Fluency Is Not Accuracy

AI Essentials for the Busy Clinician · Brian Whittaker · RRM Academy

Orientation

The most important seven minutes of the course: the one failure to understand, and the fast habit that makes the tool safe.

Lesson

The failure, in two lines

The AI predicts text; it does not look things up. When the true answer is not solidly in its patterns, it produces something that sounds right instead, with identical confidence. That is **hallucination**. Newer models do it less. None do it never, and nothing in the writing warns you.

> **Term: Hallucination.** Confident, fluent, false. Normal behavior at the edge of its knowledge, with no tell in the prose.

The one hazard to know by name

Ask for a reference and the AI will produce one: authors, title, journal, DOI, all perfectly formatted, and possibly nonexistent. It builds what a citation looks like; whether the paper exists was never part of the process. Lawyers have been sanctioned in court for filing AI-invented cases. Do not be the clinical version.

The rule: an AI citation is not real until you have found the paper yourself. PubMed, thirty seconds, every time.

> **Term: Fabricated citation.** A perfect-looking reference that may not exist. Retrieve it yourself before it goes anywhere.

The checking habit, sized to stakes

Verification here does not mean re-researching everything:

- **Rewrites of your own material, tone, formatting:** you reading it is the check.

- **Facts going to a patient or into a decision:** check the numbers and claims against a source you trust.
- **Citations:** always. No exceptions for impeccable-looking ones; impeccable-looking is what fabrication does.

Two helpers: ask the AI "what are you least sure about here?" (often genuinely useful), and remember that grounded work (Module 2) needs far less checking. That is another reason to give it the document.

> **The verification habit.** Check output against a trusted source before use, scaled to stakes. Citations always, no exceptions.

The frame

Treat it like a bright, tireless helper who is sometimes confidently wrong and never knows when. You would not let that helper send things to patients unreviewed. You would also never fire them. Review is how the value gets captured safely.

Scenario: try it yourself

Setup. A patient hands you a chatbot printout recommending a supplement, citing "a 2023 randomized trial in Fertility and Sterility." You search: no such trial. The supplement has some real, modest evidence.

Model answer. "I am glad you brought this in before starting anything, that was exactly right. I looked for the study it cites and it does not appear to exist; these tools sometimes invent convincing references, so I check every one. The supplement does have some real, more modest evidence, so let us look at what is actually known and decide together." Affirm, state the finding without ridicule, pivot to the real evidence.

Terms from this module

Term	What it means, and why you care
Hallucination	Confident, fluent, false; no tell in the prose. *Your review captures the value safely.*
Fabricated citation	Perfect-looking, possibly nonexistent. *Retrieve every one yourself. Thirty seconds.*

Term	What it means, and why you care
The verification habit	Check against trusted sources, scaled to stakes. *Own-material rewrites: light. Patient-facing facts: always. Citations: no exceptions.*