

My View on AI and University-Level English Literature

Martin Paul Eve

Professor of Literature, Technology and Publishing
Birkbeck, University of London

Associate Director of Knowledge Commons
Michigan State University

5 October 2026




Image: *The AI Race* (Gloria Mendoza / <https://betterimagesofai.org> / <https://creativecommons.org/licenses/by/4.0/>)

Cite as: Martin Paul Eve, "My View on AI and University-Level English Literature", *eve.gd: Martin Paul Eve*, October 05, 2026,
<https://doi.org/10.59348/3kh8w-q1d96>.

DOI: <https://doi.org/10.59348/3kh8w-q1d96>

Available at: <https://eve.gd/2026/10/05/my-view-on-ai-and-university-level-english-literature/>

It is bad that I have started writing this post because, for reasons you will see below, I am worried about how much time we spend talking about AI, even as I recognise that politics requires work and attention. And AI is certainly political. I am also aware that many great minds have already put their heads together on the matter of artificial intelligence, at the moment powered by LLMs (Large Language Models). Groups such as [the MLA Task Force On AI](#) reads like an all-star lineup. In the UK, the [English Association](#) and [University English](#) have also been active.

English Literature is (in part) About Teaching How to Take Intellectual Pleasure from Literature

But I nonetheless have my own views. These perhaps are probably just repetitions and assemblages of things many others have thought and said. But I want to articulate them (again), particularly in the context of UK higher education's discipline of English literature. In this discipline we teach people how to read literature, which sounds ridiculous because everybody thinks we all already know how to read novels and poems and so on. But we teach them how to think about, with and through these texts. We teach our students how to see what the text is doing at different layers and then, most importantly, how to derive *an intellectual satisfaction or pleasure* from working deeply on the interpretation, lineage, history, or other scholarly facet of a text. That's the dream.

(Of course, there are many other defences of what English literature does. We give people transferable skills, we teach them critical thinking, we provide them with tools that will gain employment, etc, etc. These are all true to some extent and I think Rick Rylance does quite a good job of defending these defences (Rylance, Rick, *Literature and the Public Good*, The Literary Agenda (Oxford University Press, 2016)). But these are defences, apologies for the discipline, ways of justifying our existence in the contemporary financialised university. What we really just want is a polity of educated readers who think critically about the texts that they encounter and know how to formulate arguments around texts. And we want people to enjoy that critical reading experience. Anyway, here I'm going to fancifully entertain the dream world where this pure state of scholarly bliss is what we actually pursue.)

But, primarily, it still boils down to the fact that we are teaching people how to understand, historicize, de-historicize, read – and in creative writing, write – literature. Literary works **mean** nothing to a machine. And the output of a machine doesn't have that kind of meaning for us, although I am very aware of the longstanding study of machine-generated texts. And **I strongly defend the continuation of the essay** as a valuable form for teaching thinking. Yes, AI can produce an essay that looks like literary criticism. A human that can produce one, though, is the worthwhile part of this. I think we should tell students this very much up front: the distinction between the output object and the process.

Indeed, my view on AI is based on many factors, but, in this context, it hinges on the fact that one of the things that AI does is to produce a facsimile output object without a person learning a process. But learning how to read, understand, and, in CW, write literature at university level is about mastering a kind of process/skillset/disciplinary mode, not really about producing an artifact. As others on the internet have already said (with apologies, I cannot find the person who put this in my feed): We don't ask undergraduates for essays because we expect to learn something astonishing from their work. We want this work to determine whether they have learned the mode of our discipline and then demonstrated that through co-incidental production of an output. Have they learned the process? The output is supposed to be merely the proxy manifestation of a learned process. Of course, it actually loops back. Producing the object also teaches the process as we learn through writing and as writing can modify our thought processes. Inscription inscribes. But AI breaks all this. AI in these contexts is total intellectual atrophy.

If we do not teach this process and allow AI to just produce the end point, the output, without students having to go through the process of learning how to do this, then we might as well give up on our discipline now.

Unfortunately, though, as Adorno noted, and with whom some students agree: "The pleasure of thinking is not to be recommended". So if a student just wants a degree certificate, it is somewhat difficult for us to stop them short-circuiting our assessment mechanism, though I have some ideas.

But I do not think that AI use should be encouraged or permitted in much of our disciplinary practice. We are talking after all about interacting with literature and understanding how to think about, through and with it. It's supposed to be a human activity. What are we doing by ceding our ability to take intellectual pleasure to a machine?

Two Important Upfront Caveats

But: an important note. Students often now **will not know when they are using AI**. It is now so baked into major Microsoft and Apple software as the default way of working that their spelling and grammar tools, even, are all now based on AI. And they will not know that even they should disclose this via our policy. Interestingly, LibreOffice, the major free open source competitor to Microsoft Word, has a stance of no AI by default.

Another similar case: many disability assistance tools use large language models underneath. So, for example, I use a voice to text system to help me type when my rheumatoid arthritis symptoms are too severe for my hands to work. This synthesizes my voice and then uses the likelihood prediction of the LLM based on the sound input to work out what word I probably said. This feels slightly different because, yes, it's a generative tool and it can make corrections. So it is using an AI system to change things. It also inserts punctuation. But it is a different way of working that tries to express what you said, rather than expressing something for you.

Specific Uses of AI in our Discipline

Some activities that students are being encouraged to do with AI by university leadership:

Gathering information using these systems is potentially dangerous. LLMs are not search engines or truth machines; they are systematic prediction engines based on language within a training corpus, distilled down to a set of probabilities or “weights” that guide the most likely next token (sometimes: a word). They are technologically amazing. That simply by capturing so much

data you can make, quite frankly extraordinarily plausible, next token predictions. Certain types of AI called RAG, or retrieval-augmented generation, can provide references to external texts, but they do not work systematically as a library search would and they can potentially mislead students. It would be better for students to learn traditional library research and information retrieval techniques via the services provided there. If students are permitted to use LLMs to gather information, they must track down the actual source itself, read and understand it themselves, and confirm that their citation is correct and theoretically/contextually accurate.

Relatedly, **generating citations or references** with AI (even with RAG [retrieval-augmented generation] systems) is extremely dangerous and will lead to citations that do not resolve to real works. People call these “fake references” because they are adhering to the “truth”/search engine episteme/paradigm of AI. But really, it’s just the most plausible (not precisely the same as “likely”) linguistic representation in the training data for the prompt (the document), and that plausible output does not have to be a real citation. I would worry that we will see a swelling of unread secondary citations coupled with citations to material that does not exist. As the UCB statement says: [a reference to a fake citation should raise the presumption of AI use](#).

Furthermore, to **understand key theories and references**, you need a reliable system that outputs consistently accurate text. AI LLM systems cannot do this, and they are likely to mischaracterise or gloss key aspects of our various fields (they produce so-called “hallucinations” – and I hate that term but won’t go into why here). After all, as the famous statistical aphorism has it: all models are wrong. Students will end up with a theoretical mess that they cannot properly understand because it has not been properly explained or read (by them).

Disciplinary Legacy

Obviously, I am a big fan of digital access to research. The only way that I can read and write interesting academic material in my current medical hospital admission is via digital access to research. But that is very different to using LLMs or “AI” to structure that interaction. I have spent

years learning research techniques to find the correct path through our secondary literature.

That is something that we need to teach our students. I strongly feel that permitting AI, even in these limited areas, is not the way to pass on the skills and knowledge of our discipline to a cohort of students.

Politics (Critiquing AI) and Pleasure (Literature)

There is, of course, a much broader political context to this. AI seems to be the way to “work smarter with fewer people” (urgh). But UK universities are under severe financial threat at the moment, as we know all too well at my UK institution. If the UK university sector contracts violently, we will be left with a set of elite institutions that can afford to teach privileged students properly and pass on this knowledge of intellectual pleasure and critical engagement, while the satellite institutions, the poorer second-class university cousins, will use AI to give a degree certificate. Academics will be encouraged to grade papers using AI for speed and efficiency. Students will then realise that they could get the same effect by not going to university and just using the AI themselves (or not bothering).

The institutions that will survive, I think, are the ones that will outlaw AI by promising students they will be taught by people and that they will be expected to work without AI. They will be taught a process or tradition, or mode by people who have themselves learned it. Critical thinking and intellectual pleasure will be the preserve, in the space of English Literature, of a small elite.

But I also wanted to talk politics in a slightly broader context with respect to AI. As I stated at the top of this post, one of the things that annoys me is that so much energy and time – as evidenced by the anecdotal contexts of my Mastodon and Bluesky feeds – is spent arguing against AI. There are even posts complaining about how much material there is about AI. And this post is a prime hypocritical example.

Every time we write about AI, we feed its hype machine simply by dint of volume. But this talking about AI displaces the things we could be talking about. We could instead be talking and writing about a recent novel we read. Our hobbies. The activities we enjoy. Our accomplishments and failures. How we fell in love. Comic books, novels, recent musical concerts, enjoyable travel, works of art, our own attempts at painting, playing guitar, The Beatles, playing sports, observing sports, making music. We could just have time to think. We could talk about the things that give us joy, that give us pleasure. But talking about AI displaces this pleasure time. We have to write about what we consider important.

But you cannot give up on politics (by which, here, I probably mean openly critical dissenting speech). The people running AI systems have more money than it's possible to imagine. This gives them the ability to damage the world more than it's possible to imagine. We must continue to criticize AI and the plutocrats. The volume of critique does make a difference on occasion; although in many cases it does not (see: 1Password and Omarchy). But it sucks away our lives and time for interesting things, having the constant need to counter these oppressive systems and people. In English Literature departments, so much departmental time and administration has been spent not really discussing how we will teach our discipline and its practices, but what we do about this bold, intrusive and ultimately damaging new technology.

This is the political economic trade-off in our space: politics (time devoted to critiquing AI, developing policies for handling it etc.) vs. pleasure (reading and teaching literature). And that's what AI, along with every other bureaucratic incursion, ends up taking away from us.

References

This reference list has been generated, in many cases, through an automated process and may therefore contain errors or incomplete metadata. It is provided purely as assistance for finding the items mentioned within this post.

‘Action on AI’, *University English* <<https://universityenglish.ac.uk/action-on-ai/>>

‘AI in English Studies’, *The English Association* <<https://englishassociation.ac.uk/ai-in-english-studies/>>

‘Artificial Intelligence Policy’, *Berkeley Law*, UC Berkeley School of Law
<<https://www.law.berkeley.edu/academics/registrar/academic-rules/artificial-intelligence-policy/>>

Eve, Martin Paul, ‘The Great Automatic Grammatizator: writing, labour, computers’, *Critical Quarterly*, Wiley, October 2017 <<https://doi.org/10.1111/criq.12359>>

‘MLA Task Force on AI in Research and Teaching’, *Modern Language Association*
<<https://www.mla.org/About-Us/Governance/Committees/Committee-Listings/Professional-Issues/MLA-Task-Force-on-AI-in-Research-and-Teaching>>

Priego, Ernesto, ‘On Taking the Time’, *Ernesto Priego*, 16 March 2018
<<https://ernestopriego.com/2018/03/16/on-taking-the-time/>>

Priego, Ernesto, ‘Inauguration Day’, *Ernesto Priego*, 20 January 2017
<<https://ernestopriego.com/2017/01/20/inauguration-day/>>

Rylance, Rick, *Literature and the Public Good*, *The Literary Agenda*, Oxford University Press, 2016 <<https://global.oup.com/academic/product/literature-and-the-public-good-9780199654390>>

Vignoli, Italo, ‘Yes, no AI is now a feature’, *TDF Community Blog*, 3 September 2026
<<https://blog.documentfoundation.org/blog/2026/09/03/yes-no-ai-is-now-a-feature/>>