

JUSTIN SEELEY

BUILDING LEARNING

THAT AI CAN'T COMPLETE

A PRACTICAL GUIDE TO BUILDING CAPABILITY
INSTEAD OF DEPENDENCY WITH AI.

Introduction: The Paradox We're Not Talking About

Something strange is happening in corporate learning right now, and most organizations haven't fully grasped what it means.

According to recent research from the UK, 92% of students are using AI tools in their education. In corporate settings, the number is likely higher—there's less stigma, more pressure to be efficient, and fewer restrictions. AI has become as ubiquitous as Google, as automatic as spell-check, as integrated into daily work as email.

But here's what should concern every learning leader: the widespread adoption of AI isn't just changing how people complete training. It's revealing that much of what we've been calling "learning" was never really learning at all.

When a language model can ace your leadership development program without ever having led anyone, when ChatGPT can pass your compliance certification without understanding a single regulation, when employees can complete entire learning paths while building zero actual capability—that's not an AI problem. That's a learning design problem that's been hiding in plain sight for years.

AI didn't break corporate learning. It just made it impossible to ignore what was already broken.

For decades, L&D has optimized for the wrong outcomes. We measured completion rates instead of capability development. We tested knowledge recall instead of judgment. We designed for information transfer instead of skill building. We built systems that rewarded getting through content rather than actually changing what people could do.

This worked reasonably well—or at least, we told ourselves it did—when accessing information was hard and demonstrating knowledge required actual understanding. But AI demolished those assumptions overnight.

Now we're faced with a choice that will define the future of corporate learning: do we double down on surveillance and restrictions, trying to preserve a system that was never serving us that well? Or do we use this disruption as an opportunity to rebuild learning around what actually matters—developing capabilities that AI can't replicate and judgment that AI can't automate?

Most organizations are choosing the first path without realizing it. They're implementing AI detection software, requiring proctored exams, banning tools, and creating policies that position learners and organizations in an adversarial relationship. This approach fails on multiple levels: it's practically unenforceable, pedagogically counterproductive, and strategically shortsighted. You can't prepare people for an AI-saturated world by forbidding them from using AI.

But the alternative—blindly embracing AI without understanding its implications for learning—is equally problematic. I've seen organizations rush to AI-generate courses, implement AI grading, and celebrate efficiency gains while their people become less capable, not more.

The issue isn't whether to use AI in learning. That ship has sailed. The issue is understanding the fundamental difference between AI that builds capability and AI that creates dependency—and redesigning our entire approach to learning accordingly.

This is where learning science becomes essential. Research into how humans actually learn—cognitive load theory, metacognition, scaffolded instruction, deliberate practice—provides a foundation for making smart choices about when and how to integrate AI. The same tool can either support learning or undermine it, depending on what cognitive work is happening in the learner's mind.

Throughout this piece, I'll explore what that redesign looks like in practice. We'll examine why efficiency can be the enemy of learning, how to distinguish between scaffolds and crutches, what assessment should measure in an AI age, and how to build learning experiences that develop the distinctly human capabilities that matter more than ever.

But first, we need to understand what actually works in learning—and why so much of what we've been doing doesn't.

The stakes couldn't be higher. Get this right, and learning becomes a genuine competitive advantage, developing people who are exceptional at working with AI rather than replaced by it. Get it wrong, and L&D becomes an increasingly irrelevant compliance function while the real learning happens elsewhere.

Let's start by going back to the fundamentals of how humans learn, and why those fundamentals matter more in an AI age, not less.

What Actually Works in Learning

Before we can talk about AI's role in learning, we need to be brutally honest about what learning actually is and how it happens.

Learning is not information transfer. It's not content consumption. It's not clicking through modules or watching videos or reading articles. Those things might be part of learning, but they're not learning itself.

Learning is the acquisition of capability through effortful cognitive processing. It's doing something you couldn't do before because you've built new mental models, developed new skills, or changed how you think about problems.

This matters because AI is exceptionally good at information transfer and terrible at the effortful processing that creates actual learning. Which means if your learning design is primarily about getting information into people's heads, you've already been disrupted. AI does that better, faster, and more personalized than any training program ever could.

So what does the research tell us about how learning actually happens?

Cognitive load is the limiting factor. John Sweller's decades of research on cognitive load theory shows that human working memory is severely constrained. We can hold roughly 4-7 pieces of information in conscious attention at once. When we overload working memory with too much information or too many simultaneous demands, learning collapses. Effective instructional design manages cognitive load carefully, removing extraneous demands so learners can focus on the essential work of building understanding.

Active processing is non-negotiable. Learning requires the learner to actively work with information—connecting it to prior knowledge, applying it to problems, explaining it in their own words, testing it against examples. Passive consumption might feel like learning, but it rarely produces transferable capability. This is why lectures are less effective than problem-solving, why reading about leadership is different from actually leading, why watching someone code doesn't make you a programmer.

Struggle is productive (to a point). When learning feels easy and efficient, it's often not happening. The research on "desirable difficulties" shows that introducing certain kinds of challenges—spacing practice over time, interleaving different concepts, requiring retrieval from memory—makes learning harder in the moment but stronger in the long

run. The key is productive struggle: challenging enough to require effort, but not so overwhelming that learners give up.

Metacognition amplifies everything. The ability to think about your own thinking—to monitor whether you actually understand something, to recognize when you need different strategies, to evaluate the quality of your reasoning—is perhaps the most powerful learning skill. Students with strong metacognitive skills learn faster and more deeply across domains. And crucially, metacognition can be taught.

Context and application matter immensely. Knowledge that sits in isolation is fragile and hard to access when you need it. Knowledge that's connected to when and how you'll use it, that's been practiced in realistic contexts, that's been reflected upon—that knowledge transfers. This is why the best technical training includes realistic scenarios, why leadership development needs real leadership challenges, why you can't become good at something by only reading about it.

Here's what all of this means for AI in learning: **AI is phenomenal at reducing cognitive load, but learning requires cognitive load.** AI can provide information instantly, but learning requires effortful processing of that information. AI can make things easy and efficient, but learning often requires productive struggle.

This creates a fundamental tension. The very thing that makes AI powerful—its ability to eliminate cognitive effort—is precisely what can undermine learning if we're not careful.

Educational psychologists have long distinguished between two types of support in learning: scaffolds and crutches. A scaffold is temporary support that helps you build capability you'll eventually exercise independently. A crutch is permanent dependency that prevents capability from developing. The same tool can be either one, depending on how it's designed and what cognitive work the learner is actually doing.


This distinction is critical for everything that follows.

The Scaffold vs. Crutch Problem

When AI helps a learner understand a concept they're struggling with, explains it in multiple ways, provides examples, and allows them to practice applying it—that's a scaffold. The AI is managing cognitive load strategically, removing obstacles to understanding while the learner does the essential work of building mental models.

SCAFFOLD OR CRUTCH?

Evaluating AI's Role in Learning



When AI generates the final work product for a learner, does analysis they should be doing themselves, or makes decisions that require contextual judgment—that's a crutch. The cognitive work that would create learning isn't happening. The learner is producing output without developing capability.

Same technology. Radically different outcomes.

Here's how to distinguish between them:

AI functions as a scaffold when it:

- Explains foundational concepts the learner doesn't understand yet
- Breaks complex information into manageable pieces
- Provides examples to illustrate abstract principles
- Offers feedback on the learner's attempts
- Gradually increases challenge as competence develops
- Makes the learner less dependent over time

AI functions as a crutch when it:

- Generates the final deliverable the learner should create
- Does analysis that requires the learner's contextual knowledge
- Makes decisions that should exercise the learner's judgment
- Prevents productive struggle that builds understanding
- Becomes more necessary over time instead of less
- Allows the learner to produce work without processing

The goal of any learning intervention should be to make itself unnecessary. If your training program makes learners more dependent on support systems rather than more autonomous, it's not training—it's offloading.

Let me give you a concrete example of how this plays out.

A sales enablement team builds a comprehensive product knowledge certification. It covers technical specifications, competitive positioning, common objections, customer success stories. When completed through traditional means—reading materials, taking practice quizzes, discussing with peers—reps who passed the certification consistently outperformed those who hadn't.

Then AI tools became widely available. Completion times dropped by 60%. Reps loved being able to ask questions and get instant answers instead of searching through dense materials. Satisfaction scores went up. By traditional metrics, the program was more successful than ever.

But six months later, they noticed something troubling. New reps who'd completed the AI-enhanced certification were struggling in actual customer conversations. They could answer direct questions about product specifications, but they couldn't read the room. They couldn't adapt their pitch when they sensed resistance. They couldn't connect product features to customer pain points in real time.

The veterans who'd gone through the harder, slower certification didn't have these problems.

What changed? The old program forced reps to struggle with the material—to search for information, connect concepts, build mental models. That cognitive effort was frustrating and time-consuming, but it was creating integrated understanding that allowed for flexible, contextual application.

The AI-enhanced version eliminated that struggle. Reps could get answers instantly, but they never built the deep mental models. They had knowledge on demand, but they hadn't internalized understanding.


This is the hidden cost of efficiency in learning. Sometimes the hard way is the only way that actually works.

The Four-Level Integration Framework

Based on cognitive load theory and experimentation with AI in corporate learning, I've developed what I call the AI Integration Ladder—a framework for determining how deeply AI should be integrated at different stages of skill development.

The key insight is that the appropriate role for AI changes dramatically depending on where learners are in their capability journey. What works at mastery level can be destructive at foundation level.

THE AI INTEGRATION LADDER



Build foundational skills before advancing to higher levels of AI integration

Level 1: Foundational Skills (Minimal AI)

When learners are building foundational capabilities, AI should be minimal or absent. This is where cognitive effort matters most—where learners need to struggle productively to build the mental models that enable everything else.

If you're learning basic financial analysis, you need to work with spreadsheets yourself, make mistakes, figure out why your formulas aren't working. If you're developing critical thinking skills, you need to construct arguments without assistance, wrestle with logic, learn to spot flaws in reasoning. If you're building sales conversation skills, you need to have awkward real conversations, not rehearse smooth AI-generated scripts.

AI at this level is like using a calculator before you understand multiplication. It gives you answers without building the foundational understanding you need for more complex work.

Appropriate AI use at Level 1:

- Explaining prerequisite concepts the learner should already know but doesn't
- Providing clarification on confusing instructions
- Offering encouragement and motivation
- Nothing that does the primary cognitive work

Examples:

- Financial analyst learning modeling fundamentals: builds models manually, struggles with logic errors, develops intuition
- New manager learning feedback frameworks: practices giving real feedback to real people, learns from actual responses
- Developer learning to code: writes code themselves, debugs their own errors, builds problem-solving capability

Level 2: Guided Practice (AI as Explainer)

Once foundational understanding exists, AI can serve as an explainer and elaborator. The learner is still doing the primary cognitive work, but AI helps when they're genuinely

stuck, provides alternative explanations when the first one didn't click, and offers examples to illustrate concepts.

This is the scaffold phase. AI reduces unproductive struggle (the kind that comes from poor explanations or missing prerequisite knowledge) while preserving productive struggle (the kind that builds robust understanding).

Appropriate AI use at Level 2:

- Answering specific questions about concepts
- Providing worked examples to illustrate principles
- Offering alternative explanations when learners are confused
- Suggesting practice problems at the right difficulty level
- Explaining why a learner's attempt was incorrect

Examples:

- Marketing professional learning customer segmentation: uses AI to refresh statistical concepts, but does the actual segmentation analysis themselves
- Project manager learning risk assessment: asks AI to explain probability concepts, but identifies and evaluates risks in their own projects
- Designer learning accessibility: gets AI to explain WCAG standards, but makes their own design decisions

Level 3: Skill Integration (AI as Thought Partner)

When learners have solid understanding but are working on integration and application, AI can serve as a thought partner. It challenges assumptions, offers alternatives, plays devil's advocate, helps explore edge cases.

This is where AI adds maximum value. The learner has enough expertise to evaluate AI's contributions critically, but AI expands their thinking beyond what they'd generate alone. The learner is making final decisions, but AI is making those decisions better.

Appropriate AI use at Level 3:

- Generating alternative approaches to compare

- Challenging the learner's assumptions
- Identifying potential blind spots or risks
- Simulating stakeholder perspectives
- Stress-testing ideas against different scenarios

Examples:

- Experienced leader developing strategy: uses AI to generate alternative strategic approaches, then evaluates each against organizational culture, resource constraints, and stakeholder needs
- Senior analyst building a model: has AI suggest alternative methodologies, then makes informed decisions about which approach fits their context
- Learning designer creating a program: uses AI to brainstorm engagement approaches, then selects based on learner needs and constraints

Level 4: Mastery and Innovation (AI as Collaborator)

At mastery level, AI becomes a true collaborator. The expert uses AI to extend their capabilities, automate routine elements, and tackle problems at a scale or complexity that wouldn't be possible alone.

But this only works if the previous levels built genuine expertise. Without deep understanding, you can't distinguish good AI outputs from plausible-sounding nonsense. You can't make the judgment calls about when to trust AI, when to modify its suggestions, and when to override it entirely.

Appropriate AI use at Level 4:

- Automating routine analysis to focus on interpretation
- Generating multiple options to synthesize
- Handling scale that would be impossible manually
- Rapid prototyping and iteration
- Complex simulation and modeling

Examples:

- Data scientist uses AI to generate code, test edge cases, and explore alternative algorithms, but understands the underlying mathematics well enough to validate results
- Executive uses AI to analyze vast amounts of market data, but brings decades of industry experience to interpret what matters
- Learning architect uses AI to rapidly prototype course structures, but designs the pedagogical approach based on deep expertise

The critical insight: Most corporate learning fails because it tries to jump straight to Level 3 or 4 without building the foundational understanding in Levels 1 and 2. People end up dependent on AI without the capability to use it effectively. They can produce plausible-looking work, but they lack the judgment to know when it's actually good.

What This Means for Assessment

Traditional assessment is dying, and we need to build something better in its place.

Here's what doesn't work anymore:

- Multiple choice questions (language models ace these)
- Essays that ask for information synthesis (GPT-4 writes better than most humans)
- Knowledge checks that test recall (that was never valuable anyway)
- Time-limited exams (AI is faster)
- Any assessment that measures what someone knows rather than how they think

The fundamental problem is that we built assessment around information and knowledge, and AI is fundamentally an information technology. If your assessment can be completed by someone with zero understanding but access to ChatGPT, it's not measuring learning—it never was.

Here's what research and early implementations suggest does work:

Portfolio-based assessment showing learning progression over time. Not just final deliverables, but the evolution of thinking, the iterations, the refinements. Ethan Mollick

at Wharton has students submit logs of their AI interactions alongside their business plans, showing what questions they asked, what responses they got, what they kept, what they modified, and why.

Process documentation that makes thinking visible. Have learners explain their reasoning, document their decision-making, reflect on alternatives they considered. This can't be faked because the thinking process itself is what you're assessing.

Mastery-based progression with adaptive challenge. Rather than one-shot exams, use systems that continue probing understanding until learners demonstrate they can apply concepts in novel contexts. The cognitive co-pilot project I'm part of does this—it won't let learners progress until they've shown genuine comprehension, not just memorized responses.

Collaborative problem-solving with peer evaluation. Have learners work together on complex problems, then evaluate each other's contributions. The interpersonal dynamics, the negotiation of ideas, the building of shared understanding—that's distinctly human and can't be automated.

Applied judgment in context. Give learners realistic scenarios that require them to integrate multiple pieces of knowledge, make decisions with incomplete information, and explain their reasoning. The judgment call is what matters, not the information recall.

Let me give you a concrete example of what this looks like in practice.

Traditional approach: "Write a strategic plan for entering a new market."

Result: AI can generate a completely plausible strategic plan. You have no idea if the learner understands strategy or just knows how to prompt ChatGPT.

Redesigned approach: "Use AI to generate three different strategic approaches for entering a new market. Document the prompts you used and your reasoning. Evaluate each approach against our company's specific constraints—budget, timeline, current capabilities, organizational culture. Write an analysis explaining which elements from each approach you'd combine and why, including what the AI got wrong or couldn't know. Have a peer challenge your reasoning. Revise based on their feedback. Submit everything: your prompts, the AI responses, your evaluations, the peer feedback, and your final synthesis."

Now you're assessing:

- Prompt design (requires clarity and strategic thinking)
- Critical evaluation (can they spot AI's limitations?)
- Contextual judgment (do they understand what AI can't know?)
- Synthesis (can they integrate insights effectively?)
- Receptiveness to feedback (do they update their thinking?)
- Metacognition (are they aware of their reasoning process?)

That's not something you can outsource to AI. The AI can help with pieces, but the cognitive work that matters is distinctly human.

ASSESSMENT REDESIGN:

From Output to Thinking

TRADITIONAL ASSESSMENT

TESTING OUTPUTS

THE ASSIGNMENT

Write a strategic plan
for entering a new market

WHAT THIS MEANS

- Can they write clearly?
- Do they know the format?
- Did they use AI?
- Do they understand strategy?

All unknowable from final output

THE PROBLEM

AI can generate a perfect
strategic plan in 30 seconds.

*Zero proof of actual
learning or understanding*

REDESIGNED ASSESSMENT

TESTING THINKING

THE ASSIGNMENT

1. Use AI to generate 3 different strategies
2. Document your prompts
3. Evaluate each against constraints
4. Explain what AI got wrong
5. Get Peer Feedback
6. Submit everything

WHAT'S ACTUALLY BEING TESTED

- Prompt design
- Critical evaluation
- Contextual judgment
- Synthesis ability
- Receptiveness to feedback
- Metacognitive awareness

THE KEY SHIFT: Don't assess what they produced. Assess how they thought about it.

The Real Skills Gap

Here's what makes this urgent: the skills that matter most in an AI age are precisely the skills that traditional corporate learning has been worst at developing.

Metacognition—the ability to monitor and adjust your own thinking—is now a core workplace competency. When you're working alongside AI systems that confidently assert things that are completely wrong, you need sophisticated self-awareness. You need to know when you actually understand something versus when you're fooling yourself. You need to recognize when AI is operating within its capabilities versus when it's hallucinating. You need to catch yourself when you're accepting plausible-sounding nonsense because it's easier than thinking critically.

Critical discernment becomes essential when information is infinite but wisdom is scarce. Large language models reproduce biases from their training data. They optimize for sounding confident rather than being correct. They can't distinguish between well-established facts and fringe theories. Knowing how to evaluate sources, spot logical fallacies, recognize when something sounds right but isn't—these capabilities require deep domain expertise combined with intellectual humility.

Contextual judgment is what separates competent from exceptional performance. AI can generate five strategic approaches to a business problem in 30 seconds. But it can't know that your CEO hates a particular framework because of bad history, that your best engineer is about to go on parental leave, that your largest customer is privately considering a switch, that your organizational culture will reject certain types of change. That contextual knowledge, and the judgment to apply it, is irreducibly human.

Collaborative sense-making is what happens when humans work through complexity together. It's messy, iterative, fundamentally social. It involves reading body language, building on half-formed ideas, managing group dynamics, knowing when to push and when to let go. AI can support this process—taking notes, suggesting alternatives, organizing information. But it can't replace the human work of building shared understanding.

Ethical reasoning in ambiguous situations. AI can tell you what's legal, what's common practice, what outcomes are likely. It can't tell you what you should do when every option has trade-offs, when stakeholders have conflicting needs, when short-term and

long-term interests diverge. That requires values, experience, and the kind of practical wisdom that only comes from having real stakes in real situations.

These capabilities can't be developed through content consumption. They can't be assessed through multiple choice questions. They can't be automated or outsourced. And critically, they can't be developed in learning programs that treat AI as either forbidden or as a way to make everything faster and easier.

They require learning experiences that are challenging, reflective, collaborative, and grounded in realistic application. They require assessment that measures thinking, not just knowledge. They require time for struggle and space for failure.

In other words, they require the exact opposite of what most corporate learning has become.

The Crisis Is the Opportunity

Here's the uncomfortable truth: AI didn't create the problems in corporate learning. It exposed them.

The crisis in assessment isn't that AI can pass our tests. It's that our tests were never measuring anything that mattered.

The problem with efficiency isn't that AI makes things too easy. It's that we confused ease with effectiveness.

The challenge with detection isn't that we can't catch people using AI. It's that we built systems where using AI is indistinguishable from legitimate work.

This is actually good news. Because it means the solution isn't better surveillance or more sophisticated detection. The solution is rebuilding learning around what actually matters: developing capabilities that AI can't replicate and judgment that AI can't automate.

What would that look like?

Learning experiences designed around cognitive work that matters. Not "complete this module" but "solve this realistic problem, document your thinking, explain your

reasoning, get challenged by peers, revise your approach." Not "watch this video" but "try this in your actual work, reflect on what happened, identify what you'd do differently."

Assessment that measures thinking, not outputs. Show us your reasoning process, not just your final answer. Document how you used AI, what you kept, what you changed, and why. Demonstrate that you can evaluate information critically, not just locate it. Prove that you understand the limitations of your knowledge, not just its breadth.

Transparent integration of AI as part of the learning process. Here's when and how AI can help you. Here's what it can't do. Here's how to know if you're using it as a scaffold or a crutch. Here's how professionals in your field are using it responsibly. Now practice using it that way, get feedback, refine your approach.

Focus on capabilities that become more valuable as AI becomes more capable. The better AI gets at analysis, the more important human judgment becomes. The faster AI can generate options, the more critical evaluation skills become. The more information is available, the more essential discernment becomes.

This isn't about banning AI or embracing it uncritically. It's about using AI's disruption as an opportunity to finally build the learning systems we should have built all along—systems that develop genuine capability instead of just documenting content consumption.

What L&D Leaders Should Do Now

If you're a learning leader trying to figure out what to do about AI, here's my honest assessment based on the research and what's working in practice:

1. Stop fighting the unwinnable war.

AI detection doesn't work reliably. Bans aren't enforceable. Surveillance damages trust and doesn't improve learning. You're spending resources on an arms race while your learners develop increasingly sophisticated workarounds. More importantly, you're

positioning the most transformative technology of their careers as something forbidden in the one place designed to help them grow.

2. Audit every learning program for obsolescence.

Look at what each program requires cognitively. If it's primarily recall, reorganization, or summarization of information, that program is already obsolete. Either redesign it around higher-order skills (analysis, evaluation, creation, judgment) or kill it. Be ruthless. Half-measures won't work.

3. Teach critical AI literacy explicitly.

Not "how to write good prompts" (people figure that out quickly), but genuine literacy: When does AI hallucinate and how do you catch it? What biases are embedded in training data? What can't AI know about your specific context? How do you validate AI-generated information? What are the ethical implications of different AI applications? Make this core curriculum, not an elective.

4. Make AI use transparent and expected.

Build AI into your learning design intentionally. Require documentation of AI use as part of assignments. Create rubrics that assess both the work product and the critical thinking around AI's contributions. Normalize the idea that professionals use AI thoughtfully, and that learning to use it well is part of professional development.

5. Redesign assessment around judgment.

Every assessment should require something AI fundamentally can't do: apply tacit knowledge from experience, make decisions with real stakes, exercise contextual judgment, demonstrate metacognitive awareness, collaborate interpersonally, or reason about ethical trade-offs. If AI could theoretically complete your assessment, it's not assessing what you think it is.

6. Apply the Integration Ladder rigorously.

For each learning objective, determine what level of AI integration makes sense. Foundational skills need minimal AI. Guided practice can use AI as an explainer. Skill integration can employ AI as a thought partner. Mastery work can treat AI as a collaborator. Don't let convenience drive these decisions—let learning science drive them.

THE CHOICE POINT: Two Paths for L&D

PATH A: DEFENSIVE



Surveillance & detection
Ban AI tools
Lock down assessments
Proctored exams
Preserve existing system



**SLOW THEN
SUDDEN
IRRELEVANCE**

Learner disengagement
Disconnected from real work
Talent leaves for modern orgs
Compliance function only

PATH B TRANSFORMATIVE



Redesign for judgment
Transparent AI integration
Measure thinking not outputs
Build critical AI literacy
Focus on capability



**STRATEGIC
COMPETITIVE
ADVANTAGE**

Real capability development
Exceptional AI users
Engaged learners
Strategic asset
Future-ready workforce



**YOUR
ORGANIZATION
TODAY**

The Choice Point
In AI & Learning

Most organizations try both paths at once. That's really just Path A with better marketing.

7. Invest in your learning team's capabilities.

Your learning professionals need new skills themselves. They need to understand both how humans learn (the cognitive science) and how AI systems work (the technical literacy). They need comfort with ambiguity and permission to experiment. They need time to redesign programs, not just layer AI on top of existing approaches. This requires real investment, not just training.

8. Create space for experimentation.

Designate some programs as sandbox environments where you can test new approaches without institutional risk. Try portfolio assessment. Test mastery-based progression. Experiment with transparent AI integration. Measure what actually happens to capability, not just completion rates. Learn fast, share findings, iterate.

9. Partner across disciplines.

The challenges around AI in learning aren't purely educational—they're psychological, sociological, ethical, technical. Build partnerships with people who bring different expertise. Talk to data scientists about how models actually work. Consult ethicists about responsible use. Engage psychologists about behavior change. The best solutions will be interdisciplinary.

10. Focus ruthlessly on capability, not completion.

Stop measuring success by how many people finished something or what their satisfaction scores were. Start measuring whether they can do things they couldn't do before. Can they think more critically? Make better decisions? Use AI more effectively? Apply learning to real challenges? If you can't measure actual capability change, you're measuring the wrong things.

The Choice Point

There's a profound choice point happening in corporate learning right now, and I don't think most organizations realize how consequential it is.

Path A is defensive. Lock down assessments. Surveil learners. Ban AI tools. Try to preserve the existing system. Deploy detection software. Return to proctored exams and handwritten assignments. Treat AI as a threat to be contained.

This path leads to slow, then sudden irrelevance. Learners disengage because your programs feel disconnected from their actual work. The best people leave for organizations that embrace modern tools. Your learning function becomes a compliance necessity rather than a strategic asset. You spend increasing resources fighting an unwinnable battle while capability quietly atrophies.

Path B is transformative. Redesign learning around capabilities AI can't replicate. Make AI use visible, reflective, and integrated. Build assessment that measures thinking, not just knowledge. Teach people to be exceptional at working with AI, not just good at avoiding it. Focus on judgment, metacognition, collaboration, and contextual application.

This path is harder and more uncertain. It requires investment in new capabilities, tolerance for experimentation, and willingness to kill sacred cows. It demands that learning leaders develop new expertise and that organizations give them space to rebuild rather than just patch. It means accepting that much of what you've been doing wasn't actually working, which is uncomfortable.

But it's the only path that leads somewhere worth going.

Most organizations are trying to have it both ways—some restrictions here, some AI tools there, policies that are vague enough to avoid commitment. That's really just Path A with better marketing. It doesn't work because it doesn't address the fundamental problem: we built learning systems for a world that no longer exists.

The organizations that will thrive are those that make the hard choice: to fundamentally rethink what learning means in an AI age and rebuild from that foundation.

The Bottom Line

AI is completing your training while making your people less capable. That's the paradox we're not talking about enough.

The solution isn't better detection or stricter policies. It's building learning that AI literally cannot complete—learning that requires judgment, contextual knowledge, metacognitive awareness, collaborative sense-making, and ethical reasoning.

This means:

- Designing for cognitive work that matters, not just content delivery
- Using AI as a scaffold for building capability, never as a crutch that prevents it
- Assessing thinking and judgment, not knowledge and outputs
- Making AI use transparent and teaching critical literacy explicitly
- Rebuilding around capabilities that become more valuable as AI becomes more capable

The research is clear. The principles are established. The early implementations are showing what works.

What's missing is the institutional will to actually do the rebuilding.

Because here's the truth: you can't prepare people for an AI-saturated world by forbidding them from using AI. And you can't develop genuine capability by optimizing for efficiency and completion.

You have to make the harder choice—to build learning systems that develop the distinctly human capabilities that matter more than ever, using AI strategically to enable that development rather than undermine it.

The question isn't whether AI will transform corporate learning. It already has.

The question is whether you'll respond by defending what's broken or building what's needed.

Did you find this useful?



Like



Comment



Share



Justin Seeley

+ Follow



For more tips & career insights
for learning designers