

I have spent the last 11 years in the trenches of L&D, moving from instructional design to LMS administration, and eventually leading QA teams for enterprise-level enablement. Over the last 18 months, I have been stress-testing AI tools in our workflow. I've seen the magic—the productivity gains are undeniable—but I have also seen the insidious, subtle errors that AI inserts into otherwise polished content. If you aren't running rigorous **post-launch monitoring**, you aren't just missing bugs; you are actively eroding your team's credibility.

When we use AI to generate scripts, assessment questions, or summaries, we are essentially acting as the "human in the loop." Yet, many teams treat the launch as the finish line. In reality, with AI, the launch is just the beginning of the verification process. Here is how I manage quality, mitigate risk, and keep my "gotchas" documentation updated to ensure our learners aren't being fed hallucinations.

Redefining Validation in the AI Era

In the "Before Times," validation was about checking for brand consistency, tone, and whether the learning objective matched the assessment. Today, validation means something else entirely. It means checking for hallucinations, bias, and context-collapse.

AI is a probabilistic engine, not a knowledge repository. It predicts the next word, not the next truth. When you generate learning content, you must validate that the output is not only factually correct but also that it aligns with your internal organizational truths. My internal "gotchas" document is filled with instances where AI hallucinated internal product features or quoted nonexistent policies because it was trained on outdated or generalized public data.

The Risk-Based QA Framework

One of the most annoying things I hear in L&D is the "blanket review." You cannot treat a 30-second micro-learning tip the same way you treat a high-stakes compliance or safety module. I use a risk-based approach to determine the depth of the **post-launch monitoring**.

Content Type	Risk Level	Validation Strategy
Compliance/Safety	High	Full SME validation, source citation audit, assessment stress-testing.
Process/Technical	Medium	Peer review, logic-check, screenshot verification.
Soft Skills/General	Low	Grammar check, brand alignment, rapid learner feedback.

For high-stakes content, I don't just review it; I try to break it. I deliberately answer assessment questions in ways that mimic common misconceptions to see if the AI-generated feedback is genuinely helpful or just a generic "that's incorrect" loop. If the feedback doesn't explain **why** an answer is wrong based on our specific internal policy, the content isn't ready.

Fact-Checking and Source Tracking: The "Trust but Verify" Workflow

The most dangerous thing an AI can do is write a sentence that sounds 95% correct but is 5% wrong. I demand that all AI-generated assertions be backed by a source. If you're generating content for a product launch, every claim about a feature's capability must be traced back to a source document.

In my workflow, I force the AI to provide citations during the drafting phase. Post-launch, I treat those citations as part of the content audit. If I find an AI-generated script claiming a feature handles "data encryption at rest" when it only does "in transit," I don't just fix the line. I flag it in our error report as a logic failure. We need to know **why**

the AI drifted so we can adjust our system prompts or RAG (Retrieval-Augmented Generation) parameters for future iterations.

Targeted SME Review: Stop Wasting Their Time

Subject Matter Experts (SMEs) are busy. If you send them a 20-page document generated by AI and ask them to "look it over," you will get a vague "looks good to me" response. This is a failure of [SME sign off template](#) the L&D practitioner, not the SME. That kind of feedback is useless.

Instead, use a targeted review process. When I engage an SME, I provide a specific checklist based on the AI-generated output:



- "Are these three product benefits accurate based on the latest release notes?"
- "Does this scenario mirror the actual challenges our reps face in the field?"
- "Are there any nuances here that the AI missed which are specific to our region/market?"

By providing structured questions, you force the SME to engage critically with the content. If they identify a mistake, document it immediately. This contributes to your **continuous improvement** loop.

Establishing a Robust Learner Feedback Loop

Once the content is live, your users become your QA team. I implement a lightweight **learner feedback loop** directly within the LMS. Don't just ask "Did you like this?"—that's a vanity metric. Ask specific, trigger-based questions:

1. "Did you find this explanation confusing?"
2. "Is there a piece of information you expected to see but didn't?"
3. "Did any of the assessment questions seem ambiguous or misleading?"

When you get that feedback, don't ignore the sentiment. I categorize negative feedback by "Content Drift" (the AI hallucinated or strayed from the source) vs. "Instructional Design" (the AI wrote something that wasn't pedagogically sound). Analyzing this data is essential for **continuous improvement**. If you ignore the feedback loop, you are essentially flying blind.

The Error Reporting Process: Making it Actionable

I maintain a living document—my "gotchas" doc—that tracks every error we find. This isn't just about accountability; it's about refining our AI prompts. Every time we catch a hallucination, I trace it back to the prompt that generated it.

The **error reporting process** should be seamless. If a learner flags a piece of content, it should create a ticket in your project management system (Jira, Trello, etc.). This ticket should include:

- The exact location of the error.
- The nature of the error (factual, tone, technical).
- The source material that contradicts the error.
- A "remedy" plan (rewrite, re-prompt, or remove).

By treating AI-assisted content with the same scrutiny as human-written content, we eliminate the "AI bias" where we subconsciously give the machine a pass because it worked so fast. It didn't work fast; it just generated a draft fast. The *work* is in the vetting.

Final Thoughts: The Mindset of Continuous Improvement

Instructional design is an iterative process. Adding AI to the mix just makes the cycle move faster. My advice? Stop looking for the "perfect" AI tool and start building a perfect vetting system.

Rewrite every sentence until it is unambiguous. Test every assessment as if you are a learner trying to break the system. And for the love of all that is professional, never accept "looks good to me" as a QA result. In the AI era, our value as L&D professionals isn't in our ability to generate content—it's in our ability to distinguish between a hallucination and a truth. Stay skeptical, stay rigorous, and keep your documentation tight.

If you're interested in refining your own AI-assisted workflows, I highly recommend starting your own "gotchas" doc today. You'll be surprised at how quickly patterns emerge—and how much more confident you'll feel once you know exactly where the risks are hiding.

