RAVINDER TULSIANI
ENTERPRISE CAPABILITY EXECUTIVE

PRACTITIONER TOOL

Does It Actually Teach?

AI-Generated Learning Quality Check

Ravinder Tulsiani, DBA — Enterprise Capability Executive

A polished course is not evidence of learning.

AI can produce a course outline, lesson, quiz or activity in minutes. That speed solves a production problem. It does not answer the more important question: will the learning help someone perform better?

A course can look complete, read well and contain accurate information without giving the learner a meaningful opportunity to practise the required performance or demonstrate that they can do it. Before approving AI-generated learning, examine the path from the business need to the evidence of capability.

Six review questions

1. What must the learner be able to do differently?

Start with performance, not content. “Understand feedback” is difficult to evaluate. “Prepare for and conduct a clear performance conversation with an employee” describes something that can be practised and observed. If the required performance is unclear, AI has no reliable target. It can still generate content, but that content may be aimed at the wrong problem.

2. Does the learner actually practise that performance?

Recognition is not the same as performance. A learner may select the correct answer, remember a definition or complete a module without being able to make the decision, perform the task or handle the situation at work. Look for practice that resembles the real requirement: a decision, output, scenario, demonstration, diagnosis or explained judgment. If the performance matters, the learner should have a meaningful opportunity to practise it.

3. Do the content, practice and assessment point to the same outcome?

AI is very capable of producing individual pieces. The pieces still need to work together. Ask whether the explanations prepare the learner for the required performance, whether the practice exercises that performance, and whether the assessment tests the same capability. A well-written activity that does not support the target performance is still a design problem.

4. What evidence would convince you that capability improved?

Completion is evidence that the learning occurred. It is not evidence that capability improved. Choose evidence that matches the task: a completed work product, simulation, observed demonstration, decision with reasoning, realistic case response, or manager observation after application in the work.

5. Will the learning survive contact with the real work?

Check whether the learning reflects the conditions that make the real task difficult: ambiguity, time pressure, competing priorities, imperfect information, system constraints, approval requirements, stakeholder reactions and consequences of poor decisions. If a capable employee would still be blocked by the work system, address that barrier rather than adding more content.

6. Which decisions still require human judgment?

AI can accelerate drafting, variation, examples and first-pass analysis. It should not make consequential design decisions invisibly. Before release, a responsible human should be able to explain why the capability matters, why the practice is appropriate, what evidence is sufficient, which sources are authoritative, what AI could infer, what was checked, who approves the design and what would cause revision. AI can accelerate generation. It does not remove the need for design judgment.

Approval decision

Decision	When it fits and what to do next
Ready to test	The required performance is clear, practice is meaningful, the assessment produces credible evidence, relevant working conditions are reflected and human review is complete. Next step: test with a small representative group.
Revise	The direction is sound but one or more links are weak. Next step: repair the specific weak link and test again.
Rebuild	The learning mainly transmits information or checks recall while the required performance is unclear, unpractised or undemonstrated. Next step: return to the performance requirement and redesign the path.

Boundary: this is a practical review structure, not a validated scoring instrument.

Illustrative example

The following example is illustrative. It is not a client case or a report of measured results.

An AI-generated course for newly promoted managers contains a strong explanation of effective feedback, polished examples and a ten-question quiz. The stated goal, however, is for managers to conduct difficult performance conversations. The course asks learners to recognize good feedback but never requires them to prepare for a conversation, decide what to say, respond to an employee reaction or demonstrate

judgment in an ambiguous case. The problem is not that the course is inaccurate. The problem is that the evidence does not match the performance requirement.

A stronger design might retain the concise explanation, then add realistic cases in which managers prepare an opening, choose how to respond, explain their reasoning and receive feedback against clear criteria. The AI-generated material may still be useful. The design decision remains human.

Reusable quality-check worksheet

01	What business or performance problem is this learning intended to influence?	
02	What must the learner be able to do differently?	
03	Where does the learner practise that performance?	
04	How closely does the practice resemble the real task or decision?	
05	What content directly supports the required performance? What could be removed?	
06	What evidence will show that the learner can perform, not merely recall?	
07	Does the assessment test the same capability the learning claims to build?	
08	Which real working conditions or constraints matter to successful performance?	
09	What would still block a fully capable employee at work?	
10	Which sources are authoritative, and what information still requires verification?	
11	What consequential design decisions were made by a person rather than left to AI defaults?	
12	Decision: ready to test, revise or rebuild? What is the reason?	

Review date: _____ Learning owner: _____ Operational / subject-matter owner: _____
Final approver: _____

Evidence boundary: This tool is Ravinder Tulsiani's professional synthesis of performance-focused learning design, AI-enabled development and evidence principles. It is not a validated psychometric instrument, certification standard or guarantee of business performance.