

Young, N. and Terrone, E. (2024), “Growing the Image: Generative AI and the Medium of Gardening”, *The Philosophical Quarterly*.

# Growing the Image: Generative AI and the Medium of Gardening

Consider the following passage from an article about the artist Sean Mauss who was disqualified from a competition to design the cover of a science fiction book because he employed Midjourney, a generative AI program which transforms text prompts into images:

The same day the winner was announced, readers and fans on Twitter were questioning whether the art was created at least in part using AI tools. The incident highlighted a growing crisis of trust in science fiction and fantasy publishing: in a world where AI-generated media is common, do you know the work you’re looking at was made by a human?

[...]

Readers and community members responded negatively to the revelations about Bob the Wizard’s cover. But in the long term, it’s not clear if these tools will remain forbidden. Misrepresenting work as AI-free is one thing – but what happens if generative AI systems become more ingrained in the artistic process? (Sato 2023)

Here, we find two broad ways in which we might characterise Midjourney. The question “[how] do you know you’re looking at [an image] that was made by a human?” suggests that Midjourney is some sort of *artificial agent*: if a human didn’t make the image who did? The most obvious answer would be Midjourney itself. Alternatively, we see that Midjourney can also be thought of as a type of *tool* which has the potential to become more “ingrained in the artistic process” over time. From this perspective, Sean Mauss *used* Midjourney to create an image. We will argue that neither approach is satisfactory and we propose thinking of Midjourney, which will be our case study, and similar text to image systems as a new type of artistic medium.

## 1. Midjourney as an Agent

If Midjourney’s artistic role is that of an agent, then it could play a variety of different agentive roles. If Midjourney were sufficiently person-like then it could be thought of as an artist. If it were only somewhat person-like it might still play a simpler agentive role such as co-author, or as some sort of assistant or helper.

However, as Anscomb (2022, p. 25) points out, it is hard to think of AI systems like Midjourney as artists because they do not seem as though they could have intentions to create art. Although Midjourney can produce outputs that might be mistaken for those produced by a traditional artist,

there does not seem to be any reason to think that generative AI systems have any mental states at all (cf. Chalmers 2023).

Despite this, however, the following passage suggests that Anscomb still believes AI systems play some sort of agentive role in the creative process:

[A]lthough AI agents are not themselves artistically creative, they can work iteratively without human intervention to non-accidentally generate the formal features of an image, and thus *contribute* to the realisation of some of the salient properties of a work. Whether this *deserves some greater or lesser share of the credit* for the production of the work qua art depends upon whether one takes a monist or pluralist conception of artistic value. (ibid. p. 35, our emphasis)

Not only does Anscomb refer to text-to-image systems as “AI Agents” throughout her paper, she suggests here that even if they are not artists proper, they might still deserve some degree of “credit” for the “contribution” that they have made to an artwork’s creation by working “autonomously” and “iteratively”:

Users of Midjourney lack direct control over exactly what sort of image is produced. In this sense Midjourney is working autonomously so as to provide some of the formal features of the image. But we doubt that this is enough to think that Midjourney is “creditworthy”, or, to use a similarly agency-infused phrase, has made a “contribution” to the artwork’s features.

To see why autonomy is not sufficient for attribution of credit, consider the following example. As I pour wine into a glass you take photos of the liquid splashing and rippling as the glass is filled. The wine is autonomous in the sense that neither I nor you has direct control over exactly how the liquid will splash into the glass (e.g. the size of the ripples, how many bubbles appear) but we would not think that the wine deserves any credit for the resulting photos in any interesting sense, nor would we say it has made any sort of contribution.

A supporter of Anscomb’s view might reply that she is using ‘agent’ in the sense as used by Boden (2016, p. 45) who defines an AI agent as “a self-contained (‘autonomous’) procedure”. With this understanding of an agent, we might say that while it is not forming its own intentions, a system such as Midjourney, unlike the splashing and rippling wine, is not only working without direct human intervention but also constitutes a procedure.

However, this “non-intentional” understanding of an agent is quite different from the “intentional” sense in which human users of Midjourney are agents (cf. Moruzzi 2022). Thus, it is hard to see how such different sorts of agents can cooperate in a meaningful sense so as to achieve co-authorship. Co-authorship, indeed, seems to require not simply agency but rather the *same sort of* agency.

If, for the sake of the argument, we concede that Midjourney is an agent in Anscomb’s sense, we are left with the dilemma of ascribing the artistic merit of the resulting image either to Midjourney’s

actions or to the user's actions since there is no way to make sense of their cooperation as agents. Both options are unsatisfying. While ascribing the artistic merit to the human user would overlook Midjourney's active contribution, ascribing the merit to Midjourney would downplay the creative activity of prompt-crafting.

## 2. Midjourney as a Tool

There is an obvious alternative to thinking of Midjourney as an agent. We do not think of a painter's brush as deserving credit for its contribution to a painting but rather, the brush is a *tool* used by the artist to create images. As Hertzmann puts it: "Computers do not create art, people using computers create art" (2018, p. 2). Should we not say the same thing about Midjourney?

The idea that Midjourney and similar systems are tools we take to be quite intuitive. Wojtkiewicz (2024) and Martinez and Scarbrough (draft) draw on this intuition to argue that images produced through generative A.I. can be art. In this section we shall put some pressure on this intuition. If Midjourney is a tool, it is quite unlike other tools. While our arguments here will not be definitive they will serve as a motivation for exploring other possibilities and in particular as motivation for our preferred view which we shall present in the next section.

If Midjourney is a tool, what sort of tool is it? Lowe (2014) splits tools into two types: utensils and machines. Utensils involve their users exerting more or less continuous control to carry out their functions. A paintbrush won't paint and a pen won't write unless they are being moved by someone's hand. Machines on the other hand are autonomous, in the sense that they can carry out their function without constant user input. Clocks for example can be set up and will carry out their function of telling the time without further user input. While a great deal of artistic creation is carried out by utensils, we can still think of some artistic tools that would be categorised as machines on Lowe's account. A drum machine can be set up to play a rhythm and then left to do its thing; we can press record on a video camera and it will record whatever is put in front of it without further human input.

If Midjourney is a type of tool it is clearly a type of machine rather than a type of utensil, as it does not require constant action on the part of the user to create an image; rather, a user enters a prompt and Midjourney will set about creating an image. However, Midjourney differs from video cameras and drum machines in that it is *inherently unpredictable*. A user can specify in their prompt what they want an image to depict, its colour scheme, the style etc. But they do not have control over the exact formal properties of the image that Midjourney spits out.

One might object here that Midjourney's unpredictability is not especially unique. An old drum machine might be unpredictable in so much as its owner is never quite sure whether it will turn on when it is plugged in, and a watercolour painter, even an extremely skilled one, is not able to control exactly how the paper will absorb and distribute the paint that they apply. Yet, this is no reason to think that they are not tools.

The comparison with the drum machine has a straightforward response. ‘Unpredictable’ should not be taken to mean ‘unreliable’. An old drum machine has a proper function –producing rhythms– and when it fails to switch on, it fails to fulfil its proper function. As Esposito (2022 p. 9) puts it, “If the outcome of a traditional machine becomes unpredictable, we do not think that it is creative or original –we think that it is broken”. Midjourney however, is reliably unpredictable. A prompter can never be sure what image they are going to get, but this is a feature, not a bug.

As regards the watercolour case, it is indeed true that the painter cannot be entirely sure how paint will –atom for atom– distribute itself over the canvas, but this will only affect the fine details of the image. Even if they do not have complete control over how the picture surface is marked, a watercolour painter still has control as to what depicted objects go where in her picture, as well as their size, shape, and colour. Even a photographer still has a relevant amount of fine-grained control as to what will appear in her picture. A Midjourney prompter, on the other hand, will always lack this sort of fine-grained control. Entering the prompt ‘watercolour painting of a London street’ will produce an image of a London street that looks like it was painted in watercolours, but Midjourney will decide what this street looks like, what sort of objects and people populate the image, and what specific style of watercolour is produced.

Helliwell (2023) suggests that this unpredictability might count against the idea that generative AI systems are simply tools:

An attitude common among philosophers and computer scientists is that AI is just a tool. I would advise against forming this judgement hastily. We are certainly using AI for our purposes, but there is much about AI systems that differs from what we might consider to be a typical artistic tool [...] outputs from machine learning systems may not be fully within our control—or our understanding.

While Helliwell does not deny that users of generative AI can be given some creative credit, the more autonomous, unpredictable work is being performed by the system, the more pressure is put on the idea that a generative AI such as Midjourney is “just a tool”.

### 3. Midjourney and Gardening

We propose, perhaps surprisingly, that creating images with Midjourney is best understood through a comparison with gardening. Midjourney is not an agent which creates, or contributes to the creation of, art, neither, strictly speaking, is it a tool which artists use; rather, it is more similar to a plot of land in which a gardener might sow seeds.

Philosophers like Bruno and Spinoza have deployed the Latin expression *natura naturans* to stress nature’s creative capacity to autonomously generate and dynamically shape things in ways that are so complex and chaotic to appear unpredictable. Spinoza contrasts *natura naturans*, what nature is doing, to *natura naturata*, what nature has done. Borrowing this terminology, we might say

that a garden is a piece of *natura naturata* that the gardener has produced together with *natura naturans*.

Still, there are cases in which the gardener will have some significant control over how *natura naturans* makes her plants grow: she can cut back the branches of bushes that grow over the pathway, or cut the early blooms of her roses to facilitate the growth of more flowers later in the summer. Of course there is a reasonable chance that the bush will grow back and block the pathway in just the same way as it did previously, and pruning the roses early does not guarantee that more flowers will grow later. The gardener can coax her garden into growing the way she wants it to, but she cannot shape it with the same precision she might fold a piece of paper or shape a lump of clay. She is thus somewhat detached from the results of her actions: while she can manipulate her secateurs in real-time she will have to wait to see whether the early pruning of the roses has the desired effect. This process can be thought of in terms of iteration. The gardener acts on her garden (planting seeds, pruning flowers etc.) and then waits to see how *natura naturans* responds. How it responds will influence her next move. Doing so, the gardener can exert skillful yet incomplete control over a piece of *natura naturans* which has a significant degree of autonomy in its own development.

Drawing on this account of *natura naturans*, we introduce the category of *machina naturans* to characterise those machines that are capable of generating and shaping things in the same way as nature. Generative AI systems are paradigm cases of *machina naturans*. The practice of making images with Midjourney is in this sense strikingly similar to gardening. Both, as “*naturans*” systems, are under the incomplete control of the user.

First, just as the types of flowers that grow will depend on the seeds sown, images produced by Midjourney will depend on the prompts it receives as input. The prompt:

/imagine a sunflower

is going to produce pictures of sunflowers, not pumpkins or primroses. However, just as which flowers can be grown, and how well they grow, depends on the environment in which the gardener is working, what images can be produced, and how well they can be produced will depend on which text to image system is being used. More importantly, the user of Midjourney, like the gardener, will not be able to control every detail of the images that are created. While a painter has a very high degree of control over how her painted sunflower looks, its precise size and shape, the prompter, just like the gardener, does not: while you can ask Midjourney for a tall or short sunflower, you cannot ask it for a sunflower that takes up exactly one third of the canvas.

At the same time, like the gardener, a prompter can iterate on Midjourney’s products to get (or get closer to) the result for which they are aiming. If Midjourney produces an image with pink flowers but a user decides that they would prefer green they can update their prompt and then hit the ‘remix’

button, and then, as the gardener must wait for *natura naturans* to make its move, wait a minute or so before a new image is produced by *machina naturans*.

Arguably, neither the gardener nor the Midjourney user have what we might call *fine-grained* control over their outputs. As we have seen, a painter can paint a sunflower in the precise way that they want to, but a gardener cannot grow one, and a prompter cannot generate one, with this degree of control. At the same time, it does not seem as though gardeners and prompters have no control, or even limited control, over their output: bushes can be shaped and cut back, Midjourney prompts can be tweaked so that future iterations fall more in line with what the user wants.

As “the gardener [is] a creative agent who is nevertheless thoroughly dependent on the cooperation of natural processes” (Cooper 2009), the prompter is a creative agent who is nevertheless thoroughly dependent on the cooperation of artificial processes. Given the structural similarities between Midjourney and gardening, the reasons why the latter can be an artistic medium should also be reasons why the former can be an artistic medium as well.

#### 4 Midjourney as a Medium

These considerations give us a third way of categorising Midjourney’s role in relation to its user. Midjourney is not an agent, nor a tool, but a medium. In the first instance, we can think of an artistic medium as the stuff that constitutes a particular work of art. We describe artists as *working in* a medium: a sculptor might work in the medium of marble or bronze, a painter in the medium of oil on canvas etc.

A slightly more sophisticated way of understanding artistic media casts them as systems of resources and practices that mediate between the artist and the audience (cf. Wollheim 1980, Davies 2004, Lopes 2014, Irvin 2022). From this perspective, Wollheim introduces the term *recalcitrance* to designate the unique difficulties which “materials present that can be dealt with only in the actual working of them” (1980, p. 28). An idea which Thomson-Jones develops as follows:

“[t]he medium presents particular challenges and possibilities for artistic creativity, and the artwork makes manifest the artist’s response to these challenges and possibilities.” (2016, p. 40; see also Nguyen 2020, pp. 145–152).

If this is correct, then a gardener’s medium is not simply the plants, seeds, and soil, but the distinct *recalcitrance* –the distinct challenges– that these physical materials provide in relation to the tools that she has available. Specifically, given the autonomous and generative character of *natura naturans*, the medium of gardening exhibits what we might call a *dynamic recalcitrance* characterised by unpredictability and autonomy of biological processes. A gardener can only coax, through various iterations, how her garden grows. While painters and sculptors face distinct challenges, they do not work with materials which are intrinsically autonomous. Experts in traditional image making arts can exert a very fine-grained control over their materials: figures in

photographs or paintings do not change shape, colours do not creep across the canvas when their author is not looking.

Recalcitrance is a useful way of thinking about different artistic media, but it should be further refined. The stuff that the artist works on (marble, paper and paint) presents various forms of recalcitrance, but so too do the tools (chisels, paintbrushes) that the artist uses. In water colour painting, one form of recalcitrance comes from the fact that the paint can spread haphazardly as it soaks into the canvas, but the paintbrush brings its own recalcitrance too: the difference between a novice painter and an expert is that only one can wield the paintbrush with enough skill to mark the paper as desired.

The recalcitrance that a gardener faces seems to lie much more in the materials with which they are working than it does their tools. Given the autonomous and generative character of natural plants, the gardener must coax what she grows into growing the way she wants it to. The tools she uses to do this, spades, shears, watering cans etc., are not especially hard to use, there are no expert and novice watering can users; rather, it is the resistance presented by nature itself which she needs to overcome. While this dynamic recalcitrance is somewhat different to what we might call the *inert* recalcitrance faced by the painter or sculptor, this does not make it any less of a challenge.

Traditional media, having inert recalcitrance, inherit predictability from the tools which constitute them. A paintbrush is a predictable tool, and a canvas and paint are passive, thereby allowing the painter to have *strong* meaningful human control over what she creates. While both gardening and Midjourney involve predictable tools (such as spades or keyboards respectively), these tools are used to work on dynamic, unpredictable stuff. In both cases, it is never certain what will grow from the seeds planted or the prompts entered; while the artist has meaningful control over what they create, it will be only *weak* control. Nevertheless, they will still have sufficient control to stamp their “personal, expressive signature” over what is created. In this sense, Midjourney makes room for the sort of “meaningful human control” that, according to Epstein et al. (2023 p. 4), “is achieved if human creators can creatively express themselves through the generative system, leading to an outcome that aligns with their intentions and carries their personal, expressive signature”.

One may still object that the same sort of weak meaningful human control can be found in styles of painting such as Pollock’s ‘action painting’. However, action painting is just one way in which to engage with the medium of painting, while making unpredictable pictures is the only way to engage with the medium of generative AI. In painting unpredictability is a possibility, in Midjourney it is a necessity.

To sum up, generative AI systems such as Midjourney can be seen as an artistic medium, comparable to gardening. Characterising Midjourney as a type of agent captures its autonomous and unpredictable nature, but not its crucial dependency on a human creative project. Characterising Midjourney as a tool, on the other hand, fits well with our idea that we can use

generative AI as a means of creation, but does not capture Midjourney's autonomy: an unpredictable tool is a malfunctioning one rather than an autonomous one.

Considering Midjourney as a dynamic medium incorporates aspects of both the agency and tool views, while avoiding their shortcomings. This allows for the recognition of the system's autonomous nature, similar to the agency view, but also fits with the perception of Midjourney as a means of creation, as per the tool view, and yet acknowledges its autonomy as an integral, functional feature, not a malfunction. This approach permits Midjourney to be understood as neither fully autonomous nor completely controlled, but as a dynamically recalcitrant medium with which creators must grapple.<sup>1</sup>

## References

Anscomb, C. (2022). Creating Art with AI. *Odradek*, 8(1), 13-51.

Boden, M. (2016). *AI: Its nature and future*, Oxford, Oxford University Press.

Chalmers, D. J. (2023). Could a large language model be conscious?. arXiv preprint arXiv:2303.07103. <https://arxiv.org/abs/2303.07103>

Cooper, D. E. Gardens; in Davies, S. et al. (Eds.). (2009). *A companion to aesthetics* John Wiley & Sons.

Davies, D. (2004). *Art as performance*. New York, Blackwell.

Epstein, Z., Hertzmann, A., Investigators of Human Creativity, Akten, M., Farid, H., Fjeld, J., ... & Smith, A. (2023). Art and the science of generative AI. *Science*, 380(6650), 1110-1111.

Esposito, E. (2022). *Artificial Communication*, Cambridge, MA: MIT Press.

Helliwell, A. (2023, November 2). Contribution to "Eight Scholars on Art and Artificial Intelligence." *Aesthetics for Birds*. Retrieved from <https://aestheticsforbirds.com/2023/11/02/eight-scholars-on-art-and-artificial-intelligence/>

Hertzmann, A. (2018). Can computers create art?. In *Arts* (Vol. 7, No. 2, p. 18).

---

<sup>1</sup> This research was funded by the ERC-StG Project 101040535 ("PEA – The Philosophy of Experiential Artifacts"). Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. This work has been presented at various workshops and conferences, and we are grateful for the useful feedback received. Special thanks to Chiara Brozzo, Elisa Caldarola, Jacopo Domenicucci, John Kulvicki, Sabrina Martinez, Irene Olivero, Elizabeth Scarbrough, and Iris Vidmar-Jovanovic.



- Irvin, S. (2022). *Immaterial: Rules in Contemporary Art*. Oxford, Oxford University Press.
- Lowe, E. J. (2014). How real are artefacts and artefact kinds?. *Artefact kinds: Ontology and the human-made world*, 17-26.
- Lopes, D. (2014). *Beyond Art*. Oxford, Oxford University Press.
- Martinez, Sabrina; Scarbrough, Elizabeth. (draft). "Co-Creators or Mere Tools? AI and Art".
- Moruzzi, C. (2022). Creative agents: Rethinking agency and creativity in human and artificial systems. *Journal of Aesthetics and Phenomenology*, 9(2), 245–268.
- Nguyen, C. Thi. (2020). *Games: Agency as Art*. Oxford, Oxford University Press.
- Sato, M. (2023, June 9). How AI art killed an indie book cover contest. *The Verge*. <https://www.theverge.com/2023/6/9/23752354/ai-spfbo-cover-art-contest-midjourney-clarkesworld>
- Thomson-Jones, K. (2016). Movie Appreciation and the Digital Medium. In *Current Controversies in Philosophy of Film* (pp. 36-54). Routledge.
- Wojtkiewicz, K. (2023). How Do You Solve a Problem like DALL-E 2?. *The Journal of Aesthetics and Art Criticism*, 81(4), 454-467.
- Wollheim, R. (1980). *Art and its Objects: An Introduction to Aesthetics*. 2nd edn. Revised, New York: Harper and Row.