

## Research Article

# GenAI images of ‘Roman Britain’ as tools of reception

Lisa Maurice 

Classical Studies, Bar-Ilan University, Ramat Gan, Israel

### Abstract

This paper addresses the question of the role that artificial intelligence (AI) image generators play in the reception of the ancient world, examining the assumptions on which they draw in the generation of images, and how the creation of such images influences perceptions about the classical past. After a brief outline of how AI image generators work, highlighting the inherent presumptions and biases of the visual productions, a small case study is then presented, in which the prompt ‘Roman Britain’ was submitted to eight different free image generators. The conclusion drawn from this experiment is that while the technology is impressive, none of the image generators have managed to produce pictures that effectively conjure up Roman Britain. Although the tools may be good at creating a general impression, individual details are often incorrect. Moreover, the output depends heavily on the training data available. In the case of the ancient world, no photographs exist; only archaeological remains, fragments, and later imaginative reconstructions survive. Consequently, these limitations inevitably shape the generated images. Despite these disadvantages, it is likely that AI-generated images will become part of cultural heritage, and it is, therefore, important to consider the role that such images might play in the reception of antiquity. In recognition of the problems, and the advantages, of this technology, some suggestions are made in the final section of the paper as to how generative artificial intelligence (GenAI) images may be used in a positive manner, particularly within the classroom.

**Keywords:** AI images; ancient Rome; Roman Britain; Classical reception; Classics teaching

### Introduction

The modern world is a heavily visual one. Images, and in particular screen images, play a central role in influencing ideas and beliefs about all manner of issues and subjects (Gioia et al. 2014). The well-known adage that ‘a picture is worth a thousand words’ underscores the perceived persuasive power of visual productions (Joffe 2008). In the case of the ancient world, the images that have been influential in forming popular conception include artefacts and archaeological sites from the historical past, examples of imaginative artwork created hundreds of years after the period they depict, and screen depictions of Greece and Rome. With the advance of artificial intelligence (AI) technology, a new source of visual representations has emerged, as generative AI (GenAI) image generators are able to produce high-quality images in a manner that is likely to result in fundamental changes in how creative processes are understood. GenAI tools enable users to produce visually appealing pictures with minimal technical expertise, thereby expanding access to creative production. These tools are widely available and have been adopted across age groups, particularly by Generation Z, to which many students belong (Ahmad 2024). As such, the role these tools play in society – and are likely to play in the coming years – should concern educators.

Their capacity to shape opinion and to influence classroom practice warrants careful consideration.

This paper aims to explore how contemporary GenAI image generators visually interpret the prompt ‘Roman Britain’, considering the kinds of imagery, iconography, architectural forms, landscapes, and human figures prioritised in the production of pictures. It examines the assumptions and visual conventions underlying these representations and considers the extent to which they reflect entrenched popular stereotypes of Rome, Britain, and antiquity more broadly. The accuracy and historical plausibility of the images are judged against current knowledge of Roman Britain, in order to explore what the strengths and weaknesses of these outputs reveal about the training data and structural limitations of diffusion-based image models. Finally, the paper addresses the implications of AI-generated visualisations for the reception of the ancient world, particularly in educational and cultural heritage contexts.

### AI image generators: a brief introduction

Before considering the experiment, it is necessary to outline briefly how image generators work. Like other GenAI models, a prompt is given and a response returned, but it is important to understand the processes by which this occurs. Most of these tools use diffusion models. GenAI models are trained to recognise objects in images, using a large dataset, which they do by converting the data into mathematical representations. They then translate the maths back into pixels and, with an image generator, deliver the output in the form of a new, freshly created, visual product (Croitoru et al. 2022).

Email: [lisa.maurice@biu.ac.il](mailto:lisa.maurice@biu.ac.il)

**Cite this article:** Maurice L. (2026). GenAI images of ‘Roman Britain’ as tools of reception. *The Journal of Classics Teaching*, 1–15. <https://doi.org/10.1017/S2058631026101020>

© 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.

It should be stressed that the machine does not actually ‘see’ or ‘understand’ the picture as a human would but rather uses statistical patterns in the data to generate possible reasonable images (Picascia and Ferrara 2026, 37–38; Spennemann 2024, 3598). In the case of image generators, this process is the result of textual prompts (Zhang et al. 2023). After the machine learning model has scanned millions of images across the internet along with the text associated with them, the algorithms can identify trends and eventually begin to guess which image and text fit together. In this way, a text prompt can produce a picture, which can then be fine-tuned by additional prompts.

### The diffusion tools

Any GenAI image tool has to work from a model; various options exist that vary in technical ways. All models work in the same way, however, by scanning millions of images and text tags and being trained to recognise these through a Contrastive Language-Image Pre-training (CLIP) technique. An issue of importance, therefore, is the source of these images-text pairs used in training; the training sets utilise images in the public domain, and, in the case of commercial models like Adobe or Shutterstock, licensed data as well. Whatever specific data sources and training sets have been used by the various tools, the training information has been scraped from the vast collection of images on the internet. As a result, the datasets incorporate diverse visual materials, including photographs, drawings, paintings, news imagery, and other publicly available content. In some cases, they also use items that are held with copyright and regarded as intellectual property. Some organisations, such as Getty, have claimed that artists’ information has been scraped (Coulter 2024). In every case, however, it is vital to understand that this vast body of sources is created by *humans*, who have taken the photographs, painted, drawn, or created the images (Epstein et al. 2023, 3). Since the outputs of GenAI image generators are essentially just sophisticated predictions stemming from statistical relationships and patterns found in the tagged image data in their training sets, GenAI fundamentally lacks the capacity for intent and autonomous creative thinking. Any perceived creativity from GenAI models is purely based on an interpretation by the individual who is interacting with that model. Central to the concept of authorship is that the author initiates and conceptualises an original piece of writing. By their nature, GenAI models are, at least at this point in their development, incapable of initiation and conceptualisation of creative or research ideas without human activation through prompting and prompt manipulation (Spennemann 2024, 3601–3602).

People’s perspectives and standpoints are shaped by a whole range of influences – family, community, education, sociopolitical and historical contexts, and life experiences. GenAI language and image models are guided by human input and trained on human-produced material. The assumptions and biases embedded in that material are therefore likely to be encoded within the models themselves. As a result, it has been noted that:

The overrepresentation of white, fully-abled, Western men in images of high status categories, and the invisibility of women, people of color, and the disabled, except in low status categories, and the almost complete absence of realistic, non-sexualized images of women, plagues all text-to-image AI models (Jacobi and Sag 2024, 1).

This has been demonstrated by research studies, one of which found that both women and people of colour were under-

represented in generated images, while women were also depicted in a more sexualised manner and as younger than men. Similarly, people of colour are often portrayed as younger than white individuals. As a result, ‘the images overwhelmingly represented individuals as belonging to a higher socioeconomic class, pointing towards a systemic bias within AI systems towards privilege’ (Gengler 2024). Although the tools have attempted to correct such biases, the results have been, at best, clumsy and, at worst, bizarre, resulting in what one study calls ‘the Black Nazi Problem’, in which, in an overzealous attempt to compensate, AI image generators have produced pictures of impossibilities such as Black Nazis and female Indian popes (Jacobi and Sag 2024, 1).

### The ‘Roman Britain’ experiment

The question of how antiquity might be depicted by AI image generators is one that has relevance for classicists, a topic that has already begun to be explored by the Investigating Generative Artificial Intelligence in Ancient World Studies (iGAIAS) research project at Reading University led by Edward Ross and Jackie Baines (<https://edwardasross.wordpress.com/projects/igaia-investigating-generative-artificial-intelligence-in-ancient-world-studies/>). This project ‘aims to demystify the development, ethics, and uses of generative AI and make it accessible for students and teachers of ancient world studies and the wider public’, through a survey in which educators share their knowledge and experience of GenAI tools in teaching the Greco-Roman world. In order to deepen the understanding of the practical aspects of utilising GenAI images in the classroom, an project carried out an experiment, the primary aim of which was to examine how freely available GenAI image generators construct visual representations of the ancient past when given a simple prompt (‘Roman Britain’), and to assess what these constructions reveal about:

- the visual assumptions embedded in AI training datasets,
- the reliability of AI-generated historical imagery, and
- the potential influence of such imagery on contemporary perceptions of antiquity.

The experiment was deliberately designed to reflect real-world classroom conditions. Only free tools were selected, and the initial prompt was intentionally simple. This approach allowed the study to replicate the likely experience of educators and students who might use these tools without advanced prompt engineering. A secondary aim was to determine whether more detailed prompting significantly improves historical plausibility, thereby testing the limits of user intervention in correcting AI output. Importantly, the goal was not to ridicule or dismiss AI technology but to evaluate critically how such tools participate in the ongoing reception of the ancient world and to consider their potential role within cultural heritage formation.

For the experiment, eight different image generators were used. Chosen at random from the multiple options now available – and the number has grown exponentially since the experiment was carried out in January 2025 – these were:

- OpenArt.ai
- ImagineArt
- Microsoft Designer
- Adobe Firefly
- Ideogram
- Freepik
- Canva
- Recraft.ai

Several of these image generators are based on the same diffusion tools. Microsoft Designer, which uses DALL-E 3 via OpenAI, is integrated with Office and designed primarily for social media. OpenArt.ai is built on Stable Diffusion, DALL-E, and other open models, and is intended for general image generation and fine-tuning models, while Canva's model is also Stable Diffusion. ImagineArt does not disclose its model, but it is likely based on Stable Diffusion or Midjourney, and its output is stylised art, with an emphasis on fantasy.

Other generators, such as Freepik, Adobe Firefly, Ideogram.ai, and Recraft.ai, all use their own proprietary models, developed according to their individual needs and criteria. Freepik, originally a huge platform of images that expanded to include AI powered tools, is particularly aimed at commercial/marketing content creation. Firefly is part of the Adobe Creative Cloud suite of applications and is intended primarily for professional designers, marketers, and agencies. Recraft was the first AI model built for designers and is designed for professionals. Ideogram, designed by ex-Google Brain researchers, is best suited to logos, posters, and meme-like content.

### Analysis of the GenAI image output

#### 1. OpenArtAi

OpenArtAi (<https://openart.ai/>) produced six responses, one of which (Figure 1) is an inaccurate map-based response, featuring a map of the Roman Empire in 117 CE, with letter shapes representing the words that the image generator is unable to understand or reproduce. The others (Figures 2–6), however, are almost entirely military-centred, with some examples featuring ruined or incomplete Roman-style buildings (a temple, the Colosseum) in the background. The colours are vivid and bright, reminiscent of stereotyped movie depictions (Maurice 2016 111–115).

#### 2. ImagineArt

ImagineArt (<https://www.imagine.art/>) produced four images (Figures 7–10), but these were entirely map-based graphics, centring on the outline of Britain, either with or without the inclusion of Ireland, but with no indication of any areas that were or were not under Roman rule at any point. Two of the images have no further details, but one features a silhouette of three figures, two males and a cloaked female, and another presents a rear view of a man in a red, belted robe, but neither the figures nor any of the details in the pictures have any real elements of Rome.

#### 3. Microsoft Designer

Microsoft Designer (<https://designer.microsoft.com/>) created four images that depict vibrantly colourful and peaceful scenes of busy marketplaces, reminiscent of illustrations in a children's book. All of the images feature green rolling hills in the background, in an idealised British landscape. Only one of the pictures, Figure 14, has any kind of military presence, in the form of one soldier, in sharp contrast to the images from OpenArtAi. It is notable that the images, particularly Figures 11 and 12, feature a diverse ethnic population; while this is commendable in that the Roman Empire was indeed multi-ethnic, it does somehow give rather an impression of a stereotyped Eastern market, not least because of the glorious sunshine beating down. An interesting mixture of buildings also features; in Figure 13 there is a colonnade of pillars topped with a pediment, possibly representing an aqueduct, but the structure is a ruin rather than in working condition. In Figure 11, both a building reminiscent of St Paul's cathedral and a Norman

keep seem to be depicted. In Figure 14, the scene is dominated by a three-arch structure on two stories, topped by a temple style pediment; what appears to be a version of Trajan's column also features.

#### 4. Adobe Firefly

Adobe Firefly (<https://firefly.adobe.com/>) is the first example in the study of a GenAI creator based on a proprietary model. These pictures feature buildings prominently, with two centring on ruined Greco-Roman temples (Figures 15 and 16), one (Figure 17) with a Norman-style building, possibly a convent or monastery, with Gothic elements, and the last (Figure 18) showing an 'ancient' scene of a Roman temple. Only the latter two feature humans, but these are not entirely successful elements; red cloth, reminiscent of Roman soldiers in pop culture, features in the figures' clothes in Figure 17, with some use of red, but it is unclear whether these garments are robes or cloaks. Even more problematically, although a Roman helmeted figure is seen, only the top half of the figure has been created, and it seems to be floating in mid-air. In Figure 18 the characters are in what appears to be medieval dress, although red features prominently once again, but again there are issues with representing the human figures, one of whom seems to consist only of an oversized striding leg over which a cloak is draped. The style of the Adobe images is very different from the others demonstrated so far, and seems to be much more in the style of classical painting, with the sky recalling Renaissance art, colour palettes reminiscent of Bruegel's use of earth pigments, and Figure 18 a strong resemblance to Vinzenz Fischer's *Sacrifice in Front of a Roman Temple* (1791).

#### 5. Ideogram

The style of Ideogram (<https://ideogram.ai/>) is different again, and its four images also exhibit a wider variety of images. Figure 19 is a photographic image of an ancient Roman street filled with tunic-clad, only Caucasian people. Despite the 'ancient Rome' prompt, what appears to be St Paul's cathedral again features, dominating the centre of the background, while the street itself is lined with Georgian-style buildings, in which the windows are paved with glass. Although the Romans did use glass in windows, this was generally restricted to the baths and to the homes of the wealthy, and it is unlikely that such a scene would have been seen in a Roman town in Britain. Figure 20 is of a statue, in what appears to be a modern city square, of a bearded figure, with a fillet headband round his head, and missing one arm; it looks like a mixture of the Charioteer of Delphi and an emperor, possibly Antoninus Pius. The ruins of an amphitheatre are the focus of the Figure 21, while Figure 22 shows more ruins, this time featuring pillars; both pictures situate the buildings in green countryside.

#### 6. Freepik

Freepik (<https://www.freepik.com/>), which gives the largest number of results, seven in total, creates images that look more like screenshots from a television drama. Figure 23 depicts a campfire near a castellated fort, beside which sits a bearded man of Middle Eastern appearance, dressed in long robe and boots, and wearing a turban, and a young, attractive, blonde-haired woman in a backless dress. Young men, apparently in modern dress, are seen on the far side of the fire. Figure 24 shows a procession of Roman soldiers and people in yellow tunics, parading down a street, and led by a blonde woman in a white and red sleeveless dress and Roman helmet, dancing in what seems to be some kind of festival. The word 'Roman' is visible on the bunting decorating the streets. A Roman



Figure 1. OpenArtAi 1.

military encampment during a misty dawn is the scene in Figure 25, with soldiers seen among the tents. Rather more surprisingly, Figure 26 depicts a man and a woman standing in a pool; the man seems to be wearing an army cloak and possibly full uniform, although his lower half is hidden by the water, and the woman, wearing a white dress, is smiling and looking at him. Even more strangely, the man's hands are perhaps bound with rope, although this is not entirely clear, and two more female figures, who seem to be naked from the waist up, are dimly seen in the background. In Figure 27, a group of archers are depicted beside

fortified walls and towers in green countryside; these are presumably Romans because they are wearing red tunics. Notably, these archers include a Black male centurion and also, perhaps less believably, a Caucasian female archer. The next picture, Figure 28, is a scene from a villa in the countryside, surrounded by lush greenery, in which a woman in a red dress is seen walking away from a well, at which a bearded man in a tunic is drawing water into a container that resembles a modern gas canister more than an ancient urn. Finally, the last picture, Figure 29, shows a bearded man in a tunic with a red cloak over his shoulder talking to a blonde



Figure 2. OpenArtAi 2.



Figure 4. OpenArtAi 4.



Figure 3. OpenArtAi 3.



Figure 5. OpenArtAi 5.

woman whose partially braided hair marks her out as Celtic. The location in this case is a street with timbered houses.

It is notable in several of the images in this set that emphasis is placed upon a Roman military-looking male and a Caucasian blonde female, perhaps representing the relationship between the Roman military presence and the native population. There is also more diversity in the portrayal of different ethnicities in these images. Overall, however, there is limited iconography that suggests ancient Rome.

#### 7. Canva

In Canva (<https://www.canva.com/ai-image-generator/>) the images are also stylised. The first image, Figure 30, is in a watercolour style, and shows a long, thatched building in a green enclosed area, presumably an exercise field, in which horses are practising manoeuvres and archery is also taking place. Spectators are

observing both on the field and from behind or on the wall enclosing the area. The figures wear cloaks but there is no obvious Roman iconography. Figure 31 is a market scene, but again appears more stereotypically Eastern than British, with the figures dressed in long robes and some sporting turbans or other headgear. One of the stalls is selling earthenware pots and another foodstuffs in wooden bowls. The next image, Figure 32, shows two Roman temples surrounded by crowds of soldiers. The style is of an engraving in black and white, with the soldiers' cloaks tinted in red. While this example is indisputably Roman in content, there is nothing to suggest Britain. The final image, Figure 33, also shows soldiers, this time in a watercolour style and next to a stone-arch building, the upper story of which has a temple-like structure.



Figure 6. OpenArtAi 6.



Figure 7. ImagineArt 1.

A similar temple building is seen on the next hilltop; the whole scene is set among rolling green hills and a blue sky peppered with clouds.

#### 8. Recraft

Finally, Recraft (<https://www.recraft.ai/ai-image-generator>) produces only two free images. The first of these, Figure 34, is a picture of ruined pillars set in a damp, green landscape, with a tree-covered hill and a misty, cloudy sky, while the second, Figure 35, is also set in a green, rolling landscape and features two Roman soldiers, patrolling between a ruined stone building and a stone wall, possibly recalling Hadrian's Wall, although this would hardly have been in ruins in the Roman period. Moreover, closer inspection



Figure 8. ImagineArt 2.



Figure 9. ImagineArt 3.

reveals that one of the soldiers is dressed in a long, gold-embroidered, red, fringed dress under his armour and carries a medieval kite shield, and the other appears to be wearing shorts, while both have closed shoes rather than *caligae*.

It should be stressed, however, that the prompt used in this experiment was relatively crude; perhaps a more sophisticated request could produce more accurate images. With that in mind, one final set of images were produced, this time with more specific prompts.

Using Recraft again, the prompt used was 'a scene of daily life in Roman Britain in a photographic style'; the result, Figure 36, is an



Figure 10. ImagineArt 4.



Figure 12. Microsoft Designer 2.



Figure 11. Microsoft Designer 1.



Figure 13. Microsoft Designer 3.

image that resembles a staged television set rather than an authentic ancient scene, showing market stalls and a paved street, with a ruined temple on one side and a castellated fort on the other. Two figures are in the foreground, one in a brown cloak and the other in a more Roman red cloak, and a third figure, who seems to be the salesman, dressed in an embroidered tunic and trousers, stands between them, but other figures in modern dress are visible in the rest of the scene. In other words, there are a few Roman elements, but many inaccuracies as well, and nothing really to indicate Britain. Attempting to address this, one final adjustment was made to the prompt, now asking for an image that portrayed, ‘a

scene of daily life in Roman Britain in a photographic style including recognisable British elements’ (Figure 37). This resulted in little more than the addition of a Tudor beamed building and a stone keep above which what seems to be the flag of England is flying.

### Overall results

One of the most obvious points that emerged from the experiment is that stereotyped visual patterns merged across platforms, with recurring motifs of ruined temples, red-cloaked soldiers, generic



Figure 14. Microsoft Designer 4.



Figure 16. Adobe firefly 2.




Figure 15. Adobe firefly 1.




Figure 17. Adobe firefly 3.

‘classical’ columns, marketplaces, and idealised green landscapes. These elements reflect dominant modern visual tropes of ‘Rome’ rather than specific features of Roman Britain. The findings demonstrate that GenAI tools tend to reproduce widely circulated cultural imagery rather than historically grounded reconstructions.

It was also clear that periods and geographies are conflated, with many outputs blending Roman, medieval, Renaissance, and even modern architectural elements. Norman keeps, Tudor buildings, St Paul’s Cathedral, medieval shields, and modern glass windows appeared alongside Roman iconography. This indicates that the models do not meaningfully distinguish between different periods

but instead aggregate visually associated features from their training data.

In addition, the structural limits of the source material quickly became apparent. Because diffusion models depend on existing visual material, they are constrained by the available imagery in their datasets. Since there are no photographs of ancient Roman Britain – only archaeological remains and modern reconstructions – AI models rely heavily on cinematic, artistic, and popular representations. This structural limitation helps explain both the strengths (general atmosphere) and weaknesses (incorrect detail) of the results.



Figure 18. Adobe firefly 4.



Figure 20. Ideogram 2.



Figure 19. Ideogram 1.



Figure 21. Ideogram 3.

The results of this experiment back up those found elsewhere. A similar exercise conducted by one blogger (<https://blog.myli.page/generative-archaeology-ai-created-realistic-images-of-ancient-worlds-b8bf7f53926e>) that explored some AI-generated artwork about ancient Assyria produced comparable conclusions:

The AI failed miserably at tasks that 1) require attention to too many details and 2) lack adequate training data . . . . On the other hand, it does an excellent job in themes that are well-explored by artists throughout history . . . scenes with fewer details . . . and portraits of single persons.

Although the ‘Roman Britain’ experiment was conducted almost a year and a half later – a very long time in terms of technological advancement – it is striking that the results are remarkably similar, suggesting that improvement is happening only slowly and that the technology is still far from perfect. Likewise, a recent experiment in which GenAI images were created in response to an ancient Greek text also concluded that cases in which there was a lack of digital corpora represented in AI training models resulted in inaccuracies (Kalargirou et al. 2025, 15).

Despite the drawbacks of the GenAI tools, the images are often visually compelling and stylistically polished, and their aesthetic



Figure 22. Ideogram 4.



Figure 23. Freepik 1.

persuasiveness risks granting them an unwarranted aura of authenticity. This has important implications, since if such depictions circulate widely, they may subtly shape public understanding of antiquity and become part of virtual cultural heritage. Uncritical use of GenAI images may reinforce misconceptions. Nevertheless, if used carefully, critical engagement with these outputs can foster analytical skills, encourage historical analysis,





Figure 24. Freepik 2.



Figure 25. Freepik 3.



**Figure 26.** Freepik 4.



**Figure 28.** Freepik 6.



**Figure 27.** Freepik 5.



**Figure 29.** Freepik 7.



Figure 30. Canva 1.



Figure 31. Canva 2.

and expose embedded biases related to race, gender, and power. Overall, the experiment reveals both the creative power and epistemological fragility of AI-generated historical imagery.

### Cultural heritage

The aim of this experiment was not to mock AI technology or to assert human superiority, but to examine the broader implications of the findings and their potential impact on contemporary cultural heritage. Cultural heritage refers to the culture, values, and traditions that are passed down from generation to generation.



Figure 32. Canva 3.



Figure 33. Canva 4.

It includes tangible items, which are derived from people's interaction with the physical environment, and include museums, buildings, art, and monuments etc., and intangible items such as language, music, customs, and skills (Vecco 2010, 323). There is now an emerging third domain of cultural heritage: virtual heritage. This comprises hardware heritage (i.e., computers, keyboards, storage media, printers), digital artefacts (i.e., items of material culture generated by computers, such as paper printouts or 3D products), virtual artefacts (i.e., virtual, computer-generated content that is visually or auditorily perceivable by humans via computer screens or speakers), latent digital signatures (i.e., volatile data written on ferromagnetic or optical media surfaces),



Figure 34. Recraft 1.



Figure 36. Recraft 3.



Figure 35. Recraft 2.



Figure 37. Recraft 4.

and digital ephemera (i.e., programs, interactions, and content performed and generated by computers without a tangible output) (Spennemann 2024, 3603). Whether any of the components of virtual heritage are deemed significant will depend on the cultural values ascribed to them and the role they played in culture and history.

GenAI is affecting how cultural heritage is managed and practised, primarily by providing analysis and decision-making tools (Foka and Griffin 2024; Spennemann 2024), but it is also becoming used as a 'go to' source of information by the public (Huschens et al. 2023, 11–12). As Foka and Griffin state,

The rise of generative AI models like ChatGPT and DALL-E has captured the public imagination; cultural and creative sectors increasingly turn to predictive models for analysing and categorising their materials and for performing what we would think of as creative or intellectual processes (Foka and Griffin 2024, 6128).

However, as the 'Roman Britain' experiment demonstrated, the responses provided by GenAI models have limited depth and occasionally suffer from inverted logic. Unlike many digital tools – such as word processors, spreadsheets, and image-editing

software – where processes unfold visibly and sequentially, GenAI systems operate as ‘black boxes’. The stages that lead to the final output remain opaque to the user (Spennemann 2024, 3601). The specific composition of these datasets is often uncertain or deliberately undisclosed. This opacity makes it difficult to determine how, or at what stage, inaccuracies arise. As a result, flaws can be hard to correct and, therefore, perpetuated further. It is entirely possible that over time AI images, such as those produced by the experiment, may enter the domain of cultural heritage and become accepted as an authentic representation of the past.

### GenAI and the reception of the ancient world in the classical classroom

In light of the recent spread of AI as a widely accessible tool, it is important to consider the role that GenAI in general and GenAI images in particular might play in the reception of antiquity. There is little doubt that the technology will improve rapidly, but at present, as has been demonstrated, there is a danger that misconceptions will be formed because of reliance on AI-generated material. As educators, it is vital to recognise this, and, being aware of the limitations, exercise caution, but at the same time to understand that in education, in popular culture, and in wider society such images may come to be regarded as ‘the truth’. If the images created by this technology are valid creations, and part of cultural heritage, they may take on a life of their own and influence ideas about the classical past.

This raises the issue of how such developments may be employed constructively. In this regard, the technological sophistication of these systems should not be overlooked; these images, when used within the constraints that are built into them, can produce pictures that are vivid and stimulating. One example of this is a GenAI recreation of the Seven Wonders of the Ancient World shown in the *Daily Mail* in December 2023 as the technology was becoming widespread (Fidler 2023). Another creative attempt by a blogger brought to life a relief from the Metropolitan Museum in New York, in which the AI produced a striking and colourful authoritative figure, with strong visual impact (<https://blog.myli.page/generative-archaeology-ai-created-realistic-images-of-ancient-worlds-b8bf7f53926e>).

Such images, and many others, have great potential within education, where they provide a valuable and exciting teaching opportunity. One interesting methodological idea for using AI image generators in the classroom has already been outlined in a paper in the *Journal of Classics Teaching* from November 2023 (Díaz-Sánchez and Chapinal-Heras 2023). According to this proposal, students search for ancient literary sources about a particular topic, translate them, and then, using an image generator, create and then modify an image to complement the piece before presenting it orally to the class. This process, as the authors argue, rather than acting as a replacement for the instructor, enhances the learning experience, as the versatility of AI provides an effective and dynamic resource, enabling students to develop their analytical skills.

Creating and analysing GenAI images is a valuable exercise, both in the classroom and beyond. It develops the ability to evaluate critically the quality and reliability of AI outputs. GenAI pictures can provide a thought-provoking starting point for discussion about questions such as the nature of the content incorporated, how convincing the image is, what elements may be problematic, and if any biases are revealed. Using GenAI tools also enables anyone, no matter how lacking in artistic skill, to produce

pictures. This applies to educators, who can use the tools to produce illustrations, games, and activity sheets, and also to students, who can be instructed to employ image generators to create pictures that, for example, illustrate an aspect of ancient history, or a scene from drama or literature that has been taught. If they are then encouraged to critique their own and others’ outputs, the exercise can be very productive. Moreover, because the source material available for representing the ancient world is inherently limited, GenAI outputs reveal the constraints under which these systems operate. Examining such images therefore illuminates not only the ancient past but also the mechanisms and constraints of the technology itself.

The advantages of using AI in the classroom in this manner are illustrated by a few concrete examples from personal experience. In one case, after reading several books of the *Iliad* in translation, students worked in groups using an AI tool to produce an illustration of a scene from the work. They then critiqued their own pictures and presented them to the other groups, inviting further comments. On another occasion, the class in a different course was tasked with producing posters advertising products of classical reception, which they then presented to the other students. In a third example, each group on a course produced an image from ancient history or mythology and the other groups were asked to identify the subject of the scene.

During these exercises, students, while enjoying the creativity and power of the tool, repeatedly expressed frustration at its inability to produce the kind of image they had intended. Furthermore, in the subsequent guided discussion, they noted inaccuracies, underlying biases, and assumptions of which they may not otherwise have been aware. This led not only to a deeper understanding of the subject matter being taught but also to a wider discussion of prejudice, expectation, and the limits of AI, which had relevance more broadly in their lives.

Exercises such as these have a positive impact on students’ learning. One recent study concluded that not only does the visual stimulus of an image support memorisation of information, but it also has a positive impact on motivation and student satisfaction (Berg et al. 2024, 503–505). Encouraging students to evaluate the strengths and weaknesses of such outputs fosters awareness not only of the ancient world but also of the broader uses and misuses of images. Analyses like this can then deepen sensitivity towards not just accuracy but also objectivity, as they can consider issues such as power, race, and gender that, as has been shown, underlie the GenAI images. These creations can therefore be used to highlight, and hopefully to alter, preconceptions and misconceptions; and in this way the tools may themselves make a positive contribution to the ongoing chain of receptions.

### Conclusion

This study set out to investigate how contemporary GenAI image tools visually interpret the prompt ‘Roman Britain’ and what such interpretations reveal about the reception of the ancient world in the digital age. The findings demonstrate that while GenAI technology can produce visually striking and atmospherically persuasive images, it currently lacks the capacity to generate reliably authentic representations of the ancient British province. Across multiple platforms, outputs relied heavily on generic ‘classical’ motifs – ruined temples, red-cloaked soldiers, arches, and idealised landscapes. They also frequently conflated Roman, medieval, and

modern architectural forms. Even when prompts were refined, inaccuracies persisted.

These limitations are not accidental but structural. Diffusion models generate images based on statistical associations drawn from existing internet imagery. In the case of antiquity, this imagery consists largely of archaeological remains, artistic reconstructions, cinematic depictions, and popular visual stereotypes. As a result, AI-generated images reflect modern receptions of the past rather than historically grounded reconstructions of it. Nevertheless, this does not render GenAI irrelevant or unusable. On the contrary, such images are likely to become increasingly integrated into digital culture and may themselves form part of emerging virtual heritage. Their influence on perceptions of antiquity will therefore grow.

The crucial issue is not whether GenAI images should be used, but how they should be used. When approached critically – particularly within the classroom – they can serve as powerful tools for discussion, encouraging students to question assumptions, identify inaccuracies, and explore how visual culture shapes historical understanding. In this way, GenAI images can become not passive conveyors of misinformation but active instruments in the study of reception, bias, and the construction of cultural memory. Ultimately, GenAI does not reconstruct the ancient past; it reconstructs our accumulated modern imaginings of it. Recognising this distinction is essential if educators, scholars, and the public are to engage responsibly with this rapidly developing technology.

## References

- Ahmad, I. (2024) The Growing Popularity of AI Image Generation: How Different Generations Are Learning to Prompt Creatively. *Digital Information World*, 24 November. Available at <https://www.digitalinformationworld.com/2024/11/ai-and-the-future-of-creativity-insights-from-a-new-study.html>.
- Berg, C., Omsén, L., Hansson, H. and Mozellius, P. (2024) Students' AI-generated images: Impact on motivation, learning and, satisfaction. Proceedings of the International Conference on AI Research, ICAIR 4. <https://doi.org/10.34190/icaire.4.1.3243>.
- Coulter, M. (2024) Aiming for fairness: An exploration into Getty Images v. Stability AI and its importance in the landscape of modern copyright law. *DePaul J. Art, Tech. & Intell. Prop. L.* 34, 124. Available at <https://via.library.depaul.edu/jatip/vol34/iss1/4>.
- Croitoru, F.A., Hondru, V., Ionescu, T. and Shah, M. (2022) Diffusion models in vision: A survey. *IEEE transactions on pattern analysis and machine intelligence* 45(9), 10850–10869. Available at <https://arxiv.org/pdf/2209.04747>.
- Díaz-Sánchez, C. and Chapinal-Heras, D. (2023) Use of Open Access AI in teaching classical antiquity. A methodological proposal. *Journal of Classics Teaching* 25(49). <https://doi.org/10.1017/S2058631023000739>.
- Epstein, Ziv., Hertzmann, A., Herman, L., Mahari, R., Frank, M.R., Groh, M., Schroeder, H., Smith, A., Akten, M., Fjeld, J., Farid, H., Leach, N., Pentland, A. and Russakovsky, O. (2023) Art and the science of generative AI: A deeper dive. <https://doi.org/10.48550/arXiv.2306.04141>.
- Fidler, K. (2023) AI imagines Seven Wonders of the Ancient World – and they're glorious. *The Daily Mail*, 20 December. Available at <https://metro.co.uk/2023/12/20/ai-imagines-seven-wonders-ancient-world-glorious-20004909/>.
- Foka, A. and Griffin, G. (2024) AI, cultural heritage, and bias: Some key queries that arise from the use of GenAI. *Heritage* 7, 6125–6136. <https://doi.org/10.3390/heritage7110287>.
- Gengler, E.J. (2024) Sexism, racism, and classism: Social biases in text-to-image generative AI in the context of power, success, and beauty. *Wirtschaftsinformatik* 48. Available at <https://aiselaisnet.org/wi2024/48>.
- Gioia, D.A., Hamilton, A.L. and Patvardhan, D.S. (2014) Image is everything: Reflections on the dominance of image in modern organizational life. *Research in Organizational Behavior* 34, 129–154. <https://doi.org/10.1016/j.riob.2014.01.001>.
- Huschens, M., Briesch, M., Sobania, D. and Rothlauf, F. (2023) Do you trust ChatGPT?—perceived credibility of human and AI-generated content. <https://doi.org/10.48550/arXiv.2309.02524>.
- Jacobi, T. and Sag, M. (2024) We are the AI problem. *Emory LJ Online* 74, 1. Available at <https://scholarlycommons.law.emory.edu/elj-online/50>.
- Joffe, H. (2008) The power of visual material: Persuasion, emotion and identification. *Diogenes* 55(1), 84–93. <https://doi.org/10.1177/0392192107087919>.
- Kalargiour, A., Kotsifakos, D. and Douligeris, C. (2025) The impact of Ancient Greek prompts on artificial intelligence: Image generation: A new educational paradigm. *AI* 2025, 6(4), 81. <https://doi.org/10.3390/ai6040081>.
- Maurice, L. (2016) Building a new ancient Rome. In Antony, A. and Monica, C. (eds.), *STARZ Spartacus: Reimagining an Icon on Screen*. Edinburgh University Press, pp. 111–130.
- Picascia, S. and Ferrara, A. (2026) How machines see and generate images. In Conte, P., Anna Caterina, D., Maria Giulia, D. and Andrea, P. (eds.), *Algomedia. The Image at the Time of Artificial Intelligence. Lecture Notes in Morphogenesis*. Cham: Springer. pp. 37–58. [https://doi.org/10.1007/978-3-032-08726-3\\_3](https://doi.org/10.1007/978-3-032-08726-3_3).
- Spennemann, D.H.R. (2024) Generative artificial intelligence, human agency and the future of cultural heritage. *Heritage* 7, 3597–3609. <https://doi.org/10.3390/heritage7070170>.
- Vecco, M. (2010) A definition of cultural heritage: From the tangible to the intangible. *Journal of Cultural Heritage* 11(3), 321–324. <https://doi.org/10.1016/j.culher.2010.01.006>.
- Zhang, C., Zhang, C., Zhang, M., Kweon, I.S. and Kim, J. (2023) Text-to-image diffusion models in generative AI: A survey. arXiv preprint arXiv:2303.07909. <https://doi.org/10.48550/arXiv.2303.07909>.