

AI / Large-Language-Model text generation and the writing for this course:

First things first: [read this document](#) about *kinds* of writing, before you read the paragraphs below—which are focused entirely on a particular kind of writing you'll be doing in this course. I wrote the linked document for my colleagues, but it's equally relevant to students in my classes. What follows is for you—students in my courses . . .

Free-standing text-generation products, as well as add-ins to Word and Google Docs, will expand and proliferate at a dazzling rate during your time as students and well beyond your years at this university. All of these products work, as I'm sure you know, by utilizing an enormous quantity of text that has been, in the expressive locution of software designers, "scraped" from the Web.¹ This vast harvest of language-use-samples is then used by the program to predict likely language sequences based on the questions fed to the AI and the tasks the AI is asked to perform. The more samples the program has to work with (the larger its language base) the better it will be at turning out plausible sentences—"plausible" in the sense that the sequence of words that comes out of a "Chat" program sounds OK, sounds like a sentence or a paragraph or an essay that someone might say or write if that person were unconcerned about speaking in platitudes and formulaic phrases: words following one another in expected sequences, phrases that make sense because they obey sequence-patterns that are discoverable if a program combs through enough previously articulated language samples.²

¹ See [this video](#) for an informed, informative, fascinating discussion of how Large Language Models (the engines that drive our current AI products) work. [This introductory essay](#) is longer, I'd imagine, than you'll want to read, but give it a try for at least a few screens; it's technical enough to be genuinely useful. Two books that are also useful (even if already outdated in some respects) are [The Alignment Problem](#) by Brian Christian and [Artificial Intelligence](#) by Melanie Mitchell. And it's worth considering how / whether OpenAI, Anthropic, Meta, Microsoft, X, etc. have our interests at heart; see Shoshana Zuboff's [Surveillance Capitalism](#)—written before AI began its current rise to prominence, but still profoundly relevant to a consideration of the companies that design and promote AI. Here are some [tips](#) for turning off Google's AI plug-in in Mail and Docs.

² See "[AI is homogenizing our thoughts](#)" (*New Yorker*, Summer 2025). AI as a "[mass delusion event](#)." The processing of harvested language has a sizeable carbon footprint. See [this informative piece](#) from the *NYTimes* 17 March 2025. From a [short piece](#) in *The New Yorker* by Elizabeth Kolbert (9 March 2024): "Artificial intelligence requires a lot of power for much the same reason. The kind of machine learning that produced ChatGPT relies on models that process fantastic amounts of information, and every bit of processing takes energy. When ChatGPT spits out information (or writes someone's high-school essay), that, too, requires a lot of processing. It's been estimated that ChatGPT is responding to something like two hundred million requests per day, and, in so doing, is consuming more than half a million kilowatt-hours of electricity. (For comparison's sake, the average U.S. household consumes twenty-nine kilowatt-hours a day.)" And from Scott Galloway's "No Mercy, No Malice" newsletter (10 May 2024): "Big Tech isn't just similar to the energy business, it is the new energy business. AI's growing power requirements make this concrete. AI compute requirements are doubling every 100 days, dramatically countervailing the gains in efficiency that every AI evangelist boasts about after vomiting that it will save/destroy humanity. Training the trillion-plus-parameter models being stewed in a medium-size city

When I've used these programs to generate essays on authors and texts I've taught for years, what I've gotten back has been almost instantly generated, perfectly predictable, conventionally phrased essays with all their parts in place and nothing surprising or delightful or provocative in their sentences or their larger arguments. This is exactly what one would expect, given that the very premise of the programs works against idiosyncrasy, surprise, insight, incandescent thought, linguistic novelty, creation itself. Indeed, what comes out of an AI generator looks *something like* "thought" made manifest in language, but that's an illusion. This is writing that has literally nothing to do with thinking and that is, instead, a pattern made of words that can sound or feel "thoughtful" to a human reader because we're used to such patterns being made by other humans who have worked to put ideas into language.³

This foundational difference between Chat-language and your own language use and your own thought—thought that precedes, intersects with, grows out of, and is shaped by your writing—will remain true even as the quality of prediction-generated language grows more and more impressive. We receive these newly generated texts as "writing," but we might more accurately understand them as displays of word-patterns. And in an irony so perfect that it belongs in a fairy tale or myth, these AI-products—making sentences that pass for language use but are only used language—are contaminating the source of their own facsimiles of thought: as AI-generated text floods the internet and comprises a larger and larger percentage of the language that's scraped back into the corpus upon which this AI-"writing" depends, the quality of that "writing" promises to degrade: the web is being polluted by this iterative machinery that requires our creative labors to make it work.⁴

(San Francisco) requires the energy consumption of a small country. One ChatGPT request requires 10 times the energy of a Google search. In five years, the incremental energy demand of AI will be equivalent to 40 million homes — more than California, Texas, Florida, and New York combined. Data centers make up 3% of total U.S. power demand, but that's projected to triple by 2030. BTW, 2030 is the same distance into the future as the finale of *Game of Thrones* is in the past (2019)" ("[Big Energy](#)"). See also "[The Costs of the Cloud](#)," by Ashley Dawson. AI also uses [large numbers of gig workers](#), often in the Global South, to train its models; the social and political costs of AI are hidden, of course, from users in the so-called "developed" countries. See also [this story](#) about "content moderation" from *Wired* about this work—pre GenAI but still essentially the same story, the same human costs; and [this story](#) about exploited labor that lies "behind AI."

³ See [this fine piece](#) from *The Atlantic* on the pitfalls of anthropomorphism, with particular focus on Claud's "[constitution](#)," issued in 2025–2026.

⁴ [James Somers](#) has written [a fine piece](#) in *The New Yorker* (5 October 2023) that revolves around this conundrum (and has a great deal else of importance to say as well). See this piece in the *Times* from [4/8/24](#) as well as [this piece](#) from the same day. See also [this op-ed](#) from the *Times* (29 March 2024), and [this piece](#) from the *Times* (26 August 2024); it's called "[model collapse](#)." And [this piece](#), on the difference between "Answer Machines" and "Question Machines." Another—a call for [cultivating thought](#)—also from the *Times*; see [this piece](#) by Brett Stephens, too; see also [this](#) from Ezra Klein, and [this](#) from Dan Chiasson in the *NY Review of Books*.

The writing I'm asking you to do for this course, and the writing that I hope you'll continue to do long after you've graduated from this university, is your own exploration of something that matters to you. Your sentences constitute your invention and shaping of your ideas, your discovery and creation of what it is that you think. Your essays will be as far removed from the predictive patterns of what a program might gather from the internet as is a generic floorplan of a house from anywhere you've actually lived—a place filled with particular things, saturated with your memories, registering the days and nights you've spent in those rooms, made complicated and interesting by the life you've lived there.



It's my hope that you'll take the opportunity this term to work with me on *your*

Here's a New Yorker piece, "[Will the Humanities Survive AI?](#)" that answers "yes" in ways that I find congruent with what I believe happens in good classroom gatherings and in good academic writing. Karen Hao's [new book](#), *Empire of AI* (2025), is powerful and addresses aspects of AI that are receiving far too little discussion; see also this [podcast interview](#) with her). For appropriately disturbing discussions in Hao's book of the environmental and social-psychological costs sustaining all the AI empires, see particularly part II, chapter 9, "Disaster Capitalism" and part III, chapter 12, "Plundered Earth" (these chapters are not downloadable, but you can read them, with a Mirlyn UM login [here](#)). See also Joshua Rothman, summer of 2025, a 30,000-foot view of AI and [our futures](#). See also this philosophically oriented essay by [James Gleick](#). [This op-ed](#) by Meghan O'Rourke is powerful too. UM's own, new (fall, 2025) [guidelines](#) are woefully inadequate along ethical lines, particularly. And this sentence near the end of the document—so imprecise (particularly in its final phrase), so poorly written that it was either not written by a human or written by a human who's read too much AI-generated prose, is a revealing example of the kind of slop we'll all have to contend with as we are subjected to the brayings of AI-boosters: "Understand that your usage of GenAI-based tools can give you the means to better not just yourself, but also society as a whole, and there is an ethical responsibility towards doing so." Another thought piece about AI, also from *The New Yorker*—this one about [personality\(?\) and Claude](#) (the author, Gideon Lewis-Kraus, was also interviewed on "[Fresh Air](#)." And here's something a little different—from a friend and way-long-ago former student, [Robbie Fuls](#), who's a professional musician, on [music and AI](#). An *Atlantic* piece on AI's assault on [the reader](#).

writing—on the thought that goes into your questions, your arguments, your readings; on your drafts and their revisions; on sentences that are the best version you can make of your ideas as they come to you. It will be more demanding than the generation and then, perhaps, modification of something produced for you in seconds by a [chat](#)bot, but I promise you it won't be artificial intelligence.⁵

NOTE: If you choose to use AI for any aspect of your papers for this course, in spite of my advice and pleading for you to value and cultivate your own capacities for thought rather than choosing to employ a facsimile of thinking, you are to footnote and explain that use in some detail: as in, “I used [fill in AI module name and version here] to generate ideas and an outline for this essay.” or “I used [fill in AI module name and version here] to help me write paragraph 2.” or . . . you get the idea.

⁵ [Here](#), if you want a sample of AI's not-quite-“voice,” are two ChatGPT-generated arguments against using AI for writing. And [here](#) is Gemini at work on Joyce's *Ulysses*. And here's a moving poem about interacting with an [AI entity](#).