

Notes on Writing with AI

By Matthew Pittinsky, Ph.D. · **DRAFT** · Last Revised May 8, 2026

Annette Lareau's *Home Advantage* is one of my favorite books from my doctoral studies. Partly for the topic -- social class and parental involvement in schools. But also, for the methodological appendix at the back, where she pulls back the curtain on how the ethnography was written: her false starts, the families who said no, the way her own background as a researcher was present at interviews. As a doctoral student looking ahead to a dissertation, I found that transparency and reflection demystifying and comforting.

In that spirit, I thought I would put together some notes and reflections on how I am using AI in my own writing. I recently posted several working drafts developed in active “collaboration” with Anthropic’s Claude -- *Three Books Every EdTech Entrepreneur Should Read*, *A Note on AI Language and Adoption in Education*, *The LMS as Platform*, and *Assessment: The Great AI Unlock*. I did not use AI on the *LMS at 30* essays I wrote a few months earlier. I figured if I am going to write about AI I should involve AI in the process. I doubt anything I do is particularly novel. In fact, part of the point is that it is already becoming ordinary. But ordinary does not mean settled. In the spirit of transparency here are the steps I took with some commentary.

Step 1. Dictate the Argument

I tend to write like I speak. Or, more accurately, I speak my writing in my head when I am typing. That makes the first step natural. I open SuperWhisper and talk through the piece -- argument, anchor, examples,

what to avoid, etc. For example, here is the opening of one recent voice memo (verbatim):

"I want to write an article about why assessment is the great unblock for the adoption of AI in education. The basic argument is that if we, you know, I want to start with an acknowledgement of the fears and concerns by students and faculty about AI, concerns about cheating, concerns about cognitive offloading. These concerns are one of several reasons why assessment is the great unlock. It's the great unlock because, first, it is the source of the greatest amount of fear about AI's impact in education. If you go back to the early or mid-1990s and swap internet for AI, you didn't see any of this fear. There was generally excitement about where this can all go. But now you see this fear with AI..."

I don't start until I can articulate the whole frame via a recording in one pass, albeit with pauses and starts, not some perfect dictation. Before starting an article I've usually thought about it for weeks or months or longer, although the article on AI and language came together quickly when the GSV panel discussion connected the dots between language, Postman's writings, and AI adoption. In the same memo I also dictate research-assistant type directions: pull the Gallup/WFF/GSV survey released at the conference, find a mid-1990s analog to that study about the Internet to compare, check whether Narayanan and Kapoor said *ordinary* or *normal* technology.

SuperWhisper's dictation is impressive, even when I pause for a bit or hem and haw and correct myself. Claude seems to be robust to any mistranscriptions. My experience so far is that Claude is also very good at disentangling verbatim text that is the substance of the essay from text that is direction to do research that contributes to the essay.

Step 2. Comment on the Draft

The next step is to review the draft generated from my dictation and, before doing a full direct markup, revise it through comments. The first draft is usually OK, but not great and quite different from the final. I direct revisions through both chat conversation and inline Word comments. I find chat is the better mode for directional and structural moves. Word comments are reserved for precise, paragraph-anchored ones. I think it's neat that Claude can read Microsoft Word comments as inline prompts, so to speak. Comments might be one-line objections ("too close to Postman, rewrite"), structural moves ("split this into two sections"), criticism ("you are being too wordy – I wouldn't say that"), or even full replacement text that I draft directly. Claude reads the comments and generates a full revision. I think of AI in this stage as a graduate student acting as research assistant, but as first author I am directly invested in the outcome. I usually end up doing 2-3 of these comment reviews. I stop when I feel like the structure and basic form are set as I don't want to start doing heavy revisions when there is a risk the next iteration will unravel or overwrite changes. I am explicit with Claude to not change anything except what I comment on, but that isn't foolproof (I've learned the hard way).

AI's research-assistant role, which occurs in steps 1 and 2, is an area where I feel some discomfort. In the assessment essay, three citations relate to formative assessment -- a field I do not know well. I presented Claude with my lay understanding of the literature, an understanding that I comes from osmosis being in the broader field. I asked Claude to challenge that understanding. Once I was comfortable that -- while no doubt contested and caveated, as everything in education research is to some extent -- they were correct claims to make, I then asked Claude to

surface seminal publications, summarize them, and cite. I did some cognitive offloading there, and the result is a small case of imposter syndrome, as more of my sensemaking occurred in conversation with Claude than with the works cited themselves. (There is an interesting sociology of knowledge implication in terms of the compounding advantages highly cited works benefit from and what that means for alternative ideas to break through.) I make peace with it by validating the claims myself — but less thoroughly, candidly, than if I had sourced them in the first place — and through the rigor of Step 4 below. In this case, I also reference AI's role in a direct, albeit tongue-in-cheek way, when citing the assessment works.

Step 3. Search My Files

I am lucky to be a digital pack rat: MS Office, iWork, and email files going back to 1997. AI makes that archive searchable in real time. *Pull the 1996 Blackboard business plan and confirm the wording on instructional management systems. Find the Building Blocks launch files from October 2001 and list the first cohort of partners. Find my course syllabus for ASU Sociology of Education and confirm how I phrased the meta learning objectives.* These are all things Claude does for me with little more than a directory and my vague recollections of dates and documents. The historical nuggets that I use to anchor a piece come from nearly 30 years of accumulated documents! My original dictation and notes provide the scaffolding of what to bring in and why. This has been the most magical part of my experience writing with Claude. I've received feedback that readers like when I share bits of history. Claude makes it possible, grounded in my primary source materials.

Step 4. Several Full Manuscript Edits, by Hand

Now the process turns old school. When it feels like the draft is stable, I do a full top-to-bottom revision. I take multiple passes -- in Word on screen and printed and marked up with a pen, which I then re-enter. (I would try scanning and giving the markup to Claude, but my handwriting is the great shame of my retired teacher mother.) I find that these passes are the marginal steps that make the article fully mine. In practice, I probably do at least five to seven revisions, spaced to enable fresh eyes.

Step 5. Cross-Check With Other Models

Before posting or publishing, I take the final draft manuscript and run it through Gemini and ChatGPT, as well as a new Claude session that lacks memory. I ask the models to fact-check claims and citations and suggest manuscript edits. I ask them to evaluate the strength of the argument and offer strawman counterarguments. Yes, I also score the extent to which the article presents as AI-written and why. Different models catch different things. I ask each of them to develop a prompt for Claude to make a final revision, from typos to counterarguments that need to be addressed. I give those prompts to Claude and ask Claude to engage in a discussion with me, presenting each suggestion one-by-one so I can decide yea or nay and be the one who drafts any new language being added, not Claude.

Step 6. Update the Skill File

I keep a Claude skill file that provides instructions about how I write, from my writing style to formatting preferences. I am careful about two things. First, the writing style components must stay *pre-AI*: the corpus used to shape the skill is my published writing from before any of this -- EdSurge, EDUCAUSE Review, Inside Higher Ed, Hechinger, plus the *LMS at 30* essays. That guards against model creep, where AI patterns I don't recognize as "foreign" in the articles slowly shift my writing to something that feels like me but is not. Second, the skill *does* learn from *my revisions* during AI-assisted writing -- but only the explicit ones where I rewrite sections so again, it is my word choice.

The skill file is my second area of discomfort. As a thought experiment: if the skill keeps getting better at writing to my standard by learning from my markups, the logical endpoint is that I no longer engage in the manuscript edits of Step 4. In that case, did AI just write it for me? Does the cumulative skill -- built from my own pushback over time -- still make the output mine, or does it become closer to ghostwriting in my voice? I am not sure. While the models are improving rapidly, we don't seem close to the thought experiment coming true in practice.

Up Next

One area I want to develop is linkages across my articles. Right now, they don't speak to each other as much as they could. (I've recently revised a few to introduce more linkages.) I admire how Ben Thompson at [Stratechery](#) threads previous writing into new pieces -- citing himself, evolving an argument, building a foundation a reader can follow back. I want to get better at this for two reasons. First, substantively, I think I am

working toward a broader thesis that can only emerge through linkages. Second, selfishly, I'd like a reader who lands on one piece to be one click away from the others given I have near zero distribution for my writing! AI is well-suited to identifying and recommending those linkages. A simple update to my Claude skill will encourage Claude to make connections visible by default. I am on guard to the risk that it will do so a bit overeagerly.

I am awaiting improved AI interfaces as well. On my Mac I use Claude in Word and Claude's desktop app with Cowork. Drafting and revising with Claude still produces a lot of artifacts -- multiple file versions, markdown / Word duplicates, side threads to keep track of. I no longer have to do much copy and paste, but to use a technical term, it is janky.

A final "up next" on my mind is how to address a pesky reader behavior in this AI-mediated world. I have not mastered the art of concise writing -- I tend to meander through history before getting to the point -- which makes my pieces prime candidates for getting pasted into a summarizer, which is all that is read. I am considering a few options. I could include my own AI-style summary up top, which seems too tempting for heartier readers who are up for the full monty. I could write something akin to an academic abstract, which avoids the illusion of a full read and would presumably shape the summary an AI tool produces from the full text. I could include a recommended summary prompt as a footnote. I am currently tinkering with a fourth: a brief appendix titled something like "A Note on AI-Assisted Reading," set at the end of the PDF, that names the practice and offers guidance on how I would want the historical anchors versus the argument handled in any summary. An appendix feels honest with the reader, should survive PDF text extraction better than a footnote, and sits in a position where in-document guidance tends to be

picked up reliably. Whether any of this is overthinking it, I am not sure. I could also just write my succinctly!

A Closing Note

Full disclosure, I wrote this reflection for a reason in addition to valuing transparency. I worry that by writing with such an active reliance on Claude I am one mistake away from posting something that smacks of “being written” by AI, which means different things to different people, damaging my credibility and reputation as a result. I hope sharing my process invites feedback on how to strengthen it. (Now is a good time to clarify that the actual process is not as linear as I make it seem, but more recursive and windier.) I also think sharing my process makes my authorship claims more inspectable.

By the time an article is posted, the argument has begun in as original a form as I am capable of articulating and has been hand-edited several times at the end. AI helped retrieve, organize, and draft; I remain accountable for judgment, claims, voice, and final language. I think the ideas and their presentation are mine in all meaningful senses. But I am mindful that the process is new and Claude can be confident, even seductive.