

AI Learning Claims Audit

Turn the Label Into Observable Behaviour

Ravinder Tulsiani, DBA — Enterprise Capability Executive

Do not evaluate the label. Evaluate what the system actually does.

Learning technology increasingly arrives with strong labels: adaptive, personalized, mastery-based, Socratic, AI tutor. Those labels may describe useful capabilities. They are not, by themselves, evidence of how the learning experience actually works. Before accepting a claim, translate it into something observable: what learner evidence does the system use, what changes because of that evidence, what happens when the learner struggles, what remains the same, and who decides whether the response is appropriate?

The aim is not to challenge a vendor's terminology for its own sake. It is to make the instructional mechanism visible enough to evaluate.

Five audit tests

1. State the claim precisely

Avoid evaluating a broad label. Write down what the product or design is actually claimed to do. "Personalizes feedback based on learner performance" is easier to examine than "AI-powered personalized learning."

2. Identify the learner evidence or trigger

Ask what information causes the system to respond differently: an answer, an error pattern, demonstrated performance, prior activity, stated goals, confidence, time or sequence, or instructor-defined rules. If nothing about the learner meaningfully changes the response, the claim may describe presentation rather than adaptation. That may still be useful. The important point is to know what is actually happening.

3. Observe what changes

Look for a meaningful change in the learning experience: the next task, challenge level, feedback, support, explanation, practice opportunity, sequence, pacing or requirement to demonstrate the capability again. A different wording of the same experience is still a change. It is not necessarily the same thing as changing the instructional decision.

4. Test what happens when the learner does not succeed

Many learning claims become clearer when the learner gets something wrong. Ask: what counts as sufficient performance; what happens when the learner does not demonstrate it; does the system diagnose the problem; does it provide support related to that problem; does the learner get another meaningful opportunity to demonstrate the capability; does the path change, or does the learner eventually receive the same sequence regardless?

5. Make governance visible

An AI-enabled learning system can make many small instructional decisions. Someone still needs to be responsible for the boundaries. Confirm: who defines acceptable behaviour; what sources the system may use; how incorrect or inappropriate responses are identified; whether instructors can inspect or override decisions; what learner data is used; what is retained; what happens when the system is uncertain; who owns the final instructional and governance decision. This audit does not replace privacy, security, accessibility, legal, procurement or technical due diligence. It helps make the learning claim itself easier to inspect.

Applying the audit to common claims

The following are practical questions, not formal definitions or certification standards.

Adaptive

What learner evidence causes the learning experience to change, and what changes as a result? Look beyond branching alone. Ask whether the system meaningfully changes support, challenge, practice, sequence or another instructional decision based on evidence about the learner.

Personalized

What is being personalized, for what reason, and does the change improve the learner's path toward the required capability? Possible personalization may involve timing, examples, explanations, feedback, support, practice or sequence. Do not assume all personalization is equally consequential.

Mastery-based

What constitutes sufficient performance, and what happens when the learner has not demonstrated it? Look for a clear criterion, relevant support and another meaningful opportunity to demonstrate the capability.

Socratic

Does the interaction genuinely probe the learner's reasoning, or does it simply present questions? The important behaviour is whether the system responds to the learner's thinking in a way that helps examine assumptions, reasoning or evidence.

AI tutor

What context can the tutor use, what decisions can it make, what does it remember, and what are its limits? A conversational interface alone does not tell you how much instructional context, learner evidence or governance sits behind it.

Audit decision

Decision	When it fits
Demonstrated	The claimed mechanism can be observed consistently in realistic tests and the instructional behaviour matches the stated claim closely enough for the intended use.
Partially demonstrated	Some elements are observable but important parts remain limited, inconsistent or conditional.
Not yet demonstrated	The label cannot be translated into observable behaviour, the expected mechanism does not occur in testing, or evidence is insufficient.

Boundary: these categories are a practical due-diligence aid, not a product certification or validated rating scale.

Application notes

Test with realistic learner responses, not only ideal inputs. Where possible, test contrasting cases: strong performance, partial understanding, misconception, uncertainty and repeated difficulty. Record what the system actually does rather than what a demo suggests it should do. A feature may be useful even when the broad marketing label is imprecise. The goal is not to disqualify the technology. The goal is to understand the mechanism well enough to make a responsible decision.

Reusable audit worksheet

01	What exact claim are we evaluating?	
02	Where is the claim stated, and what scope does it imply?	
03	What learner evidence or input is supposed to trigger a response?	
04	What exactly should change because of that evidence?	
05	Can we observe the change consistently in a realistic test?	
06	What happens when the learner performs well?	
07	What happens when the learner struggles or gives an incorrect response?	
08	Does the learner receive relevant support and another meaningful opportunity to demonstrate	
09	capability? What stays the same regardless of learner evidence?	
10	What instructional decision remains under human control?	
11	What data, privacy, accessibility, security, regulatory or technical review is still required	
12	outside this audit? Decision: demonstrated, partially demonstrated or not yet demonstrated? What evidence	

supports the decision?

Audit date: _____ Product / feature: _____ Claim owner / vendor contact: _____

Internal reviewer: _____

Evidence boundary: The AI Learning Claims Audit is Ravinder Tulsiani's professional synthesis for instructional and capability due diligence. It is not a formal industry definition of adaptive learning, personalization, mastery learning, Socratic teaching or AI tutoring. It is not a validated scoring instrument, vendor certification, procurement assessment, privacy review, security assessment, accessibility audit or legal opinion.