

Evidence-Based Learner Persona

AI Literacy for the Modern Workforce

Persona: The Mid-Career Knowledge Professional

The target participant uses AI tools daily, reports time savings, and lacks the foundational judgment framework to know when their outputs are reliable.¹

This persona is a research-informed composite. Every attribute reflects documented findings from the Anthropic Interviewer (N=1,250, Dec 2025)¹, the WEF Future of Jobs Report 2025², the Anthropic Economic Index series^{3,4,5}, Lee et al. (CHI 2025)⁶, and, for seniority, OpenAI's 2026 enterprise usage study⁷. It characterizes the target population's observed behaviors, documented knowledge gaps, and reported motivations.

Role Profile

The target participant is a mid-career professional working in a knowledge-intensive function such as marketing, operations, HR, finance, or product management. The same profile extends upward through senior leadership and executive roles, and nothing in the research corpus confines the gap to mid-level seniority. Enterprise usage records locate the persona by headcount and by intensity. Six months after a firm adopts ChatGPT Enterprise, managers and directors are the largest seniority group among weekly active users at about 24%⁸. Individual contributors follow at 15%, senior individual contributors at 14%, executives at 10%, and early-career workers at 7%. These user shares describe who appears in usage logs, not adoption rates by role, because the study does not observe headcount by role⁹. Message volume, however, falls with seniority: early-career workers send roughly eight to nine more weekly messages than the average active user in their firm, though the authors treat volume as a measure of intensity, not economic importance¹⁰. Mid-career professionals are therefore among the most numerous users, not the most intensive ones. That combination is what a judgment-focused program is built for.

Executives are relatively more represented in topic overviews, facts and figures, legal and regulatory work, and financial or tax tasks¹¹. Much of that work turns on the names, statistics, dates, and citations that Module 3 places in the fabrication zone. Management tasks carry the highest average labor cost in the productivity data (\$133)¹², which makes the same judgment gap most expensive at the top. The persona centers the design on the middle of the workforce; it does not cap who the program serves.

The target participant holds a degree-level qualification, has five or more years of experience in their field², and operates in an organization that has granted access to at least one AI tool. They are not a technical specialist, though using AI is increasingly a condition of their job. Experience changes what AI does for them. The studies the WEF summarizes show generative AI substituting for tacit knowledge among newer or average performers¹³. The design inference is that, for a worker with five or more years of expertise, the gain comes less from having answers supplied and more from judgment applied to the tool's output. That inference is the reason this program trains Discernment rather than tool operation.

Employers project that 39% of workers' core skills will change by 2030, with AI and big data topping the list of fastest-growing skills globally¹⁴. The share of tasks performed by humans alone is expected to decline from 47% to near-parity with machine and human-machine collaboration categories by 2030¹⁵. The pattern holds in Singapore and Canada, the two high-

usage economies named in the capability gap analysis. Employers in Singapore expect 36% of core skills to change and technology to deliver 37% of task time by 2030; 82% plan to reskill staff to work alongside AI, in an economy whose executives rate government investment in mid-career training 6 on a 7-point scale¹⁶. Canadian employers expect 37% skill change and the largest human-machine collaboration share among the Canada, Singapore, and Hong Kong profiles (38% of task time by 2030, against 33% for Singapore and 28% for Hong Kong), and 90% plan reskilling alongside AI; North American employers expect 67% of their workforce to need training by 2030¹⁷. Forty-eight percent of professionals in the Anthropic Interviewer's 1,250-person sample (spanning the general workforce, scientists, and creatives) have considered career transitions toward roles that focus on managing and overseeing AI systems¹⁸. The general-workforce interviews describe the same direction in qualitative terms: participants want to keep the tasks that define their professional identity, delegate routine work, and see their future role as overseeing AI systems¹⁹. A trucking dispatcher in the sample put the anxiety plainly, hunting for human skills AI cannot take over while doubting they will stay necessary²⁰.

Current AI Usage Patterns

This participant is already using AI, so the question is how well they use it.

Eighty-six percent of professionals in the Anthropic Interviewer sample report that AI saves them time, and 65% report satisfaction with AI's current role in their work²¹. These figures establish baseline tool engagement. Productivity data provides the task-level picture: median task-time savings of 81%, with the distribution peaking in the 80–90% range²². That observational median should be read against the more conservative 14–56% reductions reported in the randomized controlled trials referenced in the same analysis. Task-specific figures vary substantially: document drafting yields approximately 87% time savings, financial analysis approximately 80%, and compiling information from reports approaches 95%²³.

The usage pattern, however, reveals a judgment gap. Sixty-five percent of participants self-report their primary mode of AI use as augmentative: they describe AI as supporting their judgment rather than replacing it. Task-level behavioral data from Claude usage analysis shows a materially different picture: 57% augmentative, 43% automative²⁴. That gap ranges from eight percentage points against Handa et al. to eighteen against the more recent Economic Index wave the Anthropic Interviewer cites (August 2025 data, approximately 47% augmentative²⁵), an 8-to-18-point overestimate depending on the dataset^{24,26}. The discrepancy indicates that a sizable share of participants believe they are exercising oversight over AI outputs in situations where the task structure (full delegation, unverified output acceptance) shows they are not. The core design constraint is that this participant does not know they are overrelying.

Documented Misconceptions and Capability Gaps

The following gaps are documented in the research corpus as patterns observed in behavioral and survey data.

Confusing generation with retrieval. Participants interact with AI tools as though they function like search engines: they expect outputs to reflect stored, retrievable facts rather than probabilistically generated text. This misconception produces uncritical acceptance of plausible-sounding but fabricated citations, statistics, and legal or regulatory references. 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 capability¹³, a risk this misconception directly enables; the report's Box 3.1 skills analysis, run with Indeed, adds that human oversight remains crucial even where the tool can assist²⁷.

Overconfidence as a scrutiny suppressant. Lee et al. (CHI 2025), in a survey of 319 knowledge workers across 936 real-world AI-assisted tasks, find that higher confidence in generative AI is associated with less critical thinking²⁸. Among their reasons for skipping it, workers named tasks they saw as trivial, such as social posts and meeting minutes (55 of 319), and trust in the tool, often built on past success (83 of 319), precisely the conditions under which errors are least likely to be caught before they enter a deliverable. The three-month pattern is an inference from that trust finding: the participant who has used an AI tool successfully for three months is more likely to skip verification than the one who started last week.

Misreading the augmentation/automation boundary. As documented above, participants systematically overestimate the degree to which they maintain human oversight during AI-assisted work. They lack a functional vocabulary for distinguishing delegation decisions (what to hand off entirely) from collaboration decisions (what to do jointly with AI support). This is the gap the 4D framework's Delegation dimension is specifically designed to close: in the source document, Delegation involves deciding when to use AI for automation, augmentation, or independent agent-mediated work, and understanding each modality's affordances is a Task Delegation skill²⁹. This program covers the first two.

Underuse through overcaution. Some participants avoid AI for high-value tasks (complex analysis, sensitive communication, strategic drafting) out of discomfort with output variability, even in cases where AI assistance combined with human review would be both appropriate and efficient. This avoidance pattern produces no productivity gain and no skill development.³

Motivation and Engagement Barriers

What motivates engagement. This participant is time-pressured and professionally self-interested. They engage with development programs when the direct relevance to their daily work is legible within the first five minutes. The primary motivators are productivity gain (86% report time savings from current AI use³⁰; in a second sample, 80% report improved productivity and about half say AI helped them develop new skills³¹), career positioning (48% of professionals interviewed have considered moving toward AI oversight roles¹⁸), and professional legitimacy (the desire to use AI confidently and visibly, not just competently).

What blocks engagement. The most significant barrier is not skepticism about AI, and it is not personal fear of it. Just 13% of respondents in McKinsey's 2026 survey say AI makes them anxious about their career prospects³². Among midlevel managers and individual contributors the share is 18–19%³¹. It is social exposure. Sixty-nine percent of professionals report that social stigma is an active deterrent to visible AI use at work³³. The concealment dynamic is not theoretical: a fact-checker in the sample reported that when a colleague said they hated AI, they said nothing, and that they tell no one their process because they know how many people feel about AI. A participant who completes this program and returns to a team culture where AI use is seen as a shortcut or a threat to professional credibility has limited social support for applying what they learned. Program design must account for this: the shared professional vocabulary this program builds is itself a source of social permission for visible AI use.

Secondary barriers include time pressure (44 of 319 respondents in Lee et al. named lack of time as a barrier to critical engagement³⁴; MIT's 2026 AI committee heard from university learners that the temptation to offload thinking peaked under deadline pressure³⁵), perceived irrelevance of generic AI training to their specific functional context, and prior negative

experience with compliance-style technology training that was mandatory, low-quality, and outcome-irrelevant. ADP Research's employee survey, reported in the WEF's Box 4.1, points the same way: employers rank reskilling second among practices that improve talent availability (63%), while employees, asked their top reasons to stay with an employer, rank upskilling seventh, behind working hours, progression, pay, remote work, and pensions³⁶. The questions differ, so the gap is indicative, not a measure of training relevance.

Technology Access and Learning Preferences

The design assumes this participant accesses training asynchronously, primarily during business hours in blocked time or between meetings, and in a desktop browser more often than on a phone. It also assumes short windows of 15–30 minutes rather than extended sittings, and low tolerance for passive content delivery that does not connect to an immediate application. These are design assumptions, not survey findings; a deploying organization can check them against its own learning-platform data. Every module is therefore chunked into short sections, with interactive activities between reading sections and one or two knowledge-check sections per module, a design the capstone pilot adopted against cognitive overload in asynchronous delivery³⁷.

They already use AI tools (ChatGPT, Claude, Gemini, or enterprise-licensed equivalents) in their daily workflow. This program is their first structured engagement with the mechanisms and judgment frameworks that determine whether that existing use is reliable. Design implication: no time is spent convincing participants that AI matters. That case is made by the 86% time savings data the program presents in Module 1. Time is spent building the foundational competency and evaluative vocabulary that converts tool familiarity into professional fluency.

Design Implications Summary

Attribute	Research Anchor	Design Response
Already using AI; satisfaction high	86% time savings, 65% satisfaction ²¹	Skip adoption arguments; lead with the performance gap
Overestimates augmentation; underestimates automation	65% self-report vs. 57% behavioral data ^{38,24}	Build Delegation and Discernment competencies explicitly; make the gap visible
Overconfidence suppresses verification	Lee et al. (CHI 2025) ³⁹	Scenario-based practice that surfaces errors under confident conditions
Social stigma limits visible application	69% concealment ³³	Build shared vocabulary (4D framework) as a normalization instrument
Motivated by career relevance and output quality	48% considered moving toward AI oversight roles ¹⁸	Frame every module in terms of professional identity and output quality

Attribute	Research Anchor	Design Response
Time-pressured; asynchronous; browser-based	Scale of need: 59 of 100 workers need training by 2030 and 11 will not get access ⁴⁰ ; modality and session length are design assumptions, not surveyed preferences	15–30 minute modules; mobile-responsive; no live session dependency
Underuse in complex tasks (inferred)	Augmentation/automation behavioral data ³	Module 4 scope includes appropriate delegation identification, not only output verification

¹ Broad use does not mean deep use. Workplace and consumer surveys through December 2025 show weekly-use rates well above daily-use rates (Gallup⁴¹; Deloitte, June 2025⁴²). The surveys measure how often people use the tools, not why use stays shallow. This program's position is that the gap is not lack of access: people have the tools and use them, but without an understanding of how the models work they struggle to find reliable value in them. This persona is the daily user who has closed the frequency gap and still needs the understanding gap closed.

² The Lee et al. sample is younger than this persona: 71.79% of its 319 respondents were under 35 and 96.87% used ChatGPT⁴³, and all were recruited on Prolific⁴⁴. The authors name the skew as a limitation⁴⁵. Their models found no significant effect of age group on critical-thinking enactment⁴⁶, though only 28 respondents were 45 or older; that null result is the basis for extending the confidence finding to mid-career workers here.

³ Underuse through overcaution is a design-grounded inference derived from the augmentation/automation behavioral divergence³ rather than a directly observed survey item. It is included as a design premise and flagged accordingly. If future practitioner research produces a direct empirical anchor, this section should be updated.

Notes

1. Kunal Handa et al., “Introducing Anthropic Interviewer: What 1,250 Professionals Told Us About Working with AI,” Anthropic, December 2025, <https://www.anthropic.com/research/anthropic-interviewer>
2. World Economic Forum, “The Future of Jobs Report 2025,” January 2025, https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf
3. Kunal Handa et al., “Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations,” Anthropic, February 2025, <https://arxiv.org/abs/2503.04761>
4. Ruth Appel, Peter McCrory, Alex Tamkin, et al., “The Anthropic Economic Index Report: Uneven Geographic and Enterprise AI Adoption,” Anthropic, September 2025, <https://www.anthropic.com/research/anthropic-economic-index-september-2025-report>
5. Ruth Appel, Maxim Massenkoff, Peter McCrory, et al., “The Anthropic Economic Index Report: Economic Primitives,” Anthropic, January 2026, <https://www.anthropic.com/research/anthropic-economic-index-january-2026-report>
6. 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, <https://doi.org/10.1145/3706598.3713778>
7. Aaron Chatterji et al., “How Organizations Use AI: Evidence from ChatGPT,” OpenAI working paper, August 2026, <https://arxiv.org/abs/2608.12236>
8. Aaron Chatterji et al., “How Organizations Use AI: Evidence from ChatGPT,” OpenAI working paper, August 2026, p. 17, <https://arxiv.org/abs/2608.12236>
9. Aaron Chatterji et al., “How Organizations Use AI: Evidence from ChatGPT,” OpenAI working paper, August 2026, p. 16, <https://arxiv.org/abs/2608.12236>
10. Aaron Chatterji et al., “How Organizations Use AI: Evidence from ChatGPT,” OpenAI working paper, August 2026, p. 18, <https://arxiv.org/abs/2608.12236>
11. Aaron Chatterji et al., “How Organizations Use AI: Evidence from ChatGPT,” OpenAI working paper, August 2026, p. 21, <https://arxiv.org/abs/2608.12236>
12. 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>
13. 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

14. World Economic Forum, “The Future of Jobs Report 2025,” January 2025, pp. 32, 37, https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf
15. 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
16. World Economic Forum, “The Future of Jobs Report 2025,” January 2025, pp. 109, 192–193, https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf
17. World Economic Forum, “The Future of Jobs Report 2025,” January 2025, pp. 74, 124–125, 145, https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf
18. 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” (48% considered moving toward AI-oversight roles), <https://www.anthropic.com/research/anthropic-interviewer>
19. Kunal Handa et al., “Introducing Anthropic Interviewer: What 1,250 Professionals Told Us About Working with AI,” Anthropic, December 2025, research post, introduction, <https://www.anthropic.com/research/anthropic-interviewer>
20. Kunal Handa et al., “Introducing Anthropic Interviewer: What 1,250 Professionals Told Us About Working with AI,” Anthropic, December 2025, research post, introduction, general-workforce example quotations, <https://www.anthropic.com/research/anthropic-interviewer>
21. 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 survey, N = 1,000: 86% say AI saves time; 65% satisfied with AI’s role), <https://www.anthropic.com/research/anthropic-interviewer>
22. Alex Tamkin and Peter McCrory, “Estimating AI Productivity Gains from Claude Conversations,” Anthropic, November 2025, pp. 12–13, <https://www.anthropic.com/research/estimating-productivity-gains>
23. 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>
24. 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>
25. 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>
26. 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>
27. 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
28. 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. 9, 11, <https://doi.org/10.1145/3706598.3713778>
29. 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. 4, <https://ringling.libguides.com/ai/framework>
30. 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 survey, N = 1,000: 86% say AI saves time), <https://www.anthropic.com/research/anthropic-interviewer>
31. 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. 11, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
32. 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>
33. 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>
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, p. 7, <https://sites.mit.edu/ai-use/files/2026/08/AI-Committee-Final-Report-Aug-13.pdf>
36. 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
37. 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, pp. 29, 45, <https://ritchot.me/docs/capstone.pdf>

38. 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% describe their use as augmentative, 35% automative), <https://www.anthropic.com/research/anthropic-interviewer>
39. 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>
40. World Economic Forum, “The Future of Jobs Report 2025,” January 2025, p. 47, https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf
41. Benedict Evans, “AI eats the world (May 2026 edition),” May 2026, slide 32 (Gallup, June 2023–December 2025), <https://www.ben-evans.com/s/2026-Spring-AI.pdf>
42. Benedict Evans, “AI eats the world (November 2025 edition),” November 2025, slide 38 (Deloitte, June 2025), <https://www.ben-evans.com/s/2025-Autumn-AI.pdf>
43. 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, Table 3, p. 6, <https://doi.org/10.1145/3706598.3713778>
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, p. 5, <https://doi.org/10.1145/3706598.3713778>
45. 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, <https://doi.org/10.1145/3706598.3713778>
46. 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, Table 4, p. 10, <https://doi.org/10.1145/3706598.3713778>

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