

Best AI Workflow for Doctors Outside the US: Using Perplexity, Claude and Gemini Together for CPD, Exam Prep and Journal Clubs

TechRounder PDF Edition

Live article:

<https://www.techrounder.com/ai/best-ai-workflow-for-doctors-outside-the-us-using-perplexity-claude-and-gemini-together-for-cpd-exam-prep-and-journal-clubs/>

By Vipin PG | Published July 31, 2026 | Updated July 31, 2026 | Format: Explainer | 8 min read

In brief

If you practice medicine outside the US or UK, "which AI tool is best" is the wrong question. OpenEvidence needs a US National Provider Identifier and has been pulled back from the UK and EU. iatroX only covers UK guidelines. Claude for Healthcare's health-record features are US-only. That leaves general-purpose AI - Perplexity, Claude, and Gemini - as the realistic option for most of the world's doctors. The right move isn't picking a favourite. It's giving each tool the one job it's actually good at, and matching how carefully you check its answer to where that answer is going to end up: a quick curiosity check, a study note you'll memorise, or a slide your colleagues will see.

Search for "best AI tool for doctors" and you'll get one of two articles. Either it's a generic ChatGPT-vs-Claude-vs-Gemini comparison with no medical context at all, or it's a review of specialised clinical platforms, OpenEvidence, iatroX, UpToDate Expert AI, built to answer bedside questions safely. Both are useful in their own way. Neither one is written for the doctor who isn't standing at a bedside right now, and isn't in the US or UK either. That's a much bigger group than the existing articles assume. A large share of the world's physicians spend a good part of their week not treating patients, but preparing to know more: reading up for a specialist exam, building a journal club presentation, writing structured notes they'll revise before an oral exam, or just trying to keep up with a paper that changed practice last month. That work happens with general AI chatbots, whether anyone writes a guide for it or not.

Why the tools built specifically for doctors don't reach most doctors

It's worth being specific about this, because the gap isn't a vague "access issue" , it's a set of hard walls. OpenEvidence, the platform many US doctors now use daily for cited, guideline-grounded answers, is free only for verified US healthcare professionals who supply an NPI number. It has also been reported as withdrawn from the UK and the EU, with OpenEvidence's own privacy policy noting that its infrastructure and hosting sit in the US and warning that it isn't built to meet EU data-protection standards. It's an excellent tool , if you're a doctor practising in the United States. iatroX runs the same logic in reverse. It's genuinely free for anyone, no professional verification required, but it's architecturally built around UK guidance , NICE, CKS, the BNF. Ask it a question rooted in a different country's treatment protocols and you're outside what it was designed to do. Then there's Claude for Healthcare, Anthropic's healthcare-focused toolkit announced in January 2026. Its most talked-about feature , syncing lab results and health records so Claude can summarise your history in plain language , is currently available to Claude Pro and Max subscribers in the US, connecting to US-specific data sources like Apple Health and the CMS coverage database. Put those three together and a pattern shows up: the "safe," purpose-built clinical AI tools are almost all fenced to a single country. If you're practising in India, the Gulf, Southeast Asia, most of Africa, or large parts of Europe, none of them are really built with you in mind. You're left with the same general-purpose tools everyone else uses for everything else , Perplexity, Claude, and Gemini.

Doctors are already using general AI , heavily

This isn't a fringe workaround. The American Medical Association's most recent physician sentiment survey found that 81% of physicians now use AI in some part of their professional life, up from 66% in 2024 and just 38% in 2023 , one of the fastest adoption curves the AMA says it has tracked in health tech. Research summarisation is already one of the most common uses. In other words, doctors have already answered the "should I use AI" question for themselves. What's missing is guidance on doing it well when you don't have a specialised, guideline-locked tool doing the safety work for you.

Stop comparing the three tools. Start assigning them jobs.

Most comparison articles frame Perplexity, Claude, and Gemini as competitors , pick a winner, cancel the rest. For a doctor's actual weekly workload, that framing doesn't hold up. The three tools are strong at genuinely different things, and the smartest workflow uses all three in sequence rather than one in isolation. Perplexity's strength is grounded retrieval. It searches the live web and attaches citations to what it tells you, which makes it the right starting point when you need to know what the current literature or guideline actually says on a topic , the "quick literature lookup" part of the job. Treat its citations as a starting point to verify, not a finished answer. Claude's strength is structuring and reasoning over long material. Once you have the raw information , a set of papers, your rough notes, a transcript of a lecture , Claude is the stronger tool for turning that into something usable: a condensed study note organised around what an exam is likely to test, a presentation outline with a logical flow from background to critique, or a set of self-quiz questions pulled from material you've fed it. This is document-structuring work, and it's a different skill from web retrieval. Gemini's strength is sheer volume. If you're trying to work through an entire textbook chapter, a bundle of trial PDFs, or a long recorded lecture in one pass, Gemini's larger context window lets you feed in more material at once without breaking it into chunks first. None of the three is "the winner." Each one removes a different bottleneck.

The real risk isn't hallucination in general , it's where the wrong fact lands

Every article on AI in medicine repeats the same warning: verify everything, hallucination is a dealbreaker. That's true, but it treats every use of AI as equally risky, which isn't how the risk actually works. One evaluation of general-purpose chatbots found they fabricated medical citations roughly a third of the time when asked directly for references, a real and well-documented problem. But a fabricated reference matters differently depending on what happens to it next. It helps to think in four tiers, rather than one blanket rule:

- Casual curiosity. You're skimming a topic out of interest, not building anything from it. Low stakes, read it, move on, verify later if it matters.
- Study notes you'll memorise for an exam. This is a quieter risk than most guides mention. A wrong fact here doesn't get caught at the moment of use, it gets absorbed, and it can sit in your memory uncorrected for years. This deserves more scrutiny than a casual query, even though it feels like "just studying."
- Journal club or conference slides. Public and peer-reviewed in real time, if a number is wrong, a colleague in the room is likely to say so. That built-in correction mechanism makes this lower-risk than it feels, but you should still check the actual trial numbers against the source paper, not just the AI's summary of them.
- Anything touching an actual clinical decision. This is the one tier where the existing "always verify against a primary source" advice is non-negotiable and where general-purpose AI should never be the final word, cross-check against a guideline, drug reference, or senior colleague, every time, regardless of which tool you used.

A practical, step-by-step workflow

Here's what this looks like put together, using a realistic example: preparing a journal club presentation on a recent hypertension trial while also turning the same material into exam-recall notes.

Step 1: Retrieve with Perplexity

Start with a specific, narrow question, "what were the primary and secondary endpoints of [trial name], and how do the results compare with the current guideline recommendation." Let Perplexity pull the trial and any guideline commentary, and open the actual cited sources rather than trusting the summary alone. This step exists to get you to the real paper faster, not to replace reading it.

Step 2: Structure with Claude

Once you've confirmed the core facts against the source, feed the trial details, your own notes, and any related guideline text into Claude and ask it to do two separate things: build a presentation outline (background, methods, results, critique, relevance to practice), and separately, condense the same material into short recall-style notes organised by what's likely to be tested or discussed. Keep these as two distinct outputs, a slide outline and a study note are structured differently, and asking for both in one pass tends to produce a worse version of each.

Step 3: Generate a self-quiz

Ask the same tool to write ten to fifteen short-answer questions based only on the notes you just approved, not on the internet, not on anything beyond what you've verified. This keeps the quiz anchored to material you've already checked, rather than introducing new, unverified claims through the back door.

Step 4: Verify, matched to the tier

Apply the risk framework from above. For the study notes, spot-check the specific numbers , trial size, endpoint results, confidence intervals , against the original paper, since these are the details you'll carry into an exam. For the presentation slides, do the same check on anything you'll say out loud, since that's what a colleague could challenge on the spot.

Making this count toward your CPD, not just your to-do list

For doctors practising in India, this workflow has a practical bonus: it can double as documented professional development. Under the National Medical Commission's Registered Medical Practitioner regulations, doctors under 65 need to complete 30 CPD credit hours every five years to renew their licence, and journal club participation and presentations are recognised activities under that structure. A journal club talk you built using this workflow isn't just useful preparation , with the right documentation from your institution or state medical council, it's also a credit-bearing activity toward that requirement. The same logic applies under most countries' equivalent CME/CPD systems, even where the exact credit rules differ.

A short verification checklist

- Never treat an AI-generated citation as confirmed until you've opened the actual source.
- For anything you'll memorise, re-check the specific numbers, not just the general claim.
- For anything you'll present publicly, re-check whatever you plan to say out loud, since that's what gets challenged.
- For anything touching a real patient decision, AI is a starting point for research, never the final source , confirm against a guideline or clinical reference.
- Keep retrieval (Perplexity) and structuring (Claude) as separate steps rather than asking one tool to do both , it makes errors easier to catch, because you're checking a smaller thing at each stage.

Frequently asked questions

Can I use Perplexity, Claude, and Gemini for actual patient care decisions? Use them for background research and to move faster, but the final decision should always be checked against a proper clinical reference, guideline, or colleague , not left with a general-purpose chatbot, none of which are approved medical devices. Do I need to subscribe to all three tools? No. Free tiers of Perplexity, Claude, and Gemini cover this entire workflow for most doctors. A paid plan is worth it only if you're hitting daily usage limits. Is Claude for Healthcare coming to other countries? Anthropic hasn't given a public timeline for expanding Claude for Healthcare's record-syncing features beyond the US. Until it does, the general-purpose version of Claude remains what's available to most doctors outside the US.