

# Action Map

## AI Literacy for the Modern Workforce

Following Cathy Moore's Action Mapping framework, organized by the Anthropic 4D competency taxonomy of Delegation, Description, Discernment, and Diligence (Dakan & Feller, 2025).

### DELIVERY NOTE


Each practice activity below is mapped to two deliverables: a custom-coded platform component (Component 4) and a corporate Articulate (Rise 360 / Storyline) component (Component 5). The program was designed against both contexts to demonstrate that the instructional design transfers across a bespoke build and a standard authoring stack. The custom platform was built and shipped ([ai-literacy.ritchot.me](https://ai-literacy.ritchot.me)); the Articulate build was specified but not produced, a deliberate scope decision for an independent portfolio project without enterprise funding. The Component 5 mappings are retained because the interaction-fidelity tradeoffs they record are part of the design analysis.

### CENTER · MEASURABLE BUSINESS GOAL

#### MEASURABLE BUSINESS GOAL

Within 90 days of program completion, participants consistently select appropriate tasks for AI use, apply structured evaluation to AI-generated outputs before incorporation into business deliverables, and use a shared competency vocabulary to make their AI practices visible to colleagues and managers, with measurable gains in productive AI utilization and reductions in unverified AI output across participating teams.

*Note: Specific performance targets (e.g., percentage reduction in unverified outputs, increase in task range) are calibrated to organizational baselines during deployment. The directional commitment is fixed; the threshold is context-dependent.*



- CENTER**      **Measurable Business Goal** — the one outcome every element serves.

---

- RING 1**      **Observable Behaviors · D1–D12** — what the participant does, organized by the four competencies.

---

- RING 2**      **Practice Activities · P1–P12** — how each behavior is built, by module.

---

- RING 3**      **Reference Information · R1–R7** — job aids that sustain the behavior afterward.

● DELEGATION   ● DESCRIPTION   ● DISCERNMENT   ● DILIGENCE

## RING 1 · OBSERVABLE BEHAVIORS

Organized by the 4D competency framework (Delegation, Description, Discernment, Diligence) from Anthropic's AI Fluency Framework.

## ● Delegation

D1 - D3

**D1 Before using AI, the participant defines the task goal, identifies which components require human judgment, and determines which are appropriate for AI assistance.**

**SUB-COMPONENT** Problem Awareness → Task Delegation

**GAP TRACE** Misreading the augmentation/automation boundary (Persona — 65% self-report vs. 57% behavioral; participants don't recognize when they've shifted from collaboration to full delegation; [Handa et al. 2025, p. 3](#))

**MODULE TRACE** Module 2 (Live Data Dashboard — augmentation vs. automation aggregate data with self-reflection prompts that surface the documented behavioral discrepancy) → Module 4 (AI Interaction Sandbox — structured task decomposition exercises where the participant practices the delegation decision before prompting)

**D2 The participant distinguishes between tasks that benefit from AI augmentation and tasks that should remain human-only, based on stakes, verifiability, and domain expertise requirements.**

**SUB-COMPONENT** Task Delegation

**GAP TRACE** Underuse through overcaution in high-complexity tasks + overreliance in perceived-routine tasks (Persona — Lee et al. finding that reduced scrutiny is most pronounced in tasks perceived as routine; [Lee et al., CHI 2025, pp. 1–2](#)). (Note: "Underuse through overcaution" is a design-grounded inference derived from the augmentation/automation behavioral divergence in [Handa et al., Feb 2025](#) — see the persona's footnote 1. No direct survey item on task avoidance behavior is available in the current corpus.)

**MODULE TRACE** Module 2 (task-type productivity data) → Module 4 (sandbox — structured delegation exercises)

**D3 The participant identifies the capability boundaries of the AI tool they are using (what it can generate reliably vs. where it is likely to produce errors) before delegating a task.**

**SUB-COMPONENT** Platform Awareness

**GAP TRACE** Confusing generation with retrieval (Persona — interacting with AI as though it retrieves stored facts rather than generating probabilistic text); WEF adverse outcome risk from stretching AI beyond capability (WEF 2025, p. 11)

**MODULE TRACE** Module 3 (next-token prediction demo, tokenizer playground — mechanistic understanding of what the model actually does) → Module 4 (sandbox — testing boundary conditions)

## ● Description

D4 - D6

**D4 The participant provides task-relevant context, constraints, and output specifications when prompting AI, rather than issuing single-turn, underspecified requests.**

**SUB-COMPONENT** Product Description → Process Description

**GAP TRACE** Surface-level engagement producing a fraction of available efficiency ([Tamkin & McCrory, pp. 12–13](#): 81% median savings under proficient use in observational data, against a more conservative 14–56% range in the referenced randomized controlled trials); directive single-turn interactions constitute a substantial and growing share of AI usage, rising from 27% to 39% of sampled conversations between January and August 2025 ([Anthropic Economic Index, September 2025 report, p. 9](#)), indicating that many users are not iterating or providing contextual specifications.

**MODULE TRACE** Module 4 (AI Interaction Sandbox — side-by-side comparison of prompt reformulation → output quality change)

**D5 The participant specifies the desired format, audience, and quality criteria for AI output before generating it.**

SUB-COMPONENT Product Description

GAP TRACE The augmentation/automation behavioral discrepancy (65% self-report augmentative vs. 57% behavioral, Handa et al. 2025, p. 3) suggests that participants who believe they are augmenting but are actually automating are not setting output specifications because they are not planning to review. (*Design-grounded inference from Handa et al. behavioral data; no direct survey item on output specification behavior is available in the current corpus.*)

MODULE TRACE Module 4 (sandbox — structured prompt exercises with scaffolded product description)

**D6 The participant iterates on AI outputs through multi-turn refinement rather than accepting or discarding first-generation results.**

SUB-COMPONENT Process Description → Performance Description

GAP TRACE Lee et al. — overconfidence predicts reduced critical engagement (Lee et al., CHI 2025, pp. 1–2); behavioral pattern of single-turn accept/reject rather than iterative collaboration (Handa et al. task iteration patterns, Feb 2025)

MODULE TRACE Module 4 (sandbox — feedback simulator interaction flow; structured roleplay with iterative refinement)

**Discernment**

D7 - D9

**D7 The participant verifies factual claims, citations, and statistics in AI-generated output against independent sources before incorporating them into business deliverables.**

SUB-COMPONENT Product Discernment

GAP TRACE Confusing generation with retrieval (Persona — uncritical acceptance of plausible-sounding but fabricated citations, statistics, regulatory references); Lee et al. overconfidence → reduced scrutiny (Lee et al., CHI 2025, pp. 1–2)

MODULE TRACE Module 3 (understanding why models confabulate — next-token prediction, not information retrieval) → Module 4 (sequential classification scenario — evaluating AI output with planted errors)

**D8 The participant identifies when an AI-generated output is operating outside the model's reliable capability range, recognizing hallucination indicators, reasoning failures, and confidence without competence.**

SUB-COMPONENT Product Discernment → Process Discernment

GAP TRACE WEF — adverse outcomes where users unknowingly stretch technology beyond capability (WEF 2025, p. 11); generation-vs.-retrieval misconception (Persona)

MODULE TRACE Module 3 (mechanistic understanding — tokenization, attention, probability-based generation, hallucination mechanics) → Module 4 (sandbox — annotation layer where learner tags outputs as reliable/uncertain/fabricated)

**D9 The participant evaluates whether the AI's reasoning process, not just the final output, is sound, checking for logical gaps, unsupported assumptions, and circular reasoning.**

SUB-COMPONENT Process Discernment

GAP TRACE Lee et al. — scrutiny reduction is most pronounced in routine tasks, meaning process evaluation is the first thing dropped (Lee et al., CHI 2025, pp. 1–2); Handa et al. — validation is the smallest behavioral category at 2.8% of all interactions (Handa et al., Feb 2025, p. 10)

MODULE TRACE Module 4 (sequential classification scenario — consequence-based feedback when process evaluation is skipped)

**Diligence**

D10 - D12

**D10** The participant documents the role AI played in producing a deliverable, including which components were AI-generated, AI-assisted, or human-authored.

SUB-COMPONENT	Transparency Diligence
GAP TRACE	69% social stigma / concealment dynamic (Anthropic Interviewer, Dec 2025 — workers using AI, observing gains, and concealing both from colleagues and managers)
MODULE TRACE	Module 1 (stigma data in interactive data narratives — normalizing disclosure) → Module 4 (sandbox — practicing AI diligence statements)

**D11** The participant takes full accountability for the accuracy and appropriateness of any AI-assisted output before it is shared, submitted, or published, treating AI-generated content with the same review standard as work produced by a junior colleague.

SUB-COMPONENT	Deployment Diligence
GAP TRACE	Overconfidence as scrutiny suppressant (Lee et al., CHI 2025, pp. 1–2); the augmentation/automation behavioral discrepancy (Handa et al. 2025, p. 3) means participants are deploying outputs they believe they reviewed but did not
MODULE TRACE	Module 4 (sequential classification scenario — consequences of deploying unverified AI output in a workplace context)

**D12** The participant discusses their AI usage practices openly with colleagues and managers, using the 4D competency vocabulary to describe how they delegate, prompt, evaluate, and verify.

SUB-COMPONENT	Transparency Diligence (extended to organizational behavior change)
GAP TRACE	69% concealment dynamic (Anthropic Interviewer, Dec 2025); the Executive Problem Statement's argument that individual competency without social normalization produces no measurable business outcome
MODULE TRACE	Module 1 (stigma data, normalization framing) → Module 2 (peer usage patterns — making invisible practices visible) → Module 4 (reflection prompts using 4D vocabulary)

## RING 2 · PRACTICE ACTIVITIES

**MODULE 1** Why AI Literacy Matters Now**P1** Interactive Data Narrative: The Workforce Shift

The participant explores scrollable data stories combining WEF skill demand projections with Anthropic Economic Index adoption data. At three embedded decision points, the participant is asked to interpret a data trend and select which implication is best supported by the evidence.

BUILDS	D1, D10
COMPONENT	4C (Interactive Explorations — Module 1) / 5B (Rise 360 — Module 1)
INTERACTION	Guided data exploration with embedded interpretation checks

**P2** Stigma Data Reflection

After viewing the 69% concealment finding in context, the participant responds to a guided reflection prompt: "Have you ever used AI for a work task and chosen not to mention it? What drove that decision?" The reflection is private (not submitted), but the prompt is followed by a short content block framing the 4D vocabulary as a normalization tool.

BUILDS	D10, D12
COMPONENT	4C (Module 1 interactive) / 5B (Rise 360 — Module 1)
INTERACTION	Reflective prompt (private, ungraded)

**MODULE 2** How AI Is Actually Being Used at Work**P3** Dashboard Exploration: Augmentation vs. Automation

The participant filters the Live Data Dashboard to view the augmentation/automation split by occupation category and task type. A guided reflection prompt asks them to locate their own functional area and compare the aggregate behavioral pattern (57% augmentation / 43% automation) against the self-report data (65% / 35%). A follow-up prompt asks: "For the task types most common in your role, does the data suggest most people are augmenting or automating? What would augmenting actually look like for those tasks?"

<b>BUILDS</b>	D1, D2
<b>COMPONENT</b>	4A (Live Data Dashboard) / 5B (Rise 360 — Module 2, with static data visualization alternative)
<b>INTERACTION</b>	Filterable dashboard with guided reflection prompts

**P4** Dashboard Exploration: Productivity Distributions

The participant explores the productivity gain distribution (80–90% peak range) and task-specific breakdowns (87% document drafting, 80% financial analysis, 95% compiling). A guided prompt asks them to identify one task they currently do without AI that falls into a high-productivity-gain category, and one task they currently use AI for that might be in a lower-gain or higher-risk category.

<b>BUILDS</b>	D2, D3
<b>COMPONENT</b>	4A (Live Data Dashboard) / 5B (Rise 360 — Module 2)
<b>INTERACTION</b>	Filterable dashboard with self-assessment prompt

**MODULE 3** Understanding How Language Models Work**P5** Tokenizer Playground

The participant inputs text samples — a standard English sentence, a non-English phrase, a string of numbers, a code snippet — and observes how each is broken into tokens. They are prompted to predict the token count before seeing the result, then explain why certain inputs produce more tokens than expected.

<b>BUILDS</b>	D3, D8
<b>COMPONENT</b>	4C (Interactive Explorations — Module 3) / 5A (Storyline — Module 3 drag-and-drop interaction)
<b>INTERACTION</b>	Input-output sandbox with prediction prompts

**P6** Next-Token Prediction Demonstration

The participant views a partially generated sentence and adjusts a temperature parameter, observing how the probability distribution over candidate next tokens changes. They see that higher temperature produces more varied but less predictable outputs. A guided prompt asks: "If this model is selecting the next word based on probability, not facts, what does that mean for the reliability of a specific claim it generates?"

<b>BUILDS</b>	D3, D7, D8
<b>COMPONENT</b>	4C (Interactive Explorations — Module 3) / 5A (Storyline — Module 3)
<b>INTERACTION</b>	Parameter-adjustment interactive with guided reasoning prompt

**P7** Context Window Scenario

The participant is given a simulated task: a long document that exceeds a model's context window. They see what happens when critical information falls outside the window: the model generates a plausible but incomplete response. A debrief prompt connects this to the platform awareness component of Delegation: "How would you verify that the model actually processed the information you needed it to use?"

<b>BUILDS</b>	D3, D8, D9
<b>COMPONENT</b>	4C (Interactive Explorations — Module 3) / 5A (Storyline — Module 3)
<b>INTERACTION</b>	Simulated failure scenario with debrief

**MODULE 4** Evaluating AI Outputs and Working Responsibly**P8** Structured Task Decomposition Exercise

The participant receives a workplace scenario (e.g., preparing a quarterly business review summary). They decompose the task into components and assign each to human-only, AI-assisted, or fully-delegated, then receive consequence-based feedback showing what happens when high-judgment components are fully delegated vs. appropriately retained.

<b>BUILDS</b>	D1, D2
<b>COMPONENT</b>	4B (AI Interaction Sandbox — structured roleplay) / 5A (Storyline branching scenario — Module 4)
<b>INTERACTION</b>	Categorization exercise with consequence-based feedback

**P9** Prompt Reformulation Comparison

The participant sees a weak, single-turn prompt and the AI output it produces. They rewrite the prompt using product, process, and performance description elements, then see the improved output side-by-side. The exercise is repeated for 2–3 task types (drafting, analysis, review).

<b>BUILDS</b>	D4, D5, D6
<b>COMPONENT</b>	4B (AI Interaction Sandbox — side-by-side comparison view) / 5A (Storyline — Module 4)
<b>INTERACTION</b>	Before/after prompt workshop with side-by-side output display

**P10** Output Verification Scenario

The participant receives an AI-generated deliverable (e.g., a client-ready summary with embedded statistics and a cited source). The deliverable contains planted errors: a fabricated citation, an inflated statistic, and a plausible but logically unsupported conclusion. The participant must identify which elements require verification, select a verification method, and observe the consequence of their decision (flagging the error vs. letting it pass into a client deliverable).

<b>BUILDS</b>	D7, D8, D9, D11
<b>COMPONENT</b>	4B (AI Interaction Sandbox — annotation layer) / 5A (Storyline branching scenario — Module 4)
<b>INTERACTION</b>	Sequential classification scenario with consequence-based feedback

**P11** Iterative Refinement Exercise

The participant begins with a generated output and iterates through 3–4 turns of refinement, guided by scaffolded prompts that model the Description-Discernment loop. After each turn, they evaluate what improved, what didn't, and what to adjust next. The exercise concludes with a reflection on their own refinement pattern.

<b>BUILDS</b>	D6, D9
<b>COMPONENT</b>	4B (AI Interaction Sandbox — feedback simulator) / 5A (Storyline — Module 4)
<b>INTERACTION</b>	Multi-turn guided interaction with self-assessment

**P12** AI Diligence Statement Practice

The participant writes a brief diligence statement for a completed sandbox exercise: which components were AI-generated, which were human-modified, and what verification was performed. They see an exemplar statement and compare their own.

<b>BUILDS</b>	D10, D11, D12
<b>COMPONENT</b>	4B (AI Interaction Sandbox) / 5A (Storyline — Module 4)
<b>INTERACTION</b>	Constructed response with exemplar comparison

## RING 3 · REFERENCE INFORMATION

**R1 4D Competency Quick-Reference Card**

One-page summary of the four competencies with sub-components and action verbs. Used when the participant is making a delegation decision or writing a diligence statement and needs the vocabulary.

**SUPPORTS** D1, D2, D10, D12

**FORMAT** Downloadable PDF / in-platform reference panel

**R2 Task Delegation Decision Guide**

A short decision tree or checklist: given a task, walk through stakes (what happens if the output is wrong?), verifiability (can I independently check this?), and domain expertise requirements (do I know enough to evaluate the output?). Outputs a recommendation: human-only, AI-assisted with review, or fully delegable.

**SUPPORTS** D1, D2

**FORMAT** Downloadable one-pager / interactive checklist in custom platform

**R3 Prompt Structure Template**

A fill-in-the-blank template organized by the three Description sub-components: product description (what do I want?), process description (how should the AI approach it?), performance description (how should the AI behave?). Not a prompt library: a structural scaffold the participant applies to any task.

**SUPPORTS** D4, D5, D6

**FORMAT** Downloadable template / in-platform reference panel

**R4 Output Verification Checklist**

A short checklist for evaluating AI-generated output before use: factual claims verified against independent source? Citations checked for existence? Statistics checked against original data? Logical reasoning traced step-by-step? Tone and register appropriate for audience? Scope consistent with what was requested? Organized by Product, Process, and Performance Discernment categories.

**SUPPORTS** D7, D8, D9, D11

**FORMAT** Downloadable one-pager / in-platform reference panel

**R5 AI Capability Boundary Reference**

A concise summary of what current language models do well and where they systematically fail, organized not by technical architecture but by task type. Reliable: drafting, summarizing, reformatting, brainstorming, translation of common languages. Unreliable: precise arithmetic, real-time or post-training-cutoff facts, legal/regulatory citation, character-level text manipulation, non-English tokenization edge cases. Updated to reflect the mechanistic explanations from Module 3.

**SUPPORTS** D3, D8

**FORMAT** Downloadable reference / in-platform panel linked from Module 3

**R6 AI Diligence Statement Template**

A fill-in template for documenting AI's role in a deliverable: which sections were AI-generated, which were AI-assisted with human revision, which were human-authored; what verification was performed; which AI tool and model were used. Designed to be appended to any business deliverable or kept as an internal record.

**SUPPORTS** D10, D11, D12

**FORMAT** Downloadable template (Word/Google Doc)

**R7 Organizational AI Use Policy Starter**

A short framework (not a complete policy, but a scaffold) that a participant could bring to their manager or team lead to initiate a conversation about team-level AI use norms. Covers: what tasks are appropriate for AI use in this team, what disclosure is expected, what verification standard applies, how AI use is documented. Directly addresses the 69% concealment finding by giving participants a tool for initiating the normalization conversation.

**SUPPORTS** D12

**FORMAT** Downloadable template

**TRACEABILITY SUMMARY**

Every Ring 1 behavior (D1–D12) maps to at least one practice activity (P1–P12) and at least one reference item (R1–R7). Every practice activity maps to a specific component deliverable in both the custom platform (Component 4) and the Articulate platform (Component 5). Every reference item is designed as a job aid that sustains the behavior after the learning experience ends. The metadata (gap traces, module traces, sub-component alignment) serves as the verification scaffold and feeds directly into the program's alignment matrix.

**TRACEABILITY VERIFICATION – ALL SIX TESTS PASS**

<b>1</b>	Every behavior traces backward to a documented gap in the capability gap analysis or the learner persona	<b>12 / 12</b>
<b>2</b>	Every behavior traces forward to at least one practice activity	<b>12 / 12</b>
<b>3</b>	Every practice activity traces to a specific Component 4 and Component 5 deliverable	<b>12 / 12</b>
<b>4</b>	Every reference item supports at least one behavior	<b>7 / 7</b>
<b>5</b>	Every behavior is supported by at least one reference item	<b>12 / 12</b>
<b>6</b>	No orphaned elements in any ring	<b>PASS</b>

**SCOPE & FEASIBILITY NOTES**

- Practice activity distribution is intentionally weighted toward Modules 3 and 4 (seven of twelve activities) where skill-building occurs. Modules 1 and 2 carry awareness and motivation work.
- High-priority behaviors (D3, D8) are covered by multiple activities across modules, creating spaced repetition consistent with the transfer research.
- Each activity has both a custom platform deliverable (Component 4) and an Articulate deliverable (Component 5), though interaction fidelity will differ by platform: sandbox activities will be richer on the custom platform. (As noted above, the Articulate build was scoped but not produced.)
- If scope cuts become necessary: P7 (context window scenario) and P2 (stigma reflection) are lowest-risk cuts. P7's concept can be covered in explanatory content; P2's reflection can be embedded within P1. The remaining activities are load-bearing. R7 (Organizational AI Use Policy Starter) is cuttable without breaking the behavior-to-reference chain.

---

**REFERENCES**

---

**Anthropic Interviewer** — [Introducing Anthropic Interviewer: What 1,250 Professionals Told Us About Working with AI](#) (Dec 2025)

**Appel, R., McCrory, P. & Tamkin, A.** — [Anthropic Economic Index Report: Uneven Geographic and Enterprise AI Adoption](#) (Sep 2025)

**Dakan, R. & Feller, J.** — [The AI Fluency Framework, Version 1.5](#) (2025)

**Handa, K. et al.** — [Which Economic Tasks Are Performed with AI? Evidence from Millions of Claude Conversations](#) (Feb 2025)

**Lee, H. 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](#), CHI 2025

**Tamkin, A. & McCrory, P.** — [Estimating AI Productivity Gains from Claude Conversations](#) (Nov 2025)

**World Economic Forum** — [Future of Jobs Report 2025](#) (Jan 2025)