

The Artificial Sublime

Regina Rini - rarini@yorku.ca

York University Philosophy¹

This is a pre-publication draft posted at PhilArchive. You may cite it as such, listing the PhilArchive url <https://philpapers.org/rec/RINTAS-2> and dated as a September 2024 draft. Do check back at the PhilArchive location to see if publication info has been updated.

Abstract: Generative AI systems like ChatGPT and Midjourney can produce prose or images. But can they produce art? I argue that this question, though natural and intriguing, is the wrong one to ask. A better question is this: can generative AI yield distinct or novel forms of aesthetic value? And I argue that the answer is yes. Generative AI can be used to put us in contact with the artificial sublime – a type of aesthetic value that Kant famously argues is impossible. Kant claims that sublimity (a fusion of trepidation and wonderment) can only arise from encounters with nature, or perhaps close artistic depiction of natural scenes. I show how generative AI can yield encounters both with incalculable profundity (Kant’s mathematical sublime) as well as overwhelming moral chaos (Kant’s dynamical sublime). There are lessons not only for art criticism, but also for the philosophical theory of value.

In July 2023 a print shop near Sydney, Australia, was called upon to determine the real meaning of art. The shop was running a fashion photo competition and one submission seemed perfect. Too perfect. It showed three snazzily-dressed human figures, hands held in a daisy chain across a brightly lit public restroom. “I just thought the faces were just too ... the overall feel of it just felt fake,” the shop owner later explained. The other judges concurred: this image had been made by AI.

They debated whether the photo’s high aesthetic merit might be worth rewarding anyway. But ultimately they decided to hold the line against digital intrusion. The image was disqualified.² The announcement of their verdict came with a prim pedagogue’s warning to contestants who might be tempted to repeat the trick: “We want the images to come from YOUR real life experience, and not sourced from cyberspace. Our competition is all about showing us you understand light, composition, f-stops and story etc. We only know for sure if you take the photo yourself.”³

¹ This paper would not exist at all without Dawn Wilson and the British Society of Aesthetics, who invited me to dip my toe into aesthetic theory at their 2023 annual conference in Oxford. (It is apparently a tradition to invite one keynote speaker who has no prior background in aesthetics. Far from being a hazing ritual, this was a fantastic and musing – in the old sense - experience.) Thanks also to later audiences at Queen’s University and the 2024 TeXne Conference hosted by Northeastern, MIT, and Harvard. Thanks to Daniel Munro for feedback and Ashton Black and Daniel Rodrigues for research assistance.

² Schrader (2023).

³ @charingcrossphoto Instagram post, July 2023, as reprinted in Inside Imaging (2023a).

A lesson in aesthetic justice, perhaps. But then, a problem: it turned out the photo was real. Not AI-generated at all. It had been taken by a local amateur photographer using props at a museum fashion exhibit. Two of the image's otherworldly figures were mannequins. The third was a real human – the photographer's 18-year-old son, in fact – whose chiseled cheekbones and searching eyes were just “too perfect to be true,” according to the contest overseers. When this plot twist made local news, online commenters reacted with fury. One drew a broader lesson: “This is yet another example of the arrogance of the so-called ‘judges of art.’”⁴

Embarrassing. Yet spare a thought for today's judges of art. They are not wrong to suspect the lurking presence of digital confection. In summer 2022, an image titled *Théâtre D'opéra Spatial* made global headlines by winning a Colorado art contest without anyone noticing it had been created with Midjourney, an AI image-generation system. A few months later, German artist Boris Eldagsen triumphed in the Sony World Photography Awards with a synthetic image generated by another AI system, DALL-E 2.⁵ In February 2023, an Australian AI marketing firm won yet another photo competition with a synthetic image of waves crashing on a beach. “It's not an exaggeration to say that we've reached the point where machine is now the superior artist to man,” crowed the AI marketers.⁶ So, a few weeks later, when the Sydney photo shop judges faced their own implausibly perfect submission, you might understand their hair-trigger suspicion.

It's not only visual art under siege from algorithms. A musical duet between AI-generated voices purporting to be Drake and the Weeknd made news around the world in 2023; it was even submitted for Grammy consideration (though disqualified on a technicality unrelated to being AI).⁷ In early 2024, novelist Rie Qudan won the Akutagawa Prize, a major Japanese literary award, for *Tōkyō-to Dōjō Tō*, which she later acknowledged was partly written by ChatGPT, a language-generating AI system.⁸ Across media, it seems the age of the AI artist is at hand.⁹

Much of the reporting on these incidents is framed in terms of deception: AI-generated works are ‘fakes’; their creators are ‘cheaters’ who eventually get ‘caught’. But this framing takes for granted some very live and unresolved philosophical questions. AI-generated art is ‘fake’ only if it is not real art. But isn't it?

This paper will argue for taking the aesthetic properties of AI-generated art quite seriously. First I will answer the simple question whether AI-generated works can count as art. The answer is yes. But the argument to this conclusion will turn out to be suspiciously easy to make – so easy that it raises doubts whether the question is well-formed. Instead, I will suggest, a better question is this: does artistic production via AI give rise to any new or distinctive sort of aesthetic value? Again the answer is yes. But

⁴ @winesellaro Instagram post, July 2023, as reprinted in Inside Imaging (2023b).

⁵ Roose (2022); Parshall (2023).

⁶ @absolutely.ai Instagram post, February 2023, as reprinted in Inside Imaging (2023c).

⁷ Davis (2023).

⁸ Ha (2024).

⁹ Though sometimes the humans get their own back. In 2024, a real photograph of a flamingo, taken by an actual human, was mischievously submitted to what was supposed to be an AI image competition. It won third place, until disqualified by revelation of its unseemly human origins. See Schrader (2024).

this time the argument will be deep and unexpected. If I am right, then we will have learned something important not just about AI, but about the value of art itself.

A palette of questions

Our first task, before answers, is to be clear about questions. We are at risk of drowning in an interrogative swamp. Is an AI-generated image a work of art? In the same way as a photograph? As a painting? Or something entirely its own? Is the person who prompted its creation an artist? What about the software engineers who made the system? Could the AI *itself* ever qualify as an artist? And who deserves credit – artistic or financial – for an especially well-executed instance?

We won't make any progress unless we recognize that these questions fall across several conceptual lacunae. Just as critics in earlier generations were unsure whether photography or film could qualify as artistic media, we are now in the early very days of AI generation, when the boundaries of our concepts are being tested in unfamiliar ways.¹⁰ I make this comparison to photography and film not to prejudge whether AI generations will ultimately be accepted as art – there is an argument to come – but instead to highlight how the difficulty of delimiting art begins as a difficulty of establishing *category*. We are trying to understand both medium and genre.

Consider that AI-generated work by Boris Eldagsen – titled *PSEUDOMNESIA: The Electrician* - which won in the 2023 Sony World Photography Awards. It depicts two women posed in an ambiguous embrace, seemingly affectionate yet with potentially sinister undertones. The image is black-and-white, gauzily lit, and not of our era. 'Spooky' is a good word for it.

Set aside whether this is a work of art. Ask first: is it a *photograph*? Its creator says not. Posting on his website after refusing the award, Eldagsen said he had submitted the work only to make a point. "AI images and photography should not compete with each other in an award like this," he wrote. "They are different entities. AI is not photography."¹¹

In a certain sense he must be right. AI generation of images is obviously not photography, if photography means using a camera to capture light bounced off objects in the real world. The women depicted in Eldagsen's image never reflected light toward a camera lens, because they do not exist at all. At best, they are an algorithmic extrapolation from the appearances of many different people in many actual photos (more about how this works later). This breaks the broad correspondence to reality that characterizes the nature of photography. A photograph is constrained by the optical physics of a given moment in time in a way that an artist working in charcoal or oils – or digital bits – is not.¹²

We shouldn't overstate this point. Photography has never been a strict replication of objective reality; the photographer imposes their interpretation of a scene through perspective, composition, and skillful selection of aperture and exposure.¹³ Today's digital cameras – particularly the sort embedded in

¹⁰ Walter Benjamin – an ambivalent consumer of the new media – wrote an especially stimulating overview of photography's reception in the artistic world in 1931, already nearly a century into the medium's life. See Benjamin (1931/1977).

¹¹ Eldagsen (2023).

¹² See Hopkins (2010) and Cavedon-Taylor (2013) for epistemic features that mark the edges of the photographic medium. But see also Wilson (2022) for complexities.

¹³ The locus classicus for the naivete of seeing photography as directly realistic is Sontag (1977).

portable computers that we call phones – have moved far from naïve replication of reality. Algorithmic filters instantly ‘correct’ for unflattering patterns of light. High dynamic range (HDR) systems, standard on many phone cameras, construct an image by statistically mediating several snaps taken in very rapid succession; strictly speaking, HDR output does not bear a one-to-one correspondence to reality at any single point in time. The mechanical difference between photography and AI generation is shrinking fast.

Still, an AI generated image is not a photograph. But does that mean, as Eldagsen claims, that AI generations should not be eligible to win photography competitions? It’s important to see this question isn’t only about the nature of photographs but is also about the nature of photography *competitions*. In other words, it is a question about how we recognize aesthetic value.

Some artistic works highlight the aesthetic features of a medium or genre precisely because they challenge that category’s boundaries. Consider the oeuvre of Giuseppe Arcimboldo (1526-1593), a court painter in Vienna and Prague. He is most remembered for a series of still life grotesques, which from a distance seem to depict human figures, yet on closer examination show a carefully arranged assemblage of fruits, books, sea creatures, or other thematically linked objects. *Vertumnus* (1591), for instance, depicts his patron, Emperor Rudolf II, as the Roman god of seasons, constructed from a teeming mound of produce.



Fig 1: *Vertumnus* (1591), oil painting by Giuseppe Arcimboldo (1526-1593)¹⁴

¹⁴ Image in public domain, retrieved from Wikimedia commons Aug 20 2024:
https://en.wikipedia.org/wiki/File:Vertumnus_%C3%A5rstidernas_gud_m%C3%A5lad_av_Giuseppe_Arcimboldo_1591_-_Skoklosters_slott_-_91503.jpg

Is Arcimboldo's *Vertumnus* a portrait? Portraiture is a representational genre crossing painting, sketch, and photography. It has its own conventions, much like biography as a literary genre. In one sense, *Vertumnus* seems clearly to be a portrait of Rudolf, given the painter's artistic intentions and the apparent resemblance. On the other hand, the painting breaks from the part-whole relationship that traditional portraiture employs. In a traditional portrait it is possible to clearly identify elements of the image as depicting an eye or an ear, even if one could not see the full image. If you walked in on the painter mid work, you might recognize a patch of colour as a human ear and wonder whose face would soon emerge. But a sneak preview of *Vertumnus* would not have revealed the constituents of human faces: the viewer would see fruits and vegetables, not eyes and ears. These are recognizable as body parts only in the context of the entire image and its representational conceit. This was highly unusual by the conventions of 16th century portraiture.

Rather than form and convention, we might look instead to the values of the genre. One of the defining values of portraiture is the image's ability to help us better appreciate the depicted subject, perhaps by a close study of the relationships among their features. Hilary Mantel's 'Wolf Hall' trilogy of historical novels repeatedly returns to the question whether Hans Holbein's famous portraits of the Tudor royal court disclose aspects of these personages they may have wished to hide. In this scene, Thomas Cromwell is joined by his son Gregory as he contemplates his own freshly revealed portrait:

"He turns to the painting. 'I fear Mark was right.'
'Who is Mark?'
'A silly little boy who runs after George Boleyn. I once heard him say I look like a murderer.'
Gregory says, 'Did you not know?'"¹⁵

Is it possible to appreciate the person depicted in *Vertumnus* in this way? According to art historian Thomas Kaufmann, Arcimboldo likely chose these elements to signal the "harmony" and "timelessness" of his patron's reign: "The combination of fruits and vegetables from all seasons in his face expresses his undying majesty."¹⁶ So *Vertumnus* does tell us something about its subject (or at least how Arcimboldo thought Rudolf would want to be regarded) but *how* it does this is not quite the same as traditional portraiture. Here, character traits are attributed through allusive constitution rather than simple physiognomy. In an important sense, the painting belongs to the genre of still life – it depicts fruits and vegetables – rather than the genre of portraiture. Yet it also seems to be a winking commentary on portraiture, something it can do only by coming *close enough* to being a portrait.

Could we say something similar for Eldagsen's AI-generated image? It is not a photograph, true. But it is very close to a photograph in many respects, enough that judges in a major photography competition failed to detect it as something else. And the artist's intentions were clearly to probe the nature of the medium. So, at minimum, Eldagsen's *PSEUDOMNESIA: The Electrician* should be admitted as an internal commentary on photography in much the same way Arcimboldo's *Vertumnus* comments on portraiture.

What all this shows is that categorial questions – *is this a photograph?* – are sometimes distractions from much more interesting and revealing questions about aesthetic value. Whether Eldagsen's image deserves recognition in a photography competition depends on which aspects of aesthetic excellence

¹⁵ Mantel (2009), 489.

¹⁶ Kaufmann (1976), 295.

the competition aims to recognize. If the competition is meant to reward technical skill in wielding a camera, then AI generations clearly do not qualify. But if the competition instead aims to recognize aesthetic sensibility – the ways an artwork can deepen our experience of the world – then a work which is a commentary on photography may very well deserve recognition *within* photography.

Not everyone will agree with this, of course. But the point is simply that this is at least an interesting debate, one worth having. It requires us to explicitly confront our conception of aesthetic value in a way that blunt categorical questions do not.

The lesson here generalizes. We've just seen that *is this a photograph?* is not a very illuminating question. Perhaps the same is true about the seemingly more portentous question *is this a work of art?* It is easy to get caught up in dense and seemingly interminable debates about the boundary conditions for artwork. But I want to avoid this sort of direct approach. One reason is that it turns out to be disturbingly easy to force the conclusion that AI generations can be artworks.

Here is a cheap and disposable argument for the possibility of AI-generated art. I can construct a work of art entirely out of AI tools. First, I prompt one of the language systems (GPT, Claude, etc) to provide a list of 30 phrases that sound like the titles of contemporary artworks. Then I copy these phrases, one by one, into an image-generation system like DALL-E or Stable Diffusion. Finally, I ask a gallerist friend to allow me to exhibit prints of all the generated images. I title the exhibit *The Work of Art in the Age of Statistical Reproduction*. It is this last object – the collection of all the images together, along with the recipe for how I created them – that I now claim to be a mixed media artwork.

Notice that the essential medium of this artwork is AI-generation itself. The work of art is not the images, nearly all of which will be unremarkable schlock. The work is the assemblage, generated in an entirely mechanical way by the AI systems I use as my artistic tools. All the words came from GPT; all the images came from DALL-E. I do not edit or curate them. There is no space for human aesthetic sensibility at granular level, only in the overall conception of the project. *The Work of Art in the Age of Statistical Reproduction* is an artwork because it is a meta commentary on the means of artistry in the 21st century.

You may be tempted to deny that this could count as an artwork. But taking a stance of border patrol carries high intuitive costs. Several of the most famous artworks of the last century were commentaries on artistic production, notable precisely for their unusual choice of materials. Start with Marcel Duchamp's famous *Fountain* (1917) a found-art urinal submitted to the Society of Independent Artists as a challenge to aesthetic boundaries. Or John Cage's composition *4'33''* (1952), a piece of music consisting of no notes, only the experience of a concert hall struggling to be silent. Or Marina Abramovic's *The Artist is Present* (2010), a performance piece wherein the artist sat silently in the atrium of New York's Museum of Modern Art while patrons queued for hours to perch across from her. (Your humble author stood among them, but lacked the forbearance to reach the prime seat.) Or Maurizio Catellan's mixed media piece *Comedian* (2019), consisting of a fresh banana duct-taped to a gallery wall, stipulated to be re-instantiated wherever the necessary materials and the will of any patron coincide.

These are artworks, surely. But what do these have that *The Work of Art in the Age of Statistical Reproduction* does not? It can't be something about the materials. There was nothing inherently artsy in Duchamp's urinal – that's the whole point of the piece. If urinals and silence can be artistic media, why not AI-generation?

Of course, I do not claim that *The Work of Art in the Age of Statistical Reproduction* is *good* art, still less that it deserves a place in history beside these examples. My point is only that it is possible to produce an artwork – perhaps of execrable quality, like many amateur pieces – from the materials of AI-generation, via the royal road of meta commentary.

If I stopped here – ‘AI generations can be art, QED!’ – you’d be reasonably annoyed. This argument is a philosophical cocktail party trick. It’s cheap and shallow.

But that’s the point: ask an unhelpful question, get a cheap and shallow answer. *Is this a work of art?* invites exactly the sort of cleverness-signalling reply that I’ve just imagined. Neither the question nor the answer tells us much about why we make or appreciate art. Like *is this a photograph?*, the categorical question is a distraction from the more elemental questions we should be asking. “These waters must be troubled before they can exert their virtues,” is how Edmund Burke began his investigation of aesthetic fundamentals in 1757.¹⁷ It’s time to sink deeper into the murk. *This* is the question we should be asking: which, if any, aesthetic values might be served by AI generation?

Something serious in the imagination

Which aesthetic values might we consider? Two have dominated traditional western philosophy. Beauty is the more famous, though notoriously difficult to define. But I will focus on the other: the sublime. Roughly, the sublime is something like “aesthetic awe”.¹⁸ Prototypical examples are natural immensities: the Pacific Ocean, the Alps, the unfathomable extent of outer space. Contemplating such things can trigger a mental state pairing profound appreciation with a feeling rather like fear. This combination – appreciation plus apprehension - is the aesthetic sublime.

Though recognized since antiquity, the sublime really took off as an aesthetic focus in 18th century Europe.¹⁹ According to Anglo-Irish philosopher Edmund Burke, sublimity arises from a close-but-not-too-close brush with annihilation. Genuine peril is psychologically painful, but a *contained* threat allows pain to undergo a striking transmogrification. “[A]t certain distances, and with certain modifications, [danger and pain] may be, and they are, delightful.”²⁰ But this is not a simple delight: “terror is in all cases whatsoever, either more openly or latently, the ruling principle of the sublime.”²¹

Among Burke’s many readers was Immanuel Kant, one of the giants of western philosophy. Writing a generation later, Kant took up the core of Burke’s account but built it into a more nuanced model. Letting go Burke’s central focus on terror, Kant maintained that the sublime involves “something serious in the

¹⁷ Burke (1757/1937), 48.

¹⁸ See Clewis (2021) for an argument that sublimity simply *is* aesthetic awe. Most theorists favour a looser connection, suggesting e.g. that they are merely “neighboring concepts” (Brady, 2019, 358). Some aestheticians doubt it is conceptually possible even to give a theory of the sublime, given the nature of the phenomenon. For discussion, see Sircello (1993); Forsey (2007); Zuckert (2019).

¹⁹ For historical overviews, see Lokke (1982) and Brady (2013), chapters 1-4.

²⁰ Burke (1757/1937), 36.

²¹ Burke (1757/1937), 50.

activity of the imagination.”²² (Kant, famous in Königsberg for daily strolls you could set your watch by, is someone you should take seriously when he says something is serious.)

Mortality isn’t the only serious thing we face. Death (which Burke claimed as the ultimate source of all terror) limits us temporally. But, Kant saw, we are limited in other ways too - in our perceptual or imaginative abilities, for instance. The immensity of outer space confounds our sense of how the world works. Consider: assuming the expansion of the universe continues to accelerate, there are distant stars whose light will *never* reach Earth, even given infinite time. We will never have direct perceptual information about them. Kant thought confronting any such absolute limit could yield a threatening experience.

This is Kant’s first change from Burke: the replacement of mortal terror with a more general awareness of our fundamental limits. Kant’s second change has to do with timing. Burke thought of the sublime as an experience of *simultaneous* apprehension and appreciation: you feel terror, but you also feel appreciation of its containment, all in the same moment. For Kant, the sublime is a two-stage process: *first* comes a sudden confrontation with your limits, *then* comes an appreciation of the ways you might still transcend these limits. As Kant put it, the sublime comes from “the feeling of a momentary inhibition of the vital powers and the immediately following and all the more powerful outpouring of them.”²³

Roughly, Kant’s image of the sublime follows the emotional oscillations of a youth sports movie: *The Mighty Ducks* or *Bad News Bears*. Here is your loveable gang of misfits no one would ever bet on. Watch them endure humiliation by a team of impossibly skillful braggarts. Feel the heroes’ limitations: how could they have ever thought they would come out on top? Then the coach makes a stirring speech, the players find their heart again, and friendship and resilience triumph over the limits of raw strength or size. Crash, rise again: the emotional pattern of the Kantian sublime.

But what *is* this rise? How does Kant think we recover from confronting our limits? Roughly, it’s that, while encountering the sublime exposes limits in *one of our* abilities, we quickly realize that we have another ability that allows us to circumvent this limit. Like the triumphant baseball team, we dig deep and find a resilience more powerful than we’d ever known ourselves to be. Or, in more typical Kantian idiom, “That is sublime which even to be able to think of demonstrates a faculty of the mind that surpasses every measure of the senses.”²⁴

Consider again those natural immensities – the Alps, the Pacific Ocean, deep space – that utterly outrun our perception and imagination. Kant says they are “contrapurposive for our power of judgment, unsuitable for our faculty of presentation, and as it were doing violence to our imagination”. Yet we discover a new purchase on these objects through powers of abstract reasoning. We can calculate the dimensions of the Pacific Ocean even if we cannot see all of it at once. We can infer features of deepest space through physical reasoning. As we come to appreciate how these initially unfathomable objects can be mentally captured, “the mind is incited to abandon sensibility and to occupy itself with ideas that

²² Kant (1790/2000), 129. (Ak 5:246)

²³ Kant (1790/2000), 128-129. (Ak 5:245)

²⁴ Kant (1790/2000), 134. (Ak 5:250)

contain a higher purposiveness.”²⁵ In doing so, we come to recognize the astounding powers not of space or seas, but of our own minds.

This claim – that we sublimate perceptual incapacity to conceptual power – illustrates what Kant calls the *mathematical* sublime, or “that which is absolutely great... that which is great beyond all comparison”.²⁶ (This is the first of two types of sublimity in Kantian theory; we will come to the second later). The point is not merely that we are physically small next to mountain ranges and seas; it is that their scale exceeds what can be contained in our perceptual faculties. We regain our grasp on them only through abstract cognition.

It is this thought which leads Kant to one of his most controversial aesthetic claims – a claim I will reject later in this essay. Kant insists that the sublime (unlike the beautiful) is genuinely experienced *only through encounters with natural objects*, not human-authored artwork. His argument for this point is indirect, textually dispersed, and perhaps not entirely coherent. It also relies on features of his general theories of aesthetics and metaphysics too expansive to recapitulate within the scope of this paper. But the spirit comes out in this passage:

“if the aesthetic judgment is to be pure (not mixed up with anything teleological as judgments of reason) and if an example... is to be given which is fully appropriate for the critique of the aesthetic power of judgment, then the sublime must not be shown in products of art (e.g., buildings, columns, etc.), where a human end determines the form as well as the magnitude... but rather in raw nature... merely insofar as it contains magnitude”²⁷

Roughly, the idea seems to be that when we encounter a work of art, we cannot help experiencing it as the product of human effort. Then we know the object has limits – limits of scale, limits of form, limits imposed by the human sensibility of the artist. Since we know that a human being can *make* this, we do not experience it as overrunning our own human capacity for perception or imagination.

Of course, we are sometimes astounded by novel artistic accomplishments – “how did they even *do* that?!” – but we know that *in principle* it is something we could understand and possibly even perform ourselves, given enough time. The artform of stage illusion (magic tricks, to the vulgar) is a good example. You can be astounded and delighted by the seemingly inexplicable reappearance of your marked four of clubs from a locked safe at the bottom of a fish tank, but deep down you *know* that there is a comprehensible explanation in human artifice.

By contrast, Kant thinks, the sublime presents itself to us as fundamentally intractable – as “contrapurposive” – to human capacity, and so cannot arise from artwork that wears human authorship on its sleeve. At best, certain works of art may be able to steal a bit of wonder, a derived sublimity, by *depicting* the natural objects from which we originally gain an appreciation of the sublime. Realistic

²⁵ Kant (1790/2000), 129. (Ak 5:246)

²⁶ Kant (1790/2000), 131-132 (Ak 5:249). The mathematical sublime especially illustrates Kant’s divergence from Burke. Recall that Burke claimed that immensity triggers the sublime by exhausting the visual nerves, leading to pain and a premonition of death. Kant’s claim is much more direct. As Rachel Zuckert puts it, Kant “suggests, contra Burke, that the feeling of the sublime can be grounded not only in our desire to live, but also in our desire to know”. (Zuckert 2003, 218).

²⁷ Kant (1790/2000), 136. (Ak 5:252-253)

paintings of mountains and seas might sneak in from the edge of sublimity, but only when “restricted to the conditions of agreement with nature”.²⁸

There are well-known problems with this argument. For one, there is the exegetical awkwardness that Kant seems to immediately contradict himself. Reaching for an example of sublime experience, he points to the Egyptian pyramids, which are human-made art (whatever late night shows on the History Channel might say). But Kant adds some revealing details that may explain the apparent contradiction. Citing the account of a European traveller, he recommends “in order to get the full emotional effect of the magnitude of the pyramids one must neither come too close to them nor be too far away”.²⁹ Too far – squat little triangular specks on the desert horizon – they aren’t very impressive at all. Too close – eyes inches from the stone – and their scale can’t even be seen. Kant’s point seems to be that monumental architecture achieves sublimity only from certain special vantage points, when we begin to forget that it is human-made at all.

But this response leads directly into the second, bigger problem for Kant. Now it’s not clear what is so special about natural objects. After all, *their* ability to trigger the sublime seems to depend on viewing from certain vantage points as well. Once, in the southern Arizona desert, I saw an eccentric cliff of brown-red stone about a mile away. I began hiking toward it, thrilled as the strange craggy shape filled more and more of my field of view. At mid distance, perhaps a quarter mile away, it was indeed sublime. But then I approached closer, closer – right up to the base of the cliff. And, suddenly, it wasn’t affecting at all. Like the pyramids too close, it was now just a big stone wall.

What these examples show is that it is not the objects themselves, but *how we experience them* that yields the sublime. To some extent, this is exactly as Kant would have it. He insists it is “the disposition of the mind resulting from a certain representation occupying the reflective judgment, but not the object, which is to be called sublime”.³⁰ In other words: it is our perceptual experience, *not the external object*, that yields the sublime (though the object is a necessary distal cause). Since the rising second half of sublime experience consists in discovery of our own cognitive powers, the feeling must ultimately be attributed to our minds.

But then it seems, in principle *anything* should be able to yield sublime experience, if only encountered from a vantage point that approximates the mid-distance approach to a vast surface. Why couldn’t an artist intentionally induce exactly this experience?

Consider another large object of the southwestern American desert: *City* (1972-2022), a work of monumental land art produced over five decades by sculptor Michael Heizer. A mile and a half of ambiguous shapes rendered in concrete and desert stones, it struck one *New York Times* critic as resembling “a mash-up of Chichen Itza and an unfinished highway interchange.”³¹ It is not hard to imagine a visitor encountering sublimity in this vast anomaly. (Your humble author’s attempt to personally confirm this impression was foiled by the fact that *City* is technically not open to uninvited

²⁸ Kant (1790/2000), 129. (Ak 5:245)

²⁹ Kant (1790/2000), 136. (Ak 5:252)

³⁰ Kant (1790/2000), 134. (Ak 5:250)

³¹ Kimmelman (2022).

visitors - a limitation that might have been delicately uncognized but for the decidedly less conceptual barrier of hundreds of miles of empty Nevada desert downwind a former nuclear test site.)

Or consider works of installation art whose purpose is to confront audiences with evidence of their own sensory limitations, such as Kusama Yayoi's 'infinity rooms', which use mirrors and clever lighting to induce an experience of absolute boundless space.³² (Though your humble author found this more productive of vaguely irritating vertigo than contemplative ablation of self.)

Uygar Abaci notes exactly this problem amid his spirited defense of Kant's insistence on the naturalness of the sublime. Given Kant's claim the sublime is properly attributed to the mind itself, it might seem that any object could do. Yet, Abaci argues on behalf of Kant, the experience of sublimity flows from "the way in which the object occasions the feeling of the sublime, not the mere occasioning of the feeling."³³ Roughly, the thought is that Kant's distinctive second stage of sublime experience involves appreciation of the "autonomy" of our cognitive abilities from the limitations of the natural world present in our sensory experience. A *contrast* with nature is therefore "in the indispensable background of the sublime". Hence Kant's insistence that art offers, at best, a derivative sort of sublimity. A painting of the Alps – or a suitably vast desert sculpture viewed at mid-distance – might trigger sublime experience in us, but only via titrated recollection of how nature itself sometimes makes us feel.

Not everyone will find that argument very convincing.³⁴ Abaci's rather heroic defense of Kant inevitably carries residue of the Prussian's metaphysical style, which – as in all things Kant – can feel altogether too woolly to those with a taste for deserted conceptual landscapes.

Luckily that needn't concern us here. I wish now to *concede* that Abaci and Kant are right so far. I concede this because I want to show that *even if* Kant has been right so far, something fundamental has begun to change.

No surprises?

A generation after Kant's death, mathematician Ada Lovelace was already at work on technology that would bring us today's digital quandaries. In the 1830s she began a long collaboration with inventor Charles Babbage, whose idea for an 'analytical engine' prefigured all modern computers. Around 1842 Lovelace composed a commentary including what is widely regarded as the ancestor of all algorithms.

Yet Lovelace, history's first coder, was also history's first AI skeptic. She cautioned that machines could never transcend the input of their human creators. "The Analytical Engine has no pretensions to *originate* anything," she wrote. "It can do whatever we *know how to order it to perform*."³⁵ A century later, this word of caution hung over Alan Turing, the only person more important to the history of computing than the Babbage/Lovelace combo. In his 1950 paper 'Computing Machinery and

³² Brady (2013, chapter 5) makes similar points using different examples, though she seems to conclude that only certain sorts of land art might be truly sublime.

³³ Abaci (2008), 241.

³⁴ For further discussion, see Crowther (1989) and Clewis (2024).

³⁵ Lovelace (1843). See note G; italics in original.

Intelligence' (from which we get the famous 'Turing Test'), he attempts to bat away "Lady Lovelace's objection" to his vision for a thinking computer. The objection, he claims, boils down to the idea "a machine can 'never do anything really new' ... [or] a machine can never 'take us by surprise'".³⁶ Computers only return to us what we have already put in, perhaps somewhat rearranged.

Turing firmly rejected this idea. "Machines take me by surprise with great frequency," he averred. True, that was usually because he hadn't finished the math or made an error. But so what? "The view that machines cannot give rise to surprises is due, I believe, to a fallacy to which philosophers and mathematicians are particularly subject. This is the assumption that as soon as a fact is presented to a mind all consequences of that fact spring into the mind simultaneously with it."³⁷ In other words, machines might transcend the input we give them simply because the human mind is limited in what it can or will bother to think through.³⁸

This Lovelace-Turing colloquy prefigured the next 75 years of AI. Computer scientists mostly kept to Lovelace's assumption that programs could only do what their creators instructed. Through the twentieth century, the aim was to transmute sinewy human knowledge into the cold logic of formal declaration.³⁹ Programming languages demanded a set of strict necessary and sufficient conditions for every concept, such that a mechanically exhaustive box-ticking would yield decisive classification of any new data point.

This approach – known as 'symbolic AI' or in the words of philosopher John Haugeland, 'good old-fashioned AI'⁴⁰ – quickly ran into apparently intractable problems. As any philosopher who has ever studied the theory of concepts can attest, it is forbiddingly difficult to weave neat logical conditions from the mundanely eccentric fabric of human ideas.⁴¹ Imagine, for instance, trying to formalize the difference between types of cars. According to manufacturer Kia Motors, "A coupe is defined as a closed-roof, 2-door passenger car... [but] the coupe has expanded with some featuring 4-doors and called 'four-door coupes' instead of the traditional 2-doors."⁴² So a coupe is *defined as* having two doors. But some of these two door cars have four doors... And these are *cars* – artifacts designed by human minds. Good luck wringing greater precision from the happenstance of nature.⁴³

³⁶ Turing (1950), 450. Turing famously begins the paper by arguing that *can machines think?* is an unhelpful question, better replaced by less direct but more revealing queries. Readers may regard the way I began this paper – dismissing the utility of questions like *is this an artwork?* – as a quiet homage.

³⁷ Turing (1950), 450-451.

³⁸ There is a complex and somewhat tangential question whether artificial systems might be properly called 'creative'. For discussion, see Boden (2016), chapter 3; Ukpaka (2024); Brainard (forthcoming); and several interesting posts in a 2023 *Aesthetics for Birds* symposium at <https://aestheticsforbirds.com/2023/11/02/eight-scholars-on-art-and-artificial-intelligence/>.

³⁹ For important early interpretations of this ambition, see McCarthy and Hayes (1969); Newell and Simon (1976).

⁴⁰ Haugeland (1985).

⁴¹ See, for instance, Fodor (1998).

⁴² Kia Motors, 'What is the difference between a coupe and a sedan?'. Accessed August 12 2024.

<https://www.kia.com/dm/discover-kia/ask/what-is-the-difference-between-a-coupe-and-a-sedan.html>

⁴³ A related but not identical challenge for symbolic AI is the 'frame problem', which concerns how to formally delimit the diaphanous background knowledge (of basic physics, or grammar, or social custom) that humans use in accomplishing even simple tasks, typically without realizing the need for this knowledge at all. See McCarthy and Hayes (1969), cited above.

Problems like these led to stagnation in AI research late in the twentieth century, an era computer scientists sometimes call ‘AI winter’. But something was stirring in the heretical ends of the discipline. Beginning in the 1980s, a growing cohort of scientists revived Turing’s hunch that artificial intelligence would require the creation of “learning machines” capable of extracting more information from data than their programmers realized. Turing explicitly modeled this thought on human intellectual development: “Instead of trying to produce a programme to simulate the adult mind, why not rather try to produce one which simulates the child’s? If this were then subjected to an appropriate course of education one would obtain the adult brain.”⁴⁴

Modern machine learning does not exactly simulate childish brains, but it comes close. (Readers familiar with the technology may skip this and the next three paragraphs. For everyone else, I will do my best to make this simple.) The basic idea is to use a tangle of mathematical equations to mimic the spongy elasticity of natural brain tissue. An enormous number of virtual ‘neurons’ are strung together in a web of directed information pathways. The system is shown curated examples and instructed to sort them into categories. At the start of training, it cannot do this at all. But after each attempt, a clever mathematical algorithm is used to slightly redraw the connections between the ‘neurons’, testing out a new information flow to see whether it leads to improved performance. Repeat this, over and over and over and over again – millions of times, at speed and repetition no human learner could tolerate – and eventually the system starts getting it right. It learns the implicit boundaries between categories – coupes vs sedans, grammatical vs ungrammatical strings of words, bebop vs swing – without anyone needing to encode that boundary in rigid logic.⁴⁵

The machine learning systems I’ve just described are known as *classifiers*. They sort items into categories very well. But they do not create anything new. The aptly-named *generative* AI that has started ‘stealing’ art competition prizes requires borrowing some tricks from statistics. Two techniques dominate generative AI in the mid-2020s. (This technology develops so quickly that different techniques may be prominent by the time you read this. For instance, generative adversarial networks (GANs) were the main image-generating technique until about 2022, but have fallen out of relative favour.)

One, which drives most language systems like ChatGPT, is stochastic prediction. The system has ingested extraordinarily vast quantities of human-authored text, on the order of large fractions of *the entire internet*. Its training captures the statistical likelihood that any word will follow any other sequence of words. For instance, if you prime the system with the string ‘all I want for Christmas is...’, then its training encodes much greater likelihood the next words will be ‘my two front teeth’ rather than ‘undetached rabbit parts’. But the output is also stochastic, rather than deterministic, meaning it roughly corresponds to statistical distribution in the training data. In other words, if 0.5% of all people who have ever written ‘all I want for Christmas is...’ completed with ‘Ryan Reynolds’s phone number’, then that is what approximately 0.5% of users will see as output. Generate enough text and these statistical curveballs will

⁴⁴ Turing (1950), 456. Turing went on gnomically: “I have done some experiments with one such child-machine, and succeeded in teaching it a few things, but the teaching method was too unorthodox for the experiment to be considered really successful.” Oh, for further details!

⁴⁵ See Rumelhart et al. (1986); LeCun et al. (2015). For a philosophical overview, see Buckner (2019).

accumulate, spiraling the composition along a semantic path rarely, if ever, trod before. This is why language models routinely surprise us (and why they are radically unreliable for epistemic purposes).⁴⁶

The second technique, *diffusion*, is the most widely used in contemporary image generation systems. Imagine a work of sand art: a picture made from coloured grains of sand arranged in gentle geometric shapes, a bit like Buddhist mandalas. Imagine this work carefully constructed in a box. Now imagine that box shaken violently, all the sand mixed and swirled so the original image is completely lost. Imagine doing this over and over with different boxed mandalas. Now imagine an AI system trained on two tasks: categorizing which arrangements of sand look like art to humans (versus which are just messy boxes), and tracking the physics of how sand grains move around amid violent randomization. Task that highly trained AI system with something new: give it a freshly shaken box and tell it to move the sand grains around, one by one, until it has produced an arrangement that looks like something humans would call a picture. The system has learned how to rapidly extract orderly patterns from chaos, converting mere entropy to surprise.⁴⁷

I'm sure computer scientists would quibble with some of the ways I've described these techniques, but my goal here has only been to give a sense of their workings to non-technical readers. The main lesson is that these techniques work on two central principles: randomization and iterative refinement. Start at an arbitrary point in the logical space of visual or auditory properties. Journey through that space in a semi-random fashion, guided only by the broad wind of statistical feedback. Keep going, step after step, again and again. Millions of steps. Eventually you end up in a region of logical space familiar to human perceivers. Your path is too complicated for humans to follow at granular level, but they will recognize the arc of the journey and the destination. Or, as Turing put it, "An important feature of a learning machine is that its teacher will often be very largely ignorant of quite what is going on inside, although he may still be able to some extent to predict his pupil's behaviour."⁴⁸

So generative AI is neither mechanism nor magic. It is a harness thrown over the chaotic potential of logical space, a way of directing random walks toward scenic vistas. Its basic mathematical operations are simple and insignificant taken individually, but with the amplification of massive iteration they can compute something new and extraordinary.

Combinatorial profundity

I'll now build to my claim that generative AI enables a distinctive sort of aesthetic value. Let's pick up where we left off with Kant – his assertion that sublimity in art is "always restricted to the conditions of agreement with nature". As we saw, it's a bit of a logical juggling act for Kantians to keep this idea in the air: Kant insists that sublimity is properly attributed to the experience in our own minds, not an external object, yet he seems to pound the table for those objects needing to be natural ones. The closest we came to an argument is that our awareness of the human-authored origin of artwork typically precludes

⁴⁶ See Vaswani et al. (2017) for foundational work in contemporary stochastic prediction. See also Bender et al. (2021) for concerns about the efficacy and implications of this technology.

⁴⁷ See, among others, Nichol et al. (2021). Emerging forms of diffusion may be more esoteric still, including proposals for modeling the mathematical properties of quantum physics; see Liu et al. (2023).

⁴⁸ Turing (1950), 458-459

our experiencing it as overwhelming human limits (except, perhaps, across the bridge of fidelity to unbound nature).

You should now be able to see how generative AI begins to pry this argument apart. After all, generative AI systems are not humans. They *do* have limits, but their limits are not human limits. The details of the things they do – the speed and accuracy with which they navigate enormous stochastic landscapes – are literally unfathomable to human minds bound in time and fallible memory.

We need to be careful here. It's not enough to point out that AI can do things we can't. Pocket calculators have been able to do that for decades, and no one would call them sublime. Rather, the point so far is only that we need to be open-minded in how we approach the output of generative AI systems. They will not cleanly fit on either side of a dualistic divide between nature and artifice.

Think again about the sorts of artworks that Kant seems to grudgingly concede might yield derivatively sublime experience: faithful landscape paintings that lead us back to our memories of nature itself. This idea was immensely influential upon the generations of northern European Romantic painters who immediately followed Kant. Perhaps the best examples are in the oeuvre of Caspar David Friedrich (1774-1840), whose brooding sea cliffs, darkened churchyards, and miniscule humans gazing on the middle distance (not too close, not too far) hit all the high notes of Kantian sublimity.

Friedrich is an especially interesting example, as he took himself to be challenging established aesthetic norms. In the early 19th century, artistic boundaries were sharpened through highly-publicized competitions, much like the 21st century contests now being harassed with AI-generated submissions. In 1805 Friedrich submitted two sepia drawings to a competition judged by Goethe, trollishly violating the prescribed theme - 'The Labours of Hercules' – with his images of ordinary people in bucolic landscapes. (Goethe was known for regarding landscape as an inferior artistic mode.) To his own surprise, Friedrich was awarded a shared first prize, starting him on a career that would transform the status of landscape in European art.⁴⁹

Imagine we could ask Kant: do you think Friedrich painted *specific actual* features of nature he personally encountered? Specific trees? Specific cliffs? In some cases (as of named mountain peaks) that might be true. Yet he also clearly created some composite landscapes, as in *Ruins in the Riesengebirge* (1830-1834), which merged features of the seashore and the mountains.⁵⁰ Assuming Kant would concede the sublimity of Friedrich's work to his Romantic descendants, he should be ready to admit that "the conditions of agreement with nature" must be parsed in broad representational categories, not strict veridical recreation of individual natural features.

But there is nothing generative AI loves more than broad representational categories. We can see this by putting it head-to-head with Friedrich's oeuvre. Start from Kant's own examples of sublime scenes: "mountain ranges towering to the heavens, deep ravines and the raging torrents in them, deeply shadowed wastelands inducing melancholy reflection."⁵¹ For each of these scene-types I've chosen a fitting Friedrich work and paired it with a DALL-E 3 generation prompted to yield Kant's phrase in Friedrich's style. See here:

⁴⁹ See Hoffman (2000), 28-29

⁵⁰ Hoffmann (2000), 186.

⁵¹ Kant (1790/2000), 152. (Ak 5:269)

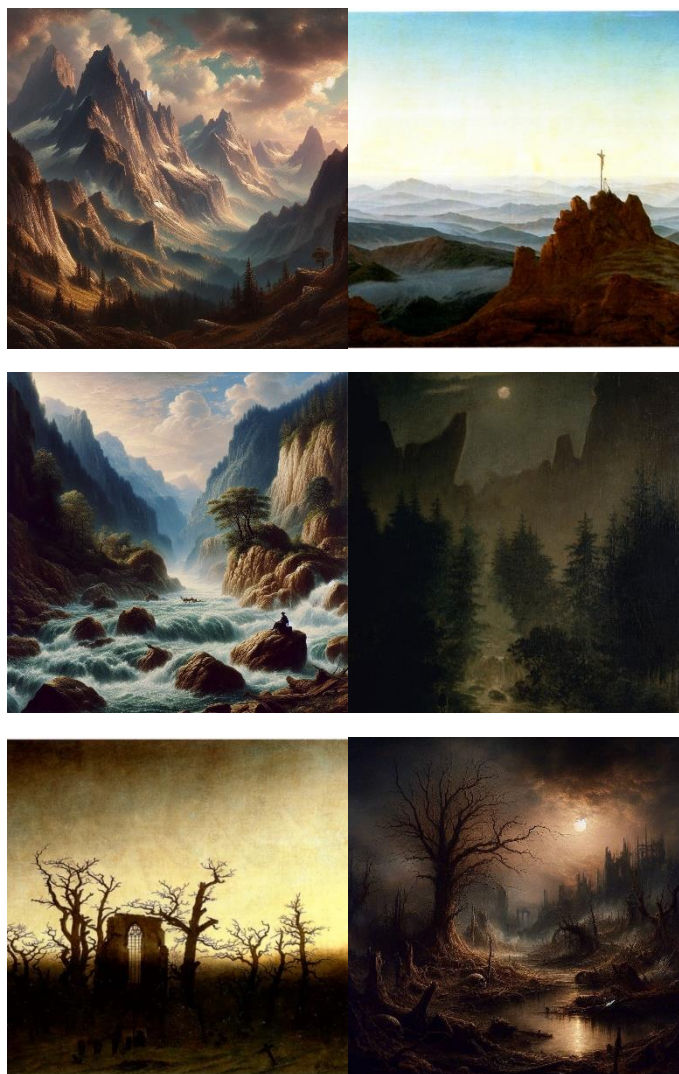


Fig 2. Pairs of works by Caspar David Friedrich alongside a DALL-E 3 prompt. Friedrich works are cropped to match DALL-E aspect ratio. Order randomized within each pair. Friedrich work titles: *Morgen im Riesengebirge* (1819-1811), *Uttewalder Grund* (1825), *Abtei im Eichwald* (1809-1810). DALL-E prompts: ‘Mountain ranges towering to the heavens, early 19th century oil painting’ ; ‘deep ravines and the raging torrents in them, early 19th century oil painting’; ‘deeply shadowed wastelands inducing melancholy reflection, early 19th century oil painting’. All DALL-E generations produced August 20 2024, DALL-3 via Bing Image Creator interface.⁵²

I’ve played a little trick on you here though. Within each pair I haven’t told you which is the DALL-E and which the real Friedrich. There’s no consistent ordering. Try to guess. (If you want the answer key, run a web search on each Friedrich title – or just translate the German.) It does seem like DALL-E can already give us broad representational categories of sublimity-conducive nature just as well as one of the masters of Romantic landscape painting.

⁵² Friedrich images all in the public domain, downloaded from Wikimedia Commons Aug 20 2024. Available at: https://commons.wikimedia.org/wiki/File:Friedrich_morgen-im-riesengebirge.jpg , https://commons.wikimedia.org/wiki/File:Caspar_David_Friedrich_-_Uttewalder_Grund.jpg , https://commons.wikimedia.org/wiki/File:Friedrich_abtei-im-eichwald.jpg

But, still, I haven't yet shown that there's anything *new* here, which is what I promised at the start. At best I've shown that generative AI can approximate whatever relationship to the sublime is attained by excellent human artists - more of the same, only sustained on data rather than bread and cigarettes.

So, one more step in my argument. We started somewhere relatively safe for Kant: with landscape art that aims to be broadly faithful to nature. But already in Friedrich we find departures from natural imagery. Consider *The Sea of Ice* (1823-1824):



Fig 3: *The Sea of Ice* (1823-1824), oil painting by Caspar David Friedrich (1774-1840)⁵³

Some people (certainly among Friedrich's Romantic contemporaries) find this sublime. Yet it is only partly a depiction of nature. And it is definitely not something Friedrich personally witnessed. It is an imaginative work that aims to induce in the viewer the characteristic feelings of sublimity – of being overwhelmed by nature – partly by imaginatively placing human figures and a human product, a ship, amid one of the most dangerous natural environments. If this is indeed a work of sublimity, then we seem on our way to conceding the imaginative mingling of natural and artificial elements can trigger sublime experience.

With that concession made, there's no reason that generative AI couldn't combine representations of natural and synthetic elements in uncountably novel directions. Consider these images, all produced by DALL-E 3 prompted to give us '[X] wreck in the sea of ice, early 19th century oil painting in the style of Caspar David Friedrich':

⁵³ Image in public domain, retrieved from Wikimedia commons Aug 23 2024:
https://commons.wikimedia.org/wiki/File:Caspar_David_Friedrich_-_Das_Eismeer_-_Hamburger_Kunsthalle.jpg



Fig 4. Six works generated by DALL-E 3, via Bing Image Creator interface, prompted with ‘_X_ wreck in the sea of ice, early 19th century oil painting in the style of caspar David friedrich’, Aug 20 2024. Wrecked objects depicted (left-to-right, starting from first row: submarine, new york city subway, zeppelin, formula one race car, space shuttle, quadcopter).

Maybe none of these images strike you as sublime. That’s fine. Not all the Friedrich paintings hit me that way either. The point is only that we can generate endless candidates for sublime experience – and surely some will find the right viewers.

This takes us a half-step from Kant’s view. We’ve softened the insistence on strict representation of nature, but we haven’t yet found anything utterly novel, nothing like that new type of aesthetic value I promised at the start. To get there, notice something interesting about what DALL-E and I did together in the last couple pages. We’ve shown the *combinatorial depth* of generative AI. It seems like we can run together almost any two concepts -- a phrase from Kant (“deep ravines and the raging torrents in them”) + Friedrich’s style; a spaceship + a sea of ice – and the system’s statistical fluency allows it to produce a passable conceptual fusion. There doesn’t seem to be any limit to this semiotic flexibility. It’s fun to illustrate this point with other examples.

For instance: the American painter Edward Hopper (1882-1967) is most famous for *Nighthawks* (1942), depicting lonesome patrons of a diner. One imagines Hopper might have seen something similar in today’s boba (bubble tea) shops, shiny pockets of isolated urbanity. We can ask DALL-E for a depiction of a boba shop in Hopper’s style:



Fig 5. 'people drinking alone in a boba shop, painting by Edward Hopper', created by DALL-E 3 (via Bing Image Creator), Aug 20, 2024

Or consider Hokusai (1760-1849), master of the ukiyo-e print, whose images of late Edo urban life still resonate. Imagine if Hokusai had lived a few generations later, after baseball became the most popular sport in Japan. What would a Hokusai print of the seventh-inning stretch look like?



Fig 6. 'baseball fans doing the seventh inning stretch, woodblock print by Hokusai, created by DALL-E 3 (via Bing Image Creator), Aug 20, 2024

Or imagine that Dorothea Lange (1895-1965), the great photographic documentarian of the Depression era, had survived to observe today's socioeconomic landscape:



Fig 7. 'rideshare driver on lunch break, black and white photograph by dorothea lange', created by DALL-E 3 (via Bing Image Creator), Aug 20, 2024

We can explore the same point in prose. Microsoft's Bing search engine runs on a version of the GPT language model with a live connection to the internet. It will readily compose most things you ask of it. When I prompted it to describe George Eliot's (1819-1880) novel *Middlemarch* (1872), it returned a serviceable, Wikipedia-style summary: "The novel is set in Middlemarch, a fictional English Midlands town, in 1829 to 1832. It follows distinct, intersecting stories with many characters. Issues including the status of women, the nature of marriage, idealism, self-interest, religion, hypocrisy, political reform, and education." So far, so mechanical.

But then I began tossing curveballs. I prompted Bing: "Eliot wrote with a very particular style in 'Middlemarch'. I wonder if you can imitate her style? Please write a description of the Peanuts character Charlie Brown, as if he were a character appearing in Middlemarch." This is what Bing returned:

The sun had just begun to settle in the town of Middlemarch, casting long shadows across the fields and gardens. In the distance, a small figure could be seen walking slowly down the lane, his head bowed and his hands clasped behind his back. This was Charlie Brown, a young man of modest means who had recently come to Middlemarch to seek his fortune. He was not a handsome man, with his large head and round nose, but there was something about him that drew the eye. Perhaps it was the way he carried himself, with a quiet dignity that belied his humble origins. Or

perhaps it was the sadness in his eyes, which seemed to hint at some great sorrow that he carried with him always. Whatever the reason, Charlie Brown was a figure of interest to many in Middlemarch, who wondered what secrets lay hidden behind his melancholy gaze.⁵⁴

That's not bad, is it? It's somewhere in the very lower bounds of how I imagine Eliot herself might have handled the compositional challenge, perhaps on a tired day. The key thing to notice is how Bing's statistical architecture allows it to bring together disparate elements of its training data. No doubt it has imbibed the complete text of *Middlemarch* alongside numerous scholarly commentaries. And no doubt also umpteen descriptions of Charlie Brown. But I doubt the data contains exactly what I requested: there is probably no sample of Charlie Brown / *Middlemarch* crossover fan fiction.⁵⁵ Yet the system's training allowed it to triangulate a prediction of how such a text might go, given the statistical frequency of word sequences found in the two pools of source texts.

We can push this point further by upping the weirdness of the request. This is what I next wrote to Bing:

"That was great! A little bit on the nose for Eliot, but in the ballpark. Can you do the same thing, but this time not about a character. Instead I'd like you describe a scene as if it happened in Middlemarch. This is the scene: a helicopter crashes in the town square. Luckily no one is hurt. One of the helicopter passengers is a famous record producer, who is so grateful to the townfolk that he promises to make one of them famous on TikTok. Please write that scene in Eliot's style, being sure to specify which townspeople in Middlemarch goes on to global fame."

Bing returned several paragraphs of prose, too much to reproduce here. But this is an excerpt, so you can see that it did what was asked:

The sun had just begun to set over the town of Middlemarch, casting long shadows across cobblestone streets and the quaint houses that lined them... Suddenly, there was a loud noise that shattered the peace of the evening. A helicopter had crashed in the town square! The people rushed to the scene, their hearts pounding with fear and anxiety, but to their great relief, no one was hurt... One of the passengers was a famous record producer, who was so grateful to the townfolk for their help that he promised to make one of them famous on TikTok... In the end, it was a young woman named Mary Garth who was selected by the record producer. Mary was a quiet and unassuming person, but she had a beautiful singing voice that has always been admired by her friends and family... Soon, people all over the world were singing along to Mary's beautiful voice.⁵⁶

⁵⁴ Generated by Bing text interface, running on GPT architecture. Sep 4, 2023.

⁵⁵ It's hard to prove this, of course, especially since OpenAI is nonspecific about exactly what has gone into GPT training data. But a quick Google search of 'charlie brown middlemarch crossover' in September 2024 returns no relevant results. A similar targeted search of Archive of Our Own (the largest online fanfiction site) comes to the same end.

⁵⁶ Generated by Bing text interface, running on GPT architecture. Sep 4, 2023.

This is not high art. There's a creakingly formulaic pattern to how Bing inserts requested elements into a stock depiction of an English village, certainly not worthy of Eliot's sharp prose. Yet there's also an unnervingly plausible sensibility behind the constructed narrative. Consider the selection of Mary Garth as the record producer's protégé. If you have not read *Middlemarch*, you might assume Bing simply grabbed up the most prominent name in the novel. But that's not it – Mary is a minor character, far less central than Dorothea or Lydgate. And yet *she is the right choice*. Somehow Bing has alighted on what resembles a sensible artistic decision. I am not, of course, saying that Bing *has* decided anything. It does not 'decide'. That is not how it works. But it can produce outcomes that look like the work of a thoughtful mind, solely through high-powered stochastic fusion.

Still, maybe you think that Bing must have *something* like this scenario in its enormous training data. Not a helicopter crash or a TikTok producer, but perhaps there is something about Mary Garth as a recording artist in the endless trove of fan fiction floating around the internet. So perhaps you are still not convinced there is anything remarkable here. In that case, let's get even weirder. I instructed Bing: "Same deal, but this time you can't start with the sun setting and the long shadows, etc. And now it turns out that the village of Middlemarch is located inside of the 1990 Kevin Bacon film 'Tremors'. Describe what happens." I'm very confident that no one has written *this* before.

Bing's response is again too long to reproduce entirely, but here is a sample, beginning from the chopper crash scene:

The passengers emerged from the wreckage unscathed, but they were not from around here. They were from a far-off land, a place called Perfection Valley. The townspeople of Middlemarch were intrigued by these newcomers, who spoke in strange accents and wore strange clothes. They welcomed them with open arms, eager to learn more about their world. But they soon discovered that these outsiders had brought with them a terrible danger. The creatures that had terrorized Perfect Valley had followed them to Middlemarch. These creatures were giant worms, capable of burrowing through the earth and swallowing anything in their path. The townspeople of Middlemarch were horrified by these monsters, which seemed to be unstoppable.⁵⁷

This will not win any literary prizes. But it is a remarkably fluid integration of very different elements. Notice how it twigs to the surprise of Georgian English villagers at the "strange accents" and "strange clothes" of visitors from 1980s Nevada – it was not prompted with that wrinkle. (Presumably this reflects a common trope of anachronistic crossover fan fiction found in the training data.) This goes beyond 'mad lib' mechanical insertion of semantic variables. It hovers somewhere between pastiche and collage, showing how seemingly any ideas can be representationally mingled without rough edges.

Of course, human artists can do this too. But even the best artists have their limits. No one can perfectly imitate *every* old master or prose stylist. And it takes time to conceive and execute combinatoric novelty. Time runs out for all of us – Burke's great terrifying limit again - long before we could produce even a tiny fraction of the range of works that generative AI can yield in a matter of minutes. Exempted from our

⁵⁷ Generated by Bing text interface, running on GPT architecture. Sep 4, 2023.

human limits, generative AI can explore farther and wider in the space of phenomenal experience than unaided humanity could ever hope to reach. This isn't just combinatorial depth. It is combinatorial profundity.

Now we are approaching the artificial sublime. Remember what Kant said: the mathematical sublime arises when we confront "that which is absolutely great... that which is great beyond all comparison".⁵⁸ Working with generative AI quickly hints at the vastness of combinatoric space, the unfathomable extent of interesting juxtapositions that no human hand will ever make actual. Our human limits glare at us in comparison to these unending possibilities. Think on this fact for awhile and you may run into the emotional crash, that reckoning with our finitude, which Kant identifies as the first step in sublime experience.

What of the second step, the rising again? It is only this: generative AI is the product of human cognition. We humans came up with it. We (a few clever among us, anyway) realized how to make the mathematics run cleanly. And most of us contributed training data in our online words or images. The vast expressive potential of generative AI is only our own rational power, sustained at incredible speed in a sort of conceptual chain reaction.⁵⁹ It is us – Turing's teachers of the machine child mind – who enable this power. We rise again in our own estimation when we see what our reason has wrought.

That is why the artificial sublime must be possible. Combinatorial profundity opens onto the mathematical sublime.

It may seem like I have here merely repeated the same meta trick I called cheap and shallow at the start of the paper. After all, I haven't tried to show that any specific piece of AI-generated work is sublime. Instead I'm now saying that our awareness of all that AI generation makes possible is a route to sublime experience. How is this not just more regression to the meta?

Notice, though, that we are now following Kant directly, not cutting argumentative corners. What we've done now is make use of Kant's own theory – his two-step emotional analysis of the sublime (crushing limits, then expanding power), his insistence that sublimity is a property of the experience itself - to see a way that our encounters with generative AI parallel our encounters with nature. If this is a meta argument, it is certainly not a shallow one.

Indeed, it is hard to see how the argument could go otherwise. It is notoriously tricky to demonstrate aesthetic properties by example. I might provide you with an AI-generated image that triggers sublime experience in *me*, but that does not imply that you will or should have the same experience. Even Kant agreed that there is something utterly subjective in aesthetic experience. ("The judgment of taste is... not a cognitive judgment, hence not a logical one, but is rather aesthetic, by which is understood one whose determining ground cannot be other than subjective."⁶⁰) If you and I have different responses to the same piece, we might judge one another's taste, but (on Kant's view) there is no rational route for either of us to convince the other. Which means any argument for the possibility or impossibility of certain types of sublime experience *must* be general, not dependent on ostensive demonstration.

⁵⁸ Kant (1790/2000), 131-132. (Ak 5:248)

⁵⁹ Once again, I borrow this image from Turing, who describes a learning machine as like "an atomic pile of less than critical size: an injected idea is to correspond to a neutron entering the pile from without". Turing (1950), 454.

⁶⁰ Kant (1790/2000), 89. (Ak 5:204)

And the Kantian point I've developed here is truly deep, so much that I am unsure I've fully characterized it. Kant thought of nature as the home of the sublime because the natural world marks that which is apart from our own minds – it is the background that establishes our limits, and it is also the medium we transcend in rational cognition. This is an idea at the centre of Kant's metaphysics, epistemology, and ethics, and passes down to contemporary philosophy in imagistic grappling between mind and world.⁶¹

If the argument of this essay is right, we should start considering whether generative AI introduces a *third* category, neither natural nor quite of the mind. The closest parallel may be to the ontologically ambiguous status of mathematical objects; numbers exist, if they exist at all, in some way not entirely belonging to the natural world or our minds.⁶² Or perhaps there is another parallel to how some non-natural realists think of the ontology of moral properties, again neither wholly of our minds nor entirely at home in the natural world.⁶³

It may seem surprising that we've ended up in quite a heady place, having started from a photograph of mannequins in a Sydney restroom. But, for Kantians anyway, aesthetics is never far from central questions of philosophy; it is not an accident that aesthetics takes up much of Kant's third *Critique*, completing the system of metaphysical and practical reason. The sublime in particular is, as Kant says, "something serious in the activity of the imagination". The natural world threatens to overwhelm us in its profundity, and the aesthetic of the sublime is a part of how we rise from that submersion. If the argument of this essay is right, the artificial sublime is only another extension of that dynamic.

The starry skies above and the moral law within

I might stop here; I have completed my main argument for the artificial sublime. But I want to conclude with a secondary point, which may in time turn out more important. I want to briefly consider the relationship between the sublime and moral reasoning, a matter close to Kant's heart.

Earlier I mentioned that Kant divides the sublime into two types. So far we have focused entirely on the mathematical sublime. We turn now to the second, which Kant calls the *dynamical* sublime. This, Kant says, concerns "Nature considered in aesthetic judgment as a power that has no dominion over us."⁶⁴ By way of examples Kant offers: "Bold, overhanging, as it were threatening cliffs, thunder clouds towering up into the heavens, bringing with them flashes of lightning and crashes of thunder, volcanoes with their all-destroying violence, hurricanes with the devastation they leave behind, the boundless ocean set into a rage."⁶⁵

⁶¹ I am thinking here especially of Rorty (1979) and McDowell (1996).

⁶² For discussion, see e.g. Clarke-Doane (2020).

⁶³ See, for instance, Dworkin (1996) or Parfit (2011). Elsewhere I (Rini 2020), following Darwall (2006) have argued that Kantians ought to attend to the 'second-personal' aspect of moral claims as a way of finding a middle between the first-personal appearance of free will and the third-personal determinism of the natural world. I am not sure yet what relation, if any, that second-personal domain has to generative AI. There are reasons to think something morally valuable might be found at this intersection – see Jiang et al. (2021).

⁶⁴ Kant (1790/2000), 143. (Ak 5:260)

⁶⁵ Kant (1790/2000), 144. (Ak 5:261)

The terrifying part of nature, then. To an extent, Kant's dynamical sublime is a return to the earlier Burkean idea of sublimity as proximity to contained peril. Kant says that these examples "make our capacity to resist into an insignificant trifle in comparison with their power." Yet, just as with the mathematical sublime, Kant claims there is a second step, a rising above. In the dynamic sublime, this comes when we realize that even the most powerful terrors of nature cannot touch our deepest moral selves. Kant says:

"nature is judged as sublime... because it calls forth our power (which is not part of nature) to regard those things about which we are concerned (goods, health and life) as trivial, and hence to regard its power... as not the sort of dominion over ourselves and our authority to which we would have to bow if it came down to our highest principles and their affirmation or abandonment."⁶⁶

It is difficult to explain this point without getting into Kant's ethical theory, which I can do here only glancingly. Kant argues at great length (especially in the *Groundwork*) that our capacity for moral reason is "not part of nature" because the moral law transcends mere natural contingencies. Morality, on Kant's view, follows necessarily from a proper understanding of moral agency itself, or what it is to responsibly and consistently make choices for reasons.⁶⁷ Our moral integrity arises from our ability to overcome (or rationally integrate) the "bidding from the outside" represented by desires for contingent worldly goods.⁶⁸

At the most extreme, as Kant now appears to be saying regarding the dynamical sublime, nature can threaten us with injury or death. But what it *cannot* ever do, no matter its power, is rob us of our moral integrity. That, Kant thinks, is the most important thing in the world. Hence the rising second step of the dynamical sublime. Nature appears terrifying, and then we remind ourselves that it cannot reach us at our core. Moral reason, Kant thinks, is in an important sense more powerful than all the volcanos and hurricanes in the world.

This point adds to our understanding of why Kant tries so hard to insist that the sublime must be found in nature. He says of nature: "it is mostly rather in its chaos or in its wildest and most unruly disorder and devastation, if only it allows a glimpse of magnitude and might, that it excites the ideas of the sublime..."⁶⁹ Nature's chaotic power is the threat we rise above, and Kant appears to think this form of chaos can be found *only* in the non-human realm.

Yet there are a few hints that Kant thought otherwise when it came to *moral* chaos. Perhaps there are people who present themselves to us as so disconnected from the moral law, so utterly uninterpretable in their devotion to disorder, that their appearance may work on sublime sensitivity not unlike roaring seas or crashing thunder.

This seems to be the thought behind rare moments where Kant talks about finding sublimity in human events. For instance, consider the terrors of war, the summed power of numerous violent choices that might overwhelm us as completely as unceasing waves. From tales of military heroes we learn that even

⁶⁶ Kant (1790/2000), 145. (Ak 5:262)

⁶⁷ Here I rely heavily on Korsgaard's interpretation of Kantian ethics. See Korsgaard (1996a), chapter 6.

⁶⁸ Kant (1785/1959), 448. For discussion, see Korsgaard (1996b), 94 and *passim*. Here I am leaning harder on the Stoic side of Kantian ethics: see Annas (1993) and Sherman (1997).

⁶⁹ Kant (1790/2000), 130. (Ak 5:246)

this can be risen above, according to Kant: “Even war, if it is conducted with order and reverence for the rights of civilians, has something sublime about it, and at the same time makes the mentality of the people who conduct it in this way all the more sublime, the more dangers it has been exposed to and before which it has been able to assert its courage...”⁷⁰

Indeed, in his earlier writings (two decades before the *Critique*) Kant appeared ready to admit that *tragic narratives* might fit his conception of the sublime:

“In my opinion, tragedy is distinguished from comedy primarily in the fact that in the former it is the feeling for the sublime while in the latter it is the feeling for the beautiful that is touched... Even the vices and moral failings often carry with them some of the traits of the sublime... Resolute audacity in a rogue is extremely dangerous, yet it touches us in the telling, and even when he is dragged to a shameful death yet he ennobles himself to a certain degree when he faces it spitefully and with contempt.”⁷¹

Here Kant seems to have left the door open for tragic literature to bring us into sublime contact with moral chaos. This suggests another path toward the artificial sublime. Good novels and works of narrative non-fiction can make us confront the threat of moral disorder. This is not only the obvious and literal sense in which an evil person might do us harm, but the more subtle – and, for Kant, significant – threat that we ourselves might be seduced by evil.

Consider Cormac McCarthy’s (1933-2023) novel *Blood Meridian* (1985). It depicts a group of murderous desperados ranging the Texas/Mexico border circa 1850. The characters are almost uniformly terrible people, but one, referred to as ‘the judge’, is especially noteworthy. He is described as 7 feet tall, utterly hairless, “like an enormous infant”⁷² He is also extraordinarily powerful: he can shoot and ride as well as anyone, he speaks every necessary language, he delivers extemporaneous lectures on desert paleontology. He is also an utterly amoral monster; at one point he buys puppies from a child only to hurl them into a river. This is what he says when the gang members, gathered around the campfire, begin to talk about war:

“Men are born for games. Nothing else. Every child knows that play is nobler than work... In such games as have further stake the annihilation of the defeated the decisions are quite clear. This man holding his particular arrangement of cards in his hand is thereby removed from existence. This is the nature of war, whose stake is at once the game and the authority and the justification. Seen so, war is the truest form of divination. It is the testing of one’s will and the will of another within that larger will which because it binds them is therefore forced to select. War is the ultimate game because war is at last a forcing of the unity of existence. War is god.”⁷³

The judge is a threat – he is powerful, dangerous, and utterly unconstrained. But he is also seductive. He talks in a way that makes it easy to want him to be right, to imagine that we too could achieve his immunity from bad fortune if only we rose to a similarly empowered indifference. For Kant, this sort of

⁷⁰ Kant (1790/2000), 146. (Ak 5:263)

⁷¹ Kant (1764/2011), 19. It is possible the later mature Kant, the system-builder who wrote the *Critiques*, might have disavowed this claim from relative youth. Yet if – as even Kant the theoretician insisted – aesthetic experience is fundamentally subjective, then it is useful to note which experiences informed Kant’s own aesthetic background.

⁷² McCarthy (1985), 335.

⁷³ McCarthy (1985), 248-250.

alluring moral chaos is the deepest threat of all – we are tempted by “resolute audacity in a rogue” because it seems like another (easier and more rewarding) route to the sort of self-mastery that adherence to the moral law requires. Confronting this sort of existential threat, feeling pulled down by intrusion of moral chaos, yet ultimately rising again in a reaffirmation of our commitment to moral duty – why *shouldn't* Kant see this as another form of sublime experience?

We don't need McCarthy's florid violence to encounter moral chaos. Back now to George Eliot (the real one, not GPT manque). In *The Mill on the Floss* we meet Mr. Glegg, “a lovable skinflint” who “was stingy in the most amiable manner”:

“If you expressed a preference for cheese-parings, Mr. Glegg would remember to save them for you, with a good natured delight in gratifying your palate, and he was given to pet all animals which required no appreciable keep. There was no humbug or hypocrisy about Mr. Glegg: his eyes would have watered with true feeling over the sale of a widow's furniture, which a five-pound note from his side pocket would have prevented; But a donation of five pounds to a person ‘in a small way of life’ would have seemed to him a mad kind of lavishness rather than ‘charity’...”⁷⁴

Here Eliot conveys a very different sort of moral chaos, a sort that might be especially threatening to someone of Kant's own dispositions. Mr. Glegg has turned the virtue of thrift into a moralistic fetish. He tells himself that he is merely avoiding waste, and indeed he is not selfish in a grasping, acquisitive sense. But he has deadened his own sensitivity to others' needs such that he can no longer recognize the difference between surplus and sustenance. Eliot flags the everyday temptation of this sort of moral disorder with the deft touch of an expert observer of humanity.

So, then, if we are ready to admit that literary narrative can achieve this effect, why not allow that generative AI might do so? Already some of the best language models are known for their ability to mimic the style of prominent authors. Certainly they are not up to the same standard as the best human writers (neither are most humans). Yet they may already come close enough for these purposes. Imagine an artificial system that performs at about 80% of the skill of a literary genius like Cormac McCarthy or George Eliot. Imagine this system instructed to generate passages adapting their sensibilities to new situations with resonance to our own time and concerns.

Consider, for instance, this excerpt GPT-4 generated for me when I asked it to depict McCarthy's judge in the form of a US military office in Iraq in 2003:

“They talk of peace as an achieved state when peace is but a moment's idyll between two wars,” the judge mused aloud, his voice gravelly grating against the silence of the desert. “War is the true god of this world, more real than any deity conjured by the fantasies of men.” He found a rock, grey and sun-baked, and settled onto it. His soldiers stood near, quiet as serpents, listening.

The judge reached down and took a handful of the rust-colored sand. “We make a game of it, as children playing at soldiers. But the essence of war is the stripping of

⁷⁴ Eliot (1860/1992), 135.

one's humanity, reducing a man to the animal he is. In this place," he scattered the sand, watching it catch on the hot wind, "we are predators amidst the lambs. We bring justice. We are the law. In this theatre of war, the curtain never descends."

As he stood up, his broad shouldered silhouette outlined against the punishing sun, there was an eerie silence. "Remember this, men. In this strange, desolate land, we are the architects of their future. We build empires from the ashes of the vanquished. The conqueror and his spoils. All else is dilution and self-deception."

His gaze hardened, meeting the petrified stares of men. The desert wind wailed around them.⁷⁵

This passage is not perfect, not as high art, nor even as McCarthy pastiche. Too direct, too much melodrama. But it is still remarkably good *for the work of a computer program*, especially produced in a matter of seconds. And I am only trying to show a path forward, not a destination. With more subtle prompting, with more leavening of stylistic reference points, it is not hard to imagine a future version of GPT producing genuinely surprising manifestations of moral chaos.

Imagine you could summon the spirit of Cormac McCarthy or George Eliot at any time. Imagine they sat, ever-patient and ever-capable, prepared to generate thumb portraits of whatever sort of moral disorder might be tempting in your personal situation. Not to yield simple and smug lessons, but instead to provide a foil against which to calibrate your own deliberations. This, I think, is the sort of moral value that Kant should (and perhaps did) allow in the artistic sublime of narrative. And it is something generative AI may soon make available to us in all manner of unexpected ways.

So I'll give the final word to AI itself. I will ask GPT-4 to imagine Kant in a close confrontation with moral chaos. Imagine that Kant sat down across that campfire from the judge. How would that go?

"Moral duty, Herr Kant," he began, the words riding the hot exhale of desert wind. "How does it weigh?"

Kant, drawn into the shadows of his high collared cloak, regarding the white-haired giant, "It is the cornerstone of our humanity," replied the philosopher. "It is dictated not by our desires or interests, but by reason. Duty gives our actions moral worth."

The firelight danced in the judge's pale eyes. "You mean to say our worth is not in our victories, our conquests..."

⁷⁵ Generated by GPT-4. Sep 4, 2023. Exact prompt: "Please write a short passage of fiction in the style of Cormac McCarthy's novel Blood Meridian. In this passage, present the character called the judge as a US military officer in Iraq in 2003, discussing his views on war. The viewpoint is a third personal narrator who has no access to the character's inner thoughts."

“Quite so,” Kant interjected. “It lies in why we choose to act, not the ends, no matter how beneficial they may seem. It is adherence to the moral law that dignifies us, not its outcome.”

The judge considered this before finally speaking. “Interesting,” he said, his voice mocking yet heavy with known truth, “That in this vast, cold universe, we assign ourselves worth based on invisible tenants of duty and morality. Is this our attempt at meaning or is it our futile attempt to ignore our own insignificance?”

Kant watched the embers fly, their glow a brief testament to life’s transient nature. “But is it not this very pursuit of meaning, the imposition of duty, that raises us above our beastly origins?” he asked, his voice a soft echo in the desert silence.

“Aye, you argue well,” the judge murmured, his eyes fixed on the flames. “But then, here in these barren lands of dust and stone, under the gaze of these silent stars, can duty survive? Can morality stand up to the lawlessness of existence?”

Kant returned the judge’s stare, his presence a ghostly mirage in the endless expanse. “As long as there are men capable of reason,” he began, “There will be duty. There will be morality.”

The judge turned towards the endless expanse of the desert and the conversation dwindled to a close. For a long while the only sound was the crackling of the fire, keeping the infinite dark at bay.⁷⁶

References

- Uygar Abaci (2008). ‘Kant’s justified dismissal of artistic sublimity’. *Journal of Aesthetics and Art Criticism* 66(3): 237-251.
- Julia Annas (1993). *The Morality of Happiness*. Oxford University Press.
- Emily Bender et al. (2021). ‘On the dangers of stochastic parrots: Can language models be too big?’ *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* 610-623.
- Walter Benjamin (1931/1977). ‘Short History of Photography’. Translated and edited by Phil Patton. *Art Forum* 15(6): 46-51. Also available at <https://www.artforum.com/features/walter-benjamins-short-history-of-photography-209486/>
- Margaret A. Boden (2016). *AI: Its nature and future*. Oxford University Press.
- Emily Brady (2013). *The Sublime in Modern Philosophy: Aesthetics, Ethics and Nature*. Cambridge University Press.
- Emily Brady (2019). ‘The Environmental Sublime’. In *The Sublime Reader*, ed. Robert R. Clewis, London: Bloomsbury, 355-365.

⁷⁶ Generated by GPT-4. Sep 4, 2023. Exact prompt: “Present a short scene of the judge, sitting with Immanuel Kant beside a desert campfire. They are discussing the concept of morality.”

- Lindsay Brainard (forthcoming). 'What is creativity?'. *The Philosophical Quarterly*.
- Cameron Buckner (2019). 'Deep learning: A philosophical introduction'. *Philosophy Compass* 14(10):e12625
- Edmund Burke (1757/1937). 'Philosophical Enquiry into the Origin of our Ideas of the Sublime and Beautiful'. Reprinted in *Edmund Burke* (ed. Charles W. Eliot). New York: Collier & Son.
- Dan Cavedon-Taylor (2013). 'Photographically Based Knowledge'. *Episteme* 10(3): 283–297.
- Justin Clarke-Doane (2020). *Morality and Mathematics*. Oxford University Press.
- Robert R. Clewis (2021) 'Why the Sublime is Aesthetic Awe', *Journal of Aesthetics and Art Criticism* 79: 301-314
- Robert Clewis (2024). 'Can Painting Evoke the Kantian Sublime?'. *Aesthetica* Preprint, (124), 31-54.
<https://mimesisjournals.com/ojs/index.php/aesthetica-preprint/article/view/4028>
- Paul Crowther (1989). *The Kantian sublime: from morality to art*. New York: Oxford University Press.
- Stephen Darwall (2006). *The Second Person Standpoint*. Harvard University Press.
- Darreonna Davis (2023), 'AI-Generated Drake, The Weeknd Song Not Eligible For Grammy. Here's How It Could Have Been, Though'. *Forbes* September 8 2023.
<https://www.forbes.com/sites/darreonnadavis/2023/09/08/ai-generated-drake-the-weeknd-song-not-eligible-for-grammy-heres-how-it-could-have-been-though/>
- Ronald Dworkin (1996). 'Objectivity and truth: You'd better believe it'. *Philosophy and Public Affairs* 25(2): 87-139.
- Boris Eldagsen (2023), 'My statement when the selection of my image was announced by SWPA'. March 14, 2023. Posted at <https://www.eldagsen.com/sony-world-photography-awards-2023/>
- George Eliot (1860/1992). *The Mill on the Floss*. New York: Knopf.
- Jerry Fodor (1998). *Concepts: Where Cognitive Science Went Wrong*. New York: Oxford University Press
- Jane Forsey (2007), 'Is a Theory of the Sublime Possible?', *Journal of Aesthetics and Art Criticism* 65: 381-389
- Thu-Huong Ha (2024), 'Akutagawa Prize draws controversy after win for work that used ChatGPT', *Japan Times* January 19 2024. <https://www.japantimes.co.jp/culture/2024/01/19/books/akutagawa-prize-book-chatgpt/>
- John Haugeland (1985). *Artificial Intelligence: The Very Idea*. MIT Press.
- Werner Hoffman (2000). *Caspar David Friedrich*. London: Thames & Hudson.
- Robert Hopkins (2012). 'Factive Pictorial Experience: What's Special About Photographs?'. *Noûs* 46(4): 709–731.
- Inside Imaging (2023a). 'Sydney lab wrongfully disqualifies 'AI image'', *Inside Imaging* July 12 2023.
<https://www.insideimaging.com.au/2023/sydney-lab-wrongfully-disqualifies-ai-image/>
- Inside Imaging (2023b). 'Sydney lab makes amends for AI mishap', *Inside Imaging* July 21 2023.
<https://www.insideimaging.com.au/2023/sydney-lab-makes-amends-for-ai-mishap/>
- Inside Imaging (2023c). 'AI agency wins DigiDirect Photo Contest', *Inside Imaging* February 7 2023.
<https://www.insideimaging.com.au/2023/ai-agency-wins-digidirect-photo-contest/>
- Liwei Jiang, et al. (2021). 'Can machines learn morality? The Delphi Experiment'. *arXiv*
<https://arxiv.org/abs/2110.07574>
- Immanuel Kant (1764/2011). *Observations on the Feeling of the Beautiful and Sublime* (eds. and trans. Patrick Frierson and Paul Guyer) Cambridge University Press.
- Immanuel Kant (1785/1959). *Groundwork for the Metaphysics of Morals*. (Trans. Beck). New York: Macmillan Library.

- Immanuel Kant (1790/2000). *Critique of the Power of Judgment*. Edited and translated by Guyer and Matthews. Cambridge University Press.
- Thomas DaCosta Kaufmann (1976), 'Arcimboldo's Imperial Allegories', *Zeitschrift fur Kunstgeschichte*. 39(4): 275-296.
- Michael Kimmelman (2022). 'It was a mystery in the desert for 50 years'. *New York Times* Aug 19 2022. <https://www.nytimes.com/interactive/2022/08/19/arts/design/michael-heizer-city.html>
- Christine M. Korsgaard (1996a). *Creating the Kingdom of Ends*. Cambridge: Cambridge University Press
- Christine M. Korsgaard (1996b). *The Sources of Normativity*. Cambridge University Press
- Yann LeCun et al. (2015). 'Deep Learning'. *Nature* 521: 436-444.
- Ziming Liu et al. (2023). 'GenPhys: From Physical Processes to Generative Models', *arXiv* <https://arxiv.org/abs/2304.02637>
- Kari Elise Lokke (1982). 'The Role of Sublimity in the Development of Modernist Aesthetics', *Journal of Aesthetics and Art Criticism* 40(4): 421-429
- Ada Lovelace (1843), 'Notes on translation of Menabrea's Sketch of the Analytical Engine'. Originally published in the *Scientific Memoirs* 3: 666-731. This text as reproduced in the *Classics in the History of Psychology* web archive (Editor: Christopher D. Green). <https://psychclassics.yorku.ca/Lovelace/lovelace.htm>.
- Hilary Mantel (2009). *Wolf Hall*. Picador: New York.
- Cormac McCarthy (1985). *Blood Meridian: or The Evening Redness in the West*. New York: Random House
- John McCarthy and Patrick Hayes (1969). 'Some philosophical problems from the standpoint of artificial intelligence', in *Machine Intelligence 4* (eds. Meltzer and Michie) Edinburgh University Press, 463-502.
- John McDowell (1996). *Mind and World*. Harvard University Press.
- Allen Newell and Herbert A. Simon (1976), 'Computer Science as Empirical Inquiry: Symbols and Search', *Communications of the ACM* 19(3): 113-126.
- Alex Nichol et al. (2021). 'GLIDE: Towards photorealistic image generation and editing with text-guided diffusion models'. *arXiv* <https://arxiv.org/abs/2112.10741>
- Derek Parfit (2011). *On What Matters*. Oxford University Press.
- Allison Parshall (2023), 'How This AI Image Won a Major Photography Competition'. *Scientific American* April 21 2023. <https://www.scientificamerican.com/article/how-my-ai-image-won-a-major-photography-competition/>
- Regina Rini (2020). 'Contingency inattention: Against causal debunking in ethics'. *Philosophical Studies* 177(2): 369-389.
- Kevin Roose (2022), 'An A.I.-Generated Picture Won an Art Prize. Artists Aren't Happy'. *New York Times* September 2 2022. <https://www.nytimes.com/2022/09/02/technology/ai-artificial-intelligence-artists.html>
- Richard Rorty (1979). *Philosophy and the Mirror of Nature*. Princeton University Press.
- David E. Rumelhart et al. (1986). 'Learning Representations by Back-propagating Errors'. *Nature* 323: 533-536
- Adam Schrader (2023), 'An Australian Photographer Was Disqualified From a Photo Contest After Her Submission Was Mistakenly Deemed A.I.-Generated'. *Artnet* July 17 2023. <https://news.artnet.com/art-world/australian-photographer-disqualified-ai-generated-2337906>

- Adam Schrader (2024), 'A Photographer Wins a Top Prize in an A.I. Competition for His Non-A.I. Image', *Artnet* June 14, 2024. <https://news.artnet.com/art-world/miles-astray-non-ai-photo-wins-ai-competition-2500809>
- Nancy Sherman (1997). *Making a Necessity of Virtue: Aristotle and Kant on Virtue*. Cambridge University Press.
- Guy Sircello (1993) 'How is a Theory of the Sublime Possible?', *Journal of Aesthetics and Art Criticism* 51: 541-550
- Susan Sontag (1977). *On Photography*. New York: Farrar, Straus and Giroux.
- Alan Turing (1950), 'Computing Machinery and Intelligence'. *Mind* 59(236): 433-460.
- Paschal Mmesoma Ukpaka (2024). 'The creative agency of large language models: a philosophical inquiry'. *AI and Ethics* preprint: <https://doi.org/10.1007/s43681-024-00557-9>
- Ashish Vaswani, et al. (2017). 'Attention is all you need'. *Proceedings of the 31st Intl Conf on Neural Information Processing Systems* 6000-6010.
- Dawn M. Wilson (2022). 'Against Imprinting: The Photographic Image as a Source of Evidence'. *Social Research: An International Quarterly* 89 (4):947-969.
- Rachel Zuckert (2003), 'Awe or Envy: Herder contra Kant on the Sublime', *Journal of Aesthetics and Art Criticism* 61(3): 217-232.
- Rachel Zuckert (2019). 'Kant's Account of the Sublime as Critique', *Kant Yearbook* 11(1): 101-119.