
Catch Up with AI without Starting Over

Think Like a NextGen LXD Starter Guide

What employers require now, and how to build on what you already know — the real language of current ID/LXD job listings, a ten-cluster self-check, and the honest map of where your experience already meets the AI-era market.

Brenda Bannan, PhD

Professor Emerita, George Mason University · Founder, AI4LD

AI4LD · ILLUMINATING THE CONNECTION BETWEEN LEARNING AND AI · GUIDE + VIDEO
+ SELF-CHECK

CONTENTS

Statement on the Use of AI in AI4LD Materials	
Introduction: What the Listings Already Told You	
Part 1: What the Job Listings Actually Say	
Three registers, not one	
The honest hedges	
Part 2: The 10-Cluster Self-Check	
Part 3: Why Your Foundation Still Wins	
Part 4: The Path	
Start free	
Bundle A — How AI and Learning Work · \$97 (units \$25 each)	
Bundle B — AI Affordances for Learning Design · \$147	
ILDF 2.0 Practitioner Stack · \$197	
ILDF 2.0 Full Stack — coming soon	
Where AI4LD is honest about the gap	
Sources	

Statement on the Use of AI in AI4LD Materials

At AI4LD™, we are committed to the transparent, ethical, and responsible use of artificial intelligence in learning and development.

This guide was conceptualized by Dr. Brenda Bannan and written collaboratively with AI: generative AI tools supported drafting and content development. Every sentence was reviewed by an expert with over thirty years of experience in learning science and instructional design. Human agency guided all decisions, and the author takes full responsibility for the accuracy and pedagogical soundness of the final content.

At AI4LD™, we design with integrity, teach with clarity, and put how people learn first.

Introduction: What the Listings Already Told You

If you've scrolled through instructional design and learning experience design job postings recently — really read them, not skimmed the titles — you've probably noticed something specific. Not the generic “AI is changing everything” unease. Something more precise: language sitting inside these postings that you don't fully recognize, right next to language you've known for twenty years.

A recent Thoughtful AI posting asks for proficiency in “instructional design methodologies (ADDIE, SAM, or similar) and adult-learning principles” — familiar ground — in the same listing that asks the candidate to “design modular learning paths that deepen knowledge in healthcare RCM, agentic AI, and compliance.” Harvard's HBS AI Institute wants “5 years of experience in instructional design or content development” alongside, as a preferred qualification, “working knowledge of AI/ML concepts and architectures, including foundational models, training methodologies, fine-tuning, data pipelines, and agentic systems.” Stripe's Learning Architect posting doesn't ask for a checklist of AI tools at all — it asks whether you can move “beyond simple prompting to think about how AI can be architected into the learning journey.”

Your training didn't cover any of that. Not because your program failed you — because it couldn't have. The Masters programs and certificate academies that trained most working IDs and LXDs were built around a stable set of processes: needs assessment,

ADDIE or SAM, backwards design, Kirkpatrick evaluation. Those processes are not wrong. They are the foundation nearly every listing above still requires, explicitly, even the ones written by AI-native companies. What's changed is that a second layer has been added on top, and almost nobody's formal training addressed it, because it didn't exist yet when most of us were trained.

I hit a version of this wall myself, a few years into leading research projects that paired instructional design with AI and machine learning experts. It wasn't a technical wall. I wasn't trying to become a computer scientist, and I didn't need to. It was a conceptual one: AI tools were showing up in every conversation, moving fast, and the learning science wasn't moving with them. Enthusiasm and apprehension were the two responses I kept seeing in the field — and neither one started from a rigorous model of what instructional design actually involves. I had something most of the AI enthusiasts in the room didn't: thirty years of understanding how people learn. That gap — between AI capability and learning science — is where this guide, and the rest of the AI4LD system, lives.

This guide is built differently from the version of it you may have seen before. Instead of asking you to take the premise on faith, it shows you the evidence: what real employers are actually asking for, in their own words, across roughly thirty current postings. Not because job listings are the whole truth about a profession — they aren't — but because they're one of the most concrete signals available of where the market is moving, and they let you check your own standing against real language instead of against anxiety.

The shift is real, it is uneven, and no single statistic settles the question of what it means for your career. What the evidence does show, consistently, is that the most senior, most prestigious postings are not asking candidates to abandon instructional design fundamentals for AI fluency. They're asking for both — and specifically for the connection between them. That connection is what this guide, and the rest of the AI4LD path, is built to give you.

Part 1: What the Job Listings Actually Say

This section draws on a curated review of roughly thirty current instructional design and learning experience design postings (2025–2026) — seventeen with new, AI-inflected titles (“AI Learning Experience Designer,” “Learning Architect,” “AI Instructional

Designer”) and thirteen with traditional titles that now carry AI language inside them. Some postings were fetched in full; others had expired or were blocked at the source and are represented only by search-engine snippets. Both kinds of evidence are labeled where they appear in this guide’s Sources section.

THREE REGISTERS, NOT ONE

The AI language inside these postings doesn’t say one thing. It clusters into three distinct registers, and they are not interchangeable — a candidate who is fluent in one can be badly underprepared for another, even under an identical job title.

Register one: tool fluency. This is the most common ask, and the lowest bar. Demandbase’s Senior Instructional Designer posting wants candidates to “utilize AI tools to optimize content development processes, personalize learning experiences, and deliver targeted learning solutions,” with “strong proficiency in AI tools for learning design (e.g., ChatGPT, Gemini).” Forward Financing’s Instructional Designer posting asks the candidate to “use Generative AI to rapidly build training outlines, draft scripts, and develop course assets.” This register treats AI as an efficiency multiplier bolted onto the existing authoring toolkit — draft faster, not design differently. It’s table stakes, and it’s also the layer that’s commoditizing fastest: RedThread Research’s 2026 vendor landscape study (461 vendors) found Content Creation functionality dropping from the #1 learning-tech capability to fifth, while Analytics & Measurement rose from 33% to 83% of vendors over the same three years.

Register two: AI-adoption coaching. A materially different job, often hiding under the same title. HoneyBook’s AI Instructional Designer posting asks for “a strategic mindset with a people-first approach to AI and automations adoption, focused on building confidence, reducing friction, and preparing teams for increasingly complex, high-value work” — the ID as the organization’s AI-enablement lead, not just an AI-assisted content producer. Forward Financing’s same posting also asks for “‘Just-in-Time’ guides that help employees learn how to use AI tools effectively” — designing AI literacy for others, not just using AI oneself.

Register three: judgment-grounded design competence. The rarest register, and the one that shows up disproportionately at the higher-prestige and research-adjacent employers. Harvard’s HBS AI Institute posting asks for “the ability to choose, configure, and apply AI tools responsibly to accelerate content production,” paired — as a

preferred, not required, qualification — with “demonstrated working knowledge of AI/ML concepts and architectures, including foundational models, training methodologies, fine-tuning, data pipelines, and agentic systems.” Prolific’s AI Instructional Designer posting asks for design work that ensures “adherence to ethical AI principles.” This register pairs the AI ask with real judgment language, and it’s where almost nobody currently teaches — which is also where AI4LD’s own material concentrates.

Where these registers meet — and where they diverge. Stripe’s Learning Architect posting is worth sitting with, because it shows registers two and three at their most demanding: “AI native: You move beyond simple prompting to think about how AI can be architected into the learning journey to leverage scale and improve content.” Notice what’s *not* in that sentence: no mention of ADDIE, no needs-analysis language, no “curriculum” at all — the posting explicitly reframes the role around “user friction” and “interaction models” instead. Paradox’s LXD posting for an AI x EdTech role asks candidates to be “able to build AI Agents” and to “write and optimize prompt chains, refine tone and flow until every exchange feels human and transformative” — again, without the traditional evaluation-framework language that anchors most ID postings. These are real postings, and they represent a genuine divergence: some of the market is moving away from the vocabulary you were trained in, not just adding to it.

THE HONEST HEDGES

None of this is a case for panic, and the evidence itself won’t support one. A few things need saying, because overstating this shift would be its own kind of dishonesty:

- **No source publishes an ID-specific displacement statistic.** Donald Taylor’s Global Sentiment Survey 2026 (n=3,500+, 105 countries) reports “a slew of redundancies at all levels” and job searches that are “harder, and longer” for affected professionals — a strong, well-regarded sentiment signal, but a sentiment survey, not a labor-market count. Indeed Hiring Lab’s July 2026 analysis tracked occupations whose tasks generative AI can most readily do — software development is their leading example — and found their postings declined more than average from 2022 to 2026, then rebounded fastest after 2025, led by senior roles. Indeed itself notes the decline began before ChatGPT and that correlation isn’t causation — and this is macro-level data, with no instructional design or L&D breakout.

- **The shift is uneven, and some major employers show none of it.** Capital One's Senior Associate, Learning Experience Designer posting contains no AI language at all — just Storyline, Captivate, and the standard design-tool stack. USAA's Instructional Designer II posting is the same: core ID/L&D requirements, no AI mentioned. These are not small or irrelevant employers. The accurate reading is “the leading edge of the market, not yet all of it” — not “everyone is being replaced.”
- **The 21–28% AI-proficient ID salary premium is a derived estimate, not a published ID-specific statistic.** It circulates widely in instructional design community discussion, and it originates with instructional designer and researcher Dr. Philippa Hardman, who applied Lightcast's general economy-wide finding — a 28% AI-skills salary premium, about \$18,000/year, across 1.3 billion job postings — to an ID salary baseline. That's a reasonable analytical move, and Hardman is transparent about it. But it is her estimate, not a number Lightcast published about instructional design specifically.
- **“AI Learning Architect” is a direction, not an arrived category.** No live posting was found using that exact title. It comes from Hardman's April 2025 specialization taxonomy, which proposes four emerging ID specializations (Learning Experience Designer, Learning Strategist, AI Learning Architect, Performance Analyst). The closest real-world match is Stripe's “Learning Architect” posting, described above — a genuine signal of direction, but treat the title itself as one researcher's proposed taxonomy, not an established job category you should expect to see widely advertised yet. A separate, independent signal points the same direction: “Learning Engineer,” a term dating to Herbert Simon in 1967, is now institutionalizing through IEEE ICICLE's professional standards work — a second, non-Hardman indicator that the generalist ID role is fracturing into named specializations, even though the specific titles differ.
- **What is solid:** ATD's August 2025 survey of 232 talent-development professionals found roughly two-thirds of instructional designers had adopted AI tools within the past year, and 70% of those adopters reported the AI use improved course quality. That's a real, ID-specific, named-sample data point — one of the few that exists.

Taken together, the evidence doesn't say “AI is replacing instructional designers.” It says employers who are hiring seriously for these roles are hiring for the combination of learning science and AI judgment — not for either alone. Dr. Philippa Hardman's synthesis of hiring-manager survey data found that 71.3% of hiring professionals rank instructional design theory and learning science above tool proficiency as a top-three

hiring criterion. That's her own analysis of the field, not an independently replicated statistic — but it lines up with what the postings themselves show. The next section turns that into something you can check yourself against.

Part 2: The 10-Cluster Self-Check

What follows are the ten skill clusters that recur across the postings reviewed for this guide, each stated in employers' own words, followed by what it actually asks of you and a short self-rating prompt. This isn't a test with a passing score. It's a way to see where your existing strengths already meet the market, and where a gap is worth naming.

Cluster 1 — GenAI/LLM fluency for content development *What employers say:* “use Generative AI to rapidly build training outlines, draft scripts, and develop course assets” (Forward Financing); “strong proficiency in AI tools for learning design (e.g., ChatGPT, Gemini)” (Demandbase). *What it means:* Comfortable, working use of generative AI tools for first-draft content — outlines, scripts, scenario text — with your own judgment applied to what the tool produces. *Self-rate (1–5):* How comfortable are you using a generative AI tool to draft a training outline you'd actually stand behind, without over- or under-trusting its output?

Cluster 2 — Architecting AI into the learning journey *What employers say:* “move beyond simple prompting to think about how AI can be architected into the learning journey to leverage scale and improve content” (Stripe). *What it means:* A structural design question rather than a tool-use one: can you design where and how AI operates across an entire learner experience? *Self-rate (1–5):* Could you sketch, right now, where AI would and would not belong across a learning journey you're currently designing?

Cluster 3 — AI agents / agentic workflows *What employers say:* “building or leveraging automations and agents to scale content creation and delivery processes” (Harvard); “able to build AI Agents” (Paradox). *What it means:* Understanding what an AI agent actually is — a system that takes multi-step action toward a goal with limited human intervention — well enough to design one responsibly or evaluate whether one belongs in a given learning workflow. *Self-rate (1–5):* Could you explain, to a colleague, the difference between an AI agent and a single-turn AI tool — and where each one is or isn't appropriate in a design process?

Cluster 4 — AI/ML conceptual knowledge *What employers say:* “working knowledge of AI/ML concepts and architectures, including foundational models, training methodologies, fine-tuning, data pipelines, and agentic systems” (Harvard, preferred qualification). *What it means:* Not becoming a machine learning engineer — understanding, at a conceptual level, what a large language model is doing, what reinforcement learning is, and how a neural network processes information, well enough to reason about what these systems can and can’t do for learning. *Self-rate (1–5):* If asked to explain how a large language model actually generates a response — in terms you’d trust yourself to say out loud to a client — how confident are you?

Cluster 5 — Teaching AI literacy to others / AI-enablement *What employers say:* “‘Just-in-Time’ guides that help employees learn how to use AI tools effectively” (Forward Financing); “a people-first approach to AI and automations adoption” (HoneyBook). *What it means:* Designing the training that helps *other people* in your organization adopt AI well — a distinct skill from using AI in your own work, and one that puts your existing instructional design expertise directly to work. *Self-rate (1–5):* Have you designed, or could you design today, a short AI-literacy resource for colleagues who are earlier in their AI adoption than you are?

Cluster 6 — Responsible/ethical AI use and output verification *What employers say:* “the ability to choose, configure, and apply AI tools responsibly” (Harvard); “adherence to ethical AI principles” (Prolific). *What it means:* Knowing how to critically evaluate an AI system’s output rather than accepting it at face value — what Deloitte’s 2026 Human Capital Trends report names as an “Interpretation and Verification” competency. *Self-rate (1–5):* When you use an AI tool’s output in your own design work, how systematic is your process for checking it before it goes further?

Cluster 7 — Prompt engineering / prompt systems *What employers say:* “write and optimize prompt chains” (Paradox); “AI prompt libraries” (Stripe). *What it means:* Building reusable, tested prompt structures rather than one-off requests — treating prompting itself as a design artifact. *Self-rate (1–5):* Do you have any prompts you reuse deliberately across projects, refined over time — or does every AI interaction start from scratch?

Cluster 8 — AI-assisted video/avatar production *What employers say:* “video generation and editing tools (e.g., AI-assisted scripting, voice, and visual production)” (Harvard); “demonstrated experience creating video-based learning content, including

use of AI-enabled video tools” (Metropolis Technologies). *What it means:* Hands-on skill with the production tools themselves — AI voice, avatar, and video generation platforms — as distinct from the instructional design decisions about when video is the right modality at all. *Self-rate (1–5):* Have you personally produced a piece of learning content using an AI video or avatar tool, start to finish?

Cluster 9 — Learning analytics / data-driven optimization *What employers say:* “learning analytics, data-driven optimization practices” (Coursera); “measure program impact using Kirkpatrick/Phillips or similar frameworks; iterate based on data” (Thoughtful AI). *What it means:* Using data from learner interactions — including AI-generated interaction data — to identify where a design is and isn’t working, and revising accordingly. *Self-rate (1–5):* How comfortable are you interpreting learner-interaction data (LMS, assessment, or AI-tool usage data) to justify a specific design change?

Cluster 10 — Enduring instructional design foundations *What employers say:* “instructional design methodologies (ADDIE, SAM, or similar) and adult-learning principles” (Thoughtful AI); backwards design and assessment design named explicitly (Writer); 3–10+ year experience floors appear in nearly every posting reviewed. *What it means:* Everything you were trained to do. Needs assessment, learning objectives, backwards design, formative and summative evaluation, adult learning theory. *Self-rate (1–5):* This one you can probably answer without reflecting. That’s the point of the next section.

Part 3: Why Your Foundation Still Wins

If you rated Cluster 10 a 5 and several of the others lower, you’re not behind. You’re holding the half of the requirement that’s harder to build and slower to fake — and the evidence in Part 1 says employers know it. Hardman’s synthesis puts learning science and ID theory above tool proficiency for 71.3% of the hiring professionals she surveyed. Nearly every posting reviewed for this guide, including the most AI-forward ones, still names years of ID experience as a floor requirement. The gap most mid-career professionals are worried about — “I don’t know enough about AI” — is real, but it’s the smaller half of the equation, not the whole of it.

Some of the fear in this space gets it backwards.

Myth: AI will replace instructional designers and LXDs. Not true, and not what the evidence shows. AI automates tasks — it doesn't replace human judgment, empathy, or the ability to design experiences that actually produce durable learning. The OECD's 2026 Digital Education Outlook found that AI can improve task performance without producing corresponding gains in actual learning when its use isn't deliberately designed for learning. That gap between task-completion and real learning is exactly where a trained instructional designer becomes indispensable — and it's exactly the judgment that no posting reviewed for this guide asks an AI system to provide.

Myth: you need to be a tech expert to use AI well. You don't need to code, and none of the postings above ask you to. What Cluster 4 actually asks for is conceptual fluency — understanding what these systems optimize for, well enough to reason about whether that aligns with the learning outcomes you care about. That's a learning science question dressed up in AI vocabulary, not a computer science requirement.

Myth: AI will make learning less human-centered. Used well, AI can support personalization, accessibility, and engagement at a scale no individual designer could manage alone. Used carelessly, it displaces the cognitive effort that actually produces learning — a pattern we call cognitive displacement. The difference between those two outcomes is a design decision, not a property of the technology. That decision is your expertise, not the tool's.

Myth: you have to master every AI tool to stay relevant. The postings say otherwise. Across the listings reviewed for this guide, no employer asks for a long tool inventory — tool fluency is the commodity register, assumed rather than rewarded. What the most demanding postings ask for instead is conceptual fluency: understanding what these systems do well, where they fail, and how to architect them into a learning journey. A few principles applied consistently beat fifty tools used superficially.

This is where the four affordances this guide has always centered on stop being an abstract framework and become the direct bridge between your existing foundation and what Part 1's evidence shows employers actually want:

- **Memory** — how AI stores, retrieves, and surfaces information at scale, and how that connects to retrieval practice and spaced repetition — directly relevant to Cluster 4's conceptual fluency and Cluster 9's data-driven optimization.

- **Reasoning** — what AI systems can and cannot do when analyzing patterns, generating scenarios, or supporting problem-solving, and where that differs from human reasoning — the foundation for Cluster 6’s verification competency and Cluster 3’s agentic-workflow judgment.
- **Agency** — how AI-powered systems adapt to individual learners, and where scoped autonomy belongs versus where it doesn’t — the design question at the center of Cluster 2’s “architecting AI into the learning journey.”
- **Human-AI Collaboration** — AI and human intelligence working together, each doing what it does best — the frame that runs underneath every register in Part 1, and the one most postings are gesturing toward without naming it explicitly.

Your Masters program and your ADDIE fluency are not obsolete. The work is connecting what you already have to the half of the requirement the tool-chasers in this field are skipping. The postings in Part 1 are not asking you to become someone new — they’re asking you to become explicit about a connection your training gave you the foundation for, but not the vocabulary to name.

Part 4: The Path

This guide is a start, not a finish. Below is where each cluster connects to something concrete you can do next, starting with what costs nothing and moving toward what goes deeper. This section also names where AI4LD doesn’t teach a cluster at all — an honest map is more useful than a complete-looking one that isn’t.

START FREE

Cluster 10 (foundations) you already have — the self-check in Part 2 is itself the free tool. Revisit it as your practice evolves; nothing further to go get.

For five of the other clusters, a free resource takes you further, in the order it makes sense to use them:

1. **Before anything else — the first analysis question.** [Is It a Training Problem?](#) is grounded in Gilbert’s Behavior Engineering Model, Mager & Pipe’s analysis of performance problems, and needs-assessment work — worth running before reaching for an AI-enabled training solution at all.

2. **Cluster 4, connecting how people learn to how AI works — and the start of project support.** The [AI Learning Context Lens](#) is free and deliberately open: bring a real project's context and it returns a learning-science-grounded first read on where AI fits. It's also a working taster of the eight-agent **ILDF 2.0 Practitioner Stack** described later in this guide — the stack of agentic support for your ID/LXD projects, worth experiencing before you get there.
3. **Cluster 1, GenAI/LLM tool fluency.** The [AI Tool Intelligence tracker](#) is a free, participatory tool for the AI4LD community to score AI tools and platforms against a learning-science rubric, so you can better evaluate technology tools for your own context.
4. **Cluster 2, architecting AI into the learning journey.** The [AI Affordance Mapper](#) (free to try, then paid) gives you a sense of the performance support AI4LD builds into its apps and agents, and helps you think through the four AI affordances against a design problem you're already familiar with.
5. **Cluster 5, teaching AI literacy to others / AI-enablement.** [AI Literacy Curriculum Design](#) — ebook, video, and companion agent (the AI Literacy Curriculum Auditor) — is free. HoneyBook and Forward Financing's postings describe this same register as "Just-in-Time" guides and a "people-first approach to AI adoption" for colleagues; designing training *about* AI is instructional design work you already know how to do, applied to a new subject. This free microlearning unit demonstrates the design pattern of the microlearning units listed in the following bundles, purchased separately (\$25 each) or as a four-unit bundle (\$97). Content varies by unit, but every product is grounded in current research and literature.

None of these require a purchase to be useful on their own.

BUNDLE A — HOW AI AND LEARNING WORK · \$97 (UNITS \$25 EACH)

Maps most directly to **Cluster 4 (AI/ML conceptual knowledge)** — the register that Harvard's posting names as a preferred qualification and that almost no competing resource teaches to instructional designers specifically. Four microlearning units, each pairing an ebook and video with a live companion app:

- **Unit 1: Bridging How People Learn and AI** — the pairing this whole guide is built on: how human learning actually works (dual processing, self-directed learning) set against how AI operates, the grounding every later unit and cluster draws from. Its companion, the **AI Learning Context Lens**, is the same free, open taster from the list above — a learning-science first read on where AI fits a real project.
- **Unit 2: Introduction to GenAI and LLMs** — the conceptual grounding behind Cluster 1's "GenAI/LLM fluency" ask. Its companion, the **Multimodal Assessment Lab**, works through what counts as evidence of learning inside a multimodal (video, audio, interactive) experience — assessment design for exactly the AI-era formats employers are asking learning designers to produce.
- **Unit 3: Reinforcement Learning for Learning Designers** — connects directly to Cluster 4's "training methodologies" language. Its companion, the **Reward Design Simulator**, lets you define what an adaptive system optimizes for — and see the consequences of that choice — before anything gets built.
- **Unit 4: Neural Networks for Learning Designers** — the "foundational models... architectures" half of Cluster 4. Its companion, **Why Did It Do That?**, walks you through analyzing why an AI system behaved the way it did — the unit's neural-network concepts applied to a real, puzzling output.

This is the "just enough" technical foundation — not a computer science degree, but genuine conceptual fluency you could speak to in an interview: it starts from what you already know about how people learn and connects it to AI.

BUNDLE B — AI AFFORDANCES FOR LEARNING DESIGN · \$147

Maps directly to **Cluster 2 (architecting AI into the learning journey)** and **Cluster 3 (agentic workflows)** — the registers where Stripe, Harvard, and Paradox's postings concentrate their most demanding language. Four microlearning units, each pairing an ebook and video with a live AI Design Agent, so the four affordances from Part 3 become something you integrate into your ID/LXD practice, not just read about:

- **Memory: Memory as AI Affordance** — how AI stores, retrieves, and loses information — context windows as working memory, persistence and its absence — connected to retrieval practice and spaced repetition; the depth behind Cluster 4's conceptual fluency and Cluster 9's data-driven optimization. Its agent, the **Memory Agent**, works a spaced-recall schedule into a real memory-architecture decision.

- **Reasoning: What Reasoning Agents Actually Do** — what AI systems can and cannot do when they analyze, generate, and “reason,” and where that differs from human reasoning; the foundation for Cluster 6’s verification competency. Its agent, the **Design Reasoning Partner**, asks for your reasoning before offering its own, argues the opposing conclusion at full strength, and leaves you with a co-built design decision record.
- **Agency: Designing with Agentic AI** — where scoped autonomy belongs in a learning design and where it doesn’t; the heart of Cluster 3’s agentic-workflow judgment. Its agent, the **Scoped Autonomy Coach**, is a live example of one agent working within adjustable autonomy and fixed guardrails.
- **Human-AI Collaboration** — AI and human intelligence working together, each doing what it does best — the frame running underneath every register in Part 1. Its agent, the **Co-Design Partner**, attributes every contribution in a co-authored document as yours, AI’s, or blended — Cluster 6’s verification habit, practiced live.

The four ebooks are in staged release — check the product page for which are purchasable today, worth knowing before you buy in at the module you most need. The five live agents (the Affordance Mapper plus one per affordance) are themselves worked examples of Cluster 3’s agentic-workflow competency — working with them shows you how design-focused agents behave from the practitioner’s side: what a well-scoped agent asks of you, where it defers to your judgment, and where it holds its boundaries. That’s experience *using* agents thoughtfully — a real first step toward the agentic-workflow judgment the postings describe — not agent-building itself, which is its own more technical skill.

ILDF 2.0 PRACTITIONER STACK · \$197

Eight agents that support an applied ID process from analysis through evaluation — a good place to start as an instructional design professional. [Get the Practitioner Stack →](#)

1. **Performance & Instructional Analysis** — the needs-assessment question that comes before any solution: what is actually causing the performance gap?
2. **LXD Evidence Assistant** — the research evidence behind your design choices, cited and checkable (also available as its own standalone product).
3. **Learner Analysis** — who your learners actually are: audience analysis and learner personas grounded in empathy research.

4. **Context & Task Analysis** — the setting, constraints, and the real work your design has to serve.
5. **AI Affordance Mapper** — which of the four AI affordances genuinely fits your design problem: the same Mapper from the free list above, here in its place in the full process.
6. **Learner Experience Simulator** — walk your designed flow as a learner would, before any build cost is spent.
7. **Formative Evaluation Designer** — evaluation designed in while you design, not bolted on after launch.
8. **Learning Analytics Interpreter** — the direct answer to Cluster 9: paste pilot or course data and the claim you're tempted to make from it, and it pressure-tests that interpretation — rival explanations, measurement cautions, and a flag for when a finding means the design itself needs another look.

ILDF 2.0 FULL STACK — COMING SOON

The framework behind everything above is specified in a published white paper — Bannan, B. (2026). *The integrative learning design framework revisited: AI affordances, risks, and guardrails for learning experience design research and practice* [White paper]. Zenodo. <https://doi.org/10.5281/zenodo.20970396> — a framework over twenty years in development; ILDF 2.0 is its AI-enabled version. Whether your home tradition is ADDIE or LXD/UX design, you'll recognize the cycles; ILDF 2.0 names them explicitly and pairs each stage with a purpose-built agent. It's the full answer to **Cluster 2** at process scale, and touches **Cluster 6 (responsible/ethical AI use)** and **Cluster 9 (learning analytics)** throughout. The Full Stack extends the Practitioner Stack into the complete agentic ecosystem:

- **The complete agent roster** — 20+ design agents across all five ILDF 2.0 phases, beyond the Practitioner Stack's eight
- **Design Journey persistence** — your project's decisions and artifacts carried across agents and sessions, not one-off conversations
- **Decision-log exports** — a defensible record of the design reasoning behind each project
- **Founding Practitioner cohort entry** — a capped first group with locked pricing and a named channel for design input

WHERE AI4LD IS HONEST ABOUT THE GAP

Two clusters from Part 2 are not fully covered by anything in the AI4LD system:

- **Cluster 8 (AI-assisted video/avatar production)** is a real gap. AI4LD uses tools like HeyGen for its own video production, but doesn't sell instruction in the production tools themselves — the editing, voice, and visual-generation software Harvard's and Metropolis's postings name. If that's the specific skill you're closing, the platforms' own tutorials, or a dedicated production-tools course, are the right resource — that's a production skill, not a learning-science one, and treating it as though it were would misrepresent what this guide can teach you.
- **Cluster 7 (prompt engineering as a standalone, dedicated skill)** is only partially covered. Prompt design is addressed inside the affordance ebooks and companion agents as embedded practice, but there's no dedicated, customer-facing prompt-systems resource in the current AI4LD lineup. If you need a dedicated prompt-engineering curriculum specifically, there are multiple other resources providing that.

No part of this path requires urgency. Look at where your own self-check in Part 2 landed, pick the cluster where the gap matters most to your current work, and start there.

Sources

Sources marked **[FULL]** were fetched from a live or cached full posting/report and quotes are verbatim as retrieved. Sources marked **[PARTIAL]** were recovered from search-engine snippets only, because the original posting had expired or was blocked at fetch time; quoted text in those cases is reproduced as retrieved but should be understood as snippet-level, not full-posting, evidence.

Job listings cited:

- Forward Financing, Instructional Designer — <https://www.edtech.com/jobs/instructional-designer-7329> — [PARTIAL/snippet]
- Demandbase, Senior Instructional Designer — <https://builtin.com/job/senior-instructional-designer/3969382> — [FULL]

- Harvard Business School, Instructional Designer (HBS AI Institute) — <https://careers.harvard.edu/job/instructional-designer-hbs-ai-institute-in-boston-ma-united-states-jid-1016> — [FULL]
- Stripe, Learning Architect — <https://builtin.com/job/learning-architect/9979873> — [FULL]
- Paradox, LXD – Learning Experience Designer – AI x EdTech — <https://builtin.com/job/lxd-learning-experience-designer-ai-x-edtech-full-remote-french-fluent/7452260> — [FULL]
- Thoughtful AI, Learning & Development Specialist — <https://builtin.com/job/learning-development-l-d-specialist/6213229> — [FULL]
- Coursera, Learning Designer – Technical/AI — <https://careers.coursera.com/jobs/learning-designer-technical-ai-india> — [FULL]
- HoneyBook, AI Instructional Designer, Member Care — <https://job-boards.greenhouse.io/honeybook/jobs/7206065> — [PARTIAL/snippet]
- Prolific, AI Instructional Designer — <https://job-boards.eu.greenhouse.io/prolific/jobs/4495617101> — [PARTIAL/snippet]
- Metropolis Technologies, Senior Learning Experience Designer — <https://job-boards.greenhouse.io/metropolis/jobs/7634811003> — [PARTIAL/snippet]
- Capital One, Senior Associate, Learning Experience Designer (counter-example, no AI language) — <https://www.capitalonecareers.com/job/richmond/senior-associate-learning-experience-designer-job-family-learning/1732/84767024768> — [FULL]
- USAA, Instructional Designer II (counter-example, no AI language) — [PARTIAL, per research file]
- Writer, Senior Learning Experience Designer, Generative AI — <https://www.builtinsf.com/job/senior-learning-experience-designer-generative-ai/6360044> — [FULL]

Market and research reports cited:

- Hardman, P. (2025, July 24). Blood in the Instructional Design Machine? <https://drphilippahardman.substack.com/p/blood-in-the-instructional-design> — source of the 71.3% hiring-criterion synthesis and the 21–28% derived salary-premium estimate; transparent synthesis/opinion piece, attributed throughout as Hardman’s analysis.

- Hardman, P. (2025, April 17). Instructional Design Isn't Dying — It's Specialising. <https://drphilippahardman.substack.com/p/instructional-design-isnt-dying-its> — source of the “AI Learning Architect” specialization taxonomy.
- Johns Hopkins School of Education. What Is Learning Engineering? <https://education.jhu.edu/news/what-is-learning-engineering-inside-one-of-educations-fastest-emerging-fields/> — explainer/trade press on the “Learning Engineer” title (coined by Herbert Simon, 1967), now institutionalizing via IEEE ICICLE professional standards; independent, non-Hardman corroborating signal for the specialization-direction claim — definitions-level evidence, not a prevalence statistic.
- Taylor, D. (2026, February). Global Sentiment Survey 2026. https://donaldhtaylor.co.uk/research_base/global-sentiment-survey-2026/ — n=3,500+, sentiment survey.
- Indeed Hiring Lab. (2026, July 8). AI and Job Postings: From Destruction to Creation? <https://www.hiringlab.org/2026/07/08/ai-and-job-postings-from-destruction-to-creation/> — macro postings data, no ID-specific breakout.
- ATD Research. (2025, August 4). AI in Instructional Design: Transforming Workflows and Content Creation. <https://www.prnewswire.com/news-releases/atd-research-ai-tools-are-benefitting-instructional-designers-302521023.html> — n=232, named-sample practitioner survey.
- Lightcast. (2025, July 24). Beyond the Buzz. <https://lightcast.io/resources/blog/beyond-the-buzz-press-release-2025-07-23> — economy-wide AI-skills salary premium data (1.3B postings); NOT an ID-specific figure — see honest-hedges note in Part 1.
- RedThread Research. (2026, June 10). State of L&D 2026. <https://finance.yahoo.com/sectors/technology/articles/redthread-research-state-l-d-130000028.html> — 461-vendor supply-side database.
- Deloitte. (2026). Global Human Capital Trends. <https://www.deloitte.com/us/en/insights/topics/talent/human-capital-trends.html> — large annual global HR-leader survey.
- OECD. (2026). OECD Digital Education Outlook 2026: Exploring effective uses of generative AI in education. OECD Publishing. <https://doi.org/10.1787/062a7394-en>
- Bannan, B. (2026). The integrative learning design framework revisited: AI affordances, risks, and guardrails for learning experience design research and practice [White paper]. Zenodo. <https://doi.org/10.5281/zenodo.20970396>