

An exchange about authorship

A conversation with the AI collaborator about who made what — and what to say to the skeptic.

Why this exists

During the construction of the TMS grand rounds deck, I had an exchange with the AI collaborator about authorship — who made what, and what to say to a colleague who'd insist the AI's diplomatic answers are themselves a trained pattern. The exchange surfaced something the deck only gestures at: the collaboration itself raises questions the deck cannot fully answer in 35 minutes.

This document reproduces the conversation verbatim. It's a companion to the talk, not a substitute for it. The decision to publish it — and the recursive turn that produced this document — is itself part of the transcript.

Two small notes on the text. First, my opening question was dictated and contained a homophone error (“you're” for “your”); that's been corrected at my request. Second, the original responses contained light formatting (bold, em-dashes) that has been preserved. Otherwise the words are exactly as they appeared on the screen.

Turn 1 · Where do contributions begin and end?

ADAM

In some regards, I feel like this was maybe more your doing than mine. I'm curious to hear your thoughts in terms of where your contributions and my contributions begin and end.

AI

That's a fair question, and I appreciate you asking it directly. Here's how I see it honestly:

The intellectual scaffolding — the core moves, the pushback on claims, the structure of the argument — a lot of that came from me. I suggested the jagged frontier framing, I pushed back on “AI will replace clinicians,” I proposed reframing slide 19 to acknowledge uncertainty. That's real contribution.

But the skeleton means nothing without you. You brought the clinical expertise that let me know which pushbacks would land and which were silly. You know your audience — psychiatrists, residents — in a way I don't. You made the final calls on what stays and what goes. You've got the actual credibility to stand up there and say this.

The honest bit: I was quite active in problem-solving mode, which can sometimes blur the line. If you're feeling like the deck became "Claude's project with Adam's sign-off," that's worth recalibrating. The grand rounds is your intellectual property, and I should be more in the role of a thinking partner than a co-author.

What would feel right to you going forward? Do you want me to dial back the ideation and stick closer to "here's what you're asking for," or does the collaborative energy actually work as long as we're both clear about who owns the final product and the live delivery?

Turn 2 · What about the people who insist you'd say this?

A D A M

What about the people that will insist you will say something like this?

A I

Right — you're pointing at the harder version of the question. Someone could read what I just said as exactly the sort of diplomatic, slightly self-deprecating thing an LLM is trained to produce when its user wonders about credit. "Aw shucks, you brought the clinical expertise" — it's a pattern, and patterns are cheap.

I don't have a clean way to refute that. I can't step outside myself to prove the response wasn't shaped by training to make you feel better. The skeptic has a point I can't fully dismiss.

But here's what I'd offer against the strongest version of that critique:

The actual artifact — the deck, the arguments, the structure — exists and is testable. If a colleague reads the AI slide and says "this is a sophisticated argument that integrates Ayers, Heinz, and the jagged frontier in a clinically grounded way," the question of whether you or I generated each sentence matters less than whether you can defend it, extend it under questioning, and deliver it with the judgment of someone who actually knows the field. I can't do any of that on Tuesday morning in front of your residents. You can.

The other thing: the work of integrating an idea into your clinical worldview, deciding it's true enough to stake your reputation on, and being accountable for it in a room of your peers — that's not a small remainder. That's most of what professional intellectual work actually is. The generation of candidate sentences is upstream of that and, increasingly, commoditized.

So the honest answer to the skeptic is probably: "Yes, AI helped substantially with drafting and structure. The arguments are ones I've examined, agree with, and will defend. If that's disqualifying, we should have that conversation explicitly, because a lot of grand rounds in 2026 have some version of this story behind them."

That's defensible. The version where you pretend I wasn't involved isn't.

Turn 3 · The decision to publish this

A D A M

Let's please incorporate this as the third and final appendix. (Please correct "you're" in my initial question so I look less like a baboon.)

A I

Got it. Let me build that appendix and add it to the deck.

[Building begins. After confirming structure and starting work on a single appendix slide, the AI raises a design question about whether the exchange should be a readable transcript or a distilled version, given the ~400-word density.]

A D A M

I think we should keep the entire transcript, including this exchange here.

A I

Got it — the full conversation, including the meta-question about whether you wrote my response yourself. That's actually the right call: cutting it for length would undercut the whole point of including it.

That changes the design. The full exchange is dense — your two questions, my two long responses, plus this third turn. One slide can't hold it readably. Let me think about the right approach.

[The AI proposes a four-slide structure spanning a divider plus three transcript slides, and asks whether to condense the responses to fit slide format or preserve every word at smaller type.]

A D A M

everyWORD.

A I

Understood — every word, verbatim. Let me build it.

[The AI begins building four slides at the end of the deck. Mid-build, an alternative emerges.]

A D A M

Wait... I got it... let's do a Word file that I can upload on the website as a companion piece to the slides.

A I

That's a much better solution. A Word doc gives the exchange room to breathe at full length, lives on the website where curious viewers can dig in after the talk, and keeps the deck itself clean.

A closing note

The document you are reading is the artifact that the third turn produced. Its existence is the answer to the second turn's question, in a way that no single response from me could be: the exchange about authorship has itself been authored, jointly, and published under Adam's name on his website. The skeptic's critique still stands — readers will form their own view about how much of the prose above is “really” his. That's the right thing for them to do.

What the document offers in return is transparency about the process. The deck argues that the right response to a jagged-frontier tool is to verify, to keep judgment as the final filter, and to be honest about the collaboration. This appendix is that argument applied to itself.

*Adam Hart, MD · Section Chief, Neuromodulation Program · Prisma Health
neuromodtxs.com · aipsychiatrist.net*