

Authenticity and Authorship in the LLM Era

Brett Reynolds 

Faculty of Liberal Arts & Sciences, Humber Polytechnic

May 23, 2025

About fifteen years ago, Jonathan Lethem published “*The Ecstasy of Influence*” in *Harper’s*. Page after page he argued that creativity is recombinant – quotation, bricolage, collage – then pulled the rug: a key at the end showed every paragraph was itself spliced from other writers’ prose.

The stunt landed only because we still nurse the myth of solitary genius, the fantasy that fine writing erupts fully formed from one mind.

So let me be clear upfront: for two years, nearly every sentence I publish has first passed through a large language model. They’re the text engines trained on humanity’s archive, the things we call ChatGPT, Claude, Gemini, distinct from other kinds of AI like those that win chess games, drive cars, or solve protein folding problems.

And that kind of use in writing makes some people deeply uncomfortable, for the same reason surgeons once bristled at Atul Gawande’s *The Checklist Manifesto*. The suggestion that their expertise needed scaffolding felt demeaning, insulting to years of training and hard-won skill. But **the data were brutal and undeniable**: checklists reduced complications by 35%, dropped infection rates dramatically, and cut surgical mortality across the board.

We grab microscopes for electrons, but grab our egos for prose. We readily defer to instruments when the domain feels remote, scientific, inhuman. But when the task feels close to home – diagnosing a patient, writing a poem, driving through traffic – we resist.

I understand the romance. Rush’s “Red Barchetta” still captures something essential: the thrill of human control, flesh and metal in harmony, the open road as pure freedom. I once drove 220 km/h through Algonquin Park at midnight, windows down, music loud, feeling gloriously, recklessly alive. Stupid, thrilling, unforgettable.

But **data from Waymo’s autonomous taxis** in San Francisco are stark: they have 81% fewer injury-causing crashes than human drivers.

The pattern repeats: human judgment that feels sophisticated but performs worse than systematic approaches. Stripe, the payments company, illustrates this perfectly. For years, payment processors relied on rigid, human-designed rules that declined transactions based on simple risk flags. Adding AI judgment of this one rule improved payment success rates by 1.3 percentage points, translating to billions in recovered revenue with essentially zero increase in fraud.

Even in domains that seem creative, the tools are leaping ahead. LLM-assisted coding platforms like **Cursor** and **Aider** now let people with no programming background build fully functioning web

applications during office hours. Folks are creating sophisticated portfolio sites, interactive galleries, even basic e-commerce platforms, armed with nothing more than clear descriptions of what they want.

But there's a deeper unease here, beyond questions of safety or efficiency. It touches something we think makes us fundamentally human: authenticity. Imagine receiving a love poem, not just technically competent, but genuinely moving, with images that catch you off guard and rhythms that echo in your head. Now imagine learning it was generated by an AI.

Something changes. Not because the words are different, but because you know they were intermediated. This is *Cyrano's* problem: the beauty comes from *Cyrano's* mind and heart, but Roxane falls for Christian's face.

But the AI poem faces a distinct problem: The use of AI for art can feel like sacrilege, a trespass into exclusively human territory.

In March 2025, tech writer Nabeel Qureshi ran a [simple experiment on Twitter](#). He posted four anonymous English translations of a passage from Homer's *Odyssey* and asked followers to vote for their favourite. When the poll closed, [he revealed the sources](#): three were by acclaimed human translators, and one was by GPT-4o. The AI translation won decisively.

We are about to encounter that problem everywhere, in every classroom, every inbox, every literary journal, every conversation about what counts as real human expression.

THE PACE OF CHANGE

If you haven't used the best and newest models, you may not realize how dramatically capable they've become. And you may be less aware of the sheer variety now available, each with different strengths, personalities, reasoning styles.

Let me give you a sense of the pace. [AlphaEvolve](#), released by DeepMind this month, didn't just improve existing algorithms, it invented entirely new approaches to matrix multiplication, [the first fundamental advance since 1969](#). Then it turned around and rewrote Google's chip-layout algorithms and datacenter job schedulers, cutting compute costs by nearly 1%.

This isn't a text generator getting better at essays. This is a system of discovery improving the infrastructure it runs on.

[OpenAI's o3](#) isn't a chattier ChatGPT. It's a reasoning engine: it breaks problems, plans code, cites sources, and never gets tired. Tool-aware, image-literate, and dominating every coding benchmark. The ["Deep Research" agent](#) actually goes out and checks sources, at times taking 15 minutes or more, resulting in a long, well-structured report with citations that would take a Waterloo intern weeks to produce.

[Google's NotebookLM](#) is just that: a digital notebook where you can drop in PDFs, YouTube links, lecture notes, even entire course readers, along with your writing and other outputs, and query across all of them instantly. You could upload all of Jane Austen's novels and ask it to trace how her treatment of money and social mobility evolved from *Sense and Sensibility* to *Persuasion*.

Here's a more personal example. Eight years ago, I had a paper with what I believed were highly original ideas about grammatical gender in Modern English. It had been through one round of review at a top journal and received a "revise and resubmit" – but with only one referee report that was so confused and contradictory I couldn't figure out how to address it. The paper had theoretical merit, but the structure was apparently incomprehensible to at least one expert reader. Frustrated

and demoralized, I gave up. It sat on my computer, gathering digital dust.

Last month, I fed that abandoned paper to o3, asking it to analyze the arguments and create a new outline specifically for *Language*, the journal of the Linguistic Society of America. Within minutes, it had identified the core insights, mapped the logical flow, and suggested a completely different organizational strategy that foregrounded the empirical findings rather than burying them in theoretical apparatus.

I took that outline and the original paper to Claude 3.7 Sonnet and asked it to restructure everything according to the new plan. In one session – literally one prompt – it produced a draft that wasn't camera-ready, but was finally workable. The ideas were the same, but they were now presented in a way that made intuitive sense. I spent a few days polishing the prose and clarifying a few technical points, and the paper is now under review again at *Language*.

Eight years in a drawer; one afternoon with an LLM; a manuscript now back at *Language*.

I've even used LLMs to flesh out counter-factuals I could barely sketch: What if smallpox moved west across the Atlantic, leaving Europe immunologically naïve while Indigenous nations remained intact?

The model forced me to spell out premises, then chased consequences: land as stewardship rather than property, wealth stored in relationships not metals, consensus mechanisms that scale.

It held the whole game tree in working memory long after mine collapsed, flagged missing links, and returned a coherent narrative I could interrogate line by line.

That's the real power: it turns half-formed hunches into testable worlds.

I still don't know how to provoke students to ask, "What angle could an LLM open here?" That's an experiment I'm pitching to this room.

PROMPTING

Every good prompt needs three core elements:

1. Context first. Give the model the material to work with: paste notes, slides, or excerpts. If working with students, include relevant background like age, program, language proficiency level, and native language.
2. Role and stance. Be explicit about what you want the model to be: "You are a curious but supportive coach helping me make sense of what I've learned" or "You are a writing tutor focused on clarity and organization."
3. Guided structure. Don't just ask for help, specify the process. Ask for open questions before testing, request paraphrasing before analysis, or outline the steps you want the conversation to follow.

Reusable student prompt

You are a reflective coach.

Your job is to help me think aloud about what I've just learned.

Who I am:

<<<PROVIDE BACKGROUND>>>

Material to consider:

<<<PASTE NOTES / SLIDES / TEXT>>>

Start by asking me:

1. What feels most surprising or resonant in this material?
2. Where do I see links to ideas I knew before?
3. What still feels fuzzy or unsettled?

Listen to each answer, paraphrase it briefly, then pose one deeper follow-up.

After two rounds, invite me to sketch a real-world situation where these ideas might apply. Only then, if I ask, offer a short self-check or practice exercise.

Keep the tone probing yet non-judgmental; aim for insight, not evaluation.

A few more prompting strategies

Assign a persona: Instead of generic responses, try “Respond as Geoff Pullum analyzing my prose for grammatical precision” or “Channel Patricia Williams’ approach to legal storytelling.” The model dials down its default enthusiasm and channels more specific expertise.

Watch recency bias: Models weight the most recent instruction heavily. If you want something emphasized throughout a long interaction, restate it periodically. End with your most important constraint: “Remember: keep all feedback constructive and focused on growth.”

You don’t need to become a prompt engineer, but being deliberate about context will save you enormous time and get you much better results.

ASSESSMENT

Most of us would love to offload marking. Let’s be clear about one thing: LLMs can’t own an assessment. They can’t sit through an academic integrity hearing, or feel the sting of failure. We are ultimately accountable.

But within that framework, these tools can handle much of the mechanical labor that’s been burning us out for decades. They can:

Take dictation and structure feedback: Instead of typing comments, speak them into voice mode while you read. The model transcribes, removes the “ums” and false starts, organizes thoughts into coherent paragraphs, and formats them according to your rubric. What used to take 10 minutes per paper now takes 3–4. Much of commenting time is making the comments gentle. LLMs are good at that.

Aggregate patterns over time: Keep the comments in a NotebookLM notebook. The model can track which skills each student keeps missing, surface cohort-level patterns you might not notice, suggest targeted mini-lessons for common problems, and even flag potential grade inflation by comparing your standards across different assignments.

Audit your grading pipeline ruthlessly: Identify every step where a model could draft initial responses while keeping final judgment human. Start small – maybe just the mechanical aspects of feedback first – and gradually expand to more complex tasks as you build trust in the workflow.

What else can we do?

1. Develop and share prompt libraries. Get your students to contribute to and evaluate them.
2. Experiment with setting up and sharing a NotebookLM notebook for a class you teach. Include shared links from chats with other LLMs.
3. Embed AI reflection in lessons. NotebookLM stores searchable transcripts automatically, creating a semester-long learning journal.

4. Model intellectual humility in your own practice: Use checklists publicly. Credit the tools that help you. Show your rough drafts and revision process. Let students see that expertise means knowing when to rely on external support, not pretending you don't need it.

THE ROAD AHEAD

If creativity is already recombinant, the question is no longer whether machines will join the collage but how we curate the pieces.

In the late 2010s, most people in AI thought the path to general intelligence was through agent-based learning: AIs that acted in a world, observed the consequences, and learned from them. These were the kinds of systems that could beat Garry Kasparov at chess or learn to play any video game by trial and error – fail, retry, optimize.

It turns out we may have been lucky that this isn't the version of AI that took off. Those agentic systems – powerful but fundamentally a-linguistic, illiterate, and ignorant of humans and our society – operate through brute interaction. If scaled to real-world action, their mode of learning would've been: try it, see what breaks. Not metaphorically, but literally.

Instead, what took off were the language models. They don't "act" in the world – not yet – they read it. They ingest human culture: books, code, comments, blog posts, research papers, jokes, rants. They "know" us in a way those agentic systems never could. Not perfectly, not transparently. You can't always peer inside and see why an LLM said what it did. And often, neither can the model itself.

But still: you can talk to it. That's new. And profound. This kind of AI doesn't thrash around trying to learn what a doorknob does. It's read enough to get it. And while it might not have a theory of mind, it simulates conversation as if it does, and that turns out to be enough for a lot of practical tasks.

Now, carefully laid-out projections suggest we could have artificial general intelligence, potentially superintelligence, in 2 to 5 years. Not just tools we control from our laptops, but systems that improve themselves, replicate themselves, commandeer manufacturing capacity. It's not clear we'll remain the ones deciding what work gets done. And this isn't distant science fiction. This could happen before many of us reach retirement.

If that's the case – if machines can do most of what we currently teach students to do, faster and better – then what are we really preparing students for? If we're training them for skills that help them earn a living now, we may be building bridges to nowhere.

But if we're lucky – if these systems remain aligned with human values and priorities – then what will matter most are the capacities that we've always valued: The ability to communicate clearly and listen well. The impulse to wonder and explore. The skills of creating and appreciating meaning and beauty. The courage to ask hard questions. The wisdom to strive to be good, to be kind.

I don't pretend English teachers are uniquely positioned to develop these capacities. But we are positioned. We work daily with the stuff of human meaning-making: stories, arguments, the architecture of persuasion, the music of language. We help students find their voices. In a world where machines can simulate any voice, that may be the most valuable thing we do.

"The Ecstasy of Influence" ended with a key that exposed every stolen line.

Today the key is in full public view: we all write with machines that remix the past at light speed. Our task is not to hide the seam but to show students how to stitch meaning deliberately.