

Learning Architecture

AI Literacy for the Modern Workforce

This is the learning architecture behind *AI Literacy for the Modern Workforce*: four modules on a Context → Evidence → Mechanism → Application arc, each mapped to the 4D competencies, the practice activities that build them, the performance objectives they satisfy, and the assessment tier that measures each one. The whole structure was designed to feed a [Kirkpatrick evaluation](#): every objective resolves to a tier, and two objectives produce artifacts the [30/60/90-day behavioral follow-up](#) can track.

The program was scoped for two delivery platforms in parallel: a custom-coded React environment (Component 4, the platform you are reading this on) and an Articulate Rise/Storyline build (Component 5). The custom platform shipped. The Articulate build was deferred as a budget decision, not a design one, and the [resource plan](#) gives the rationale. The architecture keeps both columns so the instructional design stays platform-agnostic: only the interactive components are platform-dependent, and the content could be rebuilt in any equivalent authoring tool from the same objectives and activities.

Module 3 also inherits four design principles from the capstone pilot that preceded this build¹. Scaffold abstract content from concrete analogies upward. Visualize the invisible with diagrams and tokenizer demos. Embed reflection and scenario-based application. Chunk content, with feedback after each part. Three Module 3 components answer that pilot's specific findings. The tokenizer playground (P5) is the pilot's first recommendation and answers its participants' request for a tokenizer sandbox². The number-tokenization comparison in §3.2 descends from one of the two visuals participants named as making the content concrete³. The playground's free mode shows the token count and compression ratio for any text, so learners can test the first half of the pilot's cheat-sheet rule, clean input means fewer tokens and better output, which every pilot participant applied in the scenario task⁴.

Reading the diagram

- Click any module card to reveal its dual-platform delivery, its practice activities with interaction types, and its performance objectives with competency tags, assessment tier, and evaluation mapping.
- Expand the Kirkpatrick lane to read what each evaluation level measures. Objectives tagged → **L3** feed the 30/60/90-day behavioral plan.

4 MODULES **17** OBJECTIVES **12** PRACTICE ACTIVITIES **2** PLATFORMS **8** COMPETENCY-ASSESSED **2** L3 TRACKED

■ CUSTOM PLATFORM • COMPONENT 4

React-based learning environment. Interactive dashboards, tokenizer playground, AI interaction sandbox, learning analytics. Primary technical differentiator.

■ ARTICULATE PLATFORM • COMPONENT 5

Rise 360 + Storyline. Responsive scroll-based content, branching scenarios with consequence-based feedback, SCORM/xAPI assessments. Ecosystem fluency proof.

LEARNING PROGRESSION ————— Context → Evidence → Mechanism → Application

1

Why AI Literacy Matters Now

Strategic context, market forces, labor market transformation

CONTEXT

35–45 min | ■ DELEGATION ■ DILIGENCE | 2 activities · 4 objectives

■ 4C • CUSTOM PLATFORM

Interactive Explorations

Interactive data narratives: scrollable data stories combining WEF and Anthropic findings with embedded reflection prompts

■ 5B • ARTICULATE PLATFORM

Rise 360

Responsive scroll-based content with embedded interactions and scenario-based knowledge checks

■ PRACTICE ACTIVITIES (2)

P1 Interactive Data Narrative

Guided data exploration

P2 Stigma Data Reflection

Reflective prompt (private)

■ PERFORMANCE OBJECTIVES (4)

1.1 Interpret adoption data; identify implications for own role

■ DELEGATION ○ T1 Knowledge check

1.2 Articulate business case using three quantified findings

■ DELEGATION ○ T1 Knowledge check

1.3 Explain 69% concealment dynamic and shared vocabulary solution

■ DILIGENCE ○ T1 Knowledge check

1.4 Locate own usage on augmentation-automation spectrum

■ DELEGATION ○ T2 Practice activity (P1)



2

How AI Is Actually Being Used at Work

Real usage patterns, productivity data, behavioral reality vs. hype

EVIDENCE

45–60 min | ■ DELEGATION ■ DISCERNMENT | 2 activities · 4 objectives

■ 4A • CUSTOM PLATFORM

Live Data Dashboard

Filterable dashboard covering augmentation/automation split, occupation-level adoption data, collaboration patterns, productivity distributions, and guided reflection prompts

■ 5B • ARTICULATE PLATFORM

Rise 360

Responsive scroll-based content with embedded data visualizations and scenario-based knowledge checks

■ PRACTICE ACTIVITIES (2)

P3 Dashboard: Augmentation vs. Automation

Filterable dashboard

P4 Dashboard: Productivity Distributions

Filterable dashboard

■ PERFORMANCE OBJECTIVES (4)

2.1 Distinguish augmentation from automation against behavioral data

■ DELEGATION ○ T2 Practice activity (P3)

2.2 Analyze task-level productivity data for high-gain vs. high-risk

■ DELEGATION ■ DISCERNMENT ○ T1 Knowledge check

2.3 Identify one task to add AI, one to increase oversight

■ DELEGATION ■ DISCERNMENT ○ T2 Practice activity (P4) → Level 3 → L3

2.4 Describe self-report/behavioral divergence and directive patterns

■ DELEGATION ○ T1 Knowledge check



3

MECHANISM

Understanding How Language Models Work

Tokenization, attention, probability-based generation, hallucination mechanics

50–65 min | ■ DESCRIPTION ■ DISCERNMENT | 3 activities · 4 objectives

■ 4C · CUSTOM PLATFORM

Interactive Explorations

Tokenizer playground, next-token prediction demonstration with parameter adjustment, context window visualization

■ 5A · ARTICULATE PLATFORM

Storyline Branching

Decision simulator scenarios with consequence-based feedback on model behavior interpretation

■ PRACTICE ACTIVITIES (3)

P5

Tokenizer Playground

Input-output sandbox

P6

Next-Token Prediction Demo

Parameter-adjustment interactive

P7

Context Window Scenario

Simulated failure scenario

■ PERFORMANCE OBJECTIVES (4)

3.1

Explain tokenization; predict unexpected results; describe systematic errors

■ DELEGATION ■ DISCERNMENT ● T3

Practice (P5) + Competency

3.2

Describe next-token prediction vs. retrieval; explain fluent fabrication

■ DELEGATION ■ DISCERNMENT ● T3

Practice (P6) + Competency

3.3

Identify three systematically unreliable task categories

■ DELEGATION ■ DISCERNMENT ● T3

Knowledge check + Competency

3.4

Determine whether output is within reliable capability range

■ DISCERNMENT ● T3

Practice (P5–P7) + Competency

4

Evaluating AI Outputs and Working Responsibly

Applied judgment, output verification, responsible delegation

APPLICATION

70-105 min | ■ DELEGATION ■ DESCRIPTION ■ DISCERNMENT ■ DILIGENCE | 5 activities · 5 objectives

■ 4B · CUSTOM PLATFORM

AI Interaction Sandbox

Structured roleplay, feedback simulator, prompt reformulation workshop, annotation layer, diligence statement builder

■ 5A · ARTICULATE PLATFORM

Storyline Branching

Multi-path branching scenarios with compounding consequences and constructed response assessment

■ PRACTICE ACTIVITIES (5)

P8 Task Decomposition Exercise

Categorization with consequences

P9 Prompt Reformulation Comparison

Before/after workshop

P10 Output Verification Scenario

Sequential element classification

P11 Iterative Refinement Exercise

Multi-turn guided interaction

P12 Diligence Statement Practice

Constructed response

■ PERFORMANCE OBJECTIVES (5)

4.1 Decompose task into human-only, AI-assisted, and delegated components

■ DELEGATION ● T3 Practice (P8) + Competency

4.2 Reformulate underspecified prompt into structured request

■ DESCRIPTION ● T3 Practice (P9) + Competency

4.3 Evaluate AI deliverable: verify, detect hallucination, classify elements

■ DISCERNMENT ■ DILIGENCE ● T3 Practice (P10) + Competency

4.4 Iterate through 3+ refinement turns using Description-Discernment loop

■ DESCRIPTION ■ DISCERNMENT ● T2 Practice (P11)

4.5 Produce AI diligence statement using 4D vocabulary

■ DILIGENCE ● T3 Practice (P12) + Competency + Level 3 → L3

ASSESSMENT ARCHITECTURE

Feeds Kirkpatrick Levels 2-4 below

○ Tier 1: Knowledge checks

Scenario-based knowledge checks with explanatory feedback, embedded within modules

5 objectives

● Tier 2: Practice

Learner choices and outputs within interactives, sandbox, and dashboards, assessed in-flow

4 objectives

● Tier 3: Competency and transfer

Competency assessment (5C) and Level 3 behavioral follow-up (6C): post-program transfer measurement

8 objectives

KIRKPATRICK EVALUATION FRAMEWORK

L1

Reaction

Level 1

L2

Learning

Level 2

L3

Behavior

Level 3

L4

Results

Level 4

L1

Reaction

Post-program survey across four actionable dimensions: perceived relevance, confidence change, content quality, and intent to apply, plus a standard NPS item

L2

Learning
Twelve-item scenario-based pre/post assessment, parallel-form to resist gaming, measuring knowledge gain across usage patterns, failure modes, mechanics, and evaluation

L3

Behavior
Manager evidence review and participant self-assessment at 30/60/90 days, rating twelve 4D-mapped behavioral indicators; objectives 2.3 and 4.5 produce the trackable artifacts

L4

Results
Business outcomes downstream of competent use: efficiency (time-to-first-draft, revision cycles), quality (AI-error rate), transparency, and time-to-competency, run through a Phillips ROI model (156% in the illustrative worked example)

L3

Transfer Measurement by Design
Two objectives produce trackable artifacts that the 30/60/90-day behavioral follow-up checks for sustained use. **Objective 2.3** is a task-level action commitment: the participant identifies a specific task they will change their approach to. **Objective 4.5** is an AI diligence statement: the participant documents AI’s role in a completed work product using 4D vocabulary. Level 3 data shows whether the program produces behavior change.

ANALYTICS INFRASTRUCTURE


Hidden from the learner; supports Levels 2 & 3

■ 4D • ANALYTICS DASHBOARD

Per-learner progress tracking, module completion rates, practice activity engagement metrics, and time-on-task data.
Portfolio-mode analytics dashboard at /#/dashboard (Ctrl/Cmd+Shift+A enters portfolio mode). Surfaces the practice-activity and KC data that feeds Kirkpatrick Levels 2 and 3.

■ 4E • PLATFORM SHELL

Sidebar navigation, the 90% scroll-completion sentinel, hash router, theme system, and the learner-progress and analytics contexts.
Wraps every module: every section renders inside the shell’s SectionContainer and consumes its shared progress and analytics providers.



Notes

1. 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. 44–45, <https://ritchot.me/docs/capstone.pdf>
2. 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. 37, 41, 43, <https://ritchot.me/docs/capstone.pdf>

3. 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. 43, <https://ritchot.me/docs/capstone.pdf>
4. 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. 38, 41, <https://ritchot.me/docs/capstone.pdf>

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