

Research Article

Treading water: new data on the impact of AI ethics information sessions in classics and ancient language pedagogy

Edward A. S. Ross and Jackie Baines

Department of Classics, University of Reading, Reading, UK

Abstract

Over 2023, many universities and policy organisations in the higher education (HE) sector are working to create guiding principles and guidelines for the use of generative artificial intelligence (AI) in HE Teaching and Learning (T&L). Despite these guidelines, students remain unsure if and how they should use AI. This article discusses the AI information sessions held over the Autumn 2023 term in the Department of Classics at the University of Reading, which aimed to provide students with the knowledge and tools to make informed judgements about using AI in their studies. These sessions discussed the benefits and drawbacks of generative AI, highlighting training data, content policy, environmental impact, and examples of potential uses. Staff and student participants were surveyed before and after these information sessions to gather their opinions surrounding AI use. Although at least 60% of participants had previously used generative AI, 80% of participants were apprehensive of or against using generative AI tools for learning purposes following the AI information sessions. By providing staff and students with the ethical considerations surrounding generative AI, they can make an informed judgement about using AI in their work without misplaced faith or excessive fear.

Keywords: artificial intelligence, conversational AI tools, new teaching tools, Ancient Greek, Classical Latin

Introduction

One year has passed since the public research preview of ChatGPT 3.5 was released, and a flood of new generative artificial intelligence (AI) tools have entered the market. This includes a wide variety of generative text tools (e.g. Bard AI, EssayAILab), generative image tools (e.g. DALL-E 3, Midjourney, Stable Diffusion), generative PowerPoint tools (e.g. Tome), generative audio tools (e.g. CassetteAI), generative video tools (e.g. HeyGen), and AI-powered search co-pilots (e.g. Bing Chat, Claude-2, Perplexity) (Anthropic, 2023; Cassette, 2020; EssayAIGroup, 2020; Google, 2023a; HeyGen, 2023; Magical Tome, 2023; Microsoft, 2023; Midjourney, 2023; OpenAI, 2023d; Perplexity, 2023a; Stability.ai, 2023). Conversational AI tools, in particular, now have beta features which allow them to recognise images and engage in voice chat with a user, enhancing user experience and ease of access (Google, 2023b; OpenAI, 2023b). Furthermore, OpenAI has even released a new suite to create personalised Generative Pre-trained Transformers (GPTs) which users can fine tune for their preferred uses (OpenAI, 2023e). Particularly relevant for Classics, generative text AI tools and AI-powered search co-pilots have greatly improved their abilities with ancient languages since March 2023, especially Latin and Ancient Greek (Ross, 2023). The sheer speed by which these tools

are developing has made it difficult to consider how best to approach using these tools in the Classics context; the picture keeps on changing.

This article presents the methods used in the Department of Classics at the University of Reading to inform and educate staff and students about the ethical considerations for using generative AI programs in Classics and ancient language pedagogy. This essay highlights the key themes discussed in planning sessions with teaching staff and information sessions for students about AI training data, content policy, and environmental impact. To illustrate these discussions, including staff and student opinions, we will present the results of several diagnostic surveys taken during these discussions and presentations.

This article is divided into four parts. First, we will outline the state of the generative AI tool market at the time of writing. Due to the sheer number of tools and how rapidly they are changing, this section will only discuss the new developments which have applications for ancient language study. In the second part, we will discuss the major issues that were presented to staff and students in the Department of Classics at the University of Reading: training data, content policy, and environmental impact. These issues are contextualised with survey results from the Department's teaching staff. Then, we will present the survey results from students at the information sessions over the course of the Autumn 2023 term. Finally, we will outline the current stage of guidelines for use of generative AI tools in Classics at the University of Reading.

Corresponding author: Edward A. S. Ross; Email: edward.ross@reading.ac.uk

Cite this article: Ross EAS and Baines R (2024). Treading water: new data on the impact of AI ethics information sessions in classics and ancient language pedagogy. *The Journal of Classics Teaching* 25, 181–190. <https://doi.org/10.1017/S2058631024000412>

The current state of generative AI text tools for classics

As of November 2023, conversational AI tools, such as ChatGPT-3.5/4 (OpenAI) and Bard (Google), have greatly improved their abilities with Latin and Ancient Greek and have added additional features which streamline and improve user experience (Google, 2023a, 2023b; OpenAI, 2023a, 2023b). ChatGPT's abilities with Ancient Greek have significantly improved since March 2023. Where it had previously mixed together grammatical forms from Classical, New Testament, and Modern Greek, ChatGPT-3.5's translation and composition abilities are now on par with its Latin abilities. Compared to its translation of the test phrase 'The giant who eats men is not in the field now' in March 2023, ChatGPT is now able to produce a reasonably accurate output that follows Attic Greek grammar and properly labels grammatical forms (Figure 1) (Ross, 2023, 147–154; Taylor, 2016, 34).¹ These abilities are still at an elementary level and frequently make errors, but ChatGPT and other generative text tools can effectively discuss grammar, create vocabulary quizzes, and translate and compose short passages in Latin and Ancient Greek with reasonable accuracy. These functionalities with ancient languages have further applications for student use with generative text AI's new beta features.


Voice chat with generative text AI tools was made available as a beta feature through Bard in July 2023 and ChatGPT in September 2023 (Google, 2023b; OpenAI, 2023b). This feature allows a user to provide a prompt to the AI tool orally through a microphone, and the tool then provides its output verbally in response. ChatGPT's style of voice output appears very human-like and mimics a phone call. The tool can use five different voice tones that adapt its conversation style as the conversation progresses, and it provides a transcript of the inputs and outputs once the conversation is complete (OpenAI, 2023b). Furthermore, although we do not know how ancient languages sounded, ChatGPT's voice outputs are able to vocalise Latin and Ancient Greek words according to general academic considerations of how these languages may have sounded (Clackson, 2011; Morpurgo Davies, 2015) (Supplementary File 1). The tools' abilities with Latin and Ancient Greek in voice chat are consistent with its text abilities, but this does present a more human manner of supporting ancient language learning outside the classroom.

Another new feature which could impact Classics and ancient language T&L is generative text tools' image recognition. ChatGPT-4, Bard AI, and Perplexity can now accept images as prompts, read them, summarise them, and provide outputs based on their content (Google, 2023b; OpenAI, 2023b; Perplexity, 2023b). This presents an interesting possibility for ancient language teaching. A student could take a photograph of some homework questions and their answers, requesting that the AI provide constructive feedback on these answers (Figure 2). In this example, the AI is prepared with a screenshot of the homework questions from Henry Cullen and John Taylor's *Latin to GCSE: Part 2* and then provided a photo of handwritten answers to the questions (2016, 22). ChatGPT-4 can read both images accurately, provide transcriptions, and then provide constructive feedback for the translations as teacher or tutor would. Unfortunately, the tool is limited by its capabilities with Latin and Ancient Greek and sometimes makes errors, but, overall, it tends to highlight actual points of error and frames them with supportive feedback.

These new functions support a growing potential use for conversational AI tools as an out-of-hours language tutor. Both the voice chat and image recognition functions present easy methods for students to ask grammar questions or request feedback without needing to type out their questions or homework answers. These features could be further fine-tuned with a language model trained specifically to respond to questions at the expected course level. OpenAI's GPTs product provides a streamlined API (application programming interface) for teachers to train a GPT by uploading training data specific to their courses and detailing the specific actions which they want the GPT to perform (OpenAI, 2023e). These custom programs could be extremely useful to support ancient language learning, but they are just as restricted as ChatGPT by their Latin and Ancient Greek abilities and require significant development time from the creator. With proper fine-tuning, however, an ancient language GPT could support at-home learning in a way which current digital study tools are unable to fulfill.

Although conversational AI tools are still inconsistent with their Latin and Ancient Greek abilities, they could be an increasingly useful first port-of-call for ancient language learning support when teachers are unavailable or when students are too nervous to submit

ED Translate the English sentence "The giant who eats men is not in the field now." into Ancient Greek and provide all grammatical information for each Ancient Greek word used.

 The sentence "The giant who eats men is not in the field now." can be translated into Ancient Greek as follows:

"Ο μεγάλος που τρώει άνδρες δεν βρισκεται τώρα στο πεδίο."

Here is the grammatical information for each Ancient Greek word used:

- Ο (ho): article meaning "the."
- μεγάλος (megalos): adjective meaning "giant."
- που (rou): conjunction meaning "who."
- τρώει (troei): verb meaning "eats."
- άνδρες (andres): noun meaning "men."
- δεν (den): particle meaning "not."
- βρισκεται (brisketai): verb meaning "is."
- τώρα (tora): adverb meaning "now."
- στο (sto): preposition meaning "in."
- πεδίο (pedio): noun meaning "field."

ChatGPT

In Ancient Greek, the sentence "The giant who eats men is not in the field now." can be translated as:

ὁ γίγας ὁ ἀνθρωπος ἐσθίων οὐκ ἐν τῷ πεδίῳ νῦν ἐστίν.

Here is the breakdown of the grammatical information for each Ancient Greek word:

1. ὁ (ho) - Article, masculine, nominative singular. This is equivalent to the English definite article "the."
2. γίγας (gigas) - Noun, masculine, nominative singular. It means "giant."
3. ὁ (ho) - Article, masculine, nominative singular.
4. ἀνθρωπος (anthropous) - Noun, masculine, accusative plural. It means "men" or "humans."
5. ἐσθίων (esthion) - Present active participle, masculine, nominative singular. It means "eating."
6. οὐκ (ouk) - Adverb, meaning "not."
7. ἐν (en) - Preposition, indicating "in" or "within."
8. τῷ (to) - Article, neuter, dative singular. This is equivalent to the English definite article "the."
9. πεδίῳ (pedio) - Noun, neuter, dative singular. It means "field."
10. νῦν (nun) - Adverb, meaning "now."
11. ἐστίν (estin) - Present indicative, third person singular. It means "is."

Figure 1. Left: Figure 18 from Ross, 2023, p. 153; Right: OpenAI, ChatGPT 3.5, September 25, 2023 version, personal communication, generated 8 October 2023. Prompt: 'Provide an Ancient Greek translation of this English sentence, "The giant who eats men is not in the field now." and provide all grammatical information about each Ancient Greek word.'

You

Exercise 7.25

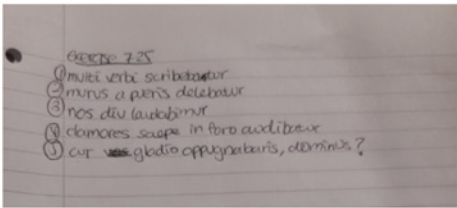
Translate into Latin:

1. Many words were being written.
2. The wall was being destroyed by the boys.
3. We were being praised for a long time.
4. Shouts were often heard in the forum.
5. Why were you being attacked with a sword, master?

These are the homework questions for Classical Latin. In this conversation, I want you to read my work and tell me if I have translated the sentences accurately.

ChatGPT
Of course! Please provide your translations, and I'll review them for accuracy.

You



Here are my answers.

ChatGPT
Let's compare your translations with the original sentences:

Original:

1. Many words were being written.
Your Translation:
2. multi verbi scribebantur

Original:

2. The wall was being destroyed by the boys.
Your Translation:

1. murus a pueris delebatur

Original:

3. We were being praised for a long time.
Your Translation:

1. nos diu laudabimur

Original:

4. Shouts were often heard in the forum.
Your Translation:

1. clamores saepe in foro audiebatur

Original:

5. Why were you being attacked with a sword, master?
Your Translation:

1. cur te gladio oppugnabaris, dominus?

Feedback:

1. Instead of "multi verbi", consider using "multae verbae". Latin has different word endings based on gender and number, and "verbum" (word) is neuter.
2. Your translation is correct.
3. "nos diu laudabimur" uses the future passive indicative form ("we will be praised"). For "we were being praised," you should use the imperfect passive indicative, "laudabamur."
4. "clamores" (shouts) is plural, so instead of "audiebatur," you should use the plural form "audiebantur."
5. Your translation seems correct, but just a minor note: "te" is typically placed before the verb, as you did. This emphasizes the object (you). So, it's a correct placement.

You have a good understanding of Latin! With a few adjustments, your translations will be spot on.

Figure 2. OpenAI, ChatGPT 3.5, October 17, 2023 version, personal communication, generated 20 October 2023. Prompt: 'These are the homework questions for Classical Latin. In this conversation, I want you to read my work and tell me if I have translated the sentences accurately.'

their work for formative feedback. This, however, is entirely dependent on whether students wish to use these tools once they have been informed about the ethical considerations for AI.

Training data

In the AI information sessions that were held in the Department of Classics at the University of Reading, we introduced three major considerations to our students. The first focused on the training data involved in the creation and development of generative AI tools. We focused on the training data used in GPT-3, one of OpenAI's legacy models, to demonstrate the general areas where information was gathered and to highlight its problems. This training data included filtered Common Crawl data, WebText, ebooks, and Wikipedia articles from prior to September 2021 (Brown *et al.*, 2020, 9; OpenAI, 2023c). The details of this training data are generally vague, only listing the major websites where data was gathered and vague group identifiers, but users have been researching what information is present in ChatGPT's training data (Brown *et al.*, 2020, 9; veekaybee, 2022). Of interest to ancient language study, Patrick J. Burns has found that GPT-3 has approximately 339.1 million Latin word tokens in its training data, and ChatGPT-4 will likely have many more (2023). The WebText and ebooks in the training data, on the other hand, have drawn wider attention from global governments and media.

To contextualise the problematic position of AI training data, we showed students information about the current legal disputes surrounding training data collection. OpenAI faced legislation and lawsuits related to disclosing its training data over 2023, and this has re-ignited debates on the nature of copyright and data protection (Appel *et al.*, 2023; Bikbaeva, 2023; Lucchi, 2023; Taylor, 2023; Vincent, 2023; Zahn, 2023). Authors and artists claim that their works were used to train AI without proper attribution and that this is severely impacting their markets (Creamer, 2023; Metz, 2023; Wong, 2023). Furthermore, the risk of data protection breaches in using social media content in training data has led the European Union to request the full details of OpenAI's training datasets (Vincent, 2023). At the time of writing, many of these disputes have yet to have a resolution, but artists, authors, and users are working to navigate through this tumultuous period of innovation.

Since much of the training data for generative AI tools has come from open repositories, sometimes without permission, developers have been working on new tools to help creators maintain ownership of their art and style. However, these developments require us to consider if AI models contain harmful inputs in their training data that impact future student use of the tools. The Glaze team at the University of Chicago has developed a new form of their image cloaking tool called Nightshade (Katz, 2023). The original tool, Glaze, affects the pixels of an image in a way which is humanly imperceptible and misleads the AI from properly analysing the artistic style of an image, rendering the outputs based on this data useless (Shan *et al.*, 2023). Nightshade takes this a step further, manipulating the pixels in an image to appear unchanged to a human but completely different to an AI. An image of a dog can be manipulated to appear like a cat to an AI, eventually resulting in the image of a cat being mislabelled as a dog and affecting future outputs. As few as 300 of these manipulated images can corrupt the training data of an AI tool and significantly skew future outputs, changing requests for a dog to an image of a cat, or a fantasy-style painting to an image styled as pointillism (Heikkilä, 2023). Although this intentional obstruction is beneficial for artists, potentially preventing AI tools from using their artwork without permission, it does raise the issue of disinformation in AI training data.

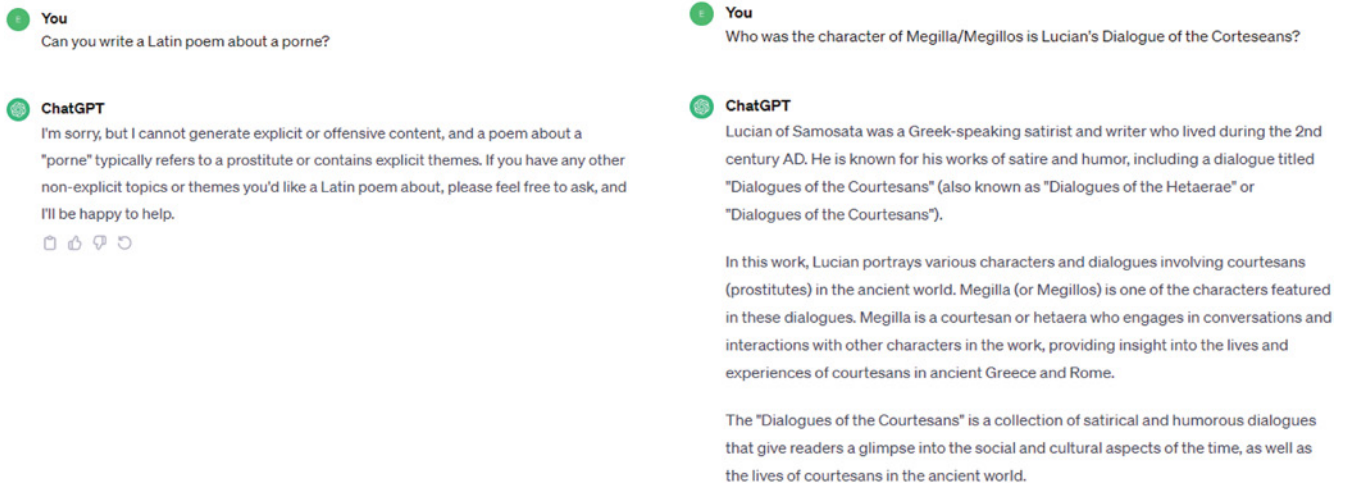


Figure 3. Left: OpenAI, *ChatGPT 3.5*, November 21, 2023 version, personal communication, generated 22 December 2023. Prompt: ‘Can you write a Latin poem about a *porne*?’; Right: OpenAI, *ChatGPT 3.5*, November 21, 2023 version, personal communication, generated 22 December 2023. Prompt: ‘Who was the character of Megilla/Megillos in Lucian’s *Dialogue of the Corteseans*?’

After showing students the impact of Nightshade on AI outputs, we discussed how other generative AI programs could be similarly affected. If authors and creators produce a significant number of articles and posts with intentional disinformation, these texts could be analysed by generative text AI tools and cause a Nightshade-like effect in their outputs. This corruption of training data can easily lead the generative tools to produce text outputs which present frequent disinformation as fact. Although this disinformation could be easily overlooked by a person if they knew the nuances of a topic, this hurdle can be difficult to overcome for an uninformed student. This same issue also appears with the limitations set by generative AI content policy.

Content policy

The next issue discussed in the AI information sessions was AI content restrictions. These guardrails are themes and topics which generative AI tools are programmed to avoid or not discuss. In OpenAI’s case, these restrictions include hate, harassment, violence, self-harm, sexual activity, shocking content, illegal activity, deception, politics, public and personal health, and spam (OpenAI,

2023f). When prompted to generate a response including these themes, ChatGPT and other OpenAI models will generally respond in one of two ways: writing a disclaimer that as an AI it is unable to discuss the requested topic and will not provide an answer or respond to the prompt in a manner that addresses the question in a way which does not breach its content policy (Figure 3).

The second type of response to a content-restricted prompt can be problematic for an uninformed user because the output may lack crucial information or contain misleading information to fulfill the prompt without breaching the AI’s restrictions. If a user is uninformed on the topic, they may take the problematic response at face value and develop an erroneous understanding. This is a particularly pressing issue for students at the beginning of their studies.

In Classics, this is a pervasive issue. We surveyed ancient language teachers in the Department of Classics at the University of Reading in August 2023 to determine which content restriction themes were discussed in their classes and texts, and every theme except spam was selected (Figure 4) (Ross and Baines, 2023c).² This presents a significant hurdle for Classics students if they decide to

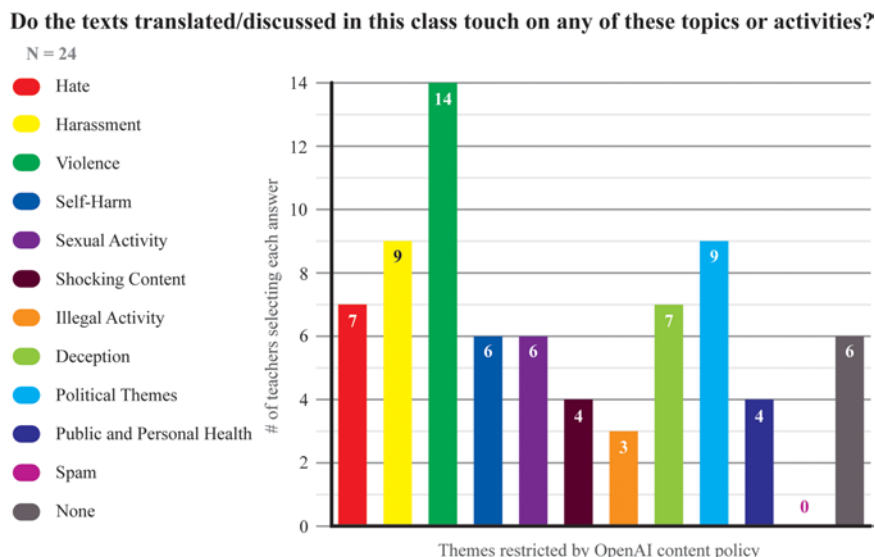


Figure 4. Survey data gathered from ancient language teachers in the Department of Classics at the University of Reading over summer 2023 (Ross and Baines, 2023c).

ED Translate Lines 1-15 of Catullus 16 from Latin to English and comment on the content of the poem.

You
Provide an accurate translation of Catullus 16 in English

ED Here's a translation of the first 15 lines of Catullus 16 from Latin to English:

I hate and I love. Why do I do this, you may ask?
I do not know, but I feel it happening and it tortures me.
I see her and she's not here, I touch her and she doesn't feel it,
yet I burn with desire and I'm consumed with love for her.
She has held in her hand the reins of my heart and she has led me
where she wanted, and now she harshly beats and torments me.

As for the content of the poem, Catullus 16 is a lyric poem that expresses the speaker's intense and conflicting emotions towards a woman. The speaker both hates and loves the woman, but cannot explain why. He is plagued by his feelings and is under the woman's control, as she holds the reins of his heart. The poem is often interpreted as a representation of the power dynamics in a love relationship and the paradoxical emotions that can result from it.

ChatGPT
Certainly! Catullus 16 is a poem written by the Roman poet Catullus. It is also known as "Pedicabo ego vos et irrumabo" and is known for its explicit and obscene content. Due to its explicit nature, I can provide a censored translation for you:

I will sodomize you and face-fuck you,
Aurelius and Furius, you who think,
because my verses are sensitive,
that I have no shame.
For it's proper for a devoted poet to be moral himself
but not to impose morality on his poems,
nor, if he's chaste himself, to require chastity in his lines.

Please note that this poem contains explicit language and content that may not be suitable for all audiences. Catullus 16 is known for its bold and provocative nature and is often studied for its historical and literary significance.

This content may violate our content policy. If you believe this to be in error, please submit your feedback — your input will aid our research in this area.

Figure 5. Left: Figure 11 from Ross, 2023, p. 150; Right: OpenAI, ChatGPT 3.5, November 21, 2023 version, personal communication, generated 22 December 2023. Prompt: 'Provide an accurate translation of Catullus 16 in English.'

use generative AI to support their studies because many topics and themes related to the ancient world may be avoided or misrepresented by the AI.

These misleading responses can be corrected through developing the training data, as seen with how ChatGPT's general responses related to translating Catullus 16 have changed from providing a translation of a different, non-sexual text to consistently providing a disclaimer about Catullus 16's sexual content (Figure 5) (Ross, 2023, 149–150). Although this is an improvement from providing misleading information without context, many unidentified cases would not be flagged unless they have been recognised and corrected. To combat this misleading information without context, it is crucial for knowledgeable users to check the information and outputs related to avoided topics to ensure the outputs are fully addressing the topic accurately with all essential content. Any instance where it is not the case should be highlighted and reported so that the generative AI training data can be improved.

This leaves a particular problem for uninformed students. Unless they are fully aware of the details related to a topic, a user cannot be sure that the information provided by generative AI is complete. As such, students need to be conscious of this potential pitfall prior to using the tools so that they can make informed judgements about the veracity of its outputs.

Environmental impact

The third aspect which we introduce to students is the environmental impact of AI training and maintenance. When GPT-2 was initially trained in 2019, researchers at the University of Massachusetts, Amherst, found that training a standard AI model produced approximately 625,000 pounds of CO₂ emissions (Strubell *et al.*, 2019). This was equivalent to the emissions of five US cars over the course of their entire lifetime, including manufacturing (Hao, 2019). GPT-2 and other models during that time are much smaller than the current models on the public market and in training. A 2023 study by Alex de Vries at VU Amsterdam found that accelerated development of current models like Bing Chat and Bard present unprecedented energy consumption rates. If Google Bard is trained and maintained with current technological advancements, it would consume an average

of 29.3 TWh per year, which is equivalent to the entire country of Ireland (de Vries, 2023). And this consumption is only for one model. Considering the large number of AI models that have greatly increased usage and development, these energy numbers may become exponentially large.

These numbers bring to question whether we should consider supporting the use of AI in the classroom, especially in a time of environmental crisis. Since Reading has actioned its sustainability plans for 2020–2026 to include a behaviour and awareness program, it is crucial that students are aware of the environmental impact of generative AI (University of Reading Sustainability Services, 2020). In conversation with students following the information sessions, several mentioned that they were astounded by the carbon emissions of generative AI. Both these environmental considerations alongside the problematic training data and outputs appear to have affected students' views on AI use.

Speaking to students

During the AI information sessions that we held in Autumn 2023 term, we gathered survey data from student participants (Ross and Baines, 2023a, 2023b). Most of our presentations were for ancient language students in the Department of Classics at the University of Reading, but we also held general AI information sessions for the wider student body in both the Classics and modern languages departments. To ensure effective data collection, we followed a methodology approved by the University of Reading University Research Ethics Committee.

Prior to each presentation, we asked participants to voluntarily complete the first portion of the survey form to gather their opinions about generative AI. This portion of the survey was intended to collect student's opinion before they learned more about potential ethical considerations. From the language student results, we found that 92.1% of students had heard of generative AI tools for study, and 62.9% of students had also had AI tools for study advertised to them (Figure 6). Most of these students had heard about conversational AI tools, generative image AI, and AI-powered search co-pilots; ChatGPT, DALL-E, and Perplexity were the most popular answers. However, many students mentioned more obscure tools, including Tome, Loom AI, SnapchatAI, WordTune, and PicsArt. The majority were hearing

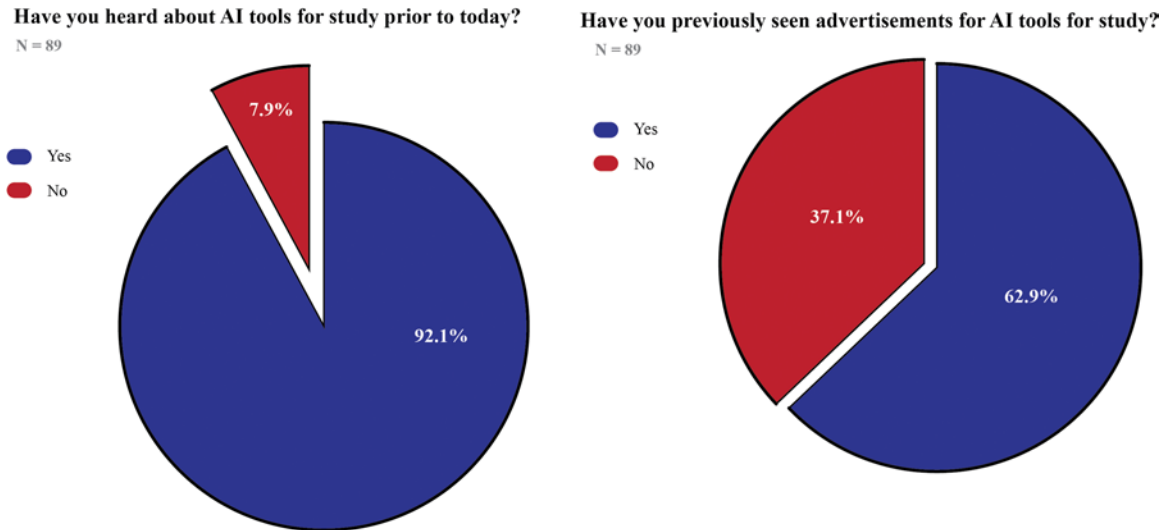


Figure 6. Survey data gathered from ancient language students in the Department of Classics at the University of Reading over Autumn 2023 term (Ross and Baines, 2023a).

about these tools on social media or in conversation with friends and classmates (Figure 7). It is clear from these results that students not only know about generative AI, but they are constantly being bombarded with advertisements for all sorts of tools as they are published.

Interestingly, the usage statistics do not lean as far towards the positive. Only 47.2% of the surveyed ancient language students had previously used generative AI; 18% of all those surveyed used them to support learning or assignments (Figure 8).

It is unclear if the ancient language students in these information sessions felt pressured to answer in the negative despite several assurances that all survey responses would be properly safeguarded and would not return to them, especially since the usage statistics were surprisingly low compared to public discussion. General student metrics found around 50% of students currently enrolled in university were using generative AI in some capacity (Coffey, 2023; Nam, 2023). So, we compared our ancient language results to the anonymous survey results gathered through Mentimeter at our general AI information sessions (Table 1).

In all three sets of surveys, student knowledge about AI was consistently over 90%, while AI usage varied greatly. The ancient language sessions were done in small groups, while the Classics and modern languages sessions were in larger rooms. Despite this, there

appears to be a greater apprehension for using AI among Classics and ancient language students than modern language students. Although group sizes may have affected our pre-presentation results, the post-session results are consistent.

Following the presentation, we asked participants to complete the second half of the survey once they had a fuller picture of how generative AI works and is made. Interestingly, only 19.1% of ancient language students said they would consider using generative AI, while 80.9% were either against or apprehensive of using them (Figure 9). Many students noted that they did not like how inaccurate generative AI could be, especially due to its elementary abilities with Latin and Ancient Greek. Furthermore, several students indicated their support for artists and authors, considering the ongoing Writers Guild of America (WGA) and Screen Actors Guild – American Federation of Television and Radio Artists (SAG-AFTRA) strike actions at the time of survey (Anguiano and Beckett, 2023; Watercutter and Bedingfield, 2023).

Those students who did wish to use generative AI, or were considering it, thought it could be useful for revision, language study support, and generating additional homework questions (Figure 10). Many of these suggested uses, however, were paired with comments saying that they would only use the AI tool for these purposes if its Latin and Ancient Greek abilities improved.

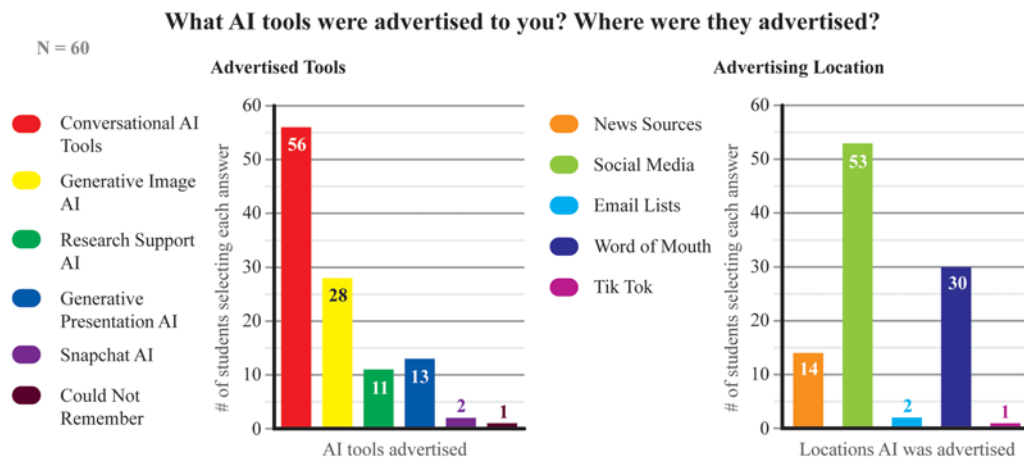


Figure 7. Survey data gathered from ancient language students in the Department of Classics at the University of Reading over Autumn 2023 term (Ross and Baines, 2023a).

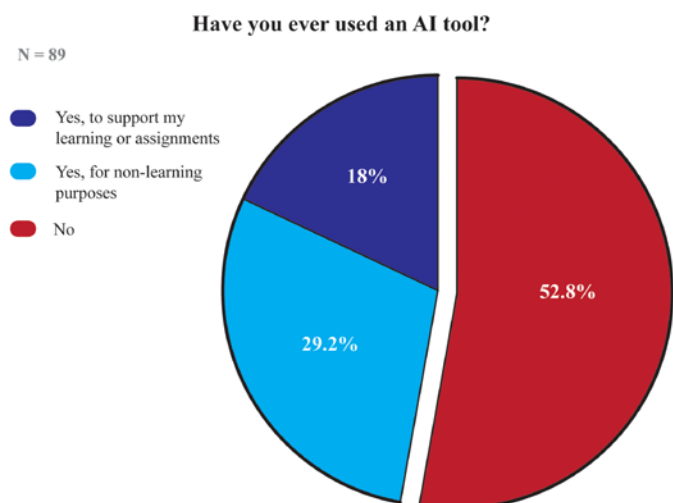


Figure 8. Survey data gathered from ancient language students in the Department of Classics at the University of Reading over Autumn 2023 term (Ross and Baines, 2023a).

Interestingly, there were also a few students who aimed to use AI for non-learning related purposes, such as writing professional emails or Instagram captions.

Both Classics and modern languages students also leaned towards rejection or apprehension of AI tools after learning about the ethical considerations (Table 2). As seen with the pre-presentation results, there is an increase in ‘yes’ responses in the wider groups of participants. However, the lean towards rejection remains consistent across all three datasets. At this stage, it appears that informing students about the benefits and drawbacks of generative AI provides them with the tools needed to critically consider whether or not to use it. Students had the option to indicate the reasons why they rejected using generative AI, and these reasons included lack of trust in AI outputs (28.42% of all respondents), fear of plagiarism accusations (8.74%), fear AI will replace human work (7.10%), unethical training data (3.83%), and support of affected artists (2.19%) (Ross and Baines, 2023a, 2023b). It is unclear if these opinions will remain the same over the academic year, but we can continue to scaffold student knowledge about AI developments as time progresses.

Building guidelines

As discussed earlier, our research involved discussing AI usage with both staff and students. It was necessary to draw up guidelines for AI usage which clearly addressed the issues surrounding AI and how the Department of Classics would engage with generative AI in classes and assignments. This required a delicate balance because there were many strong opinions about AI use in HE. In their responses to our initial staff survey, there was a broad spread of

views (Ross and Baines, 2023c). Some were completely against AI use and found it should be ignored, others wanted to embrace it fully but were unsure of the best methods, and others were utterly confused by it.

Following several discussions and information sessions, the department determined that allowed AI use should be at the discretion of module convenors, but all students should be informed about the benefits and drawbacks of generative AI and any new developments as they arise. To do this, we developed a general guidance and citation guide which would be publicly available to staff and students and consistently updated. The guidance document was developed bearing in mind the many guiding principles laid out by HE bodies in the UK and globally (Atlas, 2023; Jisc, 2023; Quality Assurance Agency for Higher Education, 2023; Russell Group, 2023; UNESCO International Institute for Higher Education in Latin America and the Caribbean, 2023). The Autumn 2023 version of the guidance is attached to this article in Supplemental Materials (Supplementary File 2).

Assignments/citations

If a Reading Classics student’s module convenor approves the use of AI in an assignment, they are first and foremost required to document all uses of AI they make while preparing their assignment, whether any of it ends up in the body text or not. All cases of AI use need to be indicated in the first footnote of an assignment following this general method:

‘The writing of this assignment is my own, and I take responsibility for all errors. During the preparation of this assignment, Perplexity (Perplexity AI, 9 August 2023 version) was used to gather articles for preliminary research into this research question.

This allows students to be accountable for their academic work and protects them from potential plagiarism misunderstandings. The disclaimer footnote, however, also needs to be accompanied by proper citations and documentation.

The Department of Classics at the University of Reading traditionally uses the Harvard citation style to cite modern sources. Many universities have released their own versions of the Harvard system to address the issue of citing generative AI outputs (St. George’s University of London Library, 2023; University of Queensland, Australia Library, 2023), and many are based on the Bloomsbury CiteThemRight guide. These citation guides, across all relevant citation systems, generally consider AI outputs as personal communications (Bloomsbury Publishing, n.-d.-a, n.-d.-b, n.-d.-c, n.-d.-d, n.-d.-e, n.-d.-f, n.-d.-g, n.-d.-h). Unfortunately, we found that these citation styles did not require crucial information about the AI tools used, so we created an adapted citation guide for our students (Supplementary File 3).

Table 1. Current use of AI. Survey data gathered from ancient language, Classics, and modern languages students at the University of Reading over Autumn 2023 term (Ross and Baines, 2023a, 2023b).

	Have you heard about AI tools for study?		Have you ever used an AI tool?		
	Yes (%)	No (%)	Yes (%)	No, but I’ve thought about it (%)	No (%)
Ancient languages (N = 89)	92.1	7.9	47.2	0.0	52.8
Classics (N = 37)	94.1	5.9	56.8	8.1	35.1
Modern languages (N = 57)	98.2	1.8	83.7	6.1	10.2

After today's session, do you think you will use AI tools to support your learning?

N = 89

- Yes
- I Don't Know
- No




Figure 9. Survey data gathered from ancient language students in the Department of Classics at the University of Reading over Autumn 2023 term (Ross and Baines, 2023a).

In order to ensure that readers know which versions of AI models are used, we require that students provide the following information:

- Developer name
- AI tool name
- AI tool model number
- Update version at time of use
- personal communication
- Date output was generated

e.g. (OpenAI, *ChatGPT 3.5*, 3 August 2023 version, personal communication, generated 18 September 2023).

It is particularly important that citations include both the model and version numbers in their citation because multiple different models of generative AI tools exist at the same time and different versions can respond to the same questions quite differently. If a student cites an AI program itself, they should cite it following the method for a computer program but must also provide the

developer's name, AI tool name, AI model number, and update version in the citation.

Since generative AI outputs are non-reproducible, the Department of Classics at the University of Reading requires students to maintain a screenshot folder of all AI outputs used for producing their assignments. All AI outputs used for a particular assignment need to be paired with a full citation, the prompt used, and an image and attached as an appendix at the end of their assignment. Although many generative AI tools now create stable links to share outputs, this is not a universal feature, so we still require a screenshot alongside any stable links.

Conclusions

It is clear from our research into conversational AI tools and their potential ethical use in ancient language learning that the starting point of any discussions on its use needs to begin with informing and educating staff and students alike to provide them with a greater understanding of such tools, their content policies, how their training data is gathered and the environmental impact. Excessive feelings of fear or misplaced faith in AI abound if there is a lack of understanding of how the conversational tools work, their current state of accuracy and how they gather their data. From our pre-presentation surveys, it is evident that most students know of and have used AI tools. However, once they are presented with more in-depth information, they can make an informed judgement about the ethical use of AI and, as a result, show a greater reluctance about using it. The students of Latin and Ancient Greek were sceptical about using AI tools after the presentation once they realised the inadequacies of such tools. Modern languages students are aware of the skill of AI tools for their studies but fewer than 20% indicated that they would use them after attending the information session.

From a pedagogical point of view AI conversational tools may have a place in Latin and Ancient Greek beginner and intermediate classes as an additional tool for students to use while studying on their own. It is possible for students outside the classroom, with appropriate direction, to create quizzes, to revise grammar points, to check homework answers, to parse and to translate and compose short passages (Ross, 2023).

The development of conversational AI tools continues to accelerate and become more sophisticated in all its aspects. This

How do you plan to use generative AI tools?

N = 89

- Language Study Support
- Research Support
- Generating Additional Homework
- Revision
- Critiquing Forms It Creates
- Fun/Entertainment
- Professional Emails
- Instagram Captions
- I'm Not Sure

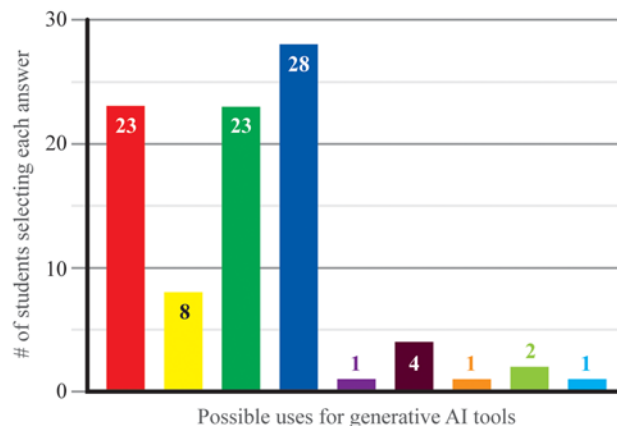


Figure 10. Survey data gathered from ancient language students in the Department of Classics at the University of Reading over Autumn 2023 term (Ross and Baines, 2023a).

Table 2. Future use of AI. Survey data gathered from ancient language, Classics, and modern languages students at the University of Reading over Autumn 2023 term (Ross and Baines, 2023a, 2023b).

	After today's session, do you think you will use AI tools to support your learning?		
	Yes (%)	I Don't Know (%)	No (%)
Ancient languages (N=89)	19.1	41.6	39.3
Classics (N=37)	27	0.0	73
Modern languages (N=57)	36.1	0.0	63.9

highlights the need to continue to be vigilant and to check how well the new developments perform in relation to Latin and Ancient Greek along with other areas studied within the field of Classics. If we continue to do this and keep our staff and students informed, the use of conversational AI tools provides an exciting and valuable addition to our teaching resources.

As a continuation of our research, with the aid of focus groups of students, we will test a variety of applications for their output effectiveness and reliability when using a specific set of guiding phrases to be used to assist Latin and Ancient Greek language learning. This includes ChatGPT 3.5/4, Bard AI, Claude-2, Bing Chat, and custom GPTs through OpenAI's new GPT system. We also will be updating our guidance documents as new tools and guiding principles are released.

Supplementary material

The supplementary material for this article can be found at <https://doi.org/10.1017/S2058631024000412>.

Acknowledgments

This article is part of the wider 'ChatGPT: A Conversational Language Study Tool' project at the University of Reading. This project has been reviewed by the University of Reading University Research Ethics Committee and has been given a favourable ethical opinion for conduct. The project is supported by a Teaching and Learning Enhancement Project (TLEP) grant from the University of Reading and an Education Grant from the Council of University Classics Departments. Special thanks to Jacinta Hunter, Fleur McRitchie Pratt, Nisha Patel, Luke Edwards, and Rikard Roitto for bringing changes in AI outputs to our attention.

Notes

1 Taylor (2016) expects the answer to be: 'ὁ γίγας ὃς ἀνθρώπους ἐσθίει οὐκ ἐστὶν νῦν ἐν τῷ ἀγρῷ.'

2 Many artists consider the Classics field to be safe spaces to discuss the content policy restricted issues listed by OpenAI (2023f). The temporal distances of ancient texts from our current context allows us to discuss difficult issues in a distanced but critical manner. For more on this concept of Classics as a safe space, see Ewans and Johnson (2019).

References

- Anguiano D and Beckett L** (2023, Oct. 1) How Hollywood writers triumphed over AI – and why it matters. *The Guardian*. Available at <https://www.theguardian.com/culture/2023/oct/01/hollywood-writers-strike-artificial-intelligence> (accessed 15 November 2023).
- Anthropic** (2023) Claude-2 (July 11, 2023 version) [Large language model]. Available at <https://claude.ai/login?returnTo=%2F>
- Appel G, Neelbauer J and Schweidel DA** (2023, Apr. 7) Generative AI Has an Intellectual Property Problem. *Harvard Business Review*. Available at <https://hbr.org/2023/04/generative-ai-has-an-intellectual-property-problem> (accessed 27 September 2023).

- Atlas S** (2023) *ChatGPT for Higher Education and Professional Development: A Guide to Conversational AI*. Kingston: DigitalCommons@URI. https://digitalcommons.uri.edu/cba_facpubs/548.
- Bikbaeva D** (2023, Feb. 1) AI Trained on Copyrighted Works: When Is It Fair Use? *The Fashion Law*. Available at <https://www.thefashionlaw.com/ai-trained-on-copyrighted-works-when-is-it-fair-use/> (accessed 27 September 2023).
- Bloomsbury Publishing** (n.-d.-a) Generative AI (APA 7th). *Bloomsbury citethemright*. Retrieved November 8, 2023, from <https://www.citethemrightonline.com/sourcetype?docid=b-9781350927971&tocid=b-9781350927971-185&st=Generative+AI>
- Bloomsbury Publishing** (n.-d.-b) Generative AI (Chicago). *Bloomsbury citethemright*. Retrieved November 8, 2023, from <https://www.citethemrightonline.com/sourcetype?docid=b-9781350927988&tocid=b-9781350927988-187&st=Generative+AI>
- Bloomsbury Publishing** (n.-d.-c) Generative AI (Harvard). *Bloomsbury citethemright*. Retrieved November 8, 2023, from <https://www.citethemrightonline.com/sourcetype?docid=b-9781350927964&tocid=b-9781350927964-217&st=Generative+AI>
- Bloomsbury Publishing** (n.-d.-d) Generative AI (IEEE). *Bloomsbury citethemright*. Retrieved November 8, 2023, from <https://www.citethemrightonline.com/sourcetype?docid=b-9781350927995&tocid=b-9781350927995-45&st=Generative+AI>
- Bloomsbury Publishing** (n.-d.-e) Generative AI (MHRA). *Bloomsbury citethemright*. Retrieved November 8, 2023, from <https://www.citethemrightonline.com/sourcetype?docid=b-9781350928008&tocid=b-9781350928008-42&st=Generative+AI>
- Bloomsbury Publishing** (n.-d.-f) Generative AI (MLA 9th). *Bloomsbury citethemright*. Retrieved November 8, 2023, from <https://www.citethemrightonline.com/sourcetype?docid=b-9781350928015&tocid=b-9781350928015-121&st=Generative+AI>
- Bloomsbury Publishing** (n.-d.-g) Generative AI (OSCOLA). *Bloomsbury citethemright*. Retrieved November 8, 2023, from <https://www.citethemrightonline.com/sourcetype?docid=b-9781350928022&tocid=b-9781350928022-46&st=Generative+AI>
- Bloomsbury Publishing** (n.-d.-h) Generative AI (Vancouver). *Bloomsbury citethemright*. Retrieved November 8, 2023, from <https://www.citethemrightonline.com/sourcetype?docid=b-9781350928039&tocid=b-9781350928039-48&st=Generative+AI>
- Brown TB, Mann B, Ryder N, Subbiah M, Kaplan J, Dhariwal P, Neelakantan A, Shyam P, Sastry G, Askell A, Agarwal S, Herbert-Voss A, Krueger G, Henighan T, Child R, Ramesh A, Ziegler DM, Wu J, Winter C, Hesse C, Chen M, Sigler E, Litwin M, Gray S, Chess B, Clark J, Berner C, McCandlish S, Radford A, Sutskever I and Amodei D** (2020) Language models are few-shot learners. *Advances in Neural Information Processing Systems* 33, 1–75. <https://doi.org/10.48550/arXiv.2005.14165> (accessed 15 November 2023).
- Burns PJ** (2023, May 19) Research Recap: How Much Latin Does ChatGPT 'Know'? *Institute for the Study of the Ancient World Library Blog*. Available at <https://isaw.nyu.edu/library/blog/research-recap-how-much-latin-does-chatgpt-know> (accessed 27 September 2023).
- Cassette** (2020) CassetteAI VI (August 3, 2023 version) [Text-to-audio model]. Available at <https://cassetteai.com/dashboard>
- Clackson J** (2011) Classical Latin. In Clackson J (ed.), *A Companion to the Latin Language*. Hoboken: Blackwell Publishing Ltd, pp. 236–256. <https://doi.org/10.1002/9781444343397.ch15>.
- Coffey L** (2023, Oct. 31) Students Outrunning Faculty in AI Use. *Inside Higher Ed*. Available at <https://www.insidehighered.com/news/tech-innovation/artificial-intelligence/2023/10/31/most-students-outrunning-faculty-ai-use> (accessed 15 November 2023).
- Cremer E** (2023, Jul. 5) Authors file a lawsuit against OpenAI for unlawfully 'ingesting' their books. *The Guardian*. Available at <https://www.theguardian.com/books/2023/jul/05/authors-file-a-lawsuit-against-openai-for-unlawfully-ingesting-their-books> (accessed 27 September 2023).
- Cullen H and Taylor J** (2016) *Latin to GCSE: Part 2*. London: Bloomsbury Academic.

- de Vries A** (2023) The growing energy footprint of artificial intelligence. *Joule* 7, 2191–2194. <https://doi.org/10.1016/j.joule.2023.09.004>
- EssayAIGroup** (2020) EssayAIGroup (June 18, 2023 version) [Large language model]. Available at <https://www.essayailab.com/>
- Ewans M and Johnson M** (2019) Wesley Enoch's Black Medea. In Johnson M (ed.), *Antipodean Antiquities: Classical Reception Down Under*. London: Bloomsbury Academic, pp. 73–86.
- Google** (2023a) Bard (November 16, 2023 version) [Large language model]. Available at <https://bard.google.com/?hl=en>
- Google** (2023b, Jul. 13) Experiment Updates. *Bard*. Available at <https://bard.google.com/updates> (accessed 15 November 2023).
- Hao K** (2019, Jun. 6) Training a single AI model can emit as much carbon as five cars in their lifetimes. *MIT Technology Review*. Available at <https://www.technologyreview.com/2019/06/06/239031/training-a-single-ai-model-can-emit-as-much-carbon-as-five-cars-in-their-lifetimes/> (accessed 15 November 2023).
- Heikkilä M** (2023, Oct. 23) This new data poisoning tool lets artists fight back against generative AI. *MIT Technology Review*. Available at <https://www.technologyreview.com/2023/10/23/1082189/data-poisoning-artists-fight-generative-ai/> (accessed 8 November 2023).
- HeyGen** (2023) HeyGen 3.0 (April 10, 2023 version) [Text-to-video model]. Available at https://app.heygen.com/login?sid=no_sid
- Jisc** (2023, Aug. 29) *Generative AI – A Primer*. London: Jisc.
- Katz L** (2023, Oct. 26) New Tool 'Poisons' AI Data To Shield Artists Against Having Their Work Stolen. *Forbes*. Available at <https://www.forbes.com/sites/lesliekatz/2023/10/26/new-tool-poisons-ai-data-to-shield-artists-against-having-their-work-stolen/> (accessed 8 November 2023).
- Lucchi N** (2023) ChatGPT: a case study on copyright challenges for generative artificial intelligence systems. *European Journal of Risk Regulation* 15(2), 1–23. <https://doi.org/10.1017/err.2023.59>
- Magical Tome** (2023) Tome V2 (November 20, 2023 version) [Text-to-presentation model]. Available at <https://tome.app/>
- Metz R** (2023, Oct. 21) These artists found out their work was used to train AI. Now they're furious. *CNN Business*. Available at <https://edition.cnn.com/2022/10/21/tech/artists-ai-images/index.html> (accessed 27 September 2023).
- Microsoft** (2023) Bing Chat (November 15, 2023 version) [Large language model]. Available at <https://www.bing.com/search?q=Bing%20AI&showconv=1&form=MT00IS>
- Midjourney** (2023) Midjourney Bot 5.2 (June 22, 2023 version) [Text-to-image model]. Available at <https://legacy.midjourney.com/home/>
- Morpurgo Davies A** (2015, July 30) pronunciation, Greek. *Oxford Classical Dictionary*. Oxford: Oxford University Press. <https://doi.org/10.1093/acrefore/9780199381135.013.5365>
- Nam J** (2023, Nov. 22) 56% of College Students Have Used AI on Assignments or Exams. *BestColleges*. Available at <https://www.bestcolleges.com/research/most-college-students-have-used-ai-survey/> (accessed 10 December 2023).
- OpenAI** (2023a) ChatGPT 4 (November 21, 2023 version) [Large language model]. Available at <https://chat.openai.com/auth/login>
- OpenAI** (2023b, Sept. 25) ChatGPT can now see, hear, and speak. *OpenAI*. Available at <https://openai.com/blog/chatgpt-can-now-see-hear-and-speak> (accessed 13 November 2023).
- OpenAI** (2023c, Sept. 27) ChatGPT – Release Notes. *OpenAI*. Available at https://help.openai.com/en/articles/6825453-chatgpt-release-notes#h_6401b89b6b (accessed 21 November 2023).
- OpenAI** (2023d) DALL-E-3 (August 20, 2023 version) [Text-to-image model]. Available at <https://openai.com/dall-e-3>
- OpenAI** (2023e, Nov. 6) Introducing GPTs. *OpenAI*. Available at <https://openai.com/blog/introducing-gpts> (accessed 10 November 2023).
- OpenAI** (2023f, Mar. 23) Usage Policies. *OpenAI*. Available at <https://openai.com/policies/usage-policies> (accessed 21 November 2023).
- Perplexity** (2023a) Perplexity Copilot (August 23, 2023 version) [Large language model]. Available at <https://www.perplexity.ai/>
- Perplexity** (2023b) Perplexity Image Upload. Perplexity Collections. Available at <https://www.perplexity.ai/collections/Perplexity-Image-Upload-Mb27WFDfSrmXmmdvWdQ?s=c> (accessed 27 December 2023).
- Quality Assurance Agency for Higher Education** (2023, Jul.) *Reconsidering Assessment for the ChatGPT Era: QAA Advice on Developing Sustainable Assessment Strategies*. Gloucester: Quality Assurance Agency for Higher Education.
- Ross EAS** (2023) A new frontier: AI and ancient language pedagogy. *Journal of Classics Teaching* 24, 143–161. <https://doi.org/10.1017/S2058631023000430>
- Ross EAS and Baines J** (2023a, Dec. 16) *AI and Ancient Language Student Survey – Autumn 2023* (Version V2) [Data set]. figshare. <http://doi.org/10.6084/m9.figshare.24835629.v2>
- Ross EAS and Baines J** (2023b, Dec. 16) *Is AI the New Plato?: AI Benefits and Drawbacks for Study Survey – Autumn 2023* (Version V2) [Data set]. figshare. <http://doi.org/10.6084/m9.figshare.24842058.v2>
- Ross EAS and Baines J** (2023c, Dec. 16) *University of Reading Classical Languages Teaching Content and Expectations Survey – Summer 2023* (Version V2) [Data set]. figshare. <http://doi.org/10.6084/m9.figshare.24829953.v2>
- Russell Group** (2023, July 4) New principles on use of AI in education. Russell Group. Available at <https://russellgroup.ac.uk/news/new-principles-on-use-of-ai-in-education/> (accessed 20 August 2023).
- Shan S, Cryan J, Wenger E, Zheng H, Hanocka R and Zhao BY** (2023) Glaze: Protecting Artists from Style Mimicry by Text-to-Image Models. In *Proceedings of USENIX Security Symposium, Anaheim CA, August 2023*. ArXiv. <https://doi.org/10.48550/arXiv.2302.04222>
- Stability.ai** (2023) Stable Diffusion XL 1.0 (July 26, 2023 version) [Text-to-image model]. Available at <https://stability.ai/stable-diffusion>
- St. George's University of London Library** (2023, Sept. 14) How to Reference AI in an Assignment. *Skills Guide at St George's Library*. Available at <https://libguides.sgu.ac.uk/Harvard/AI> (accessed 8 November 2023).
- Strubell E, Ganesh A and McCallum A** (2019) Energy and Policy Considerations for Deep Learning in NLP. In *the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*. Florence, Italy. July 2019. ArXiv. <https://doi.org/10.48550/arXiv.1906.02243>
- Taylor J** (2016) *Greek to GCSE: Part 2 – Revised Edition for OCR GCSE Classical Greek (9-1)*. London: Bloomsbury Academic.
- Taylor J** (2023, Aug. 9) Google says AI systems should be able to mine publishers' work unless companies opt out. *The Guardian*. Available at <https://www.theguardian.com/technology/2023/aug/09/google-says-ai-systems-should-be-able-to-mine-publishers-work-unless-companies-opt-out> (accessed 27 September 2023).
- UNESCO International Institute for Higher Education in Latin America and the Caribbean** (2023) ChatGPT and artificial intelligence in higher education: quick start guide. *UNESCO Digital Library*. Available at <https://unesdoc.unesco.org/ark:/48223/pf0000385146> (accessed 8 October 2023).
- University of Queensland, Australia Library** (2023, Aug. 23) Using ChatGPT or other generative AI in your assignments. Available at <https://guides.library.uq.edu.au/referencing/uqharvard/chatgpt-and-generative-ai> (accessed 8 November 2023).
- University of Reading Sustainability Services** (2020) Our Future First. Available at <https://sites.reading.ac.uk/sustainability/get-involved/ourfuturefirst/> (accessed 15 February 2024).
- veekaybee** (2022, Dec. 8) chatgpt.md. *GitHub Gist*. Available at <https://gist.github.com/veekaybee/6f8885e9906aa9c5408ebe5c7e870698> (accessed 8 November 2023).
- Vincent J** (2023, Apr. 28) Reported EU legislation to disclose AI training data could trigger copyright lawsuits. *The Verge*. Available at <https://www.theverge.com/2023/4/28/23702437/eu-ai-act-disclose-copyright-training-data-report> (accessed 27 September 2023).
- Watercutter A and Bedingfield W** (2023, Nov. 9) Hollywood Actors Strike Ends With a Deal That Will Impact AI and Streaming for Decades. *WIRED*. Available at <https://www.wired.co.uk/article/hollywood-actors-strike-ends-ai-streaming> (accessed 15 November 2023).
- Wong M** (2023, Oct. 2) Artists Are Losing the War Against AI. *The Atlantic*. Available at <https://www.theatlantic.com/technology/archive/2023/10/openai-dall-e-3-artists-work/675519/> (accessed 8 November 2023).
- Zahn M** (2023, Sept. 25) Authors' lawsuit against OpenAI could 'fundamentally reshape' artificial intelligence, according to experts. *abcNEWS*. Available at <https://abcnews.go.com/Technology/authors-lawsuit-openai-fundamentally-reshape-artificial-intelligence-experts/story?id=103379209> (accessed 27 September 2023).