

dear residents

Think. Check. Rethink.

August 16, 2026

Dear Residents,

I've been getting braver with AI asking it both mundane and profound questions. On my recent trip to Karachi, I asked ChatGPT which seats I should select on my flights. It obliged. It even helped me figure out where I needed to go when I transferred through Doha. I was worried about the trouble in the Persian Gulf, but it reassured me that airlines have become pretty savvy about rerouting flights. Sure enough, my flight from Dallas to Doha flew south before turning north, avoiding the Persian Gulf entirely. I nearly missed my connection, but fortunately the gate for the next flight was close by. These are wonderful uses of AI. I needed answers, and it gave me answers.

But I also enjoy thinking for myself. And increasingly I wonder whether we are sufficiently protective of the time it takes to do that. Our brains are not particularly well suited to the pace of the modern world. Thinking takes time. It is inefficient. It proceeds in fragments. One thought opens the door to another. We sit with incomplete information, weigh possibilities, change our minds, walk away from a problem, and then return to it.

I play the New York Times Spelling Bee most days. I could solve it much faster with ChatGPT but then why bother playing the game at all? In fact, I love getting stuck. That sounds like self-punishment, but it isn't. When I get stuck, I close the game and walk away. Yet somewhere in my mind I can still see those seven letters. Sometimes, while doing something completely unrelated, a word suddenly appears. I return to the puzzle, and that word may open an entire family of other words. Sometimes simply scrambling the letters allows me to see something that had been sitting in front of me all along. I enjoy this kind of structured procrastination. It rests the conscious thinking mind while allowing me to approach the same problem differently. And finding something that I had previously missed is deeply satisfying.

I increasingly approach AI in clinical medicine in much the same way. First, I think through the problem myself. I construct an illness script. I decide which findings matter. I look for patterns. I try to reconcile contradictory information. I pause, I ask myself what fits and what doesn't. Sometimes I notice something subtle. Sometimes I miss something obvious. Sometimes I overweight a laboratory value or a historical detail.

Only after I have formulated my own thoughts do I ask a tool such as OpenEvidence what it thinks. Most of the time, I stay close to my original formulation. Sometimes it reminds me of something I overlooked. Occasionally it introduces an angle I had not considered. I appreciate being corrected. I also click on the references and check them for myself. And then comes the most important step: What do I think now?

That third step matters because there is a difference between thinking and answering. Sometimes all I need is an answer. What is the recommended dose? What does the current guideline say? Is this a recognized adverse effect? What does this treatment cost? AI is remarkably good at retrieving and organizing this kind of information.

But a patient is not a retrieval question. Clinical reasoning requires us to sit with fragments and contradictions. We make judgments. We assign different weights to different pieces of information. We recognize patterns and occasionally realize that the pattern we initially recognized was the wrong one. Even laboratory data reward a second look. We can become blind to numbers we have already seen. A low chloride may be masking an anion gap. Checking again after a liter of normal saline may make something visible that wasn't apparent before. Dropping the entire laboratory panel into an AI tool may get us to the answer faster, but speed is not always the purpose of the exercise.

There are also things we absorb at the bedside that are difficult to put into a prompt: the look of the patient, the belongings on the bedside table, the posture of a loved one in the room, the cadence of an answer, the feel of the situation. Much of clinical expertise resides in observations and intuitions that we may not yet be able to articulate.

Our new tools are extraordinarily powerful. They work at remarkable speed. They can be impressively thorough and accurate. They suggest questions we haven't thought to ask. They retrieve information faster than any of us could. They are particularly useful when we ask, "What am I missing?" or "What could I be overlooking?" But they can also tempt us to skip the uncomfortable part.

The danger of an AI-first strategy is not simply that AI might occasionally be wrong. The greater danger for learners may be that it gives us speed over depth, reaction over reflection, and certainty over curiosity. We can become increasingly skilled at constructing prompts while becoming less practiced at constructing illness scripts. And clinical reasoning is a skill that develops through use.

That is why I have settled, at least for now, on three steps: Think. Check. Rethink. Think first. Form your own representation of the problem. Then use AI aggressively. Challenge yourself. Ask what you missed. Look for alternative explanations. Retrieve the evidence. Check the references. And then close the loop by asking yourself: What do I think now?

This sequence deliberately slows us down. It restores a little helpful friction to thinking. It allows AI to extend our cognition without replacing it. Most importantly, it preserves a sense of active participation rather than passive consumption. There is enormous pleasure in thinking and in getting stuck, noticing something you missed, changing your mind, and eventually finding your way through a difficult problem. We should not surrender that pleasure too quickly.

The tools will continue to get better. The answers will come faster. Our responsibility is to make sure that, even in an age of extraordinary assistance, we do not lose the habit or the inner authority, of forming a thought on our own.

Best regards

Dino Kazi