

Market-Level Capability Gap Analysis

AI Literacy for the Modern Workforce

Executive Summary

The global knowledge workforce is operating generative AI tools at scale without the foundational understanding required to use them accurately or evaluate their outputs reliably. Eighty-six percent of employers expect AI to transform their businesses by 2030, 63% cite workforce skill gaps as a barrier to business transformation, more than any other barrier, and half of executives name missing skills as the top barrier to AI adoption¹. One survey now shows both sides of the gap between individual and organizational return. Eighty percent of respondents to McKinsey's 2026 global survey report that AI improved their own productivity. The share reporting any earnings impact at their organization stays at 37%². At the individual level, professionals report broad time savings from AI use while simultaneously demonstrating reduced critical engagement with AI-generated outputs, a divergence with direct consequences for output quality, compliance exposure, and organizational productivity. The adoption gap is not evenly distributed: geographic disparities in AI use mean that organizations with internationally distributed teams face internal competency imbalances on top of market-level skill shortfalls. A further structural barrier compounds the skill gap: 69% of professionals who use AI tools conceal that use from colleagues³, a dynamic that individual training programs alone cannot resolve. Existing solutions (vendor-branded tool certifications, platform content libraries, and custom enterprise programs) address tool operation rather than evaluative judgment, measure completion rather than behavior change, and ignore the social conditions that determine whether individual skill translates into organizational practice. This analysis documents the market context, performance gap, business impact, change management dimension, and limitations of current solutions that define the design space for a structured, competency-based AI literacy program targeting the broad mid-career knowledge workforce.

1. Market Context

The generative AI adoption curve has moved from early-adopter experimentation to broad workforce expectation faster than organizational capability has grown. Since the release of ChatGPT in November 2022, investment flows into AI have increased nearly eightfold⁴. The supply side has kept spending. Company guidance compiled by analyst Benedict Evans puts the four largest hyperscalers' planned 2026 capital expenditure at about \$700 billion⁵. Global telecoms, by comparison, spend about \$300 billion. That spending has run well ahead of the enterprise deployment figures below, and the capability gap this program addresses is one reason. Investment intent has not slowed. Sixty percent of respondents expect their organization to raise AI spending in the coming year. The survey authors note that conviction in AI grows faster than the returns organizations can attribute to it⁶. Diffusion history supports the lag. Public cloud still held only about 30% of enterprise workloads after more than a decade, according to the Goldman Sachs CIO Survey as reproduced by Evans⁷. MGI measures the same precedent by company rather than by workload: as late as 2023, only about one in five companies ran most applications in the cloud, a technology available since the mid-2000s⁸.

The organizational response is already visible in employer planning data: 86% of employers expect AI and information processing technologies to transform their business by 2030, with Financial Services (97%) and Electronics (95%) sectors above the global average⁹. Across all the macrorends the survey covers, upskilling the workforce is the most commonly anticipated response, with 85% of surveyed employers planning to adopt this strategy over the 2025–2030 period¹⁰. Asked

about AI specifically, 77% plan to reskill existing staff to work alongside it⁹. Employers expect demand for AI-specific capability to rise fastest: AI and big data top their list of skills growing in importance through 2030, followed by networks and cybersecurity and technological literacy¹¹.

Adoption, however, is neither uniform nor matched by competency, and the disparities that exist across geographic boundaries have direct consequences for multinational organizations. Anthropic's Economic Index documents that Singapore uses Claude at 4.6 times the rate expected given its working-age population; Canada at 2.9 times. Indonesia, by contrast, sits at 0.36 times expected usage; India at 0.27 times; Nigeria at 0.2 times¹². A second dataset points the same way from the developer side. On OpenRouter, a routing platform serving 300-plus models, Asia's share of spend more than doubled from about 13% to about 31% between late 2024 and late 2025. Singapore alone accounted for 9.2% of token volume, second only to the United States¹³. The authors assign region by billing location and caution that enterprise accounts may aggregate activity across several regions under one billing entity¹⁴, so Singapore's share may reflect billing entities rather than users. Employer intent in the high-usage economies is already high. Ninety-seven percent of employers in Singapore plan to prioritize upskilling and 82% plan to reskill staff to work alongside AI, and they expect 70 of every 100 workers to need training by 2030, against 59 globally; 58% still expect to prioritize university degrees in hiring¹⁵. In Canada, 97% expect AI to transform their operations, 90% plan reskilling alongside AI, and employers expect 68 of 100 workers to need training¹⁶.

For an organization operating teams across Singapore, Jakarta, and Mumbai, the usage disparity is also an internal capability gap. Teams in the same organization, collaborating on the same deliverables, may be operating AI tools at markedly different rates and levels of fluency. Without a standardized competency baseline, distributed workforces inherit the global adoption disparity as an internal productivity and equity problem. A structured AI literacy program that establishes shared capability across geographic contexts puts internationally distributed teams on common professional footing.

Within enterprise deployment more broadly, institutional inertia and fixed adoption costs concentrate early AI use among specialized tasks where deployment is straightforward, capabilities are robust, and economic returns are clear, a pattern that systematically excludes the broad knowledge workforce from the productivity gains that AI adoption makes possible¹⁷. Morgan Stanley's CIO survey, reproduced by analyst Benedict Evans, tracks enterprise deployment across two survey waves¹⁸. In March 2025 about a quarter of CIOs had deployed at least one LLM project in production¹⁹. By September 2025 the charted share had risen to roughly a third, although the slide's headline repeats the March wording¹⁸. In both waves about 40% expected no first project in production before 2026. Six months raised the deployed share but did not shrink the share still waiting. McKinsey's mid-2026 survey, which asks about AI use in general rather than LLM projects, shows movement but no resolution. Among organizations using AI, 22% were still experimenting, 34% piloting, and 44% scaling across the enterprise, up from 38% a year earlier²⁰. Its figures are not comparable with the CIO counts above.

Enterprise usage records through March 2026 show who moved first. Among U.S. public companies, ChatGPT Enterprise adopters had a median workforce of 2,934 against 424 for non-adopters²¹. Firms in the top revenue quartile were 6.9 percentage points more likely to adopt²². Adoption tracked prior stocks of R&D, SG&A (selling, general, and administrative spending, the authors' measure of organizational capability), and capitalized software rather than scale alone²³. The authors warn that this pattern could widen productivity differences across firms even when the models are broadly available²⁴. Their estimates cover large public firms and say nothing about small private firms or mid-market companies²².

The bottleneck is not hardware, licensing, or access. It is the organizational and human infrastructure required to convert tool availability into reliable, judgment-informed use at scale. Two McKinsey publications make the same point. In commentary on the 2026 enterprise survey, one of its authors concludes that the limiting factor is increasingly the organization's ability to

absorb change²⁵. MGI argues that applying AI to individual tasks inside legacy workflows is unlikely to deliver the available productivity gains²⁶. Employers across 52 of 55 surveyed economies and 19 of 22 sectors rank workforce skill gaps as the leading barrier to business transformation, up from 60% to 63% since the 2023 edition of the same report²⁷.

2. Documented Performance Gaps

The gap between AI tool access and productive AI use is measurable at the individual level. The Anthropic Interviewer study (December 2025) draws on qualitative and quantitative data from 1,250 professionals (1,000 from the general workforce, 125 scientists, 125 creatives)²⁸. Five limits apply. Participants were recruited through crowdworker platforms, which the authors say may tilt responses toward more positive or experienced views. Claude conducted the interviews, so participants discussed AI with an AI. The post serves product development as well as research. The authors qualify the headline finding as holding 'in our sample.' And the findings are self-reports, not observed behavior.

At the aggregate level, adoption is broadly positive on self-reported outcomes: 86% of professionals report that AI saves them time, and 65% report satisfaction with the role AI plays in their work. These figures confirm that tools are being used and perceived as valuable, but not that they are being used well.

The performance gap lies in judgment quality. The Anthropic Interviewer data reveals a meaningful divergence between how professionals characterize their AI use and what task-level behavioral data shows. Sixty-five percent of participants describe AI's primary role in their work as augmentative: the model supports their judgment rather than replaces it. Anthropic's task-level analysis of actual Claude usage, however, shows a more even split: 57% augmentative, 43% automative²⁹. The size of that gap depends on the baseline: eight percentage points against Handa et al., widening to eighteen against the more recent Economic Index wave the Anthropic Interviewer cites (August 2025 data, approximately 47% augmentative)³⁰, an 8-to-18-point overestimate of augmentative use depending on the dataset^{29,31}. The discrepancy suggests that a meaningful share of users believe they are exercising oversight over AI outputs when the task structure indicates otherwise, a blind spot with direct consequences for output quality and compliance exposure.

The critical thinking research adds to this finding. Lee et al. (CHI 2025), surveying 319 knowledge workers across 936 real-world AI-assisted work tasks, found that higher confidence in generative AI is associated with less critical thinking, while higher confidence in one's own ability to do the task without AI is associated with more³². The authors measured task-specific self-confidence, not domain expertise as such, and only speculate that building expertise would help³³. The study identifies a pattern of overreliance (users report less critical thinking on tasks where they are more confident in the tool) and finds three kinds of barrier to critical engagement: awareness barriers, when tasks are perceived as trivial (55 of 319) or the tool is trusted (83 of 319); motivation barriers such as time pressure (44 of 319); and ability barriers such as difficulty inspecting AI responses, for example without enough domain knowledge (58 of 319)³⁴. A 2026 university review found the same pattern. MIT's AI committee heard from learners that the temptation to offload thinking peaked under deadline pressure, and most respondents to its campus survey said they worry about becoming dependent on AI for their academic work³⁵. The most common written-in concern was that reliance could erode critical thinking and long-term skill³⁶. This is the structural risk that tool familiarity without foundational model understanding creates: a workforce that uses AI frequently, reports satisfaction, and scrutinizes outputs less where it is more confident in the tool.

The productivity data from Tamkin and McCrory (2025) establishes what is at stake when this gap is closed. Across 100,000 real-world Claude conversations, AI reduces median task completion time by 81%, with the distribution concentrated in the 50–95% range and peaking between 80–90%³⁷. Specific task categories show sharper variation: compiling information from reports yields approximately 95% time savings; document drafting, 87%; financial analysis tasks, 80%³⁸. These figures derive from observational analysis of Claude conversations and sit at the optimistic end of the evidence; the randomized controlled trials referenced in the same work report more conservative task-time reductions in the 14–56% range. The program treats the 81% median as a capability-contingent ceiling rather than a baseline expectation: the distance between the observational median and the controlled-trial range depends in part on how much refinement and verification users must add to the model's output, which the report's in-conversation estimates do not capture³⁹. Translating task-level efficiencies to macroeconomic terms, the analysis implies a 1.8% annual increase in US labor productivity over the next decade under broad adoption, double the growth rate observed since 2019⁴⁰. Anthropic's January 2026 report replicates the 1.8% on a larger sample and then discounts each task's speedup by the model's own estimate of its success rate, which lowers the implied gain to 1.2 points for consumer use and 1.0 for enterprise API use⁴¹. That potential is contingent on workforce capability, not tool deployment alone. Context constrains sophisticated use: curating appropriate context for AI models will be important for high-impact deployments, and for some organizations the bottleneck may be less model capability than the data and organizational infrastructure needed to give the model the context complex work requires⁴²; the human side of that infrastructure is the capability this program addresses.

Tools are deployed and productivity gains are documented at the task level, but the workforce lacks the evaluative judgment to verify outputs, the conceptual foundation to understand model failure modes, and the competency vocabulary to apply these skills consistently across tasks and contexts.

3. Business Impact

The business consequences of the documented performance gap operate across three dimensions: direct productivity loss from underutilization, quality and compliance risk from uncritical output acceptance, and structural competitive disadvantage as AI-capable peers accelerate.

Productivity underutilization. Tamkin and McCrory (2025), using O*NET occupational classifications matched to BLS wage data, estimate that the median AI-assisted task would cost \$54 in professional labor to complete unassisted (the worker's hourly wage multiplied by the task's unassisted duration). By occupation category, business and financial tasks average \$69; management tasks, \$133; computer and mathematical tasks, \$82⁴³. Time savings convert directly into recovered labor cost. At the 81% observational median, a \$133 management task takes roughly \$25 of labor to complete with proficient AI use, recovering about \$108 in labor value each time the task occurs; at the more conservative 14–56% controlled-trial range, the same task still recovers between roughly \$19 and \$74. The underutilization cost is the share of that recoverable value a worker forfeits through surface-level engagement: single-turn prompting, uncritical output acceptance, or avoidance of AI for tasks where judgment concerns create hesitation. The forfeited value recurs with every occurrence of every AI-amenable task, across every knowledge worker in the organization. The 1.8% labor productivity implication from Tamkin and McCrory (2025) is a macroeconomic estimate under broad, proficient adoption⁴⁰ (1.0 to 1.2 points after the January 2026 success adjustment)⁴¹; the organizational delta between proficient and non-proficient use is the addressable gap a competency program closes.

Quality and compliance risk. Lee et al. (CHI 2025) document that confidence in AI is associated with less critical engagement, that trust in the tool can discourage critical reflection, that some users who doubted their own ability accepted outputs by default, and that a lack of domain knowledge makes outputs hard to inspect⁴⁴. For knowledge workers in functions where output accuracy carries legal, financial, or reputational stakes (compliance, finance, legal, HR, product), this reduction in scrutiny is a material risk exposure. The WEF (2025), summarizing the field experiment it cites, notes that generative AI can enhance knowledge work when applied appropriately within its capability and risks adverse outcomes where users unknowingly stretch it beyond that capability⁴; the report's Box 3.1 skills analysis, run with Indeed, adds that human oversight remains crucial even where the tool can assist⁴⁵. Without a framework for recognizing those capability boundaries (for instance, that a model generating a plausible regulatory citation may be confabulating rather than retrieving), knowledge workers in high-stakes functions have no reliable basis for determining when AI output requires independent verification and when it does not. The concern is not confined to the workplace. MIT's 2026 review names hallucinated citations as a specific risk to train against and reports hallucination as respondents' leading concern about AI⁴⁶. It states that whoever submits AI-assisted work is responsible for any inaccurate content in it⁴⁷.

Competitive positioning. Skill gaps in the labor market are the primary barrier to business transformation cited by employers for the 2025–2030 period, named by 63% of respondents, a figure that has increased since the 2023 report and ranks first across 52 of 55 economies and 19 of 22 sectors surveyed²⁷. Organizations that close the individual competency gap ahead of competitors gain an advantage that grows over time: their workforce applies AI earlier, more accurately, and across a broader range of tasks. McKinsey's 2026 survey profiles the 6% of respondents whose organizations attribute at least 5% of earnings, and significant value, to AI. Compared with the rest, they are about three times as likely to have fundamentally redesigned workflows (nearly three-quarters against one-quarter). They are twice as likely to have defined processes for measuring AI's impact and twice as likely to report visible senior-leader commitment. They are 3.3 times as likely to intend to transform their business within three years⁴⁸. Individual competency is the input those practices act on. The survey's authors expect the organizations that turn individual productivity gains into lasting enterprise returns to be those that transform their businesses, not those that only adopt tools⁴⁹. WEF respondents project that by 2030, the share of tasks performed by humans alone will decline from 47% toward near-parity with machine and human-machine collaboration categories⁵⁰, a structural shift that rewards organizations whose workforces bring sound judgment to human-machine collaboration.

The business case is a comparison, measurable in current data, between the productivity trajectory of a competency-developed workforce and one that remains at surface-level adoption.

4. Change Management Context

Training design alone cannot close the documented performance gap. The Anthropic Interviewer (2025) identifies a behavioral dynamic that operates independently of, and in some cases contrary to, individual skill acquisition: 69% of professionals surveyed report that social stigma is an active barrier to AI use at work³. The finding is not that workers reject AI tools (adoption rates and self-reported time savings contradict that reading) but that a majority of current AI users conceal their usage from colleagues. One participant's account is representative: a fact-checker reported that when a colleague expressed hostility toward AI, they said nothing and did not disclose their own practice, because they knew how many people feel about AI.

The mechanism is concealment, not resistance. Employer planning data helps explain why concealment is a rational response rather than a character flaw. Forty-one percent of employers expect to downsize where AI can replicate people's work, and 47% plan to transition staff from AI-disrupted roles to other positions⁵¹. The same employers rank organizational culture and resistance to change as the second barrier to transformation (46%)²⁷. In a separate survey of more than 11,000 executives, 43% name lack of vision among managers and leaders as a barrier to AI adoption⁹. The gap involves managers as well as the workers who report to them. Workers who have individually adopted AI tools, who report satisfaction and time savings, and who would be candidates for peer knowledge-sharing instead operate in isolation, each navigating AI use privately, without shared vocabulary, organizational norms, or social permission to make their practice visible.

This dynamic has two direct consequences for program design. First, a competency program that addresses only individual skill acquisition leaves the organizational conditions that suppress transfer entirely unaddressed. A worker who completes a program and returns to a team culture where AI use signals laziness or replaceability is no more likely to visibly apply those skills than before. Besides the gap between what workers say and do, a second perception gap runs between what employers rate as effective and what workers value. In the WEF's comparison with ADP Research's employee survey, employers rank well-being first and upskilling second among practices to increase talent availability, while employees, asked their top reasons to stay with an employer, rank them eighth and seventh and put working hours first where employers put it eighth⁵². The two surveys ask different questions, so the comparison shows a gap in priorities, not a measure of demand for training. A program that assumes workers value the training their employer plans is building on the employer's half of that gap. Second, the absence of shared professional vocabulary makes peer normalization impossible. Workers cannot calibrate their AI use against colleagues, identify best practices through informal observation, or develop team-level norms without language that makes AI work discussable in the first place.

The 4D Framework (Delegation, Description, Discernment, Diligence)⁵³ is therefore a change management instrument as well as a competency taxonomy. Its authors present it as a practical tool to inform discourse and practice, aimed primarily at higher education but expected to benefit organizations more widely, and compatible with other skills taxonomies in a variety of professional contexts⁵⁴. A workforce that shares common vocabulary for AI collaboration can name what they are doing, discuss where they draw delegation boundaries, and surface discernment concerns without the binary framing of "pro-AI" or "anti-AI" that makes current workplace conversations about AI tools socially fraught. Instead of a separate communications strategy, the program addresses the change management dimension through the design of the learning experience itself: shared language, normalized peer practice, and visible demonstration that evaluative judgment is the professional standard.

Mollick (2012, *Strategic Management Journal*), analyzing 854 PC games from 395 companies with just over \$4 billion in combined revenue, found that individual-level differences account for more variance in performance than organizational-level factors⁵⁵. Individual producers, the mid-level project managers, explained 22.3 percent of the variance in product revenue and designers 7.4 percent, against 21.3 percent for the firm itself. His applied work, *Co-Intelligence* (2024), extends this finding to AI adoption: the most advanced enterprise AI users typically operate without institutional authorization and self-direct their adoption because top-down rollouts optimize for compliance rather than capability⁵⁶. *Co-Intelligence* is a practitioner text rather than a peer-reviewed empirical study; its directional argument is consistent with both the 2012 findings and the Anthropic Interviewer's stigma data: the workers most capable of generating AI-enabled performance

returns are already using AI privately, outside organizational visibility, precisely because the organizational apparatus has not created conditions for that use to be normalized. A program that closes the individual competency gap while establishing shared professional vocabulary addresses both the skill deficit and the concealment dynamic simultaneously.

5. Existing Solutions and Their Limitations

Market Structure

The enterprise AI training market in 2025–2026 organizes into three tiers: vendor-branded programs that subsidize product adoption, platform content libraries that aggregate self-paced courses, and enterprise custom solutions deployed through organizational L&D functions. Each tier has its own design philosophy, delivery model, and structural limitation.

Tier 1 — Vendor-Branded, Tool-Centric Programs

The dominant programs in this category are produced and distributed by the technology companies whose products they promote. Three offerings define the tier.

Microsoft Career Essentials in Generative AI (LinkedIn Learning, developed with Microsoft) is, as of October 2026, a path of five short courses, about four hours in all, covering generative AI concepts, everyday AI concepts, AI literacy, Microsoft 365 Copilot, and ethics; its only tool-specific course teaches Copilot, and its description names Microsoft Copilot as the tool it covers⁵⁷. Microsoft launched it in 2023, called it the first professional certificate on generative AI in the online learning market, awarded it to learners who pass the assessment, and promised free access through 2025⁵⁸. The path is high-volume (its ethics course alone lists more than 800,000 viewers⁵⁹), and its one tool-specific course teaches Copilot. Its short introductions touch on how generative AI works and where it falls short, but at about four hours the path does not build practice in evaluating AI-generated output for accuracy, diagnosing why an LLM produces a plausible-sounding but false claim, or carrying those skills over when Copilot is replaced or supplemented by a different system.

Google AI Essentials Specialization (Coursera, five short courses that Google says take under 10 hours) covers foundational AI concepts, productivity tasks, prompting, and responsible AI, with activities that use Gemini and Google Workspace tools⁶⁰. An independent review judged it well suited to learners unfamiliar with AI but not to tech-savvy ones, and pointed experienced learners to more in-depth courses⁶¹. Like the Microsoft program, it teaches tool operation within a specific vendor ecosystem and provides no mechanism for ensuring that skills transfer to the learner's actual work context. Google subsequently released a Google AI Professional Certificate (Coursera, 2026), an eight-course series that applies AI to six work domains⁶². The foundational course still lists more enrolled learners (about 2.0 million against about 1.6 million for the certificate)^{60,62} and retains the same surface-level design.

Anthropic's AI Fluency Suite is the most conceptually sophisticated offering in this tier and the broadest in scope. The suite comprises five courses: *AI Fluency: Framework and Foundations* (2025)⁶³, organized around the 4D competency framework; *AI Capabilities and Limitations* (2026)⁶⁴, covering the behavioral properties of generative AI models; *Teaching AI Fluency* (2025)⁶⁵, designed for educators integrating AI literacy into formal instruction; *AI Fluency for Non-Profits* (2025)⁶⁶, an audience-specific adaptation of the core framework; and *Claude 101* (2025)⁶⁷, a practical tool-use guide that, like the

Microsoft and Google programs, is explicitly product-centric. The suite spans a wider range of contexts and pedagogical purposes than any competitor offering in this tier, and several of its courses are explicitly designed to transfer across tools rather than train learners on a single platform.

Within this suite, the *AI Capabilities and Limitations* course is a substantive advance. It frames generative AI around four behavioral properties (Next Token Prediction, Knowledge, Working Memory, and Steerability) and explicitly addresses hallucination mechanics: "a citation that *looks* like a real citation satisfies the pattern just as well as one pointing to a paper that actually exists"⁶⁴. The course teaches learners to locate tasks on a capability-limitation continuum and to treat specificity (names, dates, statistics, citations) as the zone where fabrication concentrates and verification is most warranted. This is real foundational content that most competitor offerings do not approach.

Two limitations remain relevant to the present analysis. First, the course addresses next-token prediction as a *behavioral* property: what AI outputs look like, when to scrutinize them, how to interpret confident-sounding errors. It does not address tokenization as a *mechanical process*: how text is parsed into subword fragments before prediction begins, why those boundary decisions produce systematic errors in arithmetic, character-level tasks, and non-English text, and why those failure modes are structurally predictable rather than random. The behavioral framing is valuable and actionable; the mechanical layer is the foundation that makes the behavioral framing explicable rather than rote. Second, the course's value is contingent on active engagement with its exercises, hands-on tasks designed to produce genuine calibration against real work. The platform does not enforce that engagement. A learner who clicks through without completing the exercises exits with terminology rather than capability, a distinction the course design cannot prevent and the assessment structure does not catch. Human-AI interaction research offers a parallel. Lee et al. (CHI 2025) summarize prior work showing that cognitive forcing functions, such as making the user wait before seeing AI output or edit it interactively, reduce overreliance more than AI explanations do⁶⁸, and their own design recommendations favor building users' ability over adding a per-task audit⁶⁹. Those studies tested interface designs, not training, so the parallel is an analogy: this program's twelve practice activities likewise require learners to act, where a video-and-quiz course asks them only to watch.

Both courses carry a 2026 copyright; the *AI Capabilities and Limitations* course published its first video content in April 2026. The skill deficiencies captured in the WEF Future of Jobs Report 2025, the Anthropic Economic Index reports, and the Anthropic Interviewer findings reflect a workforce that entered the generative AI era without access to these resources. The gap existed before the most sophisticated available solution did. Anthropic's own workforce study points the same way: after documenting skill-atrophy concerns and a 'paradox of supervision' among its engineers, the company lists establishing best practices for AI-augmented work, for example guided by its AI fluency framework, among the follow-up steps it is examining, and says it is supporting CodePath's adaptation of computer science curricula⁷⁰. The vendor names the framework as a candidate guide for its own staff's practice as well as a public course scaffold.

Tier 2 — Platform Content Libraries

Coursera, LinkedIn Learning, and DataCamp host extensive catalogs of AI-related courses from IBM, DeepLearning.AI, CompTIA, and independent instructors. DeepLearning.AI's *AI for Everyone* (Andrew Ng) provides a non-technical overview of AI terminology, of what AI can and cannot do, and of how to spot opportunities to apply it⁷¹; IBM's *AI Foundations for Everyone* covers conceptual underpinnings for learners without a technical background, with no-code exercises in each course⁷²; Google AI Essentials prioritizes practical generative AI productivity tasks for non-technical learners⁶⁰.

CompTIA's *AI Essentials* (version 2, released in December 2025) takes a more applied approach than earlier platform offerings. It targets employees in knowledge or desk-based roles and teaches prompting, use-case identification, and data security and privacy fundamentals through interactive, scenario-based instruction⁷³. CompTIA's own research, however, finds that only about one in three companies mandate AI training for staff⁷⁴. Platform availability has not translated into systematic enterprise deployment.

Organizations with mature, workforce-wide AI upskilling programs report significant positive AI ROI at twice the overall rate: 21% of surveyed leaders report significant positive returns, and the share rises to 42% among those with a mature, organization-wide program (YouGov, n = 500+ US and UK enterprise leaders)⁷⁵. That figure is from a vendor-commissioned survey and should be read accordingly; the directional finding is consistent with independent transfer research. The same DataCamp survey finds that 77% of organizations offer some kind of AI training, yet only 35% report having a mature, workforce-wide upskilling program⁷⁶.

The structural limitation of platform programs is that completion is self-directed and organizationally disconnected. A 95% completion rate confirms that employees clicked through modules; it does not confirm that they retained information, applied it on the job, or improved performance in any measurable respect.

Tier 3 — Enterprise Custom and Consulting Solutions

A smaller market segment offers custom-built programs deployed through internal L&D functions, often with vendor support. This tier produces stronger outcomes when implemented well, as the DataCamp data above suggests. The structural limitation is access: custom solutions require internal L&D capacity, budget allocation, and executive sponsorship that many organizations lack. Seventy-four percent of senior leaders say their companies still lack the skills they need to compete, despite more than \$400 billion a year in corporate training spending (Bersin)⁷⁷, a figure consistent with Hardman's (2026) synthesis of the transfer research literature, which places sustained behavior change rates from formal training at 10–20%⁷⁸. High-quality bespoke design does not scale by default.

Structural Limitations Across the Market

Three design failures appear consistently across all three tiers.

1. Tool familiarity without foundational model understanding. The dominant design paradigm across Tiers 1 and 2 trains learners to operate specific AI interfaces without establishing the conceptual model required to evaluate outputs, recognize failure modes, or transfer skills across tools and contexts. Lee et al. (CHI 2025) point the same way: generative AI creates new cognitive work, response integration above all, and shifts effort from doing tasks to stewarding them, and the authors recommend training in verification, integration, and stewardship⁷⁹. Tool-operation training does not target those skills. The WEF (2025) notes that generative AI risks producing adverse outcomes specifically where users unknowingly stretch the technology beyond its capability⁴, a risk that tool-use training does not equip learners to recognize or manage. A learner who can generate a Copilot summary or a Gemini draft cannot necessarily determine whether the output is accurate, where it may have introduced hallucination, or when a task should not be delegated to an AI system at all. Anthropic's *AI Capabilities and Limitations* course comes closest to addressing this, and is a genuine step toward transferable foundational understanding. It does not, however, address the mechanical layer (tokenization, subword parsing, the structural predictability of certain failure modes) that would allow learners to reason from first principles about where AI judgment applies and where it does not.

2. No behavior transfer mechanism. The most reliable indicator of effective AI training is task-level behavioral change, not course completion. None of the reviewed programs include a 30/60/90-day behavioral follow-up, manager observation protocol, or structured reinforcement loop. The measurement gap is not unique to corporate training. A 2023 systematic review of AI-literacy programs for pre-university learners found hardly any that assessed even whether participants understood AI concepts afterward (Casal-Otero et al., 2023)⁸⁰. Hardman's (2026) synthesis puts sustained behavior change from formal training at 10–20%⁷⁸. Peer-reviewed estimates put the level higher but show it falling with time: training professionals estimate that the share of employees applying what they learned drops from about 62% right after training to 34% a year later, and that support in the work environment before and after training matters more than activities during it⁸¹. Programs that are passive, decontextualized, and completion-tracked provide none of that support by design.

3. No change management dimension. The 69% social stigma finding from the Anthropic Interviewer (2025)³ represents an adoption barrier that no course completion strategy resolves. CompTIA's chief product officer has said publicly that companies need to invest in staff training and also consider the change management aspects of AI on staffing and job roles⁸², but program design across all three tiers addresses this concern rhetorically at best. None of the reviewed offerings include structural features (shared vocabulary, team-level norms, manager activation) that create the social conditions for normalized AI use.

Individual Competency as the Design Target

The market offers plenty of programs, but they are systematically misaligned with the performance gap the research identifies. As the change management analysis established, individual-level differences drive firm performance, and the workers most capable of generating AI-enabled returns are already self-directing their adoption privately, outside organizational visibility^{55,3}.

The structurally correct intervention targets the individual knowledge worker (their judgment, their evaluative vocabulary, their transfer of skill across tools and contexts) rather than organizational-level compliance protocols or executive-level strategy. Programs designed at the organizational level manage the risk of AI adoption, while programs designed at the individual competency level generate the performance return that makes adoption worthwhile; neither substitutes for the other. McKinsey's 2026 survey finds individual gains already widespread (80%) while the share reporting any enterprise earnings contribution stays flat (37%). It attributes the difference to workflow redesign and leadership practice rather than to tool insertion⁸³. Its authors frame growing employee AI fluency as the resource leaders should use to redesign end-to-end workflows⁸⁴. This program supplies the fluency and the shared vocabulary; the Level 3 and Level 4 instruments are where the organizational practices attach. An independent institution has arrived at a similar competency set. MIT's 2026 committee, whose report does not cite Dakan and Feller's framework, defined effective use in four parts: how to specify a prompt, how to verify an output, when a model is likely to hallucinate, and when not to reach for AI at all⁸⁵. It defined responsible use as distinguishing augmentation from automation and disclosing AI contributions honestly⁸⁵. Those capabilities map onto Description, Discernment, Delegation, and Diligence. The current market offers risk management at scale. Individual competency with an organizational transfer path remains a design gap.

Notes

1. World Economic Forum, "The Future of Jobs Report 2025," January 2025, pp. 49, 62, https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf
2. Dan Tinkoff, Lieven Van der Veken, and Michael Chui, with Tara Balakrishnan, "The state of AI in 2026: On the road to ROI," McKinsey & Company, August 2026, pp. 3–4, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>

3. Kunal Handa et al., “Introducing Anthropic Interviewer: What 1,250 Professionals Told Us About Working with AI,” Anthropic, December 2025, research post, “AI’s impact in the general workforce” (general-workforce sample, N = 1,000: 69% mentioned social stigma, a theme count from interview transcripts; fact-checker quotation), <https://www.anthropic.com/research/anthropic-interviewer>
4. World Economic Forum, “The Future of Jobs Report 2025,” January 2025, p. 11, https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf
5. Benedict Evans, “AI eats the world (May 2026 edition),” May 2026, slide 8 (company guidance, includes capital leases), <https://www.ben-evans.com/s/2026-Spring-AI.pdf>
6. Dan Tinkoff, Lieven Van der Veken, and Michael Chui, with Tara Balakrishnan, “The state of AI in 2026: On the road to ROI,” McKinsey & Company, August 2026, pp. 16, 29, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
7. Benedict Evans, “AI eats the world (November 2025 edition),” November 2025, slide 58 (Goldman Sachs CIO Survey), <https://www.ben-evans.com/s/2025-Autumn-AI.pdf>
8. Lareina Yee et al., “Agents, robots, and us: Skill partnerships in the age of AI,” McKinsey Global Institute, November 2025, p. 7, <https://www.mckinsey.com/mgi/our-research/agents-robots-and-us-skill-partnerships-in-the-age-of-ai>
9. World Economic Forum, “The Future of Jobs Report 2025,” January 2025, p. 62, https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf
10. World Economic Forum, “The Future of Jobs Report 2025,” January 2025, p. 52, Figure 4.4, https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf
11. World Economic Forum, “The Future of Jobs Report 2025,” January 2025, p. 37, Figure 3.4, https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf
12. Ruth Appel, Peter McCrory, Alex Tamkin, et al., “The Anthropic Economic Index Report: Uneven Geographic and Enterprise AI Adoption,” Anthropic, September 2025, p. 3, <https://www.anthropic.com/research/anthropic-economic-index-september-2025-report>
13. Malika Aubakirova et al., “OpenRouter State of AI: An Empirical 100 Trillion Token Study with OpenRouter,” OpenRouter and Andreessen Horowitz, December 2025, pp. 34–36, <https://openrouter.ai/state-of-ai>
14. Malika Aubakirova et al., “OpenRouter State of AI: An Empirical 100 Trillion Token Study with OpenRouter,” OpenRouter and Andreessen Horowitz, December 2025, p. 7, <https://openrouter.ai/state-of-ai>
15. World Economic Forum, “The Future of Jobs Report 2025,” January 2025, pp. 67, 192–193, https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf
16. World Economic Forum, “The Future of Jobs Report 2025,” January 2025, pp. 74, 124–125, https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf
17. Ruth Appel, Peter McCrory, Alex Tamkin, et al., “The Anthropic Economic Index Report: Uneven Geographic and Enterprise AI Adoption,” Anthropic, September 2025, pp. 30–31, <https://www.anthropic.com/research/anthropic-economic-index-september-2025-report>
18. Benedict Evans, “AI eats the world (November 2025 edition),” November 2025, slide 57 (Morgan Stanley CIO Survey, September 2025), <https://www.ben-evans.com/s/2025-Autumn-AI.pdf>
19. Benedict Evans, “AI eats the world (May 2025 edition),” May 2025, slide 44 (Morgan Stanley CIO Survey, March 2025), <https://www.ben-evans.com/s/2025-Spring-AI.pdf>
20. Dan Tinkoff, Lieven Van der Veken, and Michael Chui, with Tara Balakrishnan, “The state of AI in 2026: On the road to ROI,” McKinsey & Company, August 2026, p. 5, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
21. Aaron Chatterji et al., “How Organizations Use AI: Evidence from ChatGPT,” OpenAI working paper, August 2026, p. 11, <https://arxiv.org/abs/2608.12236>
22. Aaron Chatterji et al., “How Organizations Use AI: Evidence from ChatGPT,” OpenAI working paper, August 2026, p. 14, <https://arxiv.org/abs/2608.12236>
23. Aaron Chatterji et al., “How Organizations Use AI: Evidence from ChatGPT,” OpenAI working paper, August 2026, p. 15, <https://arxiv.org/abs/2608.12236>
24. Aaron Chatterji et al., “How Organizations Use AI: Evidence from ChatGPT,” OpenAI working paper, August 2026, p. 22, <https://arxiv.org/abs/2608.12236>
25. Dan Tinkoff, Lieven Van der Veken, and Michael Chui, with Tara Balakrishnan, “The state of AI in 2026: On the road to ROI,” McKinsey & Company, August 2026, p. 24, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
26. Lareina Yee et al., “Agents, robots, and us: Skill partnerships in the age of AI,” McKinsey Global Institute, November 2025, p. 35, <https://www.mckinsey.com/mgi/our-research/agents-robots-and-us-skill-partnerships-in-the-age-of-ai>
27. World Economic Forum, “The Future of Jobs Report 2025,” January 2025, p. 49, Figure 4.1, https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf
28. Kunal Handa et al., “Introducing Anthropic Interviewer: What 1,250 Professionals Told Us About Working with AI,” Anthropic, December 2025, research post, “Participants” and “How Anthropic Interviewer works” (1,250 via crowdworker platforms: general workforce N=1,000, scientists N=125, creatives N=125; 10–15 minute AI-run interviews), <https://www.anthropic.com/research/anthropic-interviewer>
29. Kunal Handa et al., “Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations,” Anthropic, February 2025, p. 3, <https://arxiv.org/abs/2503.04761>
30. Ruth Appel, Peter McCrory, Alex Tamkin, et al., “The Anthropic Economic Index Report: Uneven Geographic and Enterprise AI Adoption,” Anthropic, September 2025, p. 10, Figure 1.2, <https://www.anthropic.com/research/anthropic-economic-index-september-2025-report>

31. Kunal Handa et al., “Introducing Anthropic Interviewer: What 1,250 Professionals Told Us About Working with AI,” Anthropic, December 2025, research post, “Augmentation versus automation” (65% self-reported vs. 47% augmentation / 49% automation in the August 2025 Economic Index data; five candidate explanations), <https://www.anthropic.com/research/anthropic-interviewer>
32. Hao-Ping (Hank) Lee et al., “The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers,” Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems, 2025, p. 9, <https://doi.org/10.1145/3706598.3713778>
33. Hao-Ping (Hank) Lee et al., “The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers,” Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems, 2025, p. 14, <https://doi.org/10.1145/3706598.3713778>
34. Hao-Ping (Hank) Lee et al., “The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers,” Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems, 2025, p. 11 (awareness, motivation, and ability barriers; trivial task 55/319; trust 83/319; time 44/319; barriers to inspecting responses, such as limited domain knowledge, 58/319; barriers to revising queries and improving output 72/319), <https://doi.org/10.1145/3706598.3713778>
35. MIT Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training, “Report of MIT’s Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training,” Massachusetts Institute of Technology, August 2026, pp. 7, 32, <https://sites.mit.edu/ai-use/files/2026/08/AI-Committee-Final-Report-Aug-13.pdf>
36. MIT Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training, “Report of MIT’s Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training,” Massachusetts Institute of Technology, August 2026, p. 33, <https://sites.mit.edu/ai-use/files/2026/08/AI-Committee-Final-Report-Aug-13.pdf>
37. Alex Tamkin and Peter McCrory, “Estimating AI Productivity Gains from Claude Conversations,” Anthropic, November 2025, p. 12; p. 13, Figure 6, <https://www.anthropic.com/research/estimating-productivity-gains>
38. Alex Tamkin and Peter McCrory, “Estimating AI Productivity Gains from Claude Conversations,” Anthropic, November 2025, pp. 10, 12, <https://www.anthropic.com/research/estimating-productivity-gains>
39. Alex Tamkin and Peter McCrory, “Estimating AI Productivity Gains from Claude Conversations,” Anthropic, November 2025, p. 13, <https://www.anthropic.com/research/estimating-productivity-gains>
40. Alex Tamkin and Peter McCrory, “Estimating AI Productivity Gains from Claude Conversations,” Anthropic, November 2025, p. 3, <https://www.anthropic.com/research/estimating-productivity-gains>
41. Ruth Appel, Maxim Massenkoff, Peter McCrory, et al., “The Anthropic Economic Index Report: Economic Primitives,” Anthropic, January 2026, p. 48, <https://www.anthropic.com/research/anthropic-economic-index-january-2026-report>
42. Ruth Appel, Peter McCrory, Alex Tamkin, et al., “The Anthropic Economic Index Report: Uneven Geographic and Enterprise AI Adoption,” Anthropic, September 2025, pp. 5, 41, 44, <https://www.anthropic.com/research/anthropic-economic-index-september-2025-report>
43. Alex Tamkin and Peter McCrory, “Estimating AI Productivity Gains from Claude Conversations,” Anthropic, November 2025, p. 10, <https://www.anthropic.com/research/estimating-productivity-gains>
44. Hao-Ping (Hank) Lee et al., “The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers,” Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems, 2025, pp. 11, 14, <https://doi.org/10.1145/3706598.3713778>
45. World Economic Forum, “The Future of Jobs Report 2025,” January 2025, p. 44, https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf
46. MIT Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training, “Report of MIT’s Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training,” Massachusetts Institute of Technology, August 2026, pp. 22, 33, <https://sites.mit.edu/ai-use/files/2026/08/AI-Committee-Final-Report-Aug-13.pdf>
47. MIT Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training, “Report of MIT’s Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training,” Massachusetts Institute of Technology, August 2026, p. 16, <https://sites.mit.edu/ai-use/files/2026/08/AI-Committee-Final-Report-Aug-13.pdf>
48. Dan Tinkoff, Lieven Van der Veken, and Michael Chui, with Tara Balakrishnan, “The state of AI in 2026: On the road to ROI,” McKinsey & Company, August 2026, pp. 17–19, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
49. Dan Tinkoff, Lieven Van der Veken, and Michael Chui, with Tara Balakrishnan, “The state of AI in 2026: On the road to ROI,” McKinsey & Company, August 2026, p. 29, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
50. World Economic Forum, “The Future of Jobs Report 2025,” January 2025, p. 26, https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf
51. World Economic Forum, “The Future of Jobs Report 2025,” January 2025, p. 63, Figure 4.14, https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf
52. World Economic Forum, “The Future of Jobs Report 2025,” January 2025, pp. 53–54, https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf
53. Rick Dakan and Joseph Feller, “Framework for AI Fluency: Practical Overview Document (Version 1.1),” Ringling College of Art and Design; University College Cork, January 2025, p. 3, <https://ringling.libguides.com/ai/framework>
54. Rick Dakan and Joseph Feller, “Framework for AI Fluency: Practical Overview Document (Version 1.1),” Ringling College of Art and Design; University College Cork, January 2025, pp. 1, 3, <https://ringling.libguides.com/ai/framework>

55. Ethan Mollick, “People and Process, Suits and Innovators: The Role of Individuals in Firm Performance,” *Strategic Management Journal*, 33(9), 1001–1015, 2012, <https://doi.org/10.1002/smj.1958>
56. Ethan Mollick, “Co-Intelligence: Living and Working with AI,” *Portfolio* (Penguin Random House), April 2024, <https://www.penguinrandomhouse.com/books/741805/co-intelligence-by-ethan-mollick/>
57. Microsoft and LinkedIn, “Career Essentials in Generative AI by Microsoft and LinkedIn (learning path),” LinkedIn Learning, accessed October 2026, path description and course list (path updated October 2, 2026), <https://www.linkedin.com/learning/paths/career-essentials-in-generative-ai-by-microsoft-and-linkedin>
58. Kate Behncken, “Microsoft Launches New AI Skills Training and Resources as part of Skill for Jobs Initiative,” LinkedIn, June 2023, section “New, free skilling content”, <https://www.linkedin.com/pulse/microsofts-launches-new-ai-skills-training-resources-part-behncken>
59. Microsoft and LinkedIn, “Career Essentials in Generative AI by Microsoft and LinkedIn (learning path),” LinkedIn Learning, accessed October 2026, course list, “Ethics in the Age of Generative AI” (829,611 viewers), <https://www.linkedin.com/learning/paths/career-essentials-in-generative-ai-by-microsoft-and-linkedin>
60. Google, “Google AI Essentials Specialization,” Coursera, accessed October 2026, <https://www.coursera.org/specializations/ai-essentials-google>
61. Hannah Knight, “I Took Google’s AI Essentials Course, and This Is What I Learned,” *MakeUseOf*, June 2024, <https://www.makeuseof.com/google-ai-essentials-course/>
62. Google, “Google AI Professional Certificate,” Coursera, accessed October 2026, <https://www.coursera.org/professional-certificates/google-ai>
63. Rick Dakan, Joseph Feller, and Anthropic, “AI Fluency: Framework & Foundations (course),” Anthropic Academy, 2025, <https://anthropic.skilljar.com/ai-fluency-framework-foundations>
64. Anthropic, “AI Capabilities and Limitations (course),” Anthropic Academy, 2026, <https://anthropic.skilljar.com/ai-capabilities-and-limitations>
65. Anthropic, “Teaching AI Fluency (course),” Anthropic Academy, 2025, <https://anthropic.skilljar.com/teaching-ai-fluency>
66. Anthropic, “AI Fluency for Nonprofits (course),” Anthropic Academy, 2025, <https://anthropic.skilljar.com/ai-fluency-for-nonprofits>
67. Anthropic, “Claude 101 (course),” Anthropic Academy, 2025, <https://anthropic.skilljar.com/claude-101>
68. Hao-Ping (Hank) Lee et al., “The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers,” *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 2025, p. 3, <https://doi.org/10.1145/3706598.3713778>
69. Hao-Ping (Hank) Lee et al., “The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers,” *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 2025, pp. 14–15, <https://doi.org/10.1145/3706598.3713778>
70. Saffron Huang et al., “How AI Is Transforming Work at Anthropic,” Anthropic, December 2025, research post, “Looking forward”, <https://anthropic.com/research/how-ai-is-transforming-work-at-anthropic/>
71. Andrew Ng (DeepLearning.AI), “AI For Everyone,” Coursera, accessed October 2026, <https://www.coursera.org/learn/ai-for-everyone>
72. IBM, “AI Foundations for Everyone Specialization,” Coursera, accessed October 2026, <https://www.coursera.org/specializations/ai-foundations-for-everyone>
73. CompTIA, “CompTIA AI Essentials delivers effective, efficient employee training on practical workplace AI skills,” December 2025, <https://www.comptia.org/en-us/about-us/news/press-releases/comptia-ai-essentials-delivers-effective-efficient-employee-training-on-practical-workplace-ai-skills/>
74. CompTIA, “AI’s Impact on Productivity and the Workforce,” November 2025, section “Trends” (about 1 in 3 companies mandate AI training for staff), <https://www.comptia.org/en-us/resources/research/ai-impact-on-productivity-and-the-workforce/>
75. Lynn Heidmann, “The State of Data and AI Literacy in 2026: Definitions, Statistics, and the AI Skills Gap,” *DataCamp*, February 2026, section “AI ROI and workforce capability: The critical link” (21% of leaders report significant positive ROI; 42% with a mature, organization-wide program), <https://www.datacamp.com/blog/the-state-of-data-and-ai-literacy-in-2026-definitions-statistics-and-the-ai-skills-gap>
76. Lynn Heidmann, “The State of Data and AI Literacy in 2026: Definitions, Statistics, and the AI Skills Gap,” *DataCamp*, February 2026, section “Why most data and AI training isn’t working” (77% offer some AI training; 35% have a mature, workforce-wide program), <https://www.datacamp.com/blog/the-state-of-data-and-ai-literacy-in-2026-definitions-statistics-and-the-ai-skills-gap>
77. Philippa Hardman, “The Course Is Dying as the Unit of Learning,” *Substack*, April 2026, citing the Josh Bersin Company, <https://drphilippahardman.substack.com/p/the-course-is-dying-as-the-unit-of>
78. Philippa Hardman, “The Course Is Dying as the Unit of Learning,” *Substack*, April 2026, <https://drphilippahardman.substack.com/p/the-course-is-dying-as-the-unit-of>
79. Hao-Ping (Hank) Lee et al., “The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers,” *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 2025, p. 15 (verification, integration, stewardship; maintain foundational skills), <https://doi.org/10.1145/3706598.3713778>
80. Michael Ritchot, “Token-Level Abstraction in Large Language Models: A Design-Based Inquiry into AI Comprehension Pathways,” master’s capstone, Western Governors University, June 2025, p. 7 (reviewing Casal-Otero et al., 2023), <https://ritchot.me/docs/capstone.pdf>
81. Alan M. Saks and Monica Belcourt, “An Investigation of Training Activities and Transfer of Training in Organizations,” *Human Resource Management*, 45(4), 629–648, 2006, <https://doi.org/10.1002/hrm.20135>

82. CompTIA, “CompTIA AI Essentials delivers effective, efficient employee training on practical workplace AI skills,” December 2025, quoting Katie Hoenicke, chief product officer, <https://www.comptia.org/en-us/about-us/news/press-releases/comptia-ai-essentials-delivers-effective-efficient-employee-training-on-practical-workplace-ai-skills/>
83. Dan Tinkoff, Lieven Van der Veken, and Michael Chui, with Tara Balakrishnan, “The state of AI in 2026: On the road to ROI,” McKinsey & Company, August 2026, pp. 4, 14, 18, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
84. Dan Tinkoff, Lieven Van der Veken, and Michael Chui, with Tara Balakrishnan, “The state of AI in 2026: On the road to ROI,” McKinsey & Company, August 2026, p. 14, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
85. MIT Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training, “Report of MIT's Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training,” Massachusetts Institute of Technology, August 2026, p. 21, <https://sites.mit.edu/ai-use/files/2026/08/AI-Committee-Final-Report-Aug-13.pdf>

© Michael Ritchot. Code licensed under the MIT License. Instructional content licensed under CC BY-NC-SA 4.0.