human coach provides depth, experience and nuance that AI cannot. But most learners in most programs will never have access to a human coach. AI coaching fills the gap between what learners need and what organizations can afford.



The Behavioral Science Behind It

AI coaching activates the learning mechanisms that research identifies as most effective. It requires retrieval and synthesis practice. It creates desirable difficulty - the cognitive effort that strengthens encoding. And it delivers immediate, specific feedback that closes the intention-behavior gap.

A comprehensive meta-analysis² examining over 200 studies with more than 2.1 million participants found that well-designed behavioral interventions in learning produce a Cohen's d of 0.43—a statistically meaningful effect that is practically significant at scale. When learners are required to articulate specific plans for how they will apply what they have learned—which is precisely what AI coaching does when it asks learners to write their own responses—the effect size increases to $d=0.65$, a medium-to-large effect.

✦ KEY TAKEAWAY

AI coaching is not about AI. It is about giving every learner what research says they need: active practice, in their own words, with immediate personalized feedback. AI is what makes it possible at scale.



03. The Industry Is Moving

The Implementation Inflexion

Two major industry surveys conducted in 2025–2026 paint a consistent picture: L&D has moved past experimentation with AI and into active implementation—and the frontier is shifting from content creation to learner experience.

The annual AI in L&D survey by Donald H Taylor and Eglė Vinauskaitė (2025)⁸, covering 606 respondents across 53 countries, found that 54% of L&D practitioners now actively use AI—up from 40% in 2024. The authors call this the “Implementation Inflexion,” a tipping point where AI moved from novelty to part of the standard toolkit. The Synthesia AI in L&D Report (2026)⁹, surveying 421 practitioners, found even higher adoption: 87% of teams now use AI in their work, with only 2% having no plans for adoption.

From Content Speed to Learner Impact

Content creation remains the most common use of AI in L&D. But the fastest-growing applications are those that touch the learner directly. The Taylor report identifies AI role plays as having gone “mainstream” in 2025—moving from isolated pilots to widespread deployment. Dynamic AI-generated feedback that adapts to learner input is identified as a key new modality, distinct from static pre-set responses.

The Synthesia report tracks the shift quantitatively: 36% of teams are piloting AI assessments and simulations, 33% are planning personalized learning pathways, and 49% are exploring AI tutors as an agentic capability. The expected value of AI is shifting from faster production (where 88% already report time savings) toward more personalized learning (72% expect this in the next two years), wider reach (65%), and clearer business impact (55%).

✦ KEY TAKEAWAY

L&D has passed the experimentation phase. AI is in the toolkit, and the frontier has shifted from making content faster to making learning better – through practice, feedback, and personalization delivered at scale. The question is no longer whether to use AI, but where it adds the most value for learners.



04. The Evidence: What Learners Told Us

Multiple Examples across Different Domains

DEPLOYMENT 1: IOWA STATE UNIVERSITY

Iowa State University deployed a university-wide training program-Cyclone Support: Creating a Culture of Care-to equip staff with skills for supporting students facing challenges. The training taught the Ask/Acknowledge, Affirm, and Connect (AAC) conversation framework and included an AI coaching component where learners wrote their own responses to student scenarios and received personalized, rubric-based feedback.

A post-course survey via Qualtrics collected 473 responses (excluding survey previews). The survey measured skill confidence, course quality, satisfaction with interactive and AI-powered elements, effectiveness, and net promoter score, along with three open-ended questions. Findings from this deployment reflect a single institutional implementation and should be interpreted within that context.

DEPLOYMENT 2: THE FRANCE FOUNDATION

The France Foundation, an accredited medical education company, partnered with Artha Learning to create an AI-powered Case Coach for an asynchronous online module training clinicians in shared decision-making and motivational interviewing for patients with obesity. Clinician learners encountered typical “curveballs” they might encounter when caring for patients with obesity. Using CaseCoach, learners applied evidence-based communication techniques to these situations, typing responses they would share with their patients.

DEPLOYMENT 3: GARDAWORLD

GardaWorld, a global security and operations organization, explored the use of AI coaching to address a persistent challenge across distributed teams: limited opportunities for safe, repeatable practice. Across a sample of frontline workers and supervisors, the organization identified a consistent gap in traditional training approaches-particularly in preparing employees for real-world interactions and decision-making under pressure.

Charlie Piché, L&D Manager at GardaWorld described the challenge:

“There is perhaps no issue more common than the lack of opportunities to practice in a safe environment. Addressing this gap requires significant resources, time, and consistently meaningful assessment-and this often stretches training managers thin.”



By integrating AI coaching into training workflows, learners were able to practice and refine both interpersonal and operational skills, including customer service, de-escalation, supervisory decision-making, and SOP adherence.

All the above deployments used AIReady™ technology to embed AI coaching directly within the learning experience—inside the SCORM course, within the existing LMS, with no external tools or prompt engineering required by learners.

THE QUANTITATIVE EVIDENCE

The Iowa State survey represents one of the larger real-world datasets of learner feedback on AI coaching embedded in a production eLearning course. The headline findings:

87%	of learners satisfied with AI-powered practice and feedback (52% extremely satisfied)
94.6%	reported confidence in the target skill after completing the course
96%	rated the course at least moderately effective
88%	rated overall course quality as Good or Excellent
92%	found the course easy to navigate—the AI did not create friction

Source: Iowa State University, Cyclone Support Learner Survey, March 2026 (n=473)

"We're seeing strong learner engagement with the AI interactions within the Cyclone Support Training and appreciate the partnership in developing and implementing this approach."

—Leif Olsen, Student Success & Retention Specialist, Office of the Senior Vice President & Provost, Iowa State University



The Unprompted Signal

The most important finding in this data is not a Likert scale response. It is what happened in the open-ended comments.

Of the 248 learners who wrote a response to “What is one part of the course you appreciated most, and why?”—an open-ended question with no mention of AI - **60% spontaneously identified the AI practice or personalized feedback as the most valuable part of the course.** At 150 out of 248 comments, this is qualitative evidence at quantitative scale—a signal too consistent and too large to ignore. While some learners expressed neutral or critical perspectives, the majority of feedback was positive.

What Learners Said.

THEME 1: “BETTER THAN MULTIPLE CHOICE”

Learners repeatedly contrasted the AI coaching experience with the traditional eLearning interactions they were accustomed to. The consistency of this theme—across learners of different roles, experience levels, and attitudes toward technology, including frontline and supervisory roles in enterprise environments—is striking.

“It was harder to come up with the right words than just to select the right words, but it was more realistic.”

“So much better to make me think of a response than to choose from a list.”

“It also wasn’t a fake choice... that feels like the ‘choices’ I give my literal toddler.”

This aligns precisely with what cognitive science predicts. Production—generating a response from memory—creates stronger learning than recognition—identifying a correct answer from options. Learners experienced this difference directly and named it without prompting.

“The AI agents provide immediate clarification on SOPs and guide real-time decision-making. This reduces hesitation and helps ensure actions remain aligned with documented procedures when time matters most.”

- Training Director, GardaWorld

This shift away from multiple-choice toward active practice is also reflected in enterprise training contexts:



"AI coaches developed through Artha provide an opportunity not only to practice, but to practice repeatedly, refining the critical skills that drive our business and position us at the leading edge of our industry.

Whether offering opportunities to build soft skills such as customer service, sales, de-escalation, empathy, difficult conversations, or to develop more technical capabilities like effective report writing and adherence to standard operating procedures, we are moving beyond multiple-choice decision making. Instead, we are enabling learners to actively practice, while freeing up trainers to focus on foundational development."

–Charlie Piché, Learning & Development Manager, GardaWorld

THEME 2: "THE FEEDBACK FELT REAL"

A common assumption is that AI-generated feedback will feel generic or robotic. The learner data challenges this directly.

"The feedback felt genuine and was helpful."

"The AI feedback was specific to my response."

"It allowed me to practice patient-centered communication in a safe, simulated environment, helping me refine how I elicit patient motivations, address barriers, and collaboratively set realistic goals."

Some of these quotes come from healthcare professionals – a population with high standards for clinical accuracy and low tolerance for superficial feedback. These professionals report that AI coaching helped them practice the specific clinical communication skills the module targeted.

This perception of feedback as actionable and contextually relevant is not limited to clinical settings. In enterprise environments, similar value is reported in operational contexts:

"The Artha team turned around a working AI coach for us at remarkable speed – the rubric is working great, and our review team thought the feedback length and tone were spot on."

–Ailene Cantelmi, Director, Educational Development, The France Foundation



THEME 3: “I DIDN’T EXPECT THIS FROM TRAINING”

The most telling theme: learners’ expectations were exceeded. They came expecting standard compliance training. They found something different.

“I was surprised how effective the AI component was. It made me really consider what precisely I might say to a student.”

“I was surprised when it provided feedback for my responses, but I loved it.”

“I’m curious how I could create a similar module for class instruction. Very cool, and made it fun.”

When required compliance training produces delight and curiosity—when learners ask how to build this for their own courses—something has fundamentally shifted in the learning experience.

Cross-Domain Validation

Iowa State trained university staff in student support conversations. The France Foundation trained clinicians in motivational interviewing. GardaWorld trained frontline workers and supervisors in customer-facing interactions, decision-making, and adherence to standard operating procedures. These three are featured here because the data is most complete, but they are a subset of a much broader portfolio – AI coaches have also been built for claims adjusters at an insurance firm, board directors at a non-profit, front-line workers using a coach for personal growth, and others. Different domains, different learner populations, different skill types. The pattern holds: learners value the active practice, respond positively to the feedback quality, and frequently report that the experience feels more effective than traditional approaches. This consistency suggests the underlying mechanism – active practice with personalized feedback – transfers across contexts. The value is not domain-specific; it is design-specific.

With AI, we can deliver the kind of personalized, real-time coaching that traditional training just can’t match. My goal is to use AI to give learners the safe, adaptive practice they need to master tough skills like negotiation. The Claims coach we created at Liberty Mutual with Artha achieves just that without losing sight of human accountability.

–Marlie Cardiff Sr Instructional Designer Specialist (Retired) Liberty Mutual

✦ KEY TAKEAWAY

When passive content is replaced with AI-powered practice and feedback, learners do not just complete the training. They engage with it, value it, and report confidence in the skill.



05. Limitations

This paper presents post-course survey data, not longitudinal behavior change measurement. Confidence is self-reported, not observed. There is no control group comparison—these were production deployments, not randomized controlled trials. These are real constraints. They also reflect the reality of measuring AI coaching in production rather than in a lab. The converging evidence from academic research, industry data, and primary learner feedback provides a strong triangulated foundation, but further research—particularly controlled studies measuring on-the-job behavior change—would strengthen the evidence base.

There also exists a resistance to this technology on various grounds. We cover some of the primary ones below:

“Learners Don’t Want AI in Their Training”

Of 473 learners at Iowa State, **11 expressed explicit discomfort with the AI component—that is 2.3%**. Of those 11 comments, the majority objected to AI as a concept (environmental concerns, philosophical opposition) rather than to the learning experience itself. Only a small handful questioned the pedagogical value.

The strongest counter-evidence comes from a learner who acknowledged the resistance and still endorsed the approach:

“Although I know some people will bristle at anything AI, I appreciated having to type in my response... because it helped me put the practices into my own words and style. Getting the immediate feedback on it was also helpful.”

In this deployment, learner response to the AI component was overwhelmingly positive, with only 2.3% expressing discomfort.

“AI Feedback Can’t Be Trusted”

Trust in AI feedback is a function of design, not technology. The France Foundation addressed this through expert-designed rubrics, clear disclaimer language about the AI coach’s purpose and limitations, and calibration against expert faculty expectations. The learner validation—“the feedback felt genuine and was helpful” – suggests that when the design is rigorous, learners trust the output. However, this also proves the critical importance of **solid instructional design** when building such bots.



“What About Hallucination and Safety?”

AI coaching in a closed sandbox is fundamentally different from open-ended ChatGPT. The AI references only approved, expert-validated content. Rubrics constrain the feedback to specific, known dimensions. Learners are not asked to trust the AI as a source of truth—they are asked to practice a skill and receive feedback on their practice. The risk profile is categorically different from general-purpose AI use.

"An AI coach isn't always the right answer. But when it is, and when it's designed well, it works beautifully. That means applying everything we know about how adults learn, a deep understanding of the learner's context, and disciplined work on the content we train the AI on before the coach goes live. We know now that learners appreciate and engage with live AI-generated contextual content. This is a new type of learning, and L&D needs to develop the craft to create it well."

—Garima Gupta, Founder & CEO, Artha Learning Inc. and AIReady



06. Design Principles for Responsible AI Coaching

The evidence in this paper makes the case that AI coaching works. But it does not work automatically. The technology is a delivery mechanism – the design and the discipline behind it are what determine whether AI coaching produces real learning, and whether it can withstand scrutiny from regulators, accreditors, and learners themselves.

The principles that follow are the ones we have developed at Artha Learning across deployments in higher education, healthcare, and enterprise. They sit at the intersection of learning science, instructional design craft, and responsible AI practice.

How We Design AI Coaching

PRINCIPLE 1 Expert-Informed Rubrics	AI feedback is only as good as the rubric it evaluates against. Subject matter experts define what “good” looks like (the dimensions, the language, the edge cases) before any AI is involved. Without expert-informed rubrics, AI coaching is just a text box with a chatbot.
PRINCIPLE 2: Closed Sandbox, Not Open Internet	The AI references only approved, vetted content. Learners stay in the flow of the course. No prompt engineering. This is what separates AI coaching from “just use ChatGPT” and what makes it safe for regulated environments.
PRINCIPLE 3 Production Over Recognition	Force learners to write, not select. This is the single most important design decision that research has proven makes a big difference in learner’s understanding of the content.
PRINCIPLE 4 Immediate, Specific Feedback	Not “good job” but “you acknowledged the student’s concern but did not connect them to a specific resource.” When feedback is deeply personalized to the actual response, learners actively engage and put in the work needed.
PRINCIPLE 5 Retry Opportunity	Learners try again after receiving feedback. This creates the practice-feedback loop that drives skill development and signals that the goal is learning, not assessment.
PRINCIPLE 6 Embedded, Not Bolted On	AI coaching lives inside the course – within the existing LMS, inside the SCORM package, in the flow of the experience. Learners do not leave, log in elsewhere, or learn a new interface. Removing friction for learners to try a new way of learning is important.



How We Build It Responsibly

The principles above describe what the learner experiences. The next four describe what it takes to get there. They are less visible in the final product but, in our experience, they are what separates an AI coach that holds up from one that breaks down in the wild.

PRINCIPLE 7

Heavy Upfront Investment in Tone and Voice

The AI coach has a personality whether we design it or not. Before a single rubric is written, we work with the client to define how the coach should sound – warm or precise, encouraging or direct, peer or expert. We calibrate against the brand voice, the learner population, and the emotional stakes of the content (a coach for student support conversations sounds different from one for SOP compliance). This work feels disproportionate to its visible footprint, but it is what makes the feedback land as genuine rather than generic – the single quality learners called out most often in our surveys.

PRINCIPLE 8:

Designing for the Unexpected

Learners do not always respond the way the rubric anticipates. They go off-topic. They disclose something sensitive. They try to trick the system. They give one-word answers. They write in their second language. Before launch, we map the response space well beyond the “happy path” – including refusals, escalations to human support, and graceful handling of off-rubric content. A coach that only works when learners cooperate is not a coach.

PRINCIPLE 9

Multi-Method Testing Before Launch

A single round of QA is not enough. We test in layers: rubric calibration against expert-scored sample responses; adversarial testing for prompts designed to break the coach; pilot cohorts of real learners under observation; SME review of edge-case outputs; accessibility and device testing inside the LMS. Each method catches a different class of failure and adds to confidence.

PRINCIPLE 10

Continuous Audit After Production

Launch is not the finish line. Logs of every learner-AI exchange should be downloaded, accessible, and reviewed on a regular and frequent cadence. Routine sampling reveals how learners are actually using the system, how the AI is responding, where outputs drift from the original design intent, and which rubric blind spots only surface at scale.



Where the Industry Is Heading

These principles were developed from learning science and field experience, not from a regulatory checklist. But the broader industry is converging on a similar standard. In early 2026, the Accreditation Council for Continuing Medical Education (ACCME) published the most detailed accrediting-body guidance to date on AI in education, with seven requirements spanning source integrity, human oversight, transparent disclosure, learner data protection, prohibited uses, internal governance, and secure infrastructure.

We read this convergence as a signal that responsible AI in learning is graduating from experimental curiosity to accountable practice - and that the organizations who will earn learner trust and regulatory confidence over the next decade are not the ones using the most AI, but the ones who can show their work.

✦ KEY TAKEAWAY

It is not the AI. It is the design - and the discipline, frameworks and the thinking behind it which creates the learner impact. Responsible AI coaching development is essential to build contextual, meaningful and consistent coaches to support the learners.



07. Conclusion: The Shift That's Already Happening


The evidence presented in this paper converges from three independent directions.

Academic research says passive learning fails. Billions are spent on content delivery that produces single-digit improvements in practice. The forgetting curve, the knowing-doing gap, and the information deficit model all point to the same conclusion: information alone does not change behavior. Practice and feedback do.

Industry data shows the profession is moving. Over half of L&D practitioners are actively using AI. AI role plays have gone mainstream. AI tutors lead the exploration of agentic capabilities. The enterprises deploying AI coaching at scale—are reporting big improvements in pass rates, dramatic reductions in time to proficiency, and qualitative shifts in learner engagement.

Primary learner data from varied learner populations provides evidence supporting the effectiveness of the approach in production settings. Many learners called out AI practice as the best part of the course. Only 2.3% of surveyed Iowa State learners expressed discomfort with the AI component, although these findings should be interpreted within the context of a single institutional deployment.

The eLearning industry built its foundation on content. The next era will be built on practice. The shift is already happening and our learners are expecting it.





References

1. Forsetlund, L., O'Brien, M.A., et al. (2021). Continuing education meetings and workshops: effects on professional practice and healthcare outcomes. *Cochrane Database of Systematic Reviews*, 9(9), CD003030.
2. Ekstrom, V. (2026). Digital choice architecture in medical education: Applying behavioral economics to online learning environments. *JMIR Medical Education*, 12, e86497. <https://doi.org/10.2196/86497>
3. Mertens, S., Herberz, M., et al. (2022). The effectiveness of nudging: A meta-analysis of choice architecture interventions across behavioral domains. *PNAS*, 119(1), e2107346118.
4. Murre, J.M.J. & Dros, J. (2015). Replication and analysis of Ebbinghaus' forgetting curve. *PLoS One*, 10(7), e0120644.
5. Davis, D.A., et al. (2006). Accuracy of physician self-assessment compared with observed measures of competence: A systematic review. *JAMA*, 296(9), 1094–1102.
6. Bieleke, M., Keller, L. & Gollwitzer, P.M. (2021). If-then planning. *European Review of Social Psychology*, 32(1), 88–122.
7. Légaré, F., et al. (2015). The majority of accredited continuing professional development activities do not target clinical behavior change. *Academic Medicine*, 90(2), 197–202.
8. Taylor, D.H. & Vinauskaitė, E. (2025). AI in L&D 2025: The Race for Impact. The L&D Survey Series. <https://donalghtaylor.co.uk/the-research-base/>
9. Synthesia. (2026). AI in Learning & Development Report 2026: From Experimentation to Everyday. With Dr. Philippa Hardman, Kevin Alster & Kristen Budd.
10. Thaler, R.H. & Sunstein, C.R. (2008). *Nudge: Improving Decisions About Health, Wealth, and Happiness*. Yale University Press.
11. Accreditation Council for Continuing Medical Education (ACCME). (2026). Guidance on the Responsible Use of Artificial Intelligence (AI) in Accredited Continuing Education (CE). <https://accme.org/resource/guidance-on-ai/>
12. Allied Market Research. (2024). Corporate training market by training program 2024–2035. Allied Market Research. <https://www.alliedmarketresearch.com/corporate-training-market-A06445>
13. Arizton. (2024). Continuing medical education market – Global outlook & forecast 2023–2028. Arizton Advisory & Intelligence. <https://www.arizton.com/market-reports/continuing-medical-education-market-size>
14. Cochrane Effective Practice and Organisation of Care (EPoC). (2021, October 11). Educational meetings and workshops for healthcare professionals: What are the effects? <https://epoc.cochrane.org/news/educational-meetings-effects>
15. BMC Public Health. Allan, J. L., Johnston, M., & Campbell, N. (2008). Why do people fail to turn good intentions into action? The role of executive control processes in the translation of healthy eating intentions into action in young Scottish adults. <https://pmc.ncbi.nlm.nih.gov/articles/PMC2346472/>
16. JAMA. Davis, D. A., Mazmanian, P. E., Fordis, M., Van Harrison, R., Thorpe, K. E., & Perrier, L. (2006). Accuracy of physician self-assessment compared with observed measures of competence: A systematic review. <https://pubmed.ncbi.nlm.nih.gov/16954489/>



About



About the Author

Garima Gupta is the Founder and CEO of Artha Learning Inc. and AIReady, based in Mississauga, Ontario, Canada. She holds an MEd in Adult Education and a CTDp designation, with over 18 years of experience in learning and development across corporate, healthcare, academic, and association sectors. She is a recognized speaker on AI in learning design and ways adults learn best.



About AIReady™

AIReady™ is an AI-powered learning platform developed by Artha Learning Inc. that embeds AI coaching, practice, and feedback capabilities directly within eLearning courses. The AI coach deployments described in this paper were powered by AIReady™ technology. For more information: aiforInd.com



About Artha Learning Inc.

Artha Learning Inc. is a learning design company specializing in custom eLearning development, AI-powered learning experiences, and the AI Accelerator Certificate program. Artha works across corporate, healthcare, academic, and association clients. arthalearning.com

This whitepaper may be shared freely with attribution. Please cite as:

Gupta, G. (2026). Practice, Not Content: The Evidence for AI Coaching in Learning and Development. Artha Learning Inc.