The Journal of Classics Teaching (2026), 1–8
doi:10.1017/S2058631026101214

CAMBRIDGE UNIVERSITY PRESS

Research Article

An experimental creative writing course: rewriting classical myths with ChatGPT

Valentin Decloquement
Université Lumière Lyon 2 - Research Unit HiSoMA, France

Abstract

This article presents a creative writing course incorporating AI, taught at Université Lumière Lyon 2 (France) in September–December 2024. The course combined exploration of LLMs with critical reflection on their use. It drew inspiration from *Si Rome n'avait pas chuté* by Raphaël Doan (2023), written with ChatGPT 3 and 3.5, which imagines a world where industrialisation began under Nero. Students used ChatGPT to write counterfactual stories or fan fictions: these literary genres are especially useful for creating a three-way relationship between the human writer, the LLM, and an external referent that drives the interaction between the two. The pedagogical progression was structured around three phases: idea generation, narrative structuration, and writing. I generated a prompt live in class: 'What if Medea had killed Jason instead of her children?' The students conceived extended chains of prompts to guide the model's responses and refine narrative variations. Some engaged with contemporary novels, others reimagined Graeco-Roman myths, which will be the scope of this article. Mythological rewritings prove particularly rich for AI-assisted storytelling, revitalising courses on classical myths in a playful, dialectical, and interactive manner. They raise a wide range of ethical issues, especially regarding the sources from which LLMs draw and the content policies that lead to the suppression of certain aspects of ancient literature. Asking ChatGPT to rework classical myths also reveals its tendency to homogenise narratives. This highlights the ongoing need for human oversight: when approached critically, generative AI can provide an effective introduction to classical mythology.

Keywords: creative writing; ChatGPT; classical mythology; prompt engineering; dialectic

Introduction

In the field of Classics, there has been a growing interest in the use of generative AI as a tool for enhancing ancient language acquisition. Some teachers employ available systems, such as OpenAI's ChatGPT (Peddar 2025; Ross 2023), which has made significant progress in recent years, especially in ancient Greek (Ross and Baines 2024, pp. 182–183). Others have even fine-tuned generative pretrained transformer (GPT) to develop Latin-speaking chatbots, for example, GPT-2 Latin in Switzerland, as a means of revitalising active learning methods (LatinIA 2025).

This article explores another potential application of Large Language Models (LLMs) in the field of Classics: creative writing inspired by Graeco-Roman mythology. For the past two years, scholars and teachers from many countries and cultural backgrounds have highlighted the pedagogical potential of AI-assisted storytelling both in secondary schools and higher education (Bariqoh 2025; Creely and Blannin 2024; Kabeer et al. 2025; Shultz Colby 2025; Tsao and Nogues 2024). My aim is to show that mythological rewritings are particularly fruitful, as they raise methodological and ethical questions about the versions retrieved and generated by LLMs. I present here a pedagogical experiment intended for non-classicist students from September to December 2024, hoping that the teaching protocol outlined here could be adapted and refined within a Classics curriculum. Besides, the fact that it was effective with non-specialist students demonstrates the potential of AI as a tool for the dissemination of classical culture.

I am myself a classicist whose research concerns the intellectual environment of the Greek sophists under the Roman Empire (1st–4th centuries), with particular attention to their rewritings of Homeric poetry. This background has made me acquainted with creative writing courses: since 2015, I have taught ancient preliminary rhetorical exercises (*progymnasmata*) in practice, adopting the persona of ancient sophists within the modern classroom. This approach aligns with a pedagogical movement that has recently gained ground in high schools and universities (Chiron and Sans 2020). I have adapted this framework to the context of an AI-assisted creative writing course.

The targeted audience was second-year undergraduate students enrolled in an 'Applied Literary Studies' programme ('Lettres appliquées', Department of Literary Studies, Lumière Lyon 2 University, France). It prepares them for careers in writing and communication, based on Humanities, including Digital Humanities. Creative writing courses are offered each semester, with an emphasis on fictional forms (narrative, description, dialogue, poetry, rewriting, and more). By the time they enrolled in my course, the students (22 in total) had already completed two or more semesters of 'traditional' courses. They were encouraged to

Email: valentin.decloquement@univ-lyon2.fr
Cite this article: Decloquement V (2026). An experimental creative writing course: rewriting classical myths with ChatGPT. *The Journal of Classics Teaching*, 1–8. https://doi.org/10.1017/S2058631026101214

© 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.

https://doi.org/10.1017/S2058631026101214 Published online by Cambridge University Press

draw on the knowledge and skills they had developed in other contexts and to adapt them within this new, experimental framework, I tried to explore with them. Even though the students who took part in this course have no background in Latin or Greek, their programme includes an introduction to ancient literature and civilisation (in the second semester of the first year), as well as an optional lecture course entitled 'Ancient Myths and Western Culture', offered in the first semester of the second year. I chose to use Greco-Roman myths as a source material for the course, thereby complementing other teachings.

Students were encouraged to write a short novel *with* AI, rather than delegating the act of writing to AI. The course had several objectives:

1. Exploring the potential of LLMs for creative writing.
2. Fostering critical distance from these tools, which are far from neutral. Since the widespread adoption of ChatGPT in 2023, many students are using AI tools for various purposes (Nugroho et al. 2025, pp. 502–506). They therefore need a methodological framework to engage with these digital tools.
3. Preparing the students for their new professional world, knowing that their field, among many others, has been particularly impacted by the growing presence of generative AIs in the job market (Shah 2023, pp. 31–39; Poth 2024, pp. 117–128).

This approach sought to strike a balance between divergent views on the pedagogical effectiveness of generative AI, which have emerged since the launch of ChatGPT in November 2022 (Imran and Almusharraf 2023). While some scholars and teachers argue that ChatGPT diminishes learning outcomes, others identify a clear added value, and still others find no significant difference between AI-assisted courses and traditional methods. For an up-to-date bibliography on this debate, one may turn to the recent meta-analysis by Wang and Fan (2025), which surveys more than 2,700 articles: the study ultimately argues for the positive impact of ChatGPT on education, though without addressing the ethical issues raised by generative AI. In our Department of Literary Studies, we believe that teachers have a responsibility to help students develop an informed, critical, and ethical use of generative AI (Shah 2023, pp. 181–197), building on similar initiatives in other universities (Ross 2023, pp. 158–160; Ross and Baines 2024, pp. 183–185; Saunders et al. 2024).

Due to budget constraints, we relied on the free version of ChatGPT (GPT-4o), which offered only limited functionality. Despite this limitation, one positive outcome can be emphasised: students were encouraged to (re-)engage with the model they themself use most frequently. The original output was produced in French; in this article, the quoted examples have been translated by ChatGPT itself (GPT-5).

An interactive course: what if Medea had killed Jason?

The literary genre chosen for this course was alternative story, which involves imagining the consequences of a hypothetical event that diverges from the established course of history or a fictional narrative. The value of this genre lies in fostering a triangular relationship between human intelligence, ChatGPT, and an external reference point, which forms the basis of a dialogue between a student and the LLM. The teaching protocol was informed by a French work of uchronia, *If Rome Had Not Fallen* (*Si Rome n'avait pas chuté*), written by the classicist Raphaël Doan (2023) with the assistance of GPT-3 and GPT-3.5. The book imagines a scenario in which the Industrial Revolution takes place under the Roman Empire in the 1st century AD. What makes this book particularly interesting is its elaboration of a dialectical relationship between human and machine. It opens with two prefaces, one generated by ChatGPT, the other written by the human author, who explains his protocol. In the main body of the text, the AI-generated and AI-illustrated fiction is systematically examined by the author through a series of 'Historian's Commentaries', each supported by authentic sources.

Students were given complete freedom in choosing their subject, on condition that it was restricted to a clearly defined topic, and they conformed to one of three narrative models:

- Historical uchronia (e.g., 'What if the Chernobyl disaster had worsened?', 'What if Victor Hugo had invented rap?').
- Fanfiction (e.g., 'What if Sherlock Holmes was a mental illusion of Dr. Watson, suffering from schizophrenia?', 'What if Harry Potter had joined Slytherin rather than Gryffindor?').
- Mythological rewriting, which will be our focus here.

After three months of work, each student produced a story (maximum 7,000 words).

Of the 22 students, 5 explored alternative versions of ancient Greek mythology. Four of them drew inspiration from feminist reinterpretations of myth, which have flourished since Margaret Atwood's *Penelopiad* (2005). Madeline Miller's *The Song of Achilles* (2011) and *Circe* (2018) are especially popular among students, many of whom first encountered the *Iliad* and the *Odyssey* through such novels.

Their projects gave rise to scenarios such as: 'What if Persephone had escaped from Hades?', 'What if Odysseus had never returned to Ithaca, leaving Penelope alone?', 'What if Ariadne had taken revenge on Theseus?', and 'What if Circe had never been exiled?'. A fifth student turned instead to the Judgement of Paris, reimagining the episode with Paris awarding the prize to Athena rather than Aphrodite.

The semester began with a general introduction, addressing the ethical issues raised by generative AI, such as copyright problems, gender stereotypes, racial bias, environmental impact, fabricated references, and hallucinations (Furze 2024, pp. 19–37; Mizrahi 2024, pp. 268–272; Poth 2024, pp. 17–46). The course then progressed through three sequences:

1. Finding Ideas (two 2-hour sessions)
2. Structuring the Story (two 2-hour sessions)
3. Writing the Text (five 2-hour sessions)

While this structure aligns with modern theories of generative active learning applied to LLM (Pratschke 2024, pp. 66–71), my primary inspiration came from ancient rhetoric (*inventio*, *dispositio*, *elocutio*).

The example I used in class was a mythological rewriting of my own invention: 'What if Medea had killed Jason instead of her own children'? Rather than providing students with readymade outcomes, I demonstrated the protocol by projecting ChatGPT dialogue boxes onto the screen and entering prompts in real time. We then took the time to analyse the results collectively. Most sessions were divided into two parts. In the first part, we worked collectively on this uchronian Medea, which served as a framework. In the second part, students were required to apply the same protocol to their own subject. Such a task, far from being mechanical, requires initiative: it involves constantly readapting the process to their individual case, and even experimenting with other approaches. We thus engaged in

https://doi.org/10.1017/S2058631026101214 Published online by Cambridge University Press

a 'digital dialectic game', as Rebekah Shultz Colby put it: a playful interaction between (collective) human intelligence and generative AI, consisting of 'questioning the algorithm to find its ethical limits and critiquing its writing but then querying it further until the writing improves' (Shultz Colby 2025, p. 3).

Before each session, I had tested a full sequence of prompts to anticipate the kind of texts ChatGPT would generate. The process was thus a form of semi-improvisation (Furze 2024, pp. 115–119). The outputs generated during class were not radically different from those produced during my preliminary tests, though they were not identical either: the predictive nature of generative AI models makes them relatively predictable. Fortunately, students are less predictable than AI: some often asked, out of curiosity, what might happen if a prompt were altered or an idea added. I always welcomed these interventions, particularly as I did not know where the experiment would lead.

The final session introduced AI-based image generation, in which students illustrated their own narratives using DALL·E (OpenAI), and Leonardo.AI. This exercise provided an opportunity to examine the biases of generative systems, which frequently reproduce visual conventions derived from Western pop culture. Such tendencies are especially evident in representations of mythological subjects (Peddar 2025, p. 11), as highlighted by the online exhibition *Distorted History* (Ure Museum 2024).

Checking the 'general knowledge' of the AI

Before engaging in the creative process proper, we conducted a critical analysis of ChatGPT's knowledge base on Greek and Roman mythology. Such an enquiry cannot be carried out fully, since the dataset on which the model was trained is not available to users (Furze 2024, pp. 11–13). However, when restricted to a carefully chosen example, this helps drawing students' attention to OpenAI's lack of transparency in this respect. It also serves as an introduction to the factual prompt method (Mizrahi 2024, p. 12), considering that formulating a prompt carefully and fact-checking the results are skills in their own that need learning (Bowen and Watson 2024, pp. 48–61; Furze 2024, pp. 13–18; Poth 2024, pp. 74–82; Pratschke 2024, pp. 45–48; Shah 2023, pp. 23–29).

To that end, I first delivered a 'human' university lecture on the myth of Medea, without relying on ChatGPT. The presentation was supported by an entry from a French mythology dictionary (Collognat 2012, pp. 581–585), which includes translations of ancient sources (Euripides, *Medea* 1242–1250; Apollonius of Rhodes, *Argonautica* 3.270–298; Seneca, *Medea* 447–489). In Euripides' version, Jason abandons Medea for the Corinthian princess Glauce: I pointed out that she is named Creusa in later versions. I explained that the episode of infanticide, central to Euripides' plot, is absent from earlier literary sources. Instead, Hesiod and Pindar highlight Medea's role as a helper in Jason's quest for the Golden Fleece.

Following this overview, I requested a summary of the myth of Medea from ChatGPT and asked the students to analyse it. They noticed that it standardises the myth by relying on its most canonical versions (Euripides for the infanticide and Apollonius of Rhodes for Jason's quest), thereby omitting the alternative versions that circulated in Antiquity. This outcome is largely attributable to the phrasing of my own prompt: I requested a 'summary' of the myth, without explicitly asking the LLM to specify its sources. In this sense, ChatGPT fulfilled the task it was given, even mimicking the approach of many mythology websites, which frequently present a simplified and homogenised account of a multifaceted tradition.

I then refined the prompt: 'Tell the myth of Medea in detail in an academic style'. The response took the form of an extended version of the earlier summary, structured into five sections: (1) Context and Meeting with Jason; (2) Exile and Jason's Betrayal; (3) Medea's Revenge; (4) The End of Medea; and (5) Themes and Interpretation. In order to align with the conventions of academic writing, ChatGPT acknowledged the existence of multiple versions and highlighted the variation in the name of Jason's new wife.

However, the archaic versions of the myth by Hesiod and Pindar remained absent. How can this be explained? Spontaneously, the students agreed on a hypothesis: these lesser-known details were perhaps underrepresented in the model's training data. Yet a subsequent prompt ('What are the different versions of the myth of Medea in ancient literary sources?') suggested a more complex reality. In response, ChatGPT summarised, in chronological order, the versions by Hesiod, Pindar, Euripides, Apollonius of Rhodes, Ovid, Seneca, Diodorus Siculus, Hyginus, and Valerius Flaccus.

This reveals a vicious circle. To obtain a response meeting academic standards, one must (1) explicitly formulate those standards, and (2) already possess disciplinary expertise. Someone unfamiliar with the Medea myth and not trained to compare divergent sources would receive from ChatGPT only a standardised version, arguably inferior to the corresponding Wikipedia entry. An ethical limitation compounds this: why use an energy-intensive tool when searching 'Medea' on Google yields often better-sourced content?

At this stage, students were required to take initiative. For example, the student who recast Penelope as the true heroine of the *Odyssey* first asked ChatGPT for a general summary of the epic, then for a book-by-book academic account, and finally for a selection of the passages that focus on Penelope. For her part, the student who reimagined the myth of Persephone used ChatGPT to identify the principal ancient sources, following my guidelines. In doing so, she encountered works unfamiliar to non-specialists, such as the *Homeric Hymn to Demeter* and Claudian's *De raptu Proserpinae*. In a subsequent prompt, she asked the model to highlight the major divergences among these accounts. The output was more sophisticated than the corresponding Wikipedia entry on Persephone, contrasting, for instance, symbolic and rationalising readings (as in Diodorus Siculus), or Greek religious interpretations with the theatrical adaptations of the myth in Latin literature.

Finding ideas

To optimise results in the context of writing, it is advisable to provide the LLM with a complete contextual framework: specifying its role, the task to be performed, and the target audience (Mizrahi 2024, pp. 14–19). This principle of prompt engineering has already proved effective in Classics classrooms, as it encourages students to synthesise their ideas, whether in translation, grammatical analysis, or image generation (Díaz-Sánchez and Chapinal-Heras 2024, p. 19; Peddar 2025, pp. 10–11). Putting the AI into a position of confident authority enhances the relevance of its output, which was experimented in the book *Si Rome n'avait pas chuté* (Doan 2023, pp. 21–22). It is more effective, for example, to write 'You are the greatest specialist in the history of Athenian democracy in Antiquity' than the more neutral 'You are a historian of Athenian democracy'. When applied to literary fiction, this standard prompting strategy involves determining in advance the genre in which the story will be situated (Mizrahi 2024, pp. 99–103), for example: 'You are the greatest Romantic prose-writer of the

https://doi.org/10.1017/S2058631026101214 Published online by Cambridge University Press

19th century. You write fantastic short stories aimed at a literate audience'. Each student was required to define a complete context by constructing a persona for their AI author-narrator and by reflecting on the literary genre they wanted to explore.

As a preliminary exercise, we gave specific personas to ChatGPT, such as: 'You are a Romantic poet of the 19th century. Summarise the myth of Medea (250 words maximum)'. This was followed by a phase of collective analysis, in which students identified the stylistic devices used by ChatGPT, in that instance apostrophes (e.g., 'Oh dark star, Medea, daughter of winds and cursed dreams!'), rhetorical questions (e.g., 'What remained for Medea but pain and anger?'), and metaphors (e.g., 'With a spectre's laugh, she plotted her revenge'). I used this as an opportunity to adopt a more playful approach, deliberately assigning ChatGPT incongruous personas. AI can indeed support gamification (Shah 2023, pp. 169–171), and role-playing is one possible approach (Bowen and Watson 2024, pp. 172–175). For instance, with the prompt 'You are a sports commentator. Summarise the myth of Medea in a maximum of 250 words', ChatGPT opened with a direct address to a sporting audience ('Ladies and gentlemen, what a twist in this mythological final!'), and rendered the summary in a more dynamic tone (e.g., 'But wait, the game isn't over yet. Jason and Medea settle in Corinth, and everything seems to be going well . . . until the plot twist!'). The experiment then moved into a collaborative phase. We tested a wide range of personas based on the students' ideas, including a mythologist, a zoologist, a comedian, Batman's Joker, Donald Trump, a medieval knight, and Mario from Nintendo. For each generated response, we took the time to analyse the language register, stylistic choices, and the aspects of the myth that ChatGPT emphasised depending on the role it had been assigned. After this, following a class vote on the narrative approach for the Medea story, most students chose the role of a young adult author for my story.

We then moved to a brainstorming phase. I sought to challenge the common tendency to accept AI's first response without question. Prompting the model to generate multiple alternatives on the same theme often proves more productive, allowing selection of the most relevant outcome: 'You are a best-selling young adult fantasy author. You are writing a new uchronian short story in which Medea kills Jason instead of her children. Propose five different step-by-step narrative scenarios for this story'. Below are the titles and summaries ChatGPT generated in response to the above example:

1. *Flight and the Shadow of the Past* (Drama & Journey): After killing Jason to prevent him from abandoning her, Medea flees with her children. But the past haunts her, and her sons eventually uncover the truth.
2. *The Throne of Blood* (Politics & Tragedy): Following Jason's murder, Medea attempts to seize power in Corinth. However, both gods and mortals refuse to tolerate a woman who has committed such a crime.
3. *The Island of Exiles* (Adventure & Rebirth): Medea escapes after killing Jason and founds a secret colony – a refuge for outcasts and exiled witches, where she attempts to forge a new identity.
4. *The Apotheosis of Medea* (Magic & Divine Vengeance): Jason's death triggers divine retribution. Medea sets out on a perilous quest to attain immortality and escape the wrath of the gods.
5. *The Pact of Shadows* (Mystery & Manipulation): Medea kills Jason – but he returns, a spectral figure who torments her conscience and manipulates her sons against her.

It is instructive to repeat the exercise while altering the role assigned to ChatGPT, for instance, casting it as an 'ancient poet' or a 'feminist author', following suggestions from the students. After generating around 15 scenarios from various contexts, I held a vote, asking the students to select their preferred version for the uchronian Medea.

Formulating a concise yet comprehensive prompt proved considerably more challenging than selecting the most promising idea. Some students struggled to define a precise authorial role ('You are a novelist. What if Paris had chosen Athena?'), whilst others preferred to follow my example directly (e.g., a young adult fantasy novel on Ariadne, aimed at an 18–25 readership). By contrast, several students explored the tension between their topic and the assigned authorial persona: for instance, imagining a 19th-century Russian novelist rewriting the Persephone myth. One student, focusing on Circe, even replicated our classroom brainstorming exercise, assigning ChatGPT contrasting roles ('humourist', 'novelist', 'medical doctor', 'rapper') before ultimately selecting the voice of a rap artist.

Structuring the story

Once they had chosen a scenario from the range of options generated by ChatGPT, the students had to develop a complete narrative structure. The constraint I introduced was to construct a story in seven parts: exposition, inciting incident, three episodes, resolution, final situation. This task provides an opportunity to revisit the fundamentals of narratology, combining literary theory, critical analysis, and creative writing practice. For instance, for each of the five scenarios quoted above, ChatGPT provided a detailed narrative outline. The example below corresponds to the first scenario, *Flight and the Shadow of the Past*:

> Pitch: After killing Jason to prevent him from abandoning her, Medea flees with her children. But the past haunts her, and her own sons discover the truth . . .
>
> Act I: Betrayal and Murder – Jason announces his intention to marry Creusa. Furious and betrayed, Medea kills him with a ritual dagger and sets fire to the palace of Corinth.
>
> Act II: Exile and Doubt – Medea flees with her children, seeking refuge far from Greece. As her sons grow older, they begin to ask questions. Can they really trust her?
>
> Act III: Revelation and Rupture – One of her sons uncovers the truth about Jason's death. Horrified, he rejects her and alerts her enemies.
>
> Climax: The Hunt and Judgment – Pursued by Jason's former allies, Medea must decide whether to flee again or confront her past. Her son stands against her – the ultimate betrayal.
>
> Denouement: The Cursed Sorceress – Alone and hunted, Medea exiles herself once more, doomed to become a wandering spectre in legend.

In this instance, ChatGPT employed theatrical conventions, possibly influenced by Euripides and Seneca, more precisely by Western views on ancient tragedies, as they did not comply with the division in Acts as in classical French theatre. This, however, aligns well with the conventions of dark fantasy, which frequently draws on elements of dramatic genres. From this perspective,

https://doi.org/10.1017/S2058631026101214 Published online by Cambridge University Press

LLMs are a valuable tool for the critical examination of narrative tropes.

For a plan, the simplest is to formulate the following prompt: 'You are X. In order to write a short story, you must outline the main stages of the following scenario: <paste the chosen scenario>. Your plan must include an exposition, an inciting incident, three episodes, a resolution, and a final situation'. The authorial persona defined at the outset has an influence on the outcome. In the rap-inspired version of Circe's story, for example, the opening episode, entitled 'Hermes' flow', staged a poetic confrontation between Hermes and Circe, reminiscent of both the tradition of the Greek *agon* and the battle rap.

After this, I asked students to critique the initial output, then prompted ChatGPT to do likewise. With the prompt 'Provide feedback in the form of a numbered bullet list', the LLM will identify the strengths and limitations of its own proposal (Mizrahi 2024, pp. 106–109), in the same way that it does with human writing (Bowen and Watson 2024, pp. 162–168). In the classroom example, ChatGPT highlighted, amongst the strengths, the dramatic tension generated by the children's trajectory from innocence to hatred, which aligns with young adult fantasy conventions. Regarding areas for improvement, the LLM noted that the magical dimension, central to Medea's mythological identity, remained underdeveloped. Additionally, ChatGPT recommended maintaining ambiguity around Jason's death, as a means of 'exploring themes of perception and manipulation'. The aim was to allow a direct comparison between the student's own critique and the model's self-assessment, within a dialectical framework that has already proved effective in other pedagogical contexts, such as Latin prose composition (Peddar 2025, pp. 9–11). The students raised objections that ChatGPT had overlooked (e.g., the Hollywood-like scenario), while, conversely, the model proposed areas for improvement that had not been identified by anyone (e.g., Jason's mysterious death). This task requires a critical and reflective engagement with the feedback:

1. Select the most relevant points (hence the usefulness of numbered lists). For example: 'Revise the previous plan by modifying it on the basis of improvement areas 1, 2, and 5'.
2. Introduce personal ideas directly into the prompt, which may significantly reshape the entire narrative framework. For example, 'Modify the initial situation: Jason is still alive and will be killed in peripeteia no. 1'.

This dialectical approach proved productive, as it enabled students to refine the authorial role, particularly for those who had not provided detailed context at the outset. For example, the student who initially asked ChatGPT to rewrite the Judgement of Paris in the voice of a 'novelist', without further specification, restructured the entire framework. Through successive prompts, they developed counterfactual account of the Trojan War, presented as Priam's personal journal, which is reminiscent of Dictys Cretensis and Dares Phrygius. In other cases, the method helped to sharpen the distinction between author and narrator. The student who had first cast ChatGPT as a 19th-century Russian novelist retelling the myth of Persephone subsequently asked it to adopt Persephone's own perspective when shaping the final scenario.

This shift allowed her to modify some narrative steps entirely: whereas ChatGPT had produced a happy ending in which Persephone became a celebrated queen of the Underworld, the student opted for a more tragic resolution, portraying her as a captive bound to her husband.

At this stage of the process, students were required to submit their first assignment, which was assessed using the following grid:

Finding ideas	**/ 5**
• Definition of the role	/ 1
• Delimitation of a literary genre	/ 1
• Definition of the target audience	/ 1
• Generation of multiple scenarios	/ 1
• Selection of a scenario	/ 1
Structuring the narrative	**/ 5**
• Delimitation of narrative stages	/ 1
• Feedback	/ 2
• Personal intervention	/ 2
Expression	**/ 5**
• Syntax	/ 1
• Spelling	/ 1
• Clarity of instructions	/ 1
• Absence of polite formulas	/ 1
• Emphasis	/ 1
Synthesis	**/ 5**
• <10 prompts	/ 1
• 1 prompt = 1 step	/ 1
• Absence of superfluous prompts	/ 1
• Absence of superfluous information	/ 1
• Positive phrasing	/ 1
Total	**/ 20**

Writing the text

At this juncture in the process, each student possessed a complete narrative framework, divided into seven sections. The final stage, the most repetitive, entails composing each narrative section individually before consolidating them within a single document to produce the uchronian short story. This involves defining the tone, vocabulary, and style of the text (Mizrahi 2024, pp. 21–24).

As a preliminary exercise, we engaged collectively in producing pastiches, prompting ChatGPT to imitate the style of a well-known author, following a protocol closely resembling step 1 (roles). Then, the students had to reuse the coaching methods studied for the roles ('You shall employ a pathetic style that will cause your reader to shudder and numerous metaphors that will establish a sinister atmosphere'). It was also recommended that students quantify their approach ('Your text shall comprise 50% descriptive passages, 25% dialogue, and 25% narrative sections'; 'You shall incorporate 10 metaphors, 7 similes, and 3 antitheses'). Subsequently, one may test alternative proportions for the same extract and select the most satisfactory result. The feedback method employed for structuring the story remains an effective instrument for enhancing the text.

Encouraging students to define their own stylistic framework led to the production of short narratives markedly distinct from the standardised outputs typically generated by ChatGPT. The student who cast Penelope as the heroine of the *Odyssey* asked the model to adopt a Homeric register, complete with an invocation of the Muses at

https://doi.org/10.1017/S2058631026101214 Published online by Cambridge University Press

the opening: 'O Muses, daughters of Mnemosyne, guide my pen towards the tale of Penelope, the unbowed, the faithful wife of far-seeing Odysseus'. Similarly, the student who reimagined the Judgement of Paris from Priam's perspective requested a bombastic and pathos-laden style, worthy of the epic tradition yet blended with the intimacy of a personal journal: 'I, Priam, king of Troy, bear witness to the splendours and sorrows of a past now gone. What I record here are not merely memories, but testimonies I feel compelled to share, so that the choices my son once made may be remembered – choices that still echo in the hearts of mortals and immortals, shaping the destiny of Troy'. By contrast, the rap version of Circe's story had a completely different tone: 'Yo, Hermes, like a whisper he slips through,/In the shadows, his steps are smooth'. Each stage of this narrative was punctuated by a recurring chorus, repeated for emphasis.

Only at this point did I permit them to intervene personally and manually in the texts generated by ChatGPT. This final exercise constitutes corrective work. At the macro-structural level, one could not simply copy and paste the results into a word-processing file to obtain a coherent product: since each narrative stage was generated individually, there were inevitably flaws, repetitions, or minor inconsistencies between the conclusion of one stage and the commencement of another. For instance, in the example we generated in classroom, peripeteia No. 1 concluded with the following sentence:

> Medea whispered, tears in her eyes, her voice laden with hatred and pain: "Jason . . . May your choice be made. You will soon discover that the woman you betrayed is worth more than a thousand princesses with silver laughter."

Peripeteia No. 2, conversely, begins with this paragraph:

> The twilight wove its grey shadow over the palace of Corinth, stifling the gleaming marble in a shroud of silence. Medea walked through the gardens, each step an ominous echo of her past with Jason, who awaited her, unperturbed, with an almost suffocating calm.

The initial task, therefore, involved reworking the transitions between each section. To encourage student vigilance, I imposed the constraint of producing a continuous narrative, without ellipses or chapter breaks, whilst granting them freedom in their approach:

1. Manually composing the transition. One must ensure the maintenance of stylistic coherence.
2. Requesting ChatGPT, in a separate dialogue, to generate a transition. In this case, the context must be re-established in the prompt. A new dialectical process within a fresh prompt sequence may then be initiated (feedback, stylistic modification, etc.).

In practical terms, the second method, demonstrated directly during class, may yield this type of result:

> . . . with silver laughter." The echo of Medea's words resonated in the hall, mingling with the faint breath of the wind that slipped through the tall columns. Jason, frozen in an immobility he wanted to be impassive, yet felt a shiver run down his spine. She had left, carrying with her a rage he had underestimated.
>
> The silence grew heavier, as if the walls themselves held their breath in the face of the storm brewing in the sorceress's soul. Hours stretched on, and with them, the tension in the palace's air.
>
> Medea, in her apartments, prepared her departure with methodical care, her hands trembling with a contained fury. But at the moment of crossing the threshold to leave this place that had never become a home, a cold and sharp thought crept in: fleeing would not be enough. Jason had to pay. Not just for his infidelity, but for every broken promise, every stolen moment of her life.
>
> When Medea finally left her apartments, night had given way to a hesitant twilight, an in-between where shadow stretched, heralding the impending drama. She walked slowly, almost in a trance, to the gardens, where she knew he awaited her.
>
> The wind had calmed, as if the world itself held its breath. Each step she took echoed down the rose-lined path, the red petals quivering in the dying light, and she drew closer to the man who had made her a lover, a mother, and soon, a murderer.
>
> The twilight wove its grey shadow . . .

The final task was to harmonise the result by tracking down heaviness and artificial semantic associations on the stylistic level, as well as the last inconsistencies on the narrative level (e.g., a character leaves a room but is still in that room in the next scene).

At the end of the semester, students submitted their final projects, which were evaluated using the following assessment grid:

Prompt chains	**/ 5**
• Sequencing of prompt chains	/ 3
• Consistency of prompt chains history	/ 2
Individual dialogue boxes	**/ 5**
• Consistency of context (role, literary genre, style)	/ 3
• Feedback and personal intervention	/ 2
Narrative consistency of the output	**/ 5**
• Single, coherent title	/ 1
• Continuity between narrative steps	/ 2
• Consistency of the fictional world	/ 2
Style of the output	**/ 5**
• Spelling	/ 1
• Syntax	/ 1
• Absence of clumsy phrasing	/ 1
• Absence of repetitions	/ 1
• Absence of artificial semantic associations	/ 1
Total	**/ 20**

Content policy and the limits of fiction

The fictional mode often comes into conflict with ChatGPT's content filters. This applies to depictions of violence (murder, torture, suicide) and sexual content, among other themes, even

https://doi.org/10.1017/S2058631026101214 Published online by Cambridge University Press

when they are implicit, fictionalised, or expressed in poetic form. Until early 2025, OpenAI displayed an orange error message: 'This content may violate our terms of use or usage policies'. When this dialogue box appeared within a prompt sequence, one could continue generating texts from one's personal account. However, it became impossible to share the dialogue in hyperlink form, which created logistical difficulties for a teacher wishing to evaluate the process. The orange warning box has been removed from ChatGPT's interface since early 2025. The difference now is that, if ChatGPT refuses or reformulates a response, it does so more discreetly: the underlying content moderation remains in place (Mauran 2025).

Given the prevalence of 'problematic' motifs in ancient Greek and Latin literature, those are likely to be flagged as problematic (Díaz-Sánchez and Chapinal-Heras 2024, p. 19; Ross and Baines 2024, pp. 184–185). In my course, mythological narratives often triggered the system's orange warning message, for example, Medea's infanticide in the traditional version, as well as her killing of Jason in my own counterfactual scenario, even though I had not requested any explicit description of the acts themselves and clearly said in my prompt that I was writing a short novel. ChatGPT failed to distinguish between a discourse advocating infanticide and a line from Medea in a fictional narrative who resolves to kill her children. A similar issue arose for the student working on Persephone, since the myth involves sexual violence, although she deliberately avoided focusing on that dimension. This also true concerning style: Aristophanes, Plautius, Apuleius, the Marquis de Sade, and many others would not pass ChatGPT's filters.

This limitation, nevertheless, has proved valuable and triggered classroom debates. Witnessing ChatGPT censoring itself provoked indignant reactions against political correctness, impediments to freedom of expression, and constraints upon creativity. From a deontological perspective, it enables students to recognise programmer biases, often invisible, and leads them to understand that, contrary to popular belief, AI systems are not neutral instruments. Concerning literary composition, this also provides an excellent case study for explaining the fictional pact, in contrast to ChatGPT's filters that struggle to recognise it.

Conclusions

Far from impoverishing teaching, the integration of generative AI into creative writing courses can revitalise them. Mythological material is particularly well suited to such a pedagogical approach, since myths are by nature plural, fluid, and open to reinterpretation. This makes them an ideal testing ground for AI-assisted storytelling, allowing students to cultivate a range of key competences:

1. Generative AI can be used as tool for helping students developing critical thinking skills. Prompting a LLM to reframe myths foregrounds its limitations: biases in its training data, stereotypical representations, and the tendency to homogenise contradictory versions, or censor culturally sensitive material. I noted that students are more inclined to criticise AI-generated contents than human-written texts (Peddar 2025, pp. 11–12). One explanation for this phenomenon is that AI-assisted creative writing unsettles the predominance of the author: this dialectical process allows students to exercise primary agency (Tsao and Nogues 2024, pp. 5–6).
2. Students are expected to show initiative, creativity and adaptability (Kabeer et al. 2025, pp. 809–811). Faced with the wide range of possibilities produced by both ancient myths and LLMs, they are compelled to make choices between competing versions, and to assess which narrative directions merit development and which should be set aside. AI serves as a dialogic partner that pushes them to refine their own creative decisions in response to machine-generated suggestions. It is worth noting that students who already used LLMs regularly, sometimes even daily, were at a disadvantage. The course required them to deconstruct their established habits and to re-appropriate the tool within a critical framework.
3. AI-assisted mythological narrative construction demands managing multiple dialogue threads, exploratory drafts, transitional passages, or even divergent prompt chains for the same mythological episode. This process develops organisational competences: curating and labelling prompt histories, discarding redundancies, and integrating disparate textual fragments into a consistent whole.

Yet, the integration of LLMs into creative writing courses must not be imposed without acknowledging the legitimate concerns of students who object on ethical grounds.

When we redeployed this course in 2025–2026, several students refused to engage with LLMs due to reservations about the business practices of AI companies and the environmental impact of the model. We made it clear that students were not required to participate in LLM-based exercises. The process described in this article was successfully implemented using traditional methods without any reliance on generative AI. Besides, to address these concerns, we have since adopted Mistral AI 3.1 Small, a French-developed model that consumes significantly less energy than ChatGPT (Mistral AI 2026) and offers more permissive policies for literary and artistic creation. In the future, we plan to further regulate the use of LLMs by limiting the length of prompt chains. For example, students might spend an entire session crafting a single, precise prompt, followed by a trial period limited to a maximum of three iterations.

References

Atwood, M. (2005) *The Penelopiad*. Edinburgh: Canongate.

Bariqoh, A. (2025) Integration of artificial intelligence in creative writing learning: an exploratory study of the impact on the quality of student narratives. *Journal of Language and Letters Education* 1(1), 9–18.

Bowen, J.A. and Watson, C.E. (2024) *Teaching with AI: A Practical Guide to a New Era of Human Learning*. Baltimore: Johns Hopkins University Press.

Chiron, P. and Sans, B. (2020) (eds.). *Les* Progymnasmata *en pratique de l'Antiquité à nos jours/Practicing the* Progymnasmata *from ancient times to present days*. Paris: Éditions Rue d'Ulm, presses de l'École normale supérieure.

Collognat, A. (2012) *Dictionnaire de la mythologie gréco-romaine illustrée par les récits de l'Antiquité*. Paris: Omnibus.

Creely, E. and Blannin, J. (2024) Creative partnerships with generative AI. Possibilities for education and beyond. *Thinking Skills and Creativity* 56, 1–15. https://doi.org/10.1016/j.tsc.2024.101727

Díaz-Sánchez, C. and Chapinal-Heras, D. (2024) Use of open access AI in teaching classical antiquity. A methodological proposal. *Journal of Classics Teaching* 25(49), 17–21. https://doi.org/10.1017/S2058631023000739

Doan, R. (2023) *Si Rome n'avait pas chuté*. Paris: Passés composés.

Furze, L. (2024) *Practical AI Strategies: Engaging with Generative AI in Education*. Melbourne: Ambo Press.

https://doi.org/10.1017/S2058631026101214 Published online by Cambridge University Press

Imran, M. and Almusharraf, N. (2023) Analyzing the role of ChatGPT as a writing assistant at higher education level: a systematic review of the literature. *Contemporary Educational Technology*, 15(4), 1–14. https://doi.org/10.30935/cedtech/13605

Kabeer, A., Bhat, R.A., Antony, S. and Tramboo, I.A. (2025) Enhancing creative writing skills in secondary school students through prompt engineering and artificial intelligence. *Forum for Linguistic Studies* 7(3), 800–815. https://doi.org/10.30564/fls.v7i3.8511

LatinIA (2025) Specification of the Employed Models. Available at https://latinia.help.ch/more (accessed 23 August 2025).

Mauran, C. (2025) OpenAI strips warnings from ChatGPT, but its content policy hasn't changed. Available at https://mashable.com/article/openai-chatgpt-warnings-removed (accessed 28 August 2025).

Miller, M. (2011) *The Song of Achilles*. New York: Ecco Press.

Miller, M. (2018) *Circe*. London: Bloomsbury.

Mistral AI (2026) Mistral AI Small 3.1. Available at https://mistral.ai/news/mistral-small-3-1 (accessed 19 April 2026).

Mizrahi, G. (2024) *Unlocking the Secrets of Prompt Engineering: Master the Art of Creative Language Generation to Accelerate Your Journey from Novice to Pro*. Birminghan: Packt Publishing.

Nugroho, A., Andriyanti, E., Widodo, P. and Mutiaraningrum, I. (2025) Students' appraisals post-ChatGPT use: students' narrative after using ChatGPT for writing, *Innovations in Education and Teaching International* 62(2), 499–511. https://doi.org/10.1080/14703297.2024.2319184

Peddar, D. (2025) Utilising generative AI in the classics classroom. *The Journal of Classics Teaching* 26(51), 8–13. https://doi.org/10.1017/S2058631024001363

Poth, R.D. (2024) *How to Teach AI: Weaving Strategies and Activities Into Any Content Area*. Arlington: International Society for Technology in Education.

Pratschke, B.M. (2024) *Generative AI and Education: Digital Pedagogies, Teaching Innovation and Learning Design*. Cham: Springer.

Ross, E.A.S. (2023) A new frontier: AI and ancient language pedagogy. *Journal of Classics Teaching* 24(48), 143–161. https://doi.org/10.1017/S2058631023000430

Ross, E.A.S. and Baines, J. (2024) Treading water: new data on the impact of AI ethics information sessions in classics and ancient language pedagogy. *Journal of Classics Teaching* 25(50), 181–190. https://doi.org/10.1017/S2058631024000412

Saunders, S., Coulby, C. and Lindsay R. (2024) Pedagogy and policy in a brave new world. A case study on the development of generative AI literacy at the University of Liverpool. In Beckingham, S. *et al.* (eds.), *Using Generative AI Effectively in Higher Education Sustainable and Ethical Practices for Learning Teaching and Assessment*. Abingdon: Routlegde, pp. 11–20.

Shah, P. (2023) *AI and the Future of Education: Teaching in the Age of Artificial Intelligence*. New Jersey: John Wiley & Sons.

Shultz Colby, R. (2025) Playing the digital dialectic game: writing pedagogy with generative AI. *Computers and Composition* 75, 1–12. https://doi.org/10.1016/j.compcom.2025.102915

Tsao, J. and Nogues, C. (2024) Beyond the author: artificial intelligence, creative writing and intellectual emancipation. *Poetics* 102, 1–12. https://doi.org/10.1016/j.poetic.2024.101865

Ure Museum (2024) *Distorted History: AI's skewed visions of the Ancient World*. Available at https://collections.reading.ac.uk/ure-museum/explore/online-exhibitions/distorted-history-ais-skewed-visions-of-the-ancient-world/ (accessed 26 June 2025).

Wang, J. and Fang, W. (2025) The effect of ChatGPT on students' learning performance, learning perception, and higher-order thinking: insights from a meta-analysis. *Humanities and Social Sciences Communications* 12(1), 1–21. https://doi.org/10.1057/s41599-025-04787-y

https://doi.org/10.1017/S2058631026101214 Published online by Cambridge University Press