

Artificial Intelligence and Archaeological Illustration

Matthew Magnani  and Jon Clindaniel 

ABSTRACT

The reconstruction and representation of ancient artifacts and scenes through illustration is a cornerstone in the communication of archaeological findings. Sketches of the past have transformed over time, incorporating broader technological changes, from photography to the digital tools that have become prevalent through the twenty-first century. Most recently, developments in generative artificial intelligence (AI) promise to reshape the way we represent the past to professional and public audiences. This article shows how to use an accessible and inexpensive artificial intelligence platform to generate complex archaeological illustrations. As a case study, we create multiple scenes representing competing hypotheses about Neanderthal behavior. Using the images to visually communicate alternative hypotheses, we demonstrate how archaeological illustration using artificial intelligence promises to democratize the production of visual representations of the past.

Keywords: artificial intelligence, archaeological illustration, DALL-E 2

La reconstrucción y representación de artefactos antiguos y escenas a través de la ilustración es un pilar fundamental en la comunicación de hallazgos arqueológicos. Las ilustraciones del pasado han evolucionado con el tiempo, incorporando cambios tecnológicos más amplios, desde la fotografía hasta las herramientas digitales que se han vuelto predominantes a lo largo del siglo XXI. Más recientemente, los avances en la inteligencia artificial generativa prometen reconfigurar la forma en que representamos el pasado ante audiencias profesionales y públicas. En este artículo se muestra cómo utilizar una plataforma de inteligencia artificial accesible y económica para generar ilustraciones arqueológicas complejas. A través de un estudio de caso, creamos múltiples escenas que representan hipótesis competidoras sobre el comportamiento neandertal. Utilizado para comunicar visualmente hipótesis alternativas, demostramos cómo la ilustración arqueológica utilizando inteligencia artificial promete democratizar las representaciones visuales del pasado.

Palabras clave: inteligencia artificial, ilustración arqueológica, DALL-E 2

For centuries, the illustration of archaeological findings has been an essential part of communicating the past to other scholars and the public. Visual representations of artifacts, landscapes, and prehistoric scenes have been used to disseminate archaeological interpretations, reassembling unearthed material in scenes depicting past technologies, lifeways, and environments (Adkins and Adkins 1989).

These processes of making visual representations have drawn on the technologies, techniques, and conventions of their times, and they have played a role in the establishment of archaeology as a scientific discipline (Moser 2014). Line illustration, engravings, and painting in diverse media came to be complemented by photography in the nineteenth century (Bohrer 2011). Computer-based techniques and later three-dimensional modeling came to be layered alongside their predecessors through the late twentieth century into the twenty-first century (Magnani 2014; Moser 2012). Today, archaeologists commonly employ technologies like photogrammetry to visualize artifacts smaller than a few centimeters or as large as landscapes (Magnani et al. 2020).

Despite technological advances, visualizing the past remains a specialized undertaking, requiring high skill—whether conventional or digital—financing, and time. Illustrators are expensive. Digital techniques, from illustration to three-dimensional modeling, despite increasing in ubiquity and ease of practice, require software, know-how, and equipment. For this reason, depictions of the past have remained in the hands of just a few skilled individuals, often under the purview of large projects or institutions.

Recent innovations in artificial intelligence and machine learning have pushed the conversation in new directions. Now with simple text prompts, complex and original images can be generated within seconds. Although some have touted these transformations as detrimental to the employment of artists or as risks to the educational system (Huang 2023; Roose 2022), technological advances have the potential to democratize the imagination of the archaeological past through illustration. In archaeology, the conversation is just beginning about how generative AI will reshape education, research, and illustration (Cobb 2023). In this vein, we present a workflow for scholars and students of archaeology as an

Advances in Archaeological Practice 11(4), 2023, pp. 452–460

Copyright © The Author(s), 2023. Published by Cambridge University Press on behalf of Society for American Archaeology. 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 in any medium, provided the original work is properly cited.

DOI:10.1017/aap.2023.25

alternative past-making exercise, thinking through the impacts of technological shifts on archaeological representation. In the months and years to come, archaeological illustrations made using AI will find utility in publications and museum exhibitions. Through a case study related to Neanderthal–modern human interactions, this article presents a methodology for visualizing archaeological evidence using artificial intelligence.

METHODS

Specifically, this article shows readers how to use DALL-E 2, a generative AI model trained and released by OpenAI in 2022 to create images that correspond to textual input. DALL-E 2 has minimal user hardware requirements and can be used from any computer or smartphone with Internet access. DALL-E 2 employs the “Transformer” deep learning architecture (Vaswani et al. 2017)—a type of model that was originally designed for machine translation tasks (e.g., translating an English sentence into a corresponding French sentence of arbitrary length). In recent years, however, such Transformer models have been found to be extraordinarily effective in a wide variety of other contexts that involve generating an output sequence of data that corresponds to an input sequence of data.

In natural language processing tasks, for instance, Transformer models are now the preferred approach for generating answers to questions, condensing long texts into summaries, and a wide variety of text classification tasks (Brown et al. 2020; Devlin et al. 2019; Radford et al. 2018, 2019; Raffel et al. 2020). Despite their origin in text data processing, Transformers have found success in modeling other modalities as well, such as images, which are flattened into a sequence of image “patches” akin to words in a sentence in order to be processed in the same way as text by the Transformer (Dosovitskiy et al. 2020).

This capacity to consider images as sequences of data akin to text is what makes it possible to train Transformer models to produce images corresponding to input text. For instance, in recent years, there have been a variety of models that have been successfully trained to directly perform text-to-image translations, such as the first version of DALL-E (Ramesh et al. 2021), CogView (Ding et al. 2021), Make-A-Scene (Gafni et al. 2022), and Google’s Parti (Yu et al. 2022). More recently, diffusion models have also been successfully incorporated into text-to-image generation—DALL-E 2 (Ramesh et al. 2022) and Google’s Imagen (Saharia et al. 2022), for example. Although there are many potential text-to-image models to choose from, we employ DALL-E 2 in this study due to its widespread popularity and continued use as a benchmark for state-of-the-art performance in the text-to-image generation field. The strategies of iterative AI prompt engineering that we employ in this article, however, can be translated for use with other models as well, and we expect the underlying strategies to remain relevant for prompting text-to-image models into the future.

DALL-E 2 builds on the work of previously trained Transformer encoder and decoder models to do the following (Ramesh et al. 2022):

- (1) Encode text input as a CLIP embedding—a fixed-size, quantitative representation of the text that contains semantic and contextual information and is generated from the CLIP model.

CLIP is a separate model that was trained on 400 million pairs of images and associated captions drawn from across the Internet to produce a jointly learned image/text embedding space—a latent space that allows us to quantify the relationship between any text and image element. CLIP has been found to accurately predict known image categories out of the box, and as a result of it being trained on such a large and noisy web dataset, its performance is robust and generalizable to unseen domains (Radford et al. 2021). On the basis of a text embedding, DALL-E 2 has been trained to generate a possible CLIP image embedding that is closely related to the input text.

- (2) Stochastically generate a high-resolution output image conditioned on the CLIP image embedding generated in the first step and (optional) additional text prompts. This text-conditioned diffusion decoder model builds off of the GLIDE model (Nichol et al. 2021) and has been trained to upsample images up to 1024×1024 resolution, generate a diverse range of semantically related output images based on a single input image embedding (the decoding process is stochastic and nondeterministic), and also perform text-based photo edits.

There are currently two options for generating and editing images via DALL-E 2: an online graphic user interface (GUI) and a developer application programming interface (API). In the online GUI, users enter a text description of their desired image, and they are provided an output batch of four corresponding images. From this initial batch of photos, an individual image can be selected, edits can be made to the selected image, and more variants can be generated in an iterative process. The API provides the same generation and editing features as the GUI but makes it possible to interact with DALL-E programmatically via languages such as Python. The online GUI is credit based: each iteration (the generation of one batch of images, an edit request, or a variation request) costs one credit (Jang 2023), each of which was approximately 13 cents at the time of writing. The API, on the other hand, is charged by the number of images generated and the resolution of the images, and as of the time of writing, it costs 2 cents per 1024×1024 image generated.

Ethics of Illustration and Artificial Intelligence

As an emergent technology, ongoing discussions about the ethical dilemmas posed by artificial intelligence are nascent, momentous, and shared beyond archaeology with society at large (Hagerty and Rubinov 2019). Yet archaeologists will face unique challenges posed by AI, including not only its impact on the interpretation and depiction of the past but also the myriad of ways these reorganizations trickle into labor practices and power dynamics in our own field. Compared to other shifts in technology—for instance, digital revolutions that have changed the way we visualize archaeological subjects—artificial intelligence has the potential to reorder practice more rapidly. Just as we do not see an end to academic writing in sight, we do not anticipate a diminished need for specialized archaeological illustrators. As with other technological developments in the field, AI should complement and expand our ability to work rather than overturn the skillsets required to do so.

In particular, anthropologists and artists will need to support sustained discussion about the potential social biases and inequities that are pervasive in the field and reproduced by the technology

(see Lutz and Collins [1993] for a discussion of photographic representations for comparison). AI draws on a myriad of images and text scraped from the Internet to inform its output. The resulting images may tend to reproduce visual representations that reflect broader inequalities and troublesome trends in our society (Birhane et al. 2021). Because of this, a thoughtful review of AI-generated output, as well as a reflexive approach to AI-prompt engineering, is necessary. Just as AI has the potential to reiterate contemporary stereotypes through the presentation of the archaeological record, it also has the potential to help conceive alternative pasts. Anthropologists have the capacity to critically integrate these technological advances into practice, elucidating the risks but also the potentials of new tools.

Case Study

Significant interest in depicting the lifeways of closely related species, including Neanderthals, has occupied antiquarians and archaeologists for generations. Past sketches reflected the stereotypes and biases about the species at the time they were made. For instance, early twentieth-century illustrations portrayed Neanderthals as hairy, brutish cave dwellers but transformed over time to depict increasingly modern humanoids (Moser 1992). Later discussions about their interactions have considered exchange

(or lack of exchange) through an analysis of material culture (e.g., d’Errico et al. 1998) and skeletal morphology (Trinkaus 2007), whereas most recently, the burgeoning field of paleogenomics conclusively demonstrates admixture (Sankararaman et al. 2014). The application of intersectional Black feminist approaches to the study of Neanderthal–modern human interactions suggests a continued gendered and race-based bias against this closely related species (Sterling 2015).

The karst systems of southern France have long been a battleground for ideologies about Neanderthal behavior, and by extension, our relation to them as a species. The caves’ Neanderthal stone-tool industries figured prominently in the Binford/Bordes debate essential to the establishment of processual archaeology. Today, the region remains central to examinations of Neanderthal behavioral variability and their (dis)similarity to *Homo sapiens*. Did Neanderthals just control fire, or could they also make it (Dibble et al. 2017, 2018)? Did they bury their dead (sensu Balzeau et al. 2020) or unceremoniously die in caves (Sandgathe et al. 2011)?

This article presents a methodology using artificial intelligence to illustrate these conflicting archaeological hypotheses, visually representing alternative interpretations of Neanderthal

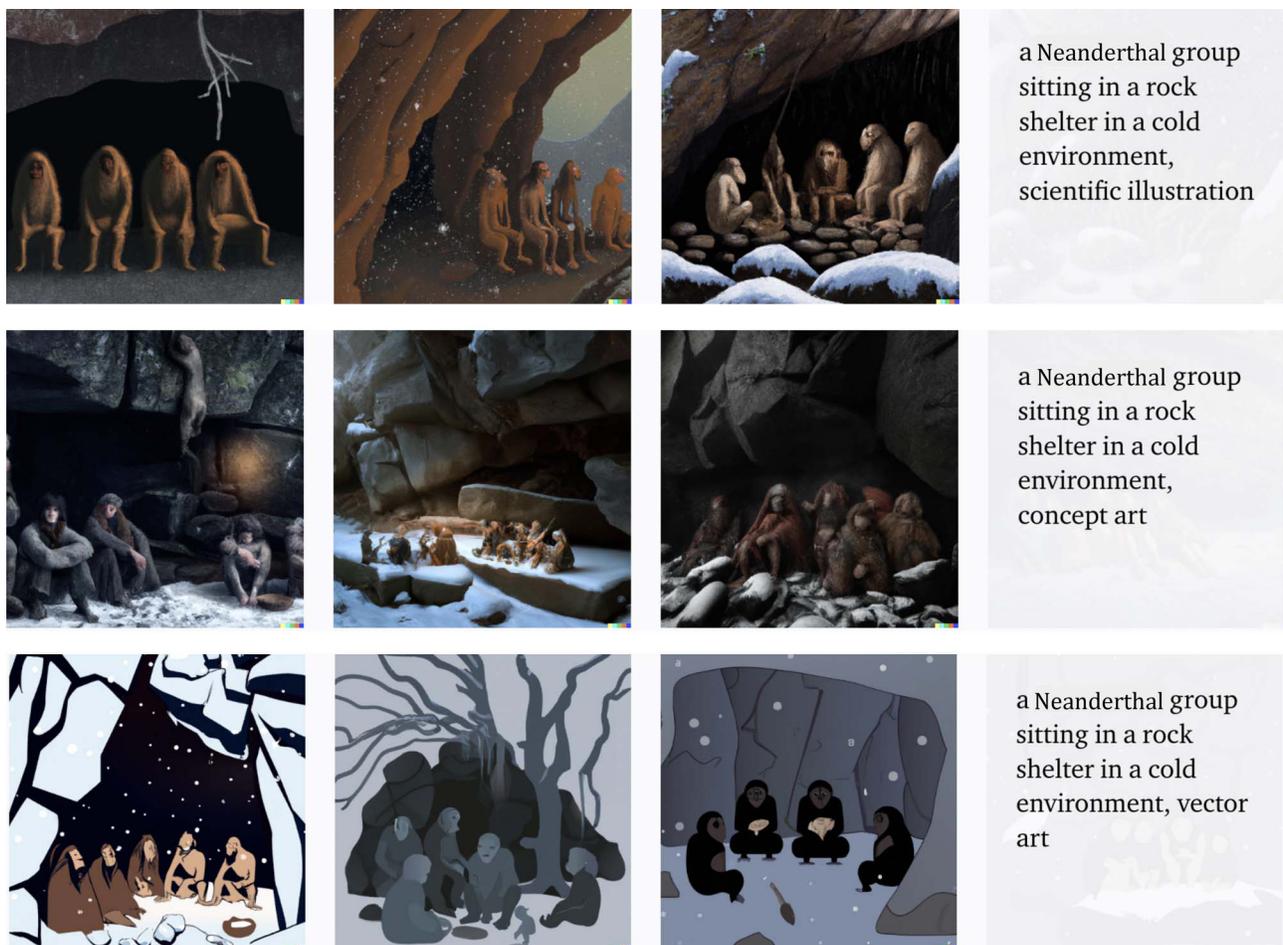


FIGURE 1. Text prompts exploring variable styles, including scientific illustration, concept art, and vector drawings. Image created by the coauthors using DALL-E 2.

behavioral variability. In the first generated image via artificial intelligence, we assume Neanderthals who neither made fire nor buried their dead. In the second, we consider a fire-making species that also buries its dead. We depict Neanderthals occupying a cave-like setting, adjusting their environmental context according to archaeological evidence from southern France. In the text that follows, we describe the process of using the DALL-E 2 GUI to generate and edit images. For more detailed information on using the API to accomplish the same tasks programmatically, consult our supplemental code (Clindaniel 2023).

RECONSTRUCTING ANCIENT SCENES

Step 1: Base Image Generation

Users begin by writing a base text prompt, encapsulating the desired image content, size, and style; larger images take longer to generate. Under the hood, DALL-E 2 converts our text input to corresponding visual information in the CLIP image/text embedding space and then stochastically decodes the resulting image

embedding as a high-resolution image corresponding to our original text input. We explored multiple styles of images prior to selecting one that appeared suitable for our purposes. For instance, a lack of a stylistic prompt generated unusable images, and a prompt indicating that the model should produce photographic realism similarly produced no usable base image. We tested additional styles featured in an online guide to DALL-E 2, including vector, concept art, and scientific illustration (see Figure 1). Through trial and error and visual inspection, we determined that a “digital art” style was conducive to visually communicating our hypotheses about Neanderthal behavioral variability. Other users will find other styles more helpful to communicate their desired scenes.

The keywords selected for our initial base illustration included a “a Neanderthal group sitting in a rock shelter in a cold environment, digital art.” Given that the DALL-E 2 decoding process (from image embeddings to output images) is stochastic and non-deterministic, entering the same keywords can produce different batches of images each time the model is instantiated. Using the aforementioned keywords, we generated three batches of images, resulting in 12 images total (see Figure 2).

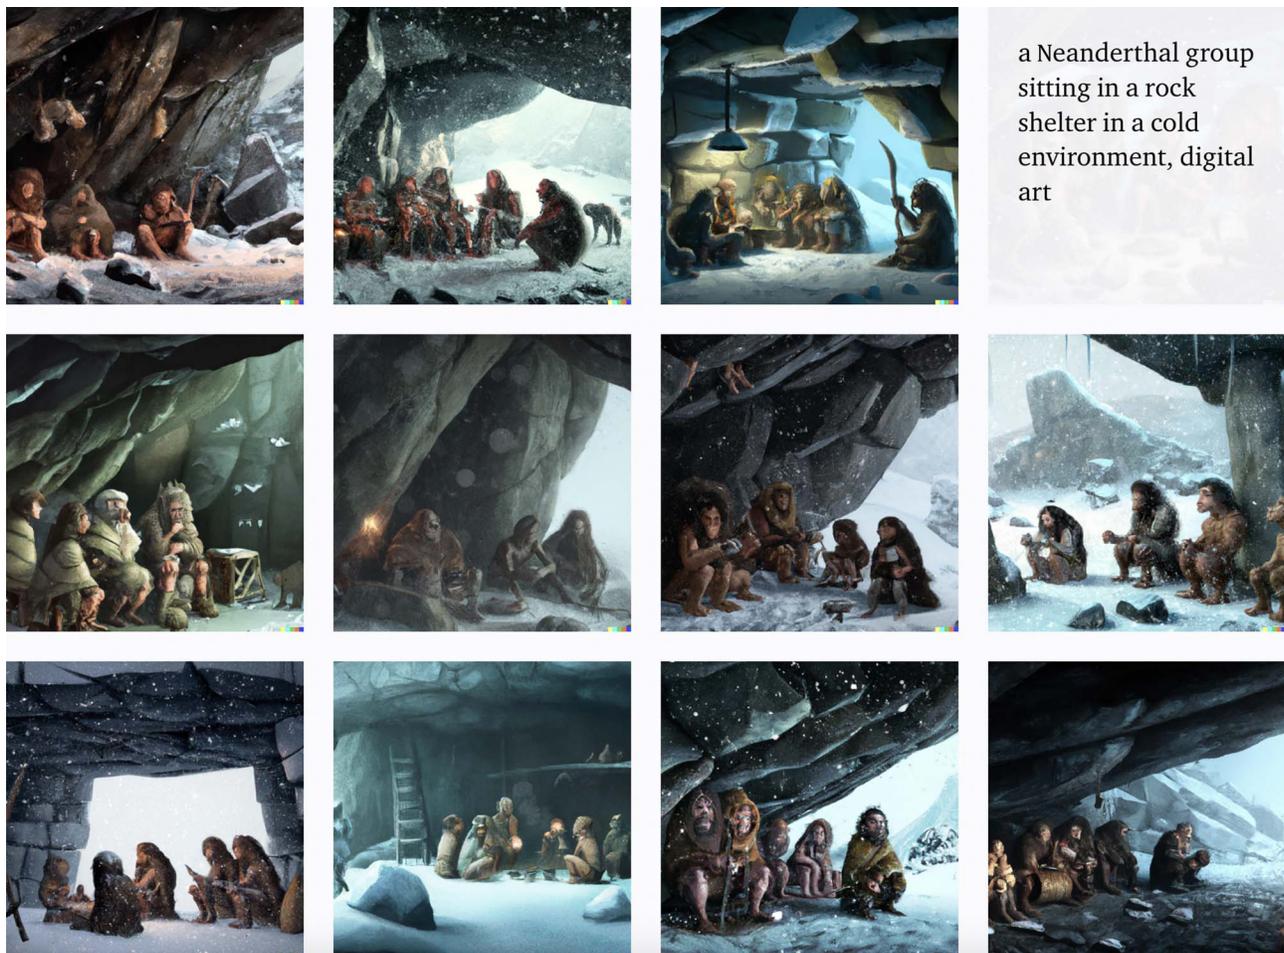


FIGURE 2. Three batches of images produced using the selected “digital art” style, from which we selected our base image. Image created by the coauthors using DALL-E 2.



FIGURE 3. Base image selected to represent a group of Neanderthals in a rock shelter in a cold-weather environment. Note the undesirable features in the image, including an anthropomorphic anomaly on the right. Image created by the coauthors using DALL-E 2.

Step 2: Base Image Selection

Given the wide range of image variants for each text prompt, we selected a base image to amend in subsequent steps through editing and integration of specific archaeological evidence.

We considered the suitability of our base images in terms of anatomical realism, cultural context, and environmental accuracy. For instance, some images returned with ape-like or stooped humanoids, whereas others incorporated groups sitting in built structures reminiscent of more recent modern-human constructions. These depictions reproduced age-old stereotypes about Neanderthals, which—although debunked through more than a century of archaeology—appear pervasive across Internet sources. Through all stages but especially during initial base-image generation, archaeologists should be critical of the kinds of implicit, past, or widespread biases that may shape the output. Our final selections presented an archaeologically sound depiction of a group occupying a cave site. Once a base image was selected as an acceptable starting point, we could begin to modify the illustration further (see Figure 3).

Step 3: Image Modification

In our third stage of archaeological illustration using artificial intelligence, we sought to modify our base image through deletion and insertion of specific archaeological evidence. We sought to create two different images, supporting hypotheses reported by different camps of archaeologists. To this end, we sought to represent the ability or lack of ability to make fire, and the burial or decomposition of Neanderthal remains in cave contexts in cold climates (marine isotope stage [MIS] 3 and 4). We integrated additional information about the primary diets of Neanderthals yielded from archaeological evidence, which was dominated by reindeer in cold climatic phases (Goldberg et al. 2012).

To do so, we modified our base images to integrate into the following prehistoric cultural environments (see Figure 4 for an example of the integration of reindeer, which required that two batches of four photos be generated before we determined that a suitable image was available). Our first image depicts the often-assumed



FIGURE 4. Two batches of photos generated using the prompt “Rangifer tarandus in the distance, digital art.” Image created by the coauthors using DALL-E 2.



FIGURE 5. Erasure of base layer section to produce mask for generating new content (i.e., “inpainting”), in this case a group of reindeer in the erased area. Image created by the coauthors using DALL-E 2.

behavioral practices of Neanderthals, occupying a cold environment rich in reindeer, practicing burial, and making fire. The second image generated represents the opposing hypothesis. Also a reindeer-rich, cold-weather environment, our second representation depicted a species that could not make fire and did not bury its dead.

DALL-E 2 was fine-tuned to learn how to fill in blank regions of an image, conditioned on optional text input (for details of this fine-tuning process, see Nichol et al. 2021). This makes it possible to mask portions of a base image that we want to stay the same and provide text prompts indicating objects that should be added in unmasked, transparent regions (in the same style as the base image). Such edits can be performed within the current borders of an image (called “inpainting”), or this masking procedure can also be used to extend images beyond their original borders (“outpainting”).

Step 3.1. In cases where it was necessary to change data with the borders of our base image, we began by deleting a fragment of the original image. Using the in-built DALL-E eraser tool in the online GUI, we created a space to add additional objects to our image using a new text prompt, following Step 1 above. For instance, to generate a group of reindeer in the background, we deleted a small section of our base-image background and used the new text prompt “Rangifer tarandus in the distance, digital art” (see Figure 5). We repeated this process serially to add or remove elements of an image—for instance, to add fire “a small campfire glowing on logs,

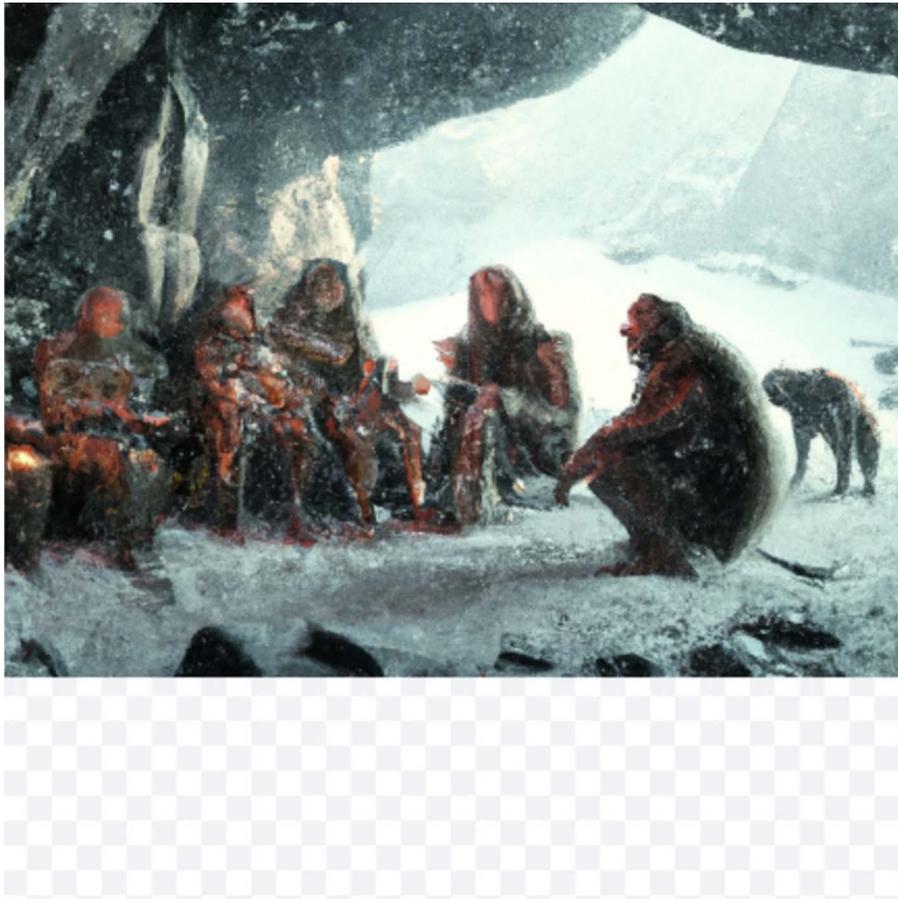


FIGURE 6. Expansion of frame downward to generate additional data representing Neanderthal remains in the cave foreground (i.e., “outpainting”). Image created by the coauthors using DALL-E 2.



FIGURE 7. A random selection of images depicting unburied Neanderthal remains. Over 75 batches of images using variable text inputs were generated before the final selection was made. Image created by the coauthors using DALL-E 2.

digital art”) or remove visual artifacts we determined were undesirable to remain in the image (deletions were replaced with prompts such as “cave wall” or “snowy background”).

Step 3.2. In cases where data needed to be added beyond the limits of the original scene, it became necessary to extend our

image. Using an add generation frame in the online GUI, we provided a new prompt and set the limits of our desired scene. For instance, when we wanted to add a scene suggesting that Neanderthals did not bury their dead, we needed to expand an activity area in the foreground of the image (see Figure 6). To do so, we used a prompt to generate unburied human remains.

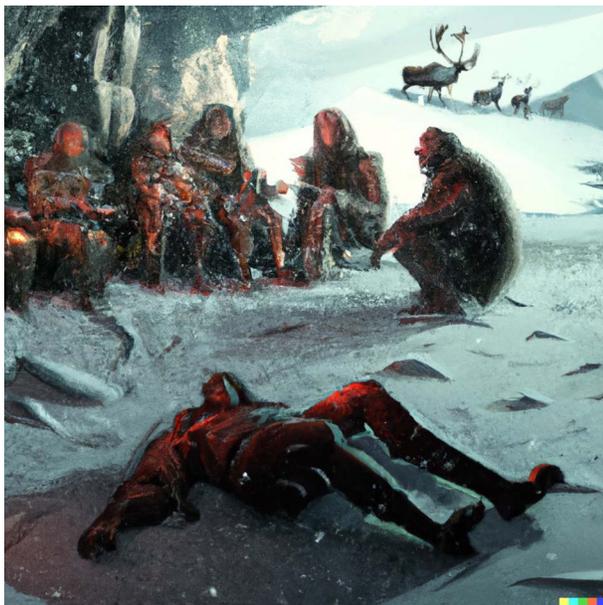


FIGURE 8. Image reflecting Neanderthals who could not create fire in cold weather climates with low environmental availability and did not inter their dead. The partially decomposed remains of a Neanderthal are in the foreground of the image. Image created by the coauthors using DALL-E 2.



FIGURE 9. Image reflecting Neanderthals who interred their dead and controlled fire during cold climatic phases. A young Neanderthal is being buried centrally in the image, with a fire to the left. Image created by the coauthors using DALL-E 2.

Again, we generated multiple batches of photos before considering which image best served the needs of the project—that is, reflected a Neanderthal consistent with current scientific understandings (e.g., anatomically correct, without buttons on clothing, metal, etc.; see [Figure 7](#), which is a random selection from more than 75 batches of images).

Step 4: Finalizing Image Selection

Following the integration of additional archaeological data and, in some cases, the expansion of the frame, we made our final selection of images portraying alternative hypotheses about Neanderthal cultural repertoires, including (1) a species that lacked the ability to make fire and that did not bury its dead, and (2) an illustration of a symbolic species that controlled fire and buried its dead (see [Figures 8 and 9](#)). Images were selected based on seamless inclusion of desired graphic elements in support of the selected hypotheses, representing behavioral repertoires reflected in competing hypotheses.

CONCLUSION

Artificial intelligence is poised to become more prevalent in our daily lives in the coming years and decades. Even in its infancy, artificial intelligence hints at its potential as a powerful tool to depict and interpret the past. Illustrations that once took weeks or months and significant capital can be rendered in hours. With nascent investments in the technology that are in the billions of dollars, the coming months and years will see unparalleled leaps forward in deep learning, artificial intelligence, and applications tailored to tackle the needs of archaeologists. As these technologies develop, archaeologists will need to be acutely aware of the ethical dilemmas that present themselves—often shared more broadly across society but in some cases unique to our field. Ultimately, the visions we have of the past will become legion in their ubiquity and ease of production, sometimes reinforcing systemic social biases, but most critically, providing the tools to present alternative interpretations of ancient life.

What does generative artificial intelligence mean for the field of archaeology? In the presented case study illustrating our methodology, we visualize multiple archaeological hypotheses about Neanderthal behavior using one program. However, the approach may be scaled to represent and interpret archaeological evidence more broadly, from any time or location, using a suite of ever-transforming algorithms. For the first time, artistic renderings of the record will become approachable by broader subsections of archaeologists and publics alike. Democratizing representations of the past, artificial intelligence will facilitate multiple perspectives and (re)interpretations to be generated and shared.

Acknowledgments

The authors would like to offer thanks to the reviewers who improved this piece through their constructive criticism.

Funding Statement

No funding was received for the research or writing of this article.

Data Availability Statement

All data and code reported in this manuscript are available on Zenodo: <https://doi.org/10.5281/zenodo.7637300>.

Competing Interests

The authors declare none.

REFERENCES CITED

- Adkins, Lesley, and Roy Adkins. 1989. *Archaeological Illustration*. Cambridge University Press, Cambridge.
- Balzeau, Antoine, Alain Turq, Saha Talamo, Camille Daujeard, Guillaume Guérin, Frido Welker, Isabelle Crevecoeur, et al. 2020. Pluridisciplinary Evidence for Burial for the La Ferrassie 8 Neandertal Child. *Scientific Reports* 10:21230. <https://doi.org/10.1038/s41598-020-77611-z>.
- Birhane, Abeba, Vinay Uday Prabhu, and Emmanuel Kahembwe. 2021. Multimodal Datasets: Misogyny, Pornography, and Malignant Stereotypes. *arXiv* 2110.01963v1. <https://arxiv.org/abs/2110.01963v1>, accessed August 24, 2023.
- Bohrer, Frederick. 2011. *Photography and Archaeology*. Reaktion Books, London.
- Brown, Tom, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, et al. 2020. Language Models Are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, edited by Hugo Larochelle, Marc'Aurelio Ranzato, Raia Thais Hadsell, Maria Florina Balcan, and H. Lin, pp. 1877–1901. Curran Associates, Red Hook, New York.
- Clindaniel, Jon. 2023. ai-illustration. Zenodo. <https://doi.org/10.5281/zenodo.7637300>, accessed October 10, 2023.
- Cobb, Peter J. 2023. Large Language Models and Generative AI, Oh My! Archaeology in the Time of ChatGPT, Midjourney, and Beyond. *Advances in Archaeological Practice* 11(3):363–369. <https://doi.org/10.1017/aap.2023.20>.
- d'Errico, Francesco, Joao Zilhão, Michèle Julien, Dominique Baffier, and Jacques Pelegrin. 1998. Neanderthal Acculturation in Western Europe? A Critical Review of the Evidence and Its Interpretation. *Current Anthropology* 39(S1):S1–S44. <https://doi.org/10.1086/204689>.
- Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding. *arXiv* 2010.11929v2. <https://doi.org/10.48550/arXiv.1810.04805>.
- Dibble, Harold L., Aylar Abodolazadeh, Vera Aldeias, Paul Goldberg, Shannon P. McPherron, and Dennis M. Sandgathe. 2017. How Did Hominins Adapt to Ice Age Europe without Fire? *Current Anthropology* 58(S16):S278–S287. <https://doi.org/10.1086/692628>.
- Dibble, Harold L., Dennis Sandgathe, Paul Goldberg, Shannon McPherron, and Vera Aldeias. 2018. Were Western European Neandertals Able to Make Fire? *Journal of Paleolithic Archaeology* 1(1):54–79. <https://doi.org/10.1007/s41982-017-0002-6>.
- Ding, Ming, Zhuoyi Yang, Wenyi Hong, Wendi Zheng, Chang Zhou, Da Yin, Junyang Lin, et al. 2021. CogView: Mastering Text-to-Image Generation via Transformers. *Advances in Neural Information Processing Systems* 34:19822–19835.
- Dosovitskiy, Alexey, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, et al. 2020. An Image Is Worth 16x16 Words: Transformers for Image Recognition at Scale. *arXiv* 2010.11929v2. <https://doi.org/10.48550/ARXIV.2010.11929>.
- Gafni, Oran, Adam Polyak, Oron Ashual, Shelly Sheynin, Devi Parikh, and Yaniv Taigman. 2022. Make-A-Scene: Scene-Based Text-to-Image Generation with Human Priors. In *Computer Vision – ECCV 2022, Lecture Notes in Computer Science* Vol. 13675, edited by Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, pp. 89–106. Springer, Cham, Switzerland.
- Goldberg, Paul, Harold Dibble, Francesco Berna, Dennis Sandgathe, Shannon J. P. McPherron, and Alain Turq. 2012. New Evidence on Neandertal Use of Fire: Examples from Roc de Marsal and Pech de l'Azé IV. *Quaternary International* 247:325–340. <https://doi.org/10.1016/j.quaint.2010.11.015>.

- Hagerty, Alexa, and Igor Rubinov. 2019. Global AI Ethics: A Review of the Social Impacts and Ethical Implications of Artificial Intelligence. *arXiv* 1907.07892v1. <https://doi.org/10.48550/arXiv.1907.07892>.
- Huang, Kalley. 2023. Alarmed by A.I. Chatbots, Universities Start Revamping How They Teach. *New York Times*, January 16, Technology section. <https://www.nytimes.com/2023/01/16/technology/chatgpt-artificial-intelligence-universities.html>, accessed February 8, 2023.
- Jang, Joanne. 2023. How DALL-E Credits Work OpenAI Help Center. *OpenAI.com*. <https://help.openai.com/en/articles/6399305-how-dall-e-credits-work>, accessed February 11, 2023.
- Lutz, Catherine A., and Jane L. Collins. 1993. *Reading National Geographic*. University of Chicago Press, Chicago.
- Magnani, Matthew. 2014. Three-Dimensional Alternatives to Lithic Illustration. *Advances in Archaeological Practice* 2(4):285–297. <https://doi.org/10.7183/2326-3768.2.4.285>.
- Magnani, Matthew, Matthew Douglass, Whittaker Schroder, Jonathan Reeves, and David R. Braun. 2020. The Digital Revolution to Come: Photogrammetry in Archaeological Practice. *American Antiquity* 85(4):737–760. <https://doi.org/10.1017/aaq.2020.59>.
- Moser, Stephanie. 1992. The Visual Language of Archaeology: A Case Study of the Neanderthals. *Antiquity* 66(253):831–844. <https://doi.org/10.1017/S0003598X0004477X>.
- Moser, Stephanie. 2012. Archaeological Visualisation: Early Artifact Illustration and the Birth of the Archaeological Image. In *Archaeological Theory Today*, 2nd ed., edited by Ian Hodder, pp. 292–322. Polity Press, Cambridge.
- Moser, Stephanie. 2014. Making Expert Knowledge through the Image: Connections between Antiquarian and Early Modern Scientific Illustration. *Isis* 105(1):58–99. <https://doi.org/10.1086/675551>.
- Nichol, Alex, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. 2021. GLIDE: Towards Photorealistic Image Generation and Editing with Text-Guided Diffusion Models. *arXiv* 2112.10741v3. <https://doi.org/10.48550/ARXIV.2112.10741>.
- Radford, Alec, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving Language Understanding by Generative Pre-Training. Electronic document, https://s3-us-west-2.amazonaws.com/openai-assets/research-covers/language-unsupervised/language_understanding_paper.pdf, accessed February 1, 2023.
- Radford, Alec, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, et al. 2021. Learning Transferable Visual Models from Natural Language Supervision. In *Volume 139: International Conference on Machine Learning*, edited by Marina Meila and Tong Zhang, pp. 8748–8763. Proceedings of Machine Learning Research Press, online.
- Radford, Alec, Jeffrey Wu, Dario Amodei, Daniela Amodei, Jack Clark, Miles Brundage, and Ilya Sutskever. 2019. Better Language Models and Their Implications. *OpenAI (blog)*. Electronic document, <https://openai.com/blog/better-language-models/>, accessed February 11, 2023.
- Raffel, Colin, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research* 21(1):1–67.
- Ramesh, Aditya, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical Text-Conditional Image Generation with CLIP Latents. *arXiv* 2206.10789. <https://doi.org/10.48550/arXiv.2206.10789>.
- Ramesh, Aditya, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. Zero-Shot Text-to-Image Generation. In *Volume 139: International Conference on Machine Learning*, edited by Marina Meila and Tong Zhang, pp. 8821–8831. Proceedings of Machine Learning Research, online. <https://proceedings.mlr.press/v139/>, accessed October 10, 2023.
- Roose, Kevin. 2022. An A.I.-Generated Picture Won an Art Prize. Artists Aren't Happy. *New York Times*, September 2, Technology section. <https://www.nytimes.com/2022/09/02/technology/ai-artificial-intelligence-artists.html>, accessed February 8, 2023.
- Saharia, Chitwan, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L. Denton, Kamyar Ghasemipour, et al. 2022. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. *Advances in Neural Information Processing Systems* 35:36479–36494.
- Sandgathe, Dennis M., Harold L. Dibble, Paul Goldberg, and Shannon P. McPherron. 2011. The Roc de Marsal Neandertal Child: A Reassessment of Its Status as a Deliberate Burial. *Journal of Human Evolution* 61(3):243–253. <https://doi.org/10.1016/j.jhevol.2011.04.003>.
- Sankararaman, Sriram, Swapan Mallick, Michael Dannemann, Kay Prüfer, Janet Kelso, Svante Pääbo, Nick Patterson, and David Reich. 2014. The Genomic Landscape of Neanderthal Ancestry in Present-Day Humans. *Nature* 507(7492):354–357. <https://doi.org/10.1038/nature12961>.
- Sterling, Kathleen. 2015. Black Feminist Theory in Prehistory. *Archaeologies* 11(1):93–120. <https://doi.org/10.1007/s11759-015-9265-z>.
- Trinkaus, Erik. 2007. European Early Modern Humans and the Fate of the Neandertals. *PNAS* 104(18):7367–7372. <https://doi.org/10.1073/pnas.0702214104>.
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. In *Advances in Neural Information Processing Systems*, edited by Ulrike von Luxburg, Isabelle Guyon, Samy Bengio, Hanna Wallach, and Rob Fergus, pp. 1–11. Curran Associates, Red Hook, New York.
- Yu, Jiahui, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, et al. 2022. Scaling Autoregressive Models for Content-Rich Text-to-Image Generation. *arXiv* 2206.10789. <https://doi.org/10.48550/arXiv.2206.10789>.

AUTHOR INFORMATION

Matthew Magnani ■ Department of Anthropology, University of Maine, Orono, ME, USA (matthew.magnani@maine.edu, corresponding author)

Jon Clindaniel ■ Computational Social Science, University of Chicago, Chicago, IL, USA