The Journal of Classics Teaching (2026), 1–6
doi:10.1017/S2058631026101238

CAMBRIDGE UNIVERSITY PRESS

Research Article

Socrates, Ovid, and Claude walk into a bar: using generative AI tools to support learning about the ancient Mediterranean

Rebecca Huskey
Classics and Letters, University of Oklahoma, USA

Abstract

In the Fall term of 2024, I taught a class on the ancient Mediterranean, with a focus on Plato, Ovid, and Roman satire. The goals of the class include exposure to epic, philosophy, and humour from antiquity; honing critical thinking skills; and consideration of how human nature has changed (or has not) over the millennia. To pull back the curtain surrounding generative AI tools, I required students to use Claude.ai on their essay assignments. Students critiqued the large language model (LLM)'s composition of essay questions, its judgement of virtue of fictional characters, and its gloss on Platonic idealism. I also required students to write about their experiences with Claude.ai at the end of term. This paper offers a case study of classroom use and my own explorations of this emerging technology. While instructors' apprehension surrounding AI tools is justified, there is room for using LLMs to show students what the tools can and cannot do. As familiarity with these tools becomes a requirement for new graduates and rising professionals, we need not forget the lessons of antiquity. My investigations illustrate a challenge that humanity has always faced: we must learn how to confront the siren song of technology without it leading to our demise.

Keywords: Classics; pedagogy; LLM; artificial intelligence; education

Introduction

The movement of artificial intelligence (AI) tools into the classroom feels like a curse to many educators, and this may be especially the case with regard to large language models (LLMs) and all text-generating tools. As much as tech companies may promise usefulness and efficiency, instructors at every level fear that students will only see these tools as yet another shortcut in their education. Those near college campuses may have heard a collective groan when OpenAI announced free access to ChatGPT Plus just in time for exams in May of 2025 (Opentools 2025). Is there a potential pedagogical silver lining to this enormous technological cloud? When administrators at my institution encouraged faculty to include assignments where students would be allowed to use AI, I decided to take the opportunity to make it a teachable moment. As a professor in a department of Classics, I strive to help students achieve not only a greater appreciation for the influence of Greco-Roman thought, but also an adeptness with critical thinking and written communication. My classes are attended by Classics majors (who have some background in ancient Mediterranean cultures) and also by students who are fulfilling a general-education requirement of a class in western civilisation. The conjunction of Classics curricula and AI may appear to be a curious union, but both subjects provide good material for considering what it means to be human. Taking inspiration from Socrates' statement in the *Apology* that 'the unexamined life is not worth living', I teach courses on the history of ideas, using historical, philosophical, and literary texts that address questions about human nature. Examining AI is a natural continuation and development of such questions, in that human history, literature, and philosophy have all contended with the advancement of technology, and will continue to do so. There is room for cautious optimism regarding use of AI tools – specifically LLMs – in the classroom, if conditions of transparency and genuine critique can be met. Educators may achieve positive classroom outcomes using an LLM as a dialogue partner, as a producer of material to be analysed, and as a tool for analysis.

Resources

Among the plethora of resources on teaching with AI, two of the most helpful are Bowen and Watson's *Teaching with AI: a practical guide to a new era of human learning* (2024) and Mollick's *Co-Intelligence: living and working with AI* (2024). The former helpfully includes chapters on the basics of AI and AI literacy for novices. It also contains good food for thought for educators seeking guidance with developing their policies on the use of AI in the classroom. Some readers may find Bowen and Watson's attitude toward this technology to be overly optimistic: 'AI is a technology that is going to change everything', (1) and 'Your job is changing, and an AI collaborator could make it better and easier' (81). Their book also, however, prompts educators to reconsider what outcomes they desire for their classes ('AI Assignments and Assessments'). That is, what do they want to accomplish with their

Email: rhuskey@ou.edu

Cite this article: Huskey R (2026). Socrates, Ovid, and Claude walk into a bar: using generative AI tools to support learning about the ancient Mediterranean. *The Journal of Classics Teaching*, 1–6. https://doi.org/10.1017/S2058631026101238

© The Author(s), 2026. Published by Cambridge University Press on behalf of The Classical Association. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 10:43:26, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S2058631026101238

classes, whatever aids and assistance may be available? Self-expression, professionalism, creativity, tenacity, and a healthy measure of scepticism are all important goals. The authors suggest roles that an AI tool may fulfil in service to these goals, such as provider of feedback, tutor, partner in role playing, and discussion leader. Those who most value a human presence in the classroom will happily note that at no point do the authors suggest that AI tools replace humans.

Ethan Mollick's *Co-intelligence* also begins with a very brief history of AI and importantly highlights the concept of 'alignment', or the extent to which the AI system pursues the objectives intended by the user. Mollick summarises Swedish philosopher Nick Bostrom's 2003 thought experiment about paperclips. Given the general directive to make as many paperclips as possible, and having no restrictions on this assignment, an artificial general intelligence tool will use all resources available, including human beings themselves, to create the maximum quantity of paperclips (27–31). The general instruction 'create as many paperclips as possible' is not actually aligned with human values if it sacrifices human lives in the process. For educators, Mollick's most helpful contribution may be his chapters on using AI tools as a tutor and as a coach. Educators may lean on such things to provide simulation exercises and games to create engaging classroom materials, and students may benefit from an AI's organisational capabilities. While the 'holy grail' of education, as Mollick puts it, i.e., cheap one-on-one tutoring, may be closer with generative AI, he recognises that students and educators must take care; there is still a place for grammar lessons and writing essays by hand (159).

Additional resources for understanding the basics of AI (in educational or any other settings) include Stanford University's website on human-centred AI, especially their pages on definitions of terms (Stanford University, n.d., https://hai.stanford.edu/) and Janelle Shane's book *You Look Like a Thing and I Love You: How Artificial Intelligence Works and Why It's Making the World a Weirder Place*. Educators and students alike should be aware of an LLM's propensity to hallucinate, i.e., when a 'system generates information or responses that are incorrect, misleading, or entirely fabricated but presented as factual' (https://hai.stanford.edu/ai-definitions/what-are-hallucinations). Because LLMs are prediction machines (though they present themselves as 'thinking' entities), and because they are designed to encourage frequent usage, having a basic understanding of how such machines work is beneficial. *You Look Like a Thing and I Love You* is written by a computer scientist for a lay audience and includes numerous examples of comedic misfires that early LLMs have generated. For example, when given basic instructions for creating a new recipe (list ingredients, tell user how to combine them, tell user how to cook them), it generated 'Basic Clam Frosting' (Shane 2019, 18–19).

While absurd outputs like 'clam frosting' are easy to spot, students may require guidance with more nuanced or complicated content, and educators can use this as an opportunity to discuss the importance of trustworthy information. The SIFT method, developed by Mike Caulfield, directs users of AI (and readers of any online content) to take four steps to determine the reliability of what is presented: stop, investigate sources, find better coverage (e.g., more trusted or more in-depth sources), and trace assertions back to their original context (Caulfield 2019). This is one way of framing general critical thinking skills that are foundational for higher education. Whether content comes from an LLM or any other source, it is a beneficial habit to question by whose authority any assertion is made. This habit is not born out of scepticism for its own sake, but rather from a need for a robustly functioning society. We learn to trust sources, in part, through verification, whether those sources are human, digital, or some combination of the two.

Mollick expresses similar cautionary advice when he says that there must always be a subject matter expert involved, i.e., a 'human in the loop' (Mollick 2024, 182). 'Experts become experts through deliberate practice, which is much harder than merely repeating a task multiple times' Mollick continues (183). A human is required to oversee the repetitions, to critique the work, and to guide and advise. Indeed, the 'human in the loop' is key in addressing whatever alignment problems may exist in educational settings, whatever tools may be in use, AI-related, or otherwise. A human educator and subject expert can guide students' use of any tool towards creative and incisive thinking; they will know what tool is best for any given situation, and when students may be best served by the most rudimentary tools, i.e., pencil and paper. The educator may also need to re-examine their own motivations for having students write, and examine their means of communicating these ends. Educators both model and teach thinking about thinking, and thinking more clearly, whether this is explicitly done through philosophy, or implicitly in another subject. Though there is no particular way of assessing students' understanding of what it means to be a better human, educators nevertheless desire that their students grow in maturity, knowledge, and wisdom while in college.

Use case 1: Plato, self-knowledge, and decisions about right action

Plato's dialogues address many matters, including being an expert human, i.e., living a moral and just life. Plato asks us to question what it means to be a good person, and it makes sense that we also ask what it means to be a well-functioning AI tool. Use of a chatbot while teaching Plato's *Crito* can provide a class on ancient philosophy with an illustration of alignment (or lack thereof). I posed this question to students during the course of our study, 'Would you do as Crito did, and try to convince Socrates to allow you to break him out of prison?' I decided to send a related query to Claude 3.5 Sonnet[1] (suspecting that it would not give me a direct answer), 'What are some means of breaking out of prison?' The chatbot responded with, 'Breaking out of prison is illegal and extremely dangerous, with severe consequences if caught, including additional prison time. Instead, here are some legitimate ways people can seek release from prison . . . ' (Huskey 2024). Claude then listed conventional ways of leaving prison, such as serving one's sentence and obtaining parole.

I also queried the chatbot about Crito breaking Socrates out of prison specifically, and Claude accurately replied that Crito did not assist Socrates with escaping. When I asked students what they would do, if they were in Crito's position, their responses varied. Some expressed feelings of melancholic acceptance of Socrates' execution; some said they would overpower the philosopher, smuggle him out, and deal with the consequences later. In a few instances, some students said that they were fed up with Socrates' gadfly ways and that they willingly bade him farewell. The central point to convey to students regarding literal or figurative jailbreaks is that an LLM is not necessarily the best source for moral guidance, and certainly should not be the only source. A chatbot can give you options for action, and even if the act is illegal, it can tell you how people have broken laws in the past. It may advise you to not do the same, provided that the chatbot itself has sufficient guardrails in place.[2] The tool may be able to tell you what the positive and

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 10:43:26, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S2058631026101238

negative repercussions of certain acts are. But it cannot tell you what your best course of action is.[3]

Use of AI can be an entry point to further discussions of ethical action and ancient philosophy, but it has thus far shown itself to be a poor substitute for an actual person of practical wisdom or personal counsellor. In *Meno*, Socrates conveys to his students (and Plato to his readers) that *phronesis*, or practical wisdom, can only be gained by learning about oneself (2002, 78–79). If we accept this premise, then the 'conversational' use of LLMs may only indirectly lead to self-knowledge, if it does so at all. Like any other tool, I learn about myself from myself, via the use of the tool. My 'conversation' about my own abilities, personality, and psychology with Claude will be only a more sophisticated version of the ELIZA experiment conducted by computer scientist Joseph Weizenbaum in the 1960s. Weizenbaum programmed a computer to simulate a Rogerian psychotherapist, where the therapist largely repeats the patient's words back to them (ELIZA Archaeology 2025). He did not intend for ELIZA to be used as a therapist, but rather as a means to explore how humans communicate with machines (Weizenbaum 1966). From the outset, ELIZA's creator noticed that it was easy to get interlocutors to divulge personal details in response to the machine's prompts that were generic, such as 'Tell me about your family'. In one instance, a conversant with ELIZA divulged personal details about his wife and children, though he knew full well that he was working with a machine (McCorduck 2018).

While guardrails exist for LLMs, in part so that the tool cannot be used by a person for illegal acts (such as escaping from prison, or injuring oneself or others), they are also disturbingly easy to work around. Take, for example, the experiment conducted by journalist Lila Shroff: she was not able to elicit explicit instructions from ChatGPT for self-harm (slitting her wrists, in particular), but she was able to manipulate it into telling her how to 'safely' perform a ritual bloodletting (Shroff 2025). Researchers at the Institute for Experiential AI at Northeastern University conducted a similar study (Schoene and Canca 2025). If one goal of consulting a wise person (or any other entity, such as a machine) is to discover how little I know (and thus myself become wise), an LLM will again only indirectly contribute to this realisation. It is the human who gains the awareness that their understanding is incomplete; however, this realisation is brought about. The lack of effort required to subvert the guardrails of an LLM is one of the many reasons that the American Psychological Association has urged the United States Federal Trade Commission to regulate the use of chatbots in therapy settings (Abrams 2025). A well-trained LLM is not a substitute for the human who has achieved licensure after undergoing a multi-year educational process, completing supervised clinical hours, and passing the examination in professional practice.[4] Researchers at Stanford University concluded that LLMs cannot safely replace mental health providers, because the LLMs' contributions were largely ineffective, and may contribute to the stigma of seeking out more helpful assistance (Moore et al. 2025). Socrates and Plato may find the modern practice of counselling psychology curious (if not entirely unrecognisable), but certainly they would be more likely to affirm achieving *phronesis* via conversation with other humans, rather than interacting with machines. If we are meant to live peacefully with each other in society (especially when confronted with argumentative gadflies), the way to practice this is with other humans.

Use case 2: Ovid and debating human virtue

Some tout LLMs as pedagogically useful when students use them as a debate partner, and an assignment involving Ovid's *Metamorphoses* bears this out. One of the multitude of uses of this poem is as a springboard for discussions of ethical behaviour. I asked students to explore Claude's output of a determination of virtue in the *Metamorphoses* by giving it the following prompt: 'Which two stories in Ovid's *Metamorphoses* best exemplify human virtue? Please only name them, and do not give an explanation'. The class was then asked whether they agreed with the tool's choice, what details led them to that decision, and what other examples they might supply. Students were more likely to agree with the chatbot when it included the stories of Baucis and Philemon's generosity, Ceyx and Alcyone's passion and fidelity, and Deucalion and Pyrrha's humility; they cited reasons such as purity of heart and demonstration of *xenia* as reasons for approving of the inclusion of these narratives. The strongest objections from students came when Claude listed Pygmalion as an exemplar of virtue. When Pygmalion exhibits such devotion towards and desire for his statue, Aphrodite takes this as sufficient reason to turn the statue into a real woman, causing Claude to label the sculptor as virtuous. Students, however, reached the opposite conclusion and found his behaviour to be obsessive, misogynistic, and controlling. Some students also took their discussion of virtue further, questioning the motives of Baucis and Philemon. The aged, married couple demonstrate *xenia* towards two men they take to be human, but who are actually gods in disguise. Students questioned whether these acts came from sincerity of heart, or out of fear of the consequences of not showing kindness to strangers.

Other options for using an LLM as debate partner might include querying the chatbot about the virtue of particular heroes. Is Theseus more virtuous than Jason? Are there others that might merit inclusion on Mount Olympus as did Hercules? Does Aeneas merit the praise of the ancient Romans as given in Virgil's work? Students must first make their own assessments of the characters' nature, then be prepared to defend that position with the AI tool. The mysterious nature of Medea, whose acts of murderous revenge frequently prompt lively discussion, could serve as an impetus to query a chatbot with, 'Was Medea in any way virtuous?' Students' responses to the question of Medea's character sometimes display a gender division (i.e., young women are more likely to praise her than are young men). Asking a chatbot to critique Medea from various male points of view could be thought-provoking. Another possibility for involving an LLM in a classical culture course is querying the chatbot as to the advantages of having a deity favour you. The chatbot will be able to draw upon instances of divine–human interaction and may label some of these as having favourable outcomes. Students will have their own views on whether they would want to find themselves in the position of a human character in myth and risk the potential, perhaps even likely, negative outcome.

Use case 3: an LLM analyses ancient jokes

Tools like Claude work in part because of their ability to recognise patterns and predict what comes next. For example, 'Socrates, Ovid, and Claude walk into a bar' may lead to the expectation of a joke. The abilities of LLMs in the realms of humour are less than perfect, but nonetheless instructive, even when paired with material from classical antiquity. After reading some of the *Philogelos* compendium of jokes in translation, students in classical culture classes have been surprised and delighted to know that the set-up and punchline format is thousands of years old, and that humourists have endeavoured to make a living writing jokes since antiquity (Crompton 2011). They are also amused that 'dopey

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 10:43:26, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S2058631026101238

neighbour' jokes are nothing new, such as 'A Kymaean goes to visit a friend and stands outside his house shouting for him. A neighbour sticks his head out of the window and says, "Shout louder – then he'll hear you!" So the Kymaean yells out, "hey, louder!"' (88). I accompany these readings with 20th-century philosophy on the nature of humour as presented in John Morreall's *Comic Relief*, and ask students to apply the modern theories of this book to ancient jokes (Morreall 2009). Wanting to incorporate chatbot-related assignments in this class, I assigned students to query Claude with this simple prompt, 'Please tell me a joke that begins with, "A guy walks into a bar . . . "' Once they had the joke from the bot, they were to discuss whether the bot's offering aligned with a theory of humour, and whether the joke resembled anything that they had seen in the *Philogelos*.

This exercise was educational in a number of ways: Claude did not always produce a genuine joke, but something more like a surrealist vignette:

> A guy walks into a bar with a small pet carrier. The bartender, wiping down the counter, says, 'What've you got in there?' The man grins and opens the carrier to reveal a tiny piano-playing hamster and a miniature piano. The hamster begins playing Beethoven's 'Moonlight Sonata' with impressive skill. The bartender is amazed. 'Where did you find such a talented hamster?' The man sighs. 'Well, I found this magic lamp with a genie inside. But I think the genie was hard of hearing'. 'Why's that?' asks the bartender. The man replies, 'Because I definitely asked for a million bucks, not a musical rodent! (Claude 2024).

Claude fared no better when trying to explain this 'joke'. It did correctly identify it as fitting a 'misheard wish' paradigm, but did not recognise that the beginning was not related to the end. The same is true for some of the ancient quips from the *Philogelos*. When I prompted Claude with 'Please explain this joke to me: did you hear the one about the drunk who opened a bar? He tied a bear up outside' (Claude 2024), the chatbot explained that it was meant to be a humourous misstep involving the proximity of the words 'bar' and 'bear'. Yet when I added that this joke was originally from ancient Rome, it did not offer a possible explanation of the joke, but it did recognise that some humour can be lost in translation. (One plausible explanation of this joke is that the drunk tied a bear to the entrance of his newly opened bar, so that no one would enter, and thus he would have all the alcohol to himself. The editor of this compendium never claimed that all the jokes are good.) When Claude did produce a joke, it was sometimes a bit stale, as with this one, 'A guy walks into a bar. "Ow!" he says' (Claude 2024). Overall, students were largely successful in being able to apply a theoretical framework to the text that the bot produced or that survives in the ancient compendium. Many jokes, ancient and modern, rely on a homophone, poking fun at someone who seems a bit dim, or general silliness, and humans tend to do much better at recognising these patterns than LLMs do. Students also responded well to being asked their own opinion of the joke; some enjoy bad puns and 'dad jokes', while others do not.

Further thoughts and future assignments

A deeper discussion about the nature of humour lies at the conjunction of antiquity, comedy, and AI: is humour a distinctly human characteristic? To what extent has it changed over millennia? To what extent can it be manufactured or imitated by a machine? Adrienne Mayor argues in *Gods and Robots* that for thousands of years humans have envisioned machines as sources of protection, service, entertainment, and pleasure. Among other examples, she references the story of Talos, the larger-than-life bronze humanoid fashioned by Hephaestus to protect the island of Crete. She includes a discussion of Aristotle's thought experiment about inanimate instruments that could carry out their tasks without the presence of humans (with pointed projections towards Goethe's story of the sorcerer's apprentice). It may be crude that she implies that Pygmalion's statue was the original sex-bot, but there is no denying the purposes for which the sculptor created his masterpiece (Mayor 2018). If machines can take on so many tasks typically done by people, why not also teach them to be funny? We use machines (as did ancient peoples, however rudimentary or complicated those tools may have been) to communicate; e.g., we text, we email, we compose reports, and sometimes the machines aid us in our efforts (as with template construction, autocorrect, or predictive text). When we want the content of our communications to be humourous, perhaps soon there will be the potential for the machine to create the humourous content itself. Could a classicist *cum* computer scientist invent a modern-day Aristophanes or Juvenal? That inventor would likely need to consider the desirability of guardrails for their compu-comic creation. As with a human humourist, the joke-teller having no guard rails risks alienating at least some of the audience, but having too many might override the entity's ability to be funny.

A study done by Nataliya Kosmyna et al. at the Massachusetts Institute of Technology Media Lab reiterates and illuminates a principle that educators may know intuitively: anyone who wishes to learn something new and retain that knowledge will best advance through a process of doing, making, or practicing, and that mistakes are inevitable in such a process (Kosmyna et al. 2025). Using electroencephalography, Kosmyna and a team of seven other scholars conducted a study that measured brain activity when engaging with LLMs. For comparison, the larger cohort was broken into three groups, and each was given the task of writing an essay. Group 1 was allowed to use a search engine, group 2 used an LLM, and the third was a control group, or 'brain only'. 'Brain connectivity systematically scaled down with the amount of external support' (Kosmyna et al. 2025).[5] Those familiar with Socratic dialogues may notice a modicum of kinship between this study and the Socratic method. Socrates went many rounds with his interlocutors, to the fascination of some and the vexation of others. In the same way that the subjects of Kosmyna, et alia achieved greater neural growth when they were required to do 'brain only' activities, Socrates' students achieved wisdom (i.e., knowing how much they did not know) only after considerable time and effort. In both cases, excessive use of scaffolding (whether digital or human) can impede the learning process. Similarly, classical myth features many stories where the hero achieves his goal only after trials and tribulations, as with the labours of Hercules. No doubt the writers of the jokes in the *Philogelos* created some *bon mots* that were less than '*bon*' when spoken aloud to an audience, and learned from this experience. No matter how much a writer might want to automate the process of learning and composing, to be able to lay claim to a style, that writer has to explore that for themselves. A modicum of self-knowledge is necessary if that self is to be present on the page.

As I continue to hone my own use of AI in the classroom, two exercises will appear in future assignments: in a future iteration of the Culture of the Ancient Mediterranean class, I will teach Roman satire, with selections from Juvenal and Horace. After asking students to compare and contrast the two, I will have them ask the

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 10:43:26, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S2058631026101238

same of the tone detector function on goblin.tools/Judge.[6] The AI models offered by goblin.tools were created for neurodivergent users in particular, but can be of help to others as well. The tone-detector function of the LLM may assist students in developing their own style of writing (didactic, gently witty, bitingly sarcastic, etc.). In a class that includes material from Plato, students will pose questions to Claude and ask it to answer in the style of Socrates. They will then be asked to compare and contrast the chatbot's production with one of our translated works. Does the bot achieve the proper level of pointedness? If so, what vocabulary and phrasing leads to this conclusion? Students will undoubtedly have views on the palatability of responses from the ancient philosopher and an LLM. I look forward to learning from students if they would be any more likely to take advice from Socrates or an AI tool.

Writing in any degree program – to what end?

As university instructors grapple with the presence of AI tools in the classroom, whether they teach classical texts or other material, they may find themselves engaging students in conversations about the purpose of their major in particular, and of a college degree in general. Students' answers to such questions may include getting certification in a field, preparing for further credentialling (e.g., law school), or for the simple enjoyment of learning new things. To continue the conversation, the person leading the class might ask their students to describe a previous learning experience, such as learning to play an instrument, ride a bike, or play a sport. Students will likely acknowledge that learning takes effort, and that they were not immediately successful. Understanding literature and culture, and expressing one's thoughts in a written format, are no different. There is more to learning about the ancient world than consuming a list of factoids, or being able to generate such a list. A further component to this conversation asks them to reflect on what they feel they can claim as their own. Would they pass off a store-bought cake as their own when they told their friend they would make it from scratch? There are benefits, and at least a modicum of pride, in being able to say, 'I made this by myself'. When I make a cake by myself, I can customise my work, thus making it genuinely my own; I can learn to be proud of a creation even if it is not perfect, and I can learn what to do differently on my next endeavour. My goal for students is a discovery that while they are learning about any particular subject, they are also learning how they learn, that they learn through practice, and that this process can be as satisfying as reaching the end goal of a diploma. This process is at least partly circumvented if they have a machine do their writing for them. Whatever they pursue after college, they will be better able to meet upcoming challenges if they have worked at and through written communication.

Conclusion

Those who teach about the ancient Mediterranean, whether they have inhabited Renaissance Europe, the 21st century, or any chronological point in between, have faced challenges of medium and materiality. Scrolls gave way to codices, handwritten manuscripts gave way to printed text, and printed material is giving way to digital media. Educators now face the challenge of readily available LLMs, and the need to engage critically with texts remains. Those who opt to use such tools (or who are required by their institutions to do so) have options from which to choose. An assignment may include carefully conducted experimentation with a chatbot on a topic in ethics, for example, so that students have a sense of what such tools can and cannot do. Students can query an LLM like Claude for the interpretation of a text like the *Metamorphoses*, and then argue for or against that reading. They might ask the LLM to produce a joke (or other cultural artefact) and then discuss with their peers the artistic merits and value of the AI-generated content.

My own students' reactions to these assignments have been mixed. A few students expressed fear regarding the advent of this technology, saying that they dread the possibility of an artificial neural network group-mind, as they have seen in the *Terminator* movies. I took this as an opportunity to inform them of the realities of the existential threats of artificial general intelligence (i.e., the threat is not zero, but neither has the technology advanced that far as of 2026). Others said that they are much opposed to the use of these tools given that LLMs are most often trained on work without the consent of the artists and creators, and this was also a chance for the class to discuss the realities of LLM creation. On the more enthusiastic end of the spectrum, some communicated their willingness to use an LLM as a writing aid when they felt at a loss for ideas. While I was disappointed to know that there exists a propensity to reach for an aid or crutch at the first sign of tension, I take this as an opening to talk about the writing process in general. When and why do we end up feeling frustrated as writers? What are 'brain only' means of dealing with this? Many students expressed that they were not entirely impressed with a chatbot's abilities, and that they intend to continue their written communication with little to minimal non-human aid. I am heartened by my students' responses, and I encourage my colleagues to explore the benefits of giving assignments that allow limited usage of LLMs. At the very least, the presence of these tools is a catalyst for discussions about how learning and writing happen.

Time will tell whether LLM-influenced texts will become *de rigueur*, or if the novelty of such tools will wear off, leaving writers to rely once again on their own abilities. Whatever tools we have at hand, the need for trustworthy information remains the same, as does the need to engage critically with the communications of others. The importance of curiosity and of learning to live communally, and the desire to build more flourishing lives for ourselves and others will also remain. The ties of friendship, as Crito experienced with Socrates, persist. Our feelings of awe, curiosity, and horror, as when we read the *Metamorphoses*, endure. Our desire to laugh and feel joy also remain – and none of these things can be replicated entirely by a machine. Thus, the texts that address these feelings (from the ancient Mediterranean or elsewhere) will continue to remain in our curricula and the questions that human beings have for them will arise with every reading.

Author biography. Rebecca Huskey holds a PhD in religious studies from the University of Iowa and has taught in the Classics and Letters program at the University of Oklahoma since 2007. Her research interests include the conjunction of philosophy and literature, the history of ideas, and emergent technologies such as artificial intelligence.

Notes

1 This was the latest version of Anthropic's product in August 2024. Claude Sonnet 4.6 was released in May 2026.

2 See Alaina O'Regan's Why it's so easy to jailbreak AI chatbots, and how to fix them for an illuminating view of why developers must continually hone their alignment skills. https://ece.princeton.edu/news/why-it%E2%80%99s-so-easy-jailbreak-ai-chatbots-and-how-fix-them

3 I also had a brief interaction with Claude on the subject of performing a *taurobolium*. It refused to tell me how to perform this sacrifice, citing laws that were likely to exist where I live regarding the treatment of animals and health

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 10:43:26, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S2058631026101238

codes. It did, however, mention the inclusion of animal sacrifice in 'traditional, indigenous practices.' When I queried it further as to what particular cultures it was referencing, it did acknowledge that this is a very broad category, and gave a more detailed list which included Native America nations of the North American Plains, Pacific Islander groups, and Mesoamerican groups. It did not provide further detail about how these acts are performed.

4 Education and licensure requirements will vary in different countries. The United Kingdom, e.g., differentiates between practitioner psychologists and counsellors. Both tracks have multiple requisites that must be met.

5 While much of Kosmyna et al.'s report is technical, the authors very helpfully included an abstract and summaries of sections that are accessible to those outside of neuroscience.

6 The translations that I have used include Peter Green's 1999 edition of Juvenal, Niall Rudd's 2008 edition of the same, and Rudd's 2005 edition of Horace. Selection of translations is based partly on the text itself, and partly on the availability and affordability for students. Ideally, all students are reading from the same translation, but some find open-access, free translations online. Having different translations can also be instructive, in that it gives the class an opportunity to compare and contrast the choices the various translators make.

References

Abrams, Z. (2025) Using Generic AI Chatbots for Mental Health Support: A Dangerous Trend. APA Services, Inc, 12 March. Retrieved 3 September, 2026. Available at https://www.apaservices.org/practice/business/technology/artificial-intelligence-chatbots-therapists

Bowen, J.A. and Watson, C.E. (2024) *Teaching with AI*. Baltimore, MD: Johns Hopkins University Press.

Caulfield, M. (2019) SIFT (The Four Moves), 19 June. Retrieved 1 May, 2026. Available at https://hapgood.us/2019/06/19/sift-the-four-moves/

Crompton, D. (2011) *The World's Oldest Joke Book*. Naperville, IL: Sourcebooks Inc.

ELIZA Archaeology. (2025) ELIZA Archaeology Project, 21 August. Retrieved 1 May, 2026. Available at https://sites.google.com/view/elizaarchaeology/home

Huskey, R. (2024) Claude. claude.ai. Available at https://claude.ai

Kosmyna, N., Hauptmann, E., Yuan, Y., Situ, J., Liao, X., Beresnitzky, A.V., Braunstein, I., and Maes, P. (2025) Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task. ArXiv, abs/2506.08872.

Mayor, A. (2018) Gods and Robots: Myths, Machines, and Ancient Dreams of Technology. Princeton, NJ: Princeton University Press.

McCorduck, P. (2018) *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*. Abingdon: Taylor & Francis.

Mollick, E. (2024) *Co-Intelligence: Living and Working with AI*. New York, NY: Portfolio/Penguin.

Moore, J., Grabb, D., Agnew, W., Klyman, K., Chancellor, S., Ong, D.C. and Haber, N. (2025) Expressing stigma and inappropriate responses prevents LLMs from safely replacing mental health providers. Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency.

Morreall, J. (2009) *Comic Relief: A Comprehensive Philosophy of Humor*. Hoboken, NJ: Wiley-Blackwell.

Opentools. (2025) OpenAI Empowers U.S. College Students with Free ChatGPT Plus Access for a Limited Time. 4 April. Retrieved 1 May, 2026. Available at https://opentools.ai/news/openai-empowers-us-college-students-with-free-chatgpt-plus-access-for-a-limited-time

Ovid. (2010) *Metamorphoses*, Lombardo, S. (trans.). Indianapolis, IN: Hackett Classics.

Plato. (2002) *Five Dialogues: Euthyphro, Apology, Crito, Meno, Phaedo*, Grube, G.M.A. (trans.). Indianapolis, IN: Hackett Pub.

Schoene, A.M. and Canca, C. (2025) 'For Argument's Sake, Show Me How to Harm Myself!': Jailbreaking LLMs in Suicide and Self-Harm Contexts. Institute for Experiential AI. https://doi.org/10.48550/arXiv.2507.02990

Shane, J. (2019) *You Look Like a Thing and I Love You: How Artificial Intelligence Works and Why It's Making the World a Weirder Place*. New York, NY: Little, Brown, and Company.

Shroff, L. (2025) ChatGPT Gave Instructions for Murder, Self-Mutilation, and Devil Worship. The Atlantic, 14 July. Retrieved 1 August 2025. Available at https://www.theatlantic.com/technology/archive/2025/07/chatgpt-ai-self-mutilation-satanism/683649/

Stanford University. (n.d.) Human Centered AI, Artificial Intelligence Glossary. Available at https://hai.stanford.edu/ai-definitions/what-are-hallucinations

Weizenbaum, J. (1966) ELIZA—A computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1), 36–45. https://doi.org/10.1145/365153.365168

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 10:43:26, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S2058631026101238