

Outsourcing Our Epiphanies

Thinking and Authorship in the Age of AI

By Jennifer Woodruff Tait

ABSTRACT: Traditionally, being an author has been considered to involve having the capacity to think. This has raised questions about the attribution of authorship to AI. In this essay, I first explain how large language models (LLMs), the technology behind popular chatbots, produce their output. I then survey recent literature which defines what thinking is with particular comparison to the way LLMs operate. This literature argues that that thinking involves opening oneself to new ideas and evaluating those ideas, always keeping alert to the possibility of unexpected discoveries. While LLMs can have new ideas put into their training data, they currently cannot evaluate those ideas correctly without human intervention, and they cannot have epiphanies. I finally argue that both of these concepts are necessary for authorship.

INTRODUCTION

In the fall of 2025, I picked up my fountain pen again. Hear me out.

For many years, I copied out by hand into a notebook thoughtful and well-written sentences from books I was reading. This is the most analog task imaginable, and one which I later discovered has been recommended to budding writers as a way to develop technique (Spencer, 2012, Joyce & Lundberg 2013; Houchin 2024). Apparently copying by hand fixes the educational target—in this case, a well-crafted sentence—more firmly in the brain. (The fountain pen, of course, is optional.)

Recently I felt drawn to begin this task again. I didn't immediately make a connection with current events until I began to notice a theme weaving through a variety of texts that I encountered discussing generative AI. As English professor Christina Bieber Lake frequently puts it to her students at Wheaton, "Writing is hard because thinking is hard" (Lake 2025). In choosing to write down others' well-crafted sentences manually, I was training my brain to think, exercising muscles that I might use later to craft a sentence of my own, to make a logical connection, to create an outline or fill in the details of a plan. There are many ways to train our brain to think, of course—I have a smartphone as well as a fountain pen and I am not suggesting a return to the days of writing out catalog cards, although studies have shown that at least some degree of handwriting in one's life and scholarship improves memory retention (Van der Weel & Van der Meer 2024; Marano 2025). I am more interested here in whether training one's brain to think is a necessary part of writing. We have generally thought this way about authorship; whether we are writing a novel, a grocery list, or a LibGuide, what we write, no matter the format we write it in or the way we write it down, is the product of thinking and we are authors when we do it.

So can generative AI think? Can it author a work? And even if it can, should it? I want to explore recent literature (including some by AI proponents) which argues that AI cannot think and unpacks why; then I want to discuss the related question of whether or not AI can author a work.

Jennifer Woodruff Tait is web editor of the Theology of Work Project.

WHAT IS A LARGE LANGUAGE MODEL?

AI in the broadest sense, of course, is behind technologies many of us already use—from Google Translate to Grammarly and autocorrect, where AI applies its text prediction models to texts that already exist. What chatbots do is take the next step and actually create text based on these models. Large Language Models (LLMs), the technology behind your average chatbot, are extremely fast and extremely complex text prediction models. They are trained to generate patterns by processing extremely vast amounts of material, noting the patterns that reoccur, and outputting those patterns: “An AI model can’t process a sentence as a continuous flow of meaning. It breaks language down into discrete, countable units called tokens, which it can analyze mathematically. Each token is converted into a string of numbers (a vector), turning language into a mathematical problem” (Robbins 2025). The output can then be tweaked and queried by humans to potentially produce better output, a process that is sometimes unintentionally humorous (Shane 2019¹; Stern 2025).

We call this process “machine *learning*” and we say that the thing doing it is “artificial *intelligence*,” but the output is not produced through actual comprehension of the material or reasoning about the material; it is produced through regurgitating more things that look as much as possible like the things the LLM has already digested. “Even the term ‘artificial intelligence’ itself is fraught with complications. It refers to systems that are artificial but not intelligent. We call them ‘intelligent’ because they can do some things that a human can do—not because they think, as the metaphor implies” (Gestwicki 2026; see also Washington 2026).

This text-prediction regurgitation is one of the major reasons chatbots hallucinate data (Kalai et. al. 2025). A citation to a scholarly work, for example, need not be *actually* correct, in terms of corresponding to the existence of such a scholarly work in the real world, as long as it follows the format the LLM has seen for scholarly citations in its training data:

LLMs learn to reason from outcomes rather than premises. Frequency substitutes for verification, narrative coherence substitutes for causality, and portability substitutes for accuracy. Process-level evidence that resists compression is underweighted, not absent, and thus rarely governs inference. The result, as any LLM user knows, is fluent, confident reasoning that is deductively unreliable (Robbins 2026).

The process performed by an LLM is perceived by us as thought for several reasons. First, it is fast. Secondly, it usually produces the sort of bland, featureless and yet authoritatively stated prose we are used to hearing in educational and academic writing. No doubt this is partially because, for AI chatbots people encounter every day such as Gemini and Claude, this kind of prose constitutes a very large portion of their training data (Robbins 2026). It is also tied into the nature of the model itself: “Computing technologies embed the values of homogenization and reductionism, and replacing wholes by their parts” (Gestwicki 2026).

So, is this kind of fast and authoritative regurgitation *thinking*? There are a fair number of people currently saying no. Naturally, some of them are English professors:

AI cannot teach thinking because AI cannot think. This common confusion—that what AI is doing can be called thinking—is causing most of our pain at this crucial juncture in higher education. Clearing up this confusion is more important than it seems. Simply telling students not to use AI without explaining why feels to them like someone told them to do dishes by hand when the dishwasher is right there (Lake 2025).

So why should we metaphorically do the dishes by hand (especially when our administrators are frequently investing in very new shiny dishwashers?) What does it actually mean to think, and furthermore to write about what we think?

HOW WE THINK (AND LLMS PROBABLY DON'T)

For one thing, thinking involves exposing yourself to new ideas, testing them out, and evaluating them—emotionally, morally, and socially as well as intellectually (Lake 2025; Waters 2025). This is why some educators who are interested in preparing students to both cope with and improve the modern world have been pushing back on the use of AI in the classroom. As one college history professor recently noted:

Generative AI offers a shortcut around the hard work of asking questions, contemplating complexity, and expressing oneself—the very practices that leave students open-minded to differing perspectives and grounded in their core convictions, less fearful of failure and more humble about success, perpetually curious about the world and resilient in their commitment to bettering it (Gehrz 2025).

While it's possible to argue that LLMs fit at least part of this definition because they are exposed to new models in their training data and can test out these new models, it still takes human input to evaluate the models and tweak the prompts. The machine is incapable of telling that it's wrong (Shane 2019; Robbins 2025). It's also incapable of feeling anything about what it's produced (Gestwicki 2026)—even though it may *say* that it feels something, or shower praise on itself or on the prompter. (See Cheng et. al. 2025 and Meyer et. al. 2025 for more on the growing sycophancy problem of chatbots). One adjunct professor reported from a workshop where he was being instructed in the use of A.I.:

It was important that I use the word *innovative* in my prompt, the presenter insisted. Omit innovative and you get different, presumably more pedestrian, results. Claude spat out the paper and told me it was proud of its work, which, after all, had a “clear thesis statement.” When I said I couldn't find the statement, Claude replied, “You're right to question this. Looking at the essay more carefully, there isn't a single, explicit thesis statement that clearly states the central argument.” Thanks, genius (Malesic 2025).

For another, thinking involves making discoveries and arriving at epiphanies in the course of your evaluations. This is true whether we are researching and writing a dissertation or just trying to solve the *New York Times* crossword puzzle, as a *Times* puzzle columnist noted: “One of the great tragedies about the increasing use of A.I. to provide answers to life's puzzles, both mathematical and abstract, is that we're opting to outsource our epiphanies. When we defer to machinery for solutions, we're missing out on the statistically proven benefits of discovering it ourselves, in our dreams or in our tubs” (Corbin 2025). Not only that, the machine itself is not having an epiphany. It's just predicting what a human epiphany would look like based on its training data.

Furthermore, the intellectual act of opening yourself to new ideas and arriving at conclusions about those ideas—sometimes surprising conclusions—has a moral dimension connected to the *imago Dei* and to our training as virtuous humans and engaged citizens, as a computer science professor and game designer argues:

Generative AI systems constantly tempt students against virtue. There is no need to persevere through difficulty if AI can instantly supply an answer, no need to reflect when AI has already suggested the next step. The AI is purely utilitarian, sycophantically supporting a user's every whim. Interacting with it carries no risk of judgment, no need for empathy, no compromise for a better future together (Gestwicki 2026).

This moral dimension has traditionally been seen as a necessary component of theological education, as one software developer noted in relationship to AI-written sermons: “Meditating on divine words is what human beings do in their inner being. This technically cannot and morally should not be automated” (Coles, 2024; see Waters 2025). But this moral dimension is not limited to theological halls nor even to the stereotypical tree-lined quads of traditional liberal arts colleges and universities:

Part of a teacher's job—certainly in the humanities, but even in professional fields like business—is to help students break out of their prisons, at least for an hour, so they can see and enhance the beauty of their

own minds. It is to help them learn, together, to defend how they want to live, precisely because they, too, unlike a machine, will one day die. I will sacrifice some length of my days to add depth to another person's experience of the rest of theirs. Many did this for me. The work is slow. Its results often go unseen for years. But it is no gimmick (Malesic 2025).

Metaphorically doing the dishes by hand, in other words, is not just about getting clean dishes as fast as possible. It is about making the dish-doer into a different kind of person. Here Lake's analogy does somewhat break down, because quite a lot of the time we *do* just want clean dishes as fast as possible—doing dishes is much more like the sort of thing AI is good for than the sort of thing it is not (Gehrz 2025). AI has proved very useful in doing calculations, figuring out probabilities, and spotting certain kinds of patterns much more quickly than a human being could do, as even Lake recognizes:

You can, and should, instruct a computer to perform a series of computations or collect some research on a topic just like you can, and should, put dirty dishes in a dishwasher. By all means, call on Claude to do your grunt work if your grunt work involves responding to emails requesting information, summarizing responses to a survey, or determining whether your team should go for it on fourth down. But I hate (and love) to be the one to break it to you: thinking is just not that kind of work (Lake 2025).

CAN AN LLM BE AN AUTHOR?

So, if the premise is accepted that AI cannot think because it cannot open itself to new ideas and evaluate those ideas by more than intellectual benchmarks, while always keeping alert to the possibility of unexpected discoveries, then can AI—specifically in the form of an LLM—be an author? Is it enough to figure out probabilities and reproduce text patterns, and by doing so produce generic and authoritatively stated prose? At least one novelist thinks that this is not the case, and is worth quoting at length on the matter:

As I've thought deeply about the matter, one problem I continue to come back to is the consequences of outsourcing the work of thinking. This work, done in the abstract and often without a direct and immediate economic yield, has been reduced to a frustration and annoyance in the writing process. If I could only get words on the page, then I could sell this novel.

Generative AI solves this problem. It can organize your ideas and plot points for you. It can offer ideas to fill in plot holes and to move you to the next chapter. It can even write a full-length novel in your own voice and style. But it can't be human.

I have a lot of thoughts, but I'll leave it at this: outsourcing the work of thinking outsources the work of being human. Yes, it is hard, but I believe doing what is hard is necessary for creating anything of lasting worth and value (Radcliff 2026).

Historian Gehrz, too, has

held a firm line against using AI to do written work — which is not simply a mechanical task, but a means of discovering one's beliefs and values and finding the voice to articulate them (Gehrz 2025).

Why would we require that an author be human? This is, of course, ultimately a philosophical question, not a practical one. Should our institution knowingly collect writings created by LLMs, we will find ways to catalog these writings and make them accessible. This question takes a step back and asks why we might not want to do that, or why, if we do so, we might want to distinguish human-authored works from AI-created ones.

Perhaps, at least right now, we require human authors because we want a better check on the fact that information is correct and not hallucinated. It is true that AI programmers are working on ways to guide chatbots to not hallucinate by optimizing the rewards the chatbots receive (Kalai et. al. 2025, Robbins 2025). But even if that problem is solved, we still need human authors because

only human authors can *think* in the fullest sense of the word, even if they only think in order to prompt chatbots: “Invisible [when you look at the inner workings of AI] are the social relations that produced the tokens: the training data, the labor that produced it, the decisions embedded in the model, the infrastructure running it” (Robbins 2025).

Authorship has traditionally been about *process* and not only *product*. Authors create books and articles and plays and stories and other kinds of writing, to be sure. They also, in the process of so doing, *think* about what they are writing—the kind of thinking that, if outsourced to something which gives us superficially similar results, we run the risk of forgetting how to do. Authors outline and revise their outlines, try on different turns of phrases, make discoveries and reject errors, and measure all of this against the sum of human knowledge as known in their particular place, time, and culture. Even some of the currently most successful attempts to have chatbots create art still require extensive human collaboration in order to approach a decent result (Robbins 2026).

Authors do this all in collaboration with the past and with an eye on the future, and in so doing ultimately not only turn themselves into better authors but also culturally preserve an inclination towards the kind of thinking which will create more authors: thinking which evaluates, discerns, meditates, and discovers. Perhaps subsequent authors will study with these previous generations, or read their books. Perhaps they will even try to imitate their styles as a means of growing into their own. Though this kind of imitation is within the skillset of an LLM, the use of that imitation to grow greater on the level of culture and epiphany and not simply the level of technique is not:

Poets work inside an historical network of existing poems. A new poem resonates when it activates prior reading in the mind of the reader. Lines and images from older work resonate with each other and with new work. That resonance is the mechanism by which the particular becomes universal. LLMs are trained on existing texts. LLMs have in their training data most of the digitized poetry in existence; their output tends to reflect patterns in that data. A model asked to write a poem about grief will draw upon images and phrasings about grief in earlier poems. . . . But LLMs do not yet, at this writing, have culture. Without culture, I’m not sure there can be great poetry (Robbins 2026).

We ultimately still need authors who can think now so that we will have more authors who think, later. We don’t even know everything they will think about yet. Perhaps authors need to be human because we still need actual epiphanies and only human authors can find them.

REFERENCES

- Cheng, Myra, Cino Lee, Pranav Khadpe, Sunny Yu, Dyllan Han, and Dan Jurafsky. 2025. "Sycophantic AI Decreases Prosocial Intentions and Promotes Dependence." Cornell University arXiv:2510.01395, October 1. <https://doi.org/10.48550/arXiv.2510.01395>.
- Coles, Arlie. 2024. "ChatGPT Goes to Church." *Plough* 40, June 6. <https://www.plough.com/topics/life/technology/chatgpt-goes-to-church>.
- Corbin, Sam. 2025. "Gameplay: Eureka!" *New York Times Gameplay* email newsletter, December 22.
- Gehrz, Chris. 2025. "AI and the Christian University." *The Pietist Schoolman*, November 21. <https://chrisgehrz.substack.com/p/ai-and-the-christian-university>.
- Gestwicki, Paul. 2026. "Artificial, Not Intelligent: How Meeting Educational Goals Requires Embracing our Humanity." *The Raised Hand*, January 28. <https://theraisedhand.substack.com/p/artificial-not-intelligent-how-meeting>.
- Houchin, Jackie. 2024. "Copy Work: What Is It? Why Do It?" *The Writers in Residence*, November 27. <https://thewritersinresidence.com/2024/11/27/copy-work-what-is-it-why-do-it/>.
- Joyce, Michael, and Anita Lundberg. 2013. "Copying to Learn: Mimesis, Plagiarism and 21st Century English Language Education." Paper presented at the 2nd International Higher Education Teaching and Learning Conference, Curtin University, Sarawak Malaysia, 9-10 December 2013. https://www.researchgate.net/profile/Anita-Lundberg-2/publication/280830021_Copying_to_learn_Mimesis_plagiarism_and_21st_century_English_language_education/links/57e1d75d08ae1f0b4d93f609/Copying-to-learn-Mimesis-plagiarism-and-21st-century-English-language-education.pdf.
- Kalai, Adam Tauman, Ofir Nachum, Santosh S. Vempala, and Edwin Zhang. 2025. "Why Language Models Hallucinate." Open AI. September 4. <https://openai.com/index/why-language-models-hallucinate/>.
- Lake, Christina Bieber. 2025. "Writing is Hard Because Thinking is Hard." *The Raised Hand*, November 25. <https://theraisedhand.substack.com/p/writing-is-hard-because-thinking>.
- Malesic, Jonathan. 2025. "ChatGPT is a Gimmick." *The Hedgehog Review*, May 22. <https://hedgehogreview.com/web-features/thr/posts/chatgpt-is-a-gimmick>.
- Marano, Giuseppe, et. al. 2025. "The Neuroscience Behind Writing: Handwriting vs. Typing—Who Wins the Battle?" *Life (Basel)* 15 (3): 345. <https://doi.org/10.3390/life15030345>.
- Meyer, Erie, Stephanie Nguyen, Laura Edelson, and Jonathan Mayer. 2025. "Tech Brief: AI Sycophancy & OpenAI." Georgetown Law, July 30. <https://www.law.georgetown.edu/tech-institute/research-insights/insights/tech-brief-ai-sycophancy-openai-2/>.
- Radcliff, Kaylena. 2026. "Looking Back, Looking Ahead." *Kaylena Radcliff, Author*. January 17. <https://kaylenaradcliff.substack.com/p/looking-back-looking-ahead>.
- Robbins, Hollis. 2025. "Token-Relief." *Anecdotal Value*, October 17. <https://hollisrobbinsanecdotal.substack.com/p/token-relief>.
- Robbins, Hollis. 2026. "LLM poetry and the 'greatness' question." *Anecdotal Value*, January 7. <https://hollisrobbinsanecdotal.substack.com/p/llm-poetry-and-the-greatness-question>.

- Robbins, Hollis. 2026a. "What Escapes Containment is Less Valuable." *Anecdotal Value*, January 29. <https://hollisrobbinsanecdotal.substack.com/p/what-escapes-containment-is-less>.
- Shane, Janelle. 2019. *You Look Like a Thing and I Love You: How Artificial Intelligence Works and Why It's Making the World a Weirder Place*. Voracious.
- Spencer, Jennifer C. 2012. "Self-Made Writer: A Grounded Theory Investigation of Writing Development Without Writing Instruction in a Charlotte Mason Home School." Ed.D. diss, Gardner-Webb University. ProQuest (3541525).
- Stern, Joanna. 2025. "We Let AI Run Our Office Vending Machine. It Lost Hundreds of Dollars." *The Wall Street Journal*, December 18. <https://www.wsj.com/tech/ai/anthropic-claude-ai-vending-machine-agent-b7e84e34>.
- Van der Weel, F. R. (Ruud) and Audrey L. H. Van der Meer. 2024. "Handwriting But Not Typewriting Leads to Widespread Brain Connectivity: A High-density EEG Study with Implications for the Classroom." *Frontiers in Psychology*, January 6. <https://doi.org/10.3389/fpsyg.2023.1219945>.
- Washington, Eve. 2026. "What Do A.I. Chatbots Discuss Among Themselves? We Sent One to Find Out." *The New York Times*, February 18. <https://www.nytimes.com/2026/02/18/upshot/moltbook-artificial-intelligence-ai.html>.
- Waters, Brent. 2025. "There Is One Occasion on Which You Are Allowed to Surprise Your President or Dean," in forum "What Do You Wish You Had Known When You Got Your First Teaching Job?" *Faith and Flourishing* 4 (2025): 151.

NOTES

- 1 Shane regularly maintains a blog where she posts thinkpieces and research more recent than her book at <https://www.aiweirdness.com/>.