CHAPTER 19

# IMAGINE WATTEAU

**Oliver Wunsch**

It appears as if in a dream, emerging from a blur of pixels: Antoine Watteau's long-lost painting *Abigail Bringing Food to David*. Art historians know the work only by its title, which refers to the Old Testament story of Abigail mollifying the bellicose David and his men by offering them large amounts of food. The picture disappeared soon after Watteau submitted it to the French Royal Academy of Painting and Sculpture in 1709, leaving scholars to wonder how the French painter of amorous seduction had handled a biblical subject.[1] Yet here it is on my screen, or at least a digital simulation of it, produced by the AI image platform known as Midjourney (fig. 19.1).

Laugh, if you like, at the image's many deficiencies, its mangled hands and Cabbage Patch faces. Mock, if you must, its inscrutable details, such as the background figure who rends her garments for no discernible reason, or the bag-headed man on the right, standing alongside a server with a strangely poufy bonnet who pours a drink for no one in particular (fig. 19.2). But if you are like me, your amusement will come with a tinge of sadness. The feeling comes, I think, less from the shortcomings of the image than from the premise behind its creation, the futile impulse to recover the missing object. My response to the image therefore offers a prime example of the melancholy that Michael Ann Holly has identified behind all art historical study: "The melancholy that courses through the history of art is a product of its perhaps unconscious awareness that works that seem so present are actually absent."[2] Holly speaks here not only about missing and destroyed works of art, but also of the lost past to which surviving art once belonged, a world that art historians seek to recover while knowing the task will always remain incomplete.

The gloom that accompanies my simulated Watteau is, for the moment, an unusual case of the phenomenon that Holly describes, but it is one that may grow familiar as AI tools become available for art historical applications. From the early days of AI image generation, data scientists have cited art restoration as a

Fig. 19.1. Digital simulation of Antoine Watteau (French, 1684–1721), *Abigail Bringing Food to David* (1709), response to "Watteau's painting showing the Old Testament story of Abigail offering large quantities of food to David and his men," 2023

Fig. 19.2. Detail of fig. 19.1

potential use case for artificial intelligence, arguing that the algorithmic analysis of an artist's surviving work might allow for the recreation of what has been lost.[3] The sophistication of the technology has improved in recent years (notwithstanding the risible details of my Watteau), and collaborations between museums and computer scientists have yielded remarkably plausible results.[4]

What, though, do such images reveal? Based as they are on data about work that has endured, they fill holes with reconstituted versions of what we already possess. To convince us that they recuperate a forgotten past, they must remind us of what we know, flattering our preconceptions about how an artist's work ought to look. What they show, then, is less a lost fragment of history than a reflection of our own art historical desires. It is these desires that I wish to analyze via my artificially resuscitated Watteau. In both its glitches and strengths, it allows us to work through the sadness lurking within our discipline's emerging digital unconscious.

• • •

My image begins as text. Midjourney, like many other AI image generators, produces its results in response to "prompts" that users write in ordinary language (competitors such as Dall·E 3 and Stable Diffusion operate according to the same principle, though Midjourney's developers have sought to distinguish their product by claiming that it possesses a greater mastery of artistic terminology and effects).[5] Every Midjourney prompt opens with the command "/imagine" followed by a description of the desired image. In my case, I prompt the program with the words "/imagine Watteau's painting showing the Old Testament story of Abigail offering large quantities of food to David and his men." The program does not treat the words as search terms to retrieve an existing image. Instead, it identifies visual patterns that were correlated with each of my words in a vast trove of images that developers used to train it. Midjourney then produces my image by starting with a randomly generated grid of pixels and looking at it for any patterns that were associated with my prompt terms. It then amplifies these patterns while reducing the surrounding noise in successive stages, each time finding and accentuating anything that appears consistent with my words. The process is not so different from Leonardo's practice of searching the veins of a rock for the outlines of a battle or human face.[6] It is closer still to the methods of the eighteenth-century artist Alexander Cozens, who produced landscape paintings by refining the accidental forms within inkblots (fig. 19.3).[7] Like Cozens, Midjourney finds design within chance.

Unlike Cozens, however, Midjourney does not depend on the visual associations of an isolated person but on those made by the many individuals who unknowingly contributed to its training data. Midjourney has not publicly disclosed the content of this dataset, but court filings and leaked documents indicate that the


Fig. 19.3. Alexander Cozens (English, 1717–1786), *A New Method of Assisting the Invention in Drawing Original Compositions of Landscape* (London: J. Dixwell, 1785), plate 14. Beinecke Rare Book and Manuscript Library, Yale University, New Haven, Connecticut

program relies on the Large-Scale Artificial Intelligence Open Network (LAION), a collection of 5.85 billion images and associated text scraped from the internet.[8] If Midjourney has a "memory," then it is the memory formed by the images that humans have put online. For this reason, scholars in media studies often speak of generative AI as exhibiting a "collective imaginary" or "collective unconscious."[9] The terms are apt, but the "collective" here is an unconventional one, referring less to a social unit than to a technological one.[10] What makes the unconscious of generative AI distinctive is the very fact that it is not located within any definable group. It is not so much a shared subjectivity as it is a digital aggregate that no human community possesses.

For a well-known artist such as Watteau, this digital unconscious runs deep, though its underlying character remains somewhat difficult to probe. As of this writing, the LAION database of images on which Midjourney trained is inaccessible, having been taken offline after researchers discovered that it contained child pornography.[11] Even if public access were restored, we would remain ignorant of the way the developers of the AI model assigned weights to the dataset's contents. We can nonetheless get an approximate sense of Midjourney's reference points and priorities from the company's founder, who has stated that the model's horizon of vision is "pretty much the internet."[12] In the case of Watteau, the internet provides a comprehensive but not entirely balanced view of the artist's work. Nearly every known, surviving painting in Watteau's oeuvre is represented online, thanks to Martin Eidelberg's ongoing efforts to assemble a digital catalogue raisonné.[13] But such scholarly undertakings do not, as far as we know, carry any more importance in the training data than other online sources. The process of data aggregation blends the content of academic resources with files from other websites, ranging from image-licensing agencies to print-on-demand poster shops. Because most of these websites are oriented toward popular interest and commercial appeal, they hew closely to the artist's most canonical work. If the AI model takes the prevalence of an image on the internet as an indication of significance, then what it provides is a market cap–weighted view of an artist's oeuvre.

To be clear, I am speculating about the workings of a technology that operates largely as a black box. But a closer look at my simulated Watteau painting confirms the algorithmic tendency to reinforce an artist's best-known techniques and effects. Consider the trees. The sagging branches and thinly painted leaves could come from any number of Watteau's works, but the darkened wood that looms over the right side of the composition finds its closest match in the artist's most canonical work, *The Pilgrimage to Cythera* (fig. 19.4).[14] In both the real and virtual paintings, the shadowy trees make the illuminated figures in the foreground stand out, drawing attention to the picture's protagonists. Watteau, in both his real and virtual incarnations, then pulls our eyes through the composition using a meandering chain of supporting characters who move into the distance at left, disappearing into the vaporous haze of the background.

Other aspects of the virtual Watteau are less obviously traceable to a single known work; instead, they appear to derive from a homogenized average of multiple compositions. The strongly lit woman in the center of the picture, presumably representing Abigail, vaguely evokes the brightly illuminated figure in Watteau's *Fêtes Vénitiennes* (fig. 19.5). She may also derive from the seated bride in the *Accordée de village* (fig. 19.6), who is far less conspicuous in the original work but is dressed and posed in a similar manner. Whatever factors went into Abigail's construction, she is closer to the courtly women we associate with Watteau than

Fig. 19.4. Antoine Watteau, *The Pilgrimage to Cythera*, 1718–19. Oil on canvas, 47 ¼ × 74¾ in. (120 × 190 cm). Charlottenburg Palace, Berlin

Fig. 19.5. Antoine Watteau, *Fêtes Vénitiennes*, 1718–19. Oil on canvas, 22 × 19 in. (56 × 46 cm). Scottish National Gallery, Edinburgh

Fig. 19.6. Antoine Watteau, *L'Accordée de village*, c. 1710–15. Oil on canvas, 24¾ × 36¼ in. (63 × 92 cm). Courtesy of the Trustees of Sir John Soane's Museum, London

Fig. 19.7. Peter Paul Rubens (Flemish, 1577–1640), *The Meeting of David and Abigail*, c. 1630. Oil on panel, 11⅝ × 26⅛ in. (45 × 66 cm). National Gallery of Art, Washington, DC

she is to the person described in the biblical narrative of Abigail and the iconographic tradition surrounding her. When other artists represented the story, they typically showed Abigail kneeling or prostrate before David (fig. 19.7), emphasizing her status in the Bible as a supplicant broker of peace. It makes some sense that Watteau would have departed from this convention, reimagining an Old Testament narrative as a scene of modern conviviality. Watteau almost never took on subjects of "grand manner" history painting, and he therefore might have simply adapted a biblical narrative to his distinctive pictorial idiom. Indeed, it is hard to imagine Watteau presenting Abigail making the demonstrative gestures that we find in, for instance, Rubens's treatment of the subject. Watteau rarely produced figures whose poses are so legible, whose interior thoughts and feelings are so transparent.

Yet, what if he surprised us? Here, I think, is where the AI rendering produces its biggest disappointments and most disturbing implications. Probabilistic by design, AI image generators encode pixels according to their likelihood. If dramatic gestures do not appear in an artist's oeuvre, Midjourney has little reason to invent them. Each of its images begin with randomly generated noise, but the aleatory elements of the process remain bound by a model that prioritizes patterns and reiteration.

The system's repetitions are not easy to recognize because they are always reformulations, not exact copies. The virtual Abigail is nowhere in Watteau's existing paintings, yet she is also everywhere, living in what computer scientists call "latent space," the field of potential images that exist between the known data points. She takes the uncanny form of the familiar presented in an unfamiliar guise. The disquieting effect goes beyond the well-known glitches in early AI images, such as extra fingers and errant limbs. If anything, those visible "mistakes" provide reassurance, allowing us to securely categorize an AI image as a mirage. As the technology improves and the missteps become harder to spot, the uncanniness of the creations will only become more acute. The risk here is not simply that art historians will lose the ability to distinguish the real from the virtual, but that we will struggle to separate new insights from repackaged versions of old ones. It is worth remembering that uncanny experiences, in Freud's original formulation, originate in "unintended repetition" and the sense of déjà vu.[15] The compulsion to repeat, Freud suggested, is what made uncanny experiences neurotic and not merely strange. An eerie dream unsettles us because it recycles a primal urge or repressed experience while giving it a new face.

If there is therapeutic value to be found in my ersatz Watteau, then, it is in the chance that it provides to confront the art historical compulsion to restate what we already know. Let us therefore use the image not as a superficial means of filling a gap in the artist's oeuvre but as a way of learning to live with the hole. In this respect, the most productive aspect of the image may be found in its

Fig. 19.8. Detail of fig. 19.1

upper-right corner, where the canvas appears to sag and the darkened paint ripples with faux craquelure (fig. 19.8). The details are troubling in their counterfeit materiality, giving false substance to a work that exists only in digital code. But as signs of age and mutability, the cracks in the digital paint also contain within them an acknowledgment of loss. Only here does the image show us the decay that it otherwise masks. May we look for our future in those dark depths.

## NOTES

**1.** The documentation and scholarship surrounding the work is synthesized in Martin Eidelberg, "Abigaïl qui apporte des vivres à David," Watteau Abecedario, 2020, http://www.watteau-abecedario.org/David_Abigail.htm.

**2.** Michael Ann Holly, *The Melancholy Art* (Princeton University Press, 2013), 21.

**3.** See, for example, Bruno Cornelis et al., "Crack Detection and Inpainting for Virtual Restoration of Paintings: The Case of the Ghent Altarpiece," in "Image Processing for Digital Art Work," ed. Patrice Abry et al., special issue, *Signal Processing* 93, no. 3 (2013): 605–19; and Varun Gupta et al., "Restoration of Artwork Using Deep Neural Networks," *Evolving Systems* 12, no. 2 (June 2021): 439–46.

**4.** I am thinking especially of the Rijksmuseum's recent initiative to recreate missing sections of Rembrandt's *Night Watch* using artificial intelligence. *The Night Watch: The Missing Pieces*, Operation Night Watch, 2021, https://www.rijksmuseum.nl/en/stories/operation-night-watch/story/night-watch-the-missing-pieces.

**5.** James Vincent, "An Interview with Midjourney Founder David Holz," *The Verge*, August 2, 2022, https://www.theverge.com/2022/8/2/23287173/ai-image-generation-art-midjourney-multiverse-interview-david-holz.

**6.** See the classic essays H. W. Janson, "The 'Image Made by Chance' in Renaissance Thought," in *De Artibus Opuscula XL: Essays in Honor of Erwin Panofsky* (New York University Press, 1961), 254–66; and E. H. Gombrich, "The Image in the Clouds," in *Art and Illusion: A Study in the Psychology of Pictorial Representation* (Pantheon Books, 1961), 181–202.

**7.** Jean-Claude Lebensztejn, *L'art de la tache: introduction à la Nouvelle méthode d'Alexander Cozens* (Limon, 1990).

**8.** For evidence of Midjourney's use of the data, see Andersen v. Stability AI Ltd., no. 3:23-cv-00201 (N.D. Cal., January 13, 2023). See also Theo Belci, "Leaked: The Names of More than 16,000 Non-Consenting Artists Allegedly Used to Train Midjourney's AI," *Art Newspaper*, January 4, 2024, https://www.theartnewspaper.com/2024/01/04/leaked-names-of-16000-artists-used-to-train-midjourney-ai.

**9.** Andreas Ervik, "Generative AI and the Collective Imaginary," *IMAGE: The Interdisciplinary Journal of Image Sciences* 37, no. 1 (2023): 42–57, https://image-journal.de/generative-ai-and-the-collective-imaginary/; and Jens Schröter, "The AI Image, the Dream, and the Statistical Unconscious," *IMAGE: The Interdisciplinary Journal of Image Sciences* 37, no. 1 (2023): 112–20, https://image-journal.de/the-ai-image-the-dream-and-the-statistical-unconscious/.

**10.** In this respect, the "collective imaginary" of generative AI should be distinguished from the formulations of the concept in contemporary social psychology. See, for example, Florence Giust-Desprairies, *L'imaginaire collectif* (Érès, 2003).

**11.** David Thiel, "Identifying and Eliminating CSAM in Generative ML Training Data and Models," Stanford Internet Observatory, 2023, https://purl.stanford.edu/kh752sm9123.

**12.** Vincent, "Interview with Midjourney."

**13.** Martin Eidelberg, "A Watteau Abecedario," 2014–, http://www.watteau-abecedario.org/.

**14.** I reproduce here the Berlin version of the composition, whose woods strike me as a somewhat closer match than the Paris painting, though the overall compositions are similar.

**15.** Sigmund Freud, *The Uncanny*, trans. David McLintock (Penguin, 2003), 144–45.