Misembodied Experiences of Artificial Art

Tom McClelland*[1]

[1] *Department of History and Philosophy of Science, Free School Lane, Cambridge, CB2 3RH, UK*

Abstract

Engagement with art typically rests on the co-embodiment of the artist and the audience. The artist is an embodied being – they have a body through which they act in the world – and their embodiment is reflected in their work. A painting, for example, contains information about the specific movements the painter made with their brush, and those implied movements are in turn suggestive of particular mental states. The viewer is able to extract this information from the artwork because they too are embodied and can use their own body to simulate the movements of the artist along with the mental states indicated by those movements. AI can be used to generate sophisticated works of art without that AI being embodied. But are our responses to such artworks still embodied? And what might this mean for the value of such AI-generated art? I argue that viewers will typically misembody AI art. That is, we typically respond to these artworks as if they were produced by an embodied artist. However, I also argue that the fact we misembody such artworks need not count against their aesthetic value.

Keywords

AI art, creativity, embodiment, illusion, value

1. Introduction

Embodiment is an integral part of our engagement with art. The artist's embodiment is reflected in the artwork they produce. Each brushstroke is the product of a specific action and contains information about the precise movement of the artist's body. And insofar as the artist's movements reflect their own mental states, those brushstrokes can contain information about the artist's cognitive and affective processes. The viewer is able to extract this information from the artwork because they too are embodied. Having a body relevantly similar to that of the artist allows you to understand how the artist moved to create those brushstrokes, and to understand the mental states expressed by those brushstrokes. These layers of comprehension and appreciation are thus possible because of the *co-embodiment* of artist and audience.

AI systems for generating original works of art are becoming increasingly sophisticated and their impact on the artworld cannot be ignored. Unlike flesh-and-blood artists, these artificial systems are disembodied. Although the AI is physically realised on some computer hardware it does not have a physical body through which it acts on the environment. When an AI generates an image it does not physically put marks on a surface. And when it creates a piece of music, it does not physically strike an instrument. Rather, these systems generate strings of code. But if embodiment is so integral to our engagement with artworks, how do viewers experience the products of these disembodied systems?

In this paper I argue that our experience of AI artworks typically misembodies the originator of the artwork. That is, we experience the artworks as if they were created by an embodied artist. Even when we know that they are the product of a disembodied AI, our embodied responses to the artwork make it seem as though the artwork were created through an embodied process. This then raises another question: does it matter that our embodied responses to AI-generated artworks are in some sense erroneous? It might be argued that our embodied responses constitute a problematic illusion that counts

* Corresponding author
ICCC'22 Workshop: The Role of Embodiment in the Perception of Human & Artificial Creativity, June 27–28, 2022, Bozen, Italy
EMAIL: twm30@cam.ac.uk
ORCID: 0000-0001-5956-1425

against the value of the artworks. I argue, however, that the value of an ordinary human-made artwork has little to do with whether our embodied responses to it are accurate and that AI-generated artwork should not be held to different standards.

My paper proceeds in three stages. First, I pin down the relevant sense of the term 'embodiment' and summarise the key features that characterise our embodied engagement with artworks. Second, I consider whether our responses to AI-generated artworks are embodied and argue that we typically misembody AI art. Third, I explore what this means for the value of AI-generated art by comparing it to other cases in which our embodied responses to artwork are erroneous and argue that such responses are unproblematic.

2. Action and Imitation

Guckelsberger, Kantosalo, Negrete-Yankelevich & Takala [1] point out that the term 'embodiment' is polysemous and suggest that work on embodied and disembodied computation must be explicit in how it uses the term. I use the term to characterise '…systems with a physical body that can interact with the environment by being subjected to and by exercising physical force' ([1], p. 2). I will also talk of 'embodied responses' in which an agent uses their own embodied abilities to simulate the movements of others. Being physically embodied is a pre-requisite of having such embodied responses. This relatively modest conception of embodiment is neutral on some of the more contentious issues surrounding the term, such as whether cognition is ultimately skull-bound. Similarly, this modest conception is not bound to theoretically loaded concepts in the vicinity such as 'structural coupling' [2].

So what role does our embodiment play in our engagement with art? There are good reasons for thinking that our own embodiment plays a central role in how we respond to works of art. We understand the embodied process through which a work of art is created by utilising our own embodied nature. Art engagement is thus built on the co-embodiment of the artist and the viewer: the viewer is embodied, the artist is embodied, and this shared embodiment is integral to the viewer's experience of the artist's work.

To appreciate the role of embodiment in our engagement with art, we should first turn to the role of embodiment in our engagement with the world in general. Understanding the actions of others is far from straightforward. Our senses might tell us how an agent's body is moving, but that isn't always enough to tell us what they are doing. Consider seeing someone drinking a cup of tea. You can perceive that the hand holding the cup is moving to the agent's mouth, but why do you interpret this as the agent drinking from the cup? After all, there are any number of other actions that are consistent with what you've seen. The agent might be warming their lips with the cup, smelling the cup's content or even sending a secret signal to a friend. One promising and well-supported suggestion is that our understanding of this action relies on *simulation*. When you see the agent's hand move, your motor system 'mirrors' what you see and the motor neurons involved in lifting a cup yourself are activated [3]. This process occurs 'off-line' meaning that you don't overtly imitate the movement. However, the covert activation of this motor process can than trigger other states associated with that movement. Specifically, it recalls information about what *you* were doing when you engaged in those movements. Because you make those movements when you're drinking the simulation suggests that the agent is drinking too. So here, your understanding of another person's actions rests on your underlying ability to simulate their movements, and your ability to simulate their movements rests on you both being embodied agents.

Although there is much more that could be said about the role of motoric simulation in action understanding, the foregoing gives us enough to shed light on our engagement with art. Consider the experience you have of viewing the ballet. Far from just witnessing the movements of the dancer, your brain simulates those movements. And those simulations can feed into various aspects of your overall experience of the ballet, including your enjoyment of them. In a study by Kirsch, Dawson & Cross [4], participants were shown a video of a ballet move that they couldn't do. They were then taught how to do that movement and shown the video again. Participants' assessment of how much they enjoyed the video increased after learning the move. Kirsch *et al* interpret these results in terms of neural simulation.

Before learning how to do the move, participants were not able to neurally simulate the relevant motor processes. Once they learned how to do it, this simulation became possible. The fact that their enjoyment of the video increased after learning the move indicates that simulating a dancer's movements contributes to your enjoyment of those movements. So again, your co-embodiment with the performer plays an important role in your aesthetic appreciation of their performance.

Interestingly, this extends to static representations of bodily movement. A painting might offer a snap-shot of a body in motion, and we are able to simulate the movement of the depicted figure. When you look at a Degas painting of a ballet dancer, you might simulate their movement. And this movement can contribute to your understanding and appreciation of the artwork just as it does in the case of the real ballet-dancers. This is particularly important to your affective response to a painting, as captured by Renaissance thinker Leon Battista Alberti:

> The painting will move the soul of the beholder when the people painted there each clearly shows the movement of his own soul...we weep with the weeping, laugh with the laughing, and grieve with the grieving. These movements of the soul are known from the movements of the body. ([5], p. 80)

Contemporary neuroscience allows us to capture this thought with greater precision. When we see someone weeping we simulate the weeping behaviour and then feel the emotion that we would feel were we to be behaving that way ourselves. Our co-embodiment with a depicted figure thus plays a central role in how we respond emotionally to that figure.

Another important dimension of our engagement with paintings is our simulation of the implied actions of the painter themselves. Besides simulating the movements of any figures in a painting, we simulate the movements made by the artist to create the painting. Consider Jackson Pollock's painting 'Number 31' at the Museum of Modern Art. The painting is a huge cream canvas splattered with thick layers of black, white and grey. It is alive with movement: great grey splashes, intricate dribbles of white and black. Pollock worked by laying his canvas on the floor and dripping the paint onto it from above using dried out brushes, sticks and – on occasion – turkey basters. As we look at the painting, we simulate the movements Pollock made. Taylor, Witt and Grimaldi wryly note '…if one is to simulate Jackson Pollock's actions it must be done indirectly, as he has been quite dead since 1956, and is therefore inanimate. Fortunately, his paintings persist as a historical record of his actions.' ([6], p. 27) The paint on the canvas is inert but, through embodied simulation, we can experience the long-past movements that created the artwork. And this simulation can play a central role in our aesthetic experience of the painting. Our neural simulation of the actions that produced the Pollock painting might tell us that there was anger in those movements, leading us to experience the painting as expressive of anger. Once more, our co-embodiment with the artist is integral to our experience of their work.

As with the ballet case above, our embodied simulation of the implied actions of the painter contribute to our enjoyment of the painting. In a study by Leder, Bär and Topolinski [7] participants were presented with pointillist paintings and asked how much they like them. This was done under three conditions. In the first condition, participants looked at the paintings whilst dabbing a pencil up and down on a surface – an act congruent with the implied actions of the artist. In the second condition, participants looked at the paintings whilst making a brushing motion with a pencil – an act incongruent with the implied actions of the artist. In a control condition, participants looked at the painting without making a motion with a pencil. The congruent group appraised the painting more positively than the control group, and the incongruent group appraised the painting less positively than the control group. This can again be interpreted in terms of our embodied simulation of the artist's movements playing a central role in our enjoyment of the artwork. In the incongruent case, the motor process behind the brushing motion impeded the embodied simulation of the artist's actions which then had a negative effect on the enjoyment of the painting. In the congruent case, the motor process behind the dabbing-motion boosted the embodied simulation and had a positive effect on the enjoyment of the painting. Our co-embodiment with the artist is once more integral to our appreciation of their work.

The simulation of the implied actions of the artist is of particular interest because it captures something about our embodied responses to both figurative and non-figurative art. As we move on to a discussion of AI-generated art, this aspect of our embodied response to paintings will be given particular attention.

3. AI and Embodiment

AI programs for generating art are becoming increasingly sophisticated. Programs such as Dall-E 2 allow users to type in detailed text-prompts that rapidly generate appropriate results [8]. They can also take an image as an input and generate variations of that image within specified parameters, or alter an image in response to text prompts. Many of the images produced are of astonishing quality and align impressively with the user's commands. But unlike a flesh-and-blood artist, such programmes are not embodied. That is, they do not have a physical body through which they act on their environment. It is a piece of software that generates strings of code. In cases like these, the AI that generates a work of art is disembodied.

Some conceptual clarifications are in order here. Programs like Dall-E 2 are disembodied in the specific sense specified. Embodiment is, as we have a seen, a polysemous term so there might well be senses of the term on which the program is embodied. For instance, it is embodied in the sense that it is physically realised in some computer hardware rather than somehow existing non-physically. Furthermore, when the program produces a new image, this too is physically encoded in some hardware. But this is not the same as the program having a body through which it acts in the world to create the artwork. The image could, of course, be physically printed out. But this still would not be a case of the program using a physical body to create the artwork. What it creates is a string of code, and the image encoded can in turn be physically realised on paper.

To say that the program that generates a work of art is disembodied is not to say that the art was produced without the role of any embodied agent. Embodied agents wrote the program. Embodied agents created many of the images on which the program was trained. And embodied agents prompt the program to generate particular images. Nevertheless, each specific image is generated by an AI that is disembodied and not by an agent using their body to form that image. Here I do not commit to who (if anyone) the artist is in this process. Is it the programmer, the user, the AI itself or some combination of the three? Even if the embodied user qualifies as an artist, it remains the case that the specific image is generated by the program and not be the user deploying their body to form the image. To say that the program is disembodied is not to say that all AI-generated art is disembodied. Robot artists such as Aidan Meller's Ai-Da and Denial Berio's Graffitizer use robotic limbs to put marks on a surface. Whether a program that controls robotic arms really qualifies as embodied is a difficult question, but it is not one that need concern us here. My concern is with AI that generates art without the use of a robotic body.

There are thus AI systems that generate art and that are (in the relevant sense) disembodied. When we view images generated by such systems, we are not co-embodied with generator of the art. We are embodied but the thing that generated the image is not. Given that our co-embodiment with the artist plays such a central role in our engagement with human-made art, what happens when such co-embodiment is absent? There are two broad possibilities:

> **The Disembodied View:** Because the AI is disembodied, we do not have embodied responses in which we simulate movements used to create the artwork.
>
> **The Misembodied View:** Despite the AI being disembodied, we respond to AI art *as if* it were produced by an embodied agent.

On the Disembodied View, an AI can do many impressive things but it cannot do what an embodied artist can do. When a flesh-and-blood artist paints a picture, their embodiment plays an ineliminable role in the creation of the artwork. The viewer then responds the artwork as the product of an embodied process. When an AI generates an image, no embodiment is manifest in the artwork produced. Accordingly, the viewer does not respond to the artwork as the product of an embodied process. This is not to say that viewers will have no embodied response to the AI-generated image. If the image contains a figure of a ballet dancer, for example, the viewer might simulate the movements of the dancer. However, there is no simulation of any implied movements of an artist. The co-embodiment of viewer and artist is absent, and any viewer responses premised on that co-embodiment will be absent too. When we see a streak of colour on a work of AI art, for instance, there is no implied motion for us

to simulate – no bodily action to be mirrored by our own motor system. This then has a knock-on effect on various other dimensions of our response: our interpretation of what we see; our affective responses; our appraisals. On this view, the fact that AI is disembodied places a boundary on the artistic capacities of even the most sophisticated AI.

I argue that this Diesmbodied View is implausible. If AI systems create works of art that are fundamentally different in kind to those created by human artists, it should be easy to tell whether the originator of an artwork is human or artificial. If we look at an artwork and have the relevant embodied responses, we would be able to immediately tell that the artwork must have been created by someone embodied like us. Conversely, if we look at an artwork and the relevant embodied responses are absent, we could infer that the artwork probably was not made by an embodied artist. Put another way, the marks of embodiment in an artwork could act as a shibboleth that reliably indicates whether the artwork was created by an embodied artist or a disembodied AI. In reality though, we often cannot tell whether a work of AI art was created by a human or an AI. One of the things that is so impressive about recent art-generating software is that many of the images they generate can pass as human-made art. When giving lectures on the topic of creativity, I have presented my classes with a series of images and asked them to identify which ones were human-made and which ones were generated by AI. In many cases students identified AI-generated images as human-made. Although these informal tests were not done under experimental conditions, the results fit with a clear trend in online discussions of the quality of AI-generated art.

So why can't we always tell when an image is created by a disembodied AI? Programmes like Dall-E 2 are trained on a vast data-set of images. When you type in a prompt such as 'painting of an afternoon tea in the Impressionist style', the program is able to generate an appropriate image because it has learned the features characteristic of the Impressionist style. And the images on which it is trained are images of paintings by flesh-and-blood artists. Since those images show the marks of the embodiment, the images created by a suitably trained program will also show marks of embodiment. Consider Figure 1 created using Dall-E 2:

Figure 1: Image generated by Dall-E 2 from the prompt 'A painting of people on the banks of the river Cam in the style of Seurat'.

Although the program does not actually apply paint to canvas, the image looks as if it were made through a series of physical movements with a brush. Put another way it *looks* like it was made by an embodied artist. And the reason for this is simple: the program has learned how pointillist paintings look, and those paintings look like they were made through a series of physical applications of a brush. We can have an embodied response to those brushstrokes, simulating the actions that an embodied agent would have made to make those marks. This might then feed into the affective response that we have: our body tells us that the paint was applied with care and we feel an emotion congruent with that. It could also feed into our appraisal of the artwork: we might evaluate the artwork positively precisely

because of the way it moves us. Furthermore, knowing that the artwork was created by an AI artist wouldn't change the effect. Even if we believe it to be created by a disembodied programme, it still looks like the creation of an embodied artist and this is enough to drive our automatic simulation of the implied movements of the artist.

To respond to an image as if it were created by an embodied artist is to misembody the generator of the artwork. Your experience of the artwork is then premised on an illusory co-embodiment between you and the generator of the artwork. I call this illusion of co-embodiment with the artist 'faux-embodiment'. It is as if the image were created by something that is embodied much like you are, but in reality the AI is not embodied at all. The foregoing suggests that such misembodiment of AI-generated art is ubiquitous and straightforwardly explicable. We can now turn to the question of whether it matters.

4. The Value Question

When AI-generated artworks lead viewers to misembody the work does that constitute a shortcoming of the artwork? Do aesthetic experiences built on faux-embodiment have less value than aesthetic experiences built on genuine co-embodiment? There are again two broad possibilities:

> **The Luddite View:** Misembodiment is a bad thing that renders the artwork, and the aesthetic experiences produced by that artwork, less valuable.
>
> **The Innocent View:** Misembodiment is not a bad thing and does not render the artwork, or the aesthetic experiences produced by that artwork, less valuable.

The Luddite View can be motivated by the thought that there's something deceptive about artworks that seem to be produced by an embodied artist when they are not. This might tie into broader worries about AI being used to mislead and to obscure reality. Consider a case where you are using an online chat function and take yourself to be talking to a human when actually you are talking to a chatbot. The chatbot could be regarded as a kind of fake conversationalist and the exchange could be regarded as a troubling simulacrum of genuine human interaction. You are victim to an illusion and illusions are a bad thing. This is not to say that there is anything wrong with what the chatbot *says*. Rather, the fact that the conversation is artificial is a negative feature regardless of the quality of the words used. By the same token, when we misembody an AI-generated artwork we are victim to an illusion. Even if our aesthetic experience is positive, it has occurred in error. Our simulation of the implied movements of the artist were meant to be veridical, but since the image was created by a disembodied system those responses will inevitably be erroneous.[2]

Other cases of artworks with deceptive appearances might be used to lend further support to this Luddite verdict. Take the notorious Vermeer forgeries by Dutch artist Hans Van Meegeren. Many experts were taken in by the forgeries and prized them as great works of art. But once it emerged that they were cleverly forged, the value of the artworks was utterly undermined. The enjoyment that people had of the artworks was built on a deception, and because their experiences were illusory they were less valuable. The experience of a genuine Vermeer would be more valuable because this element of deception is absent. AI-generated artworks are analogous to forgeries in that they have a deceptive appearance. And like the forgery, this element of deception compromises their value.

Some might advocate this kind of Luddite View wholeheartedly while others may be more conflicted and have Luddite intuitions in amongst intuitions that are more positive about AI-generated art. Either way, I think the Luddite View is a mistake. While I agree that misembodiment involves a kind of illusion, I reject the idea that such illusions are problematic. To see why, we should reflect on other cases in which artworks provoke embodied responses that are illusory i.e. artworks that encourage the simulation of implied movements that the artist did not really make.

[2] In talking of embodied responses being veridical I do not mean to imply that those responses are representational states with straightforward accuracy conditions. A less committal way to put it is that they have *aptness conditions* and that embodied simulations are apt just in case they successfully mirror the bodily movements of the relevant agent.

As the study by Leder *et al* suggests, our embodied response to the pointillist painting tells us that the painting was produced by dotting the brush onto the canvas [7]. However, this embodied response is not always veridical. Pointillist paintings can include different densities of dot. Once the dots become sufficiently dense, the eye cannot discern one dot from another. This means that the impression of a dense region of dots can be achieved by doing a messy squiggle with the brush rather than by dotting the brush. In an online guide to pointillist drawing, Megan Eckman writes 'This is where I normally cheat and begin to draw squiggles instead of dots because after a certain point, the human eye can no longer distinguish individual dots' [9]. So if the viewer has an embodied response in which they simulate a dabbing motion here, that response would be non-veridical i.e. the simulation would not match the movements actually made to produce the image.

Yet here I don't think that we would generally regard it as a problem that our embodied responses are non-veridical. It does not count against the value of the artwork that it encourages this non-veridical embodied response. And it does not count against the value of the aesthetic experience of the artwork that it is premised on a non-veridical embodiment response. Why, then, should we worry about non-veridical embodied responses in the case of the AI-generated images?

A possible Luddite response here is that the error is too small to constitute a problem. After all, it is only the dense squiggles that we respond to incorrectly so our embodied response to *most* of the dots in the pointillist picture would probably be accurate. The error in the case of AI-generated images, however, is comprehensive and cannot be brushed away as unproblematic. Against this response I would argue that increasing the scale of the error in an embodied response makes no difference. To see why, let's consider some cases in which an embodied response invited by an artwork is more inaccurate. A pointillist style can be achieved by graffiti artists using a spray can. A device can be attached to the top of the can that makes it produce short sharp bursts of paint that leave a dot on the surface. This can then be used to create a larger image in the pointillist style. Presumably looking at such an image under the right conditions would produce the same kind of embodied response as a pointillist painting does i.e. one in which a dabbing motion is simulated. But such a response would be inaccurate. And, unlike in the case of the squiggle-cheat described above, the error would be comprehensive. None of the dots were made by dabbing paint onto the surface. Yet there are no worries about these non-veridical responses counting against the value of the artwork or the value of the aesthetic experience it produces.

The Luddite might respond that the error in the spray can example is still not comprehensive. Even though the marks were not made in the way our body says they were made, they were at least made by through an embodied process. So our embodied responses are at least right about *something*. But when we have embodied responses to the AI-generated image, our embodied responses are *wholly* wrong. However, cases can be found in which embodiment is removed from the process entirely. Applications such as Photoshop allow users to upload a photograph then convert the image so that it appears to have been painted. In Figure 1 the user has uploaded a photograph of some boats then used Photoshop to give it the appearance of a pointillist painting [10]. Again, our embodied response to the resulting image would tell us that the image was made through a series of dabbing motions. But in reality the image was generated through a wholly disembodied process. Here I would again suggest that the inaccuracy of our embodied responses would not count against the value of the artwork.

Figure 2: Adobe stock image altered using Photosphop by user [10].

I've focused here on images in the pointillist style. But the conclusions reached apply much more broadly. Painters use many techniques that imply particular movements without those movements actually being made. Think again of Jackson Pollock and his turkey baster! Whatever embodied effect those paint dribbles achieve, it won't involve simulating a squeeze of the baster. Similarly, a sculpture might feel like it was bent and twisted by the artist when actually it was carefully chiseled to achieve that way. And an amplifier might make it seem as if the musician is thrashing their guitar with super-human force when in reality they are gently plucking it with a plectrum. All of these cases encourage non-veridical embodied simulations, and may achieve their effect precisely because of that simulation. But in none of those cases is the non-veridicality a problem. To object to the non-veridical nature of our embodied responses to AI-generated art would be to apply a kind of double-standard: to ask something of AI-generated art that we don't ask of art created by flesh-and-blood artists.

The Luddite might argue that their standards are consistent by putting more weight on the analogy with forgery. In the case of the Van Meegeren forgeries, for example, we really do value the artwork less when we realise that it is deceptive. However, such an analogy would not be apt. Our evaluation of a painting by an Old Master involves a complex interplay of aesthetic value and historical value. The Van Meegeren forgeries appeared to have great historical value – they seemed to be the long lost works of a hugely influential artist – but this value was deceptive. The *aesthetic* value of the paintings, however, is not the same as its historical value. It might even be argued that whatever aesthetic value the painting has is intrinsic to the painting and is not affected by who painted it. So what's problematic about the forgeries is that they were used to mislead people about their historical significance. If AI was used to create images with a deliberate intent to mislead people about the historical status of that image, then we would indeed have a problem. But the kind of misembodiment discussed does not involve any historical deception. And if we move from cases in which an artwork misleads us about its historical status to cases in which an artwork misleads us about the embodied process used to create it, it is clear that the problem disappears.

5. Acknowledgements

Many thanks to Laura Herman and Caterina Moruzzi for organising the ICCC'22 Workshop 'The Role of Embodiment in the Perception of Human & Artificial Creativity' and to the participants in the workshop for their invaluable feedback.

6. References

[1] C. Guckelsberger, A. Kantosalo, S. Negrete-Yankelevich & T. Takala, Embodiment and Computational Creativity. (2021) arXiv:2107.00949

[2] F. J. Varela, E. Thompson and E. Rosch The Embodied Mind: Cognitive Science and Human Experience, 2nd ed. Cambridge MA, MIT Press, 2017

[3] V. Halász, R. Cunnington. Unconscious Effects of Action on Perception. Brain Sciences, 2(2), (2012), 130-146.

[4] L. P. Kirsch, K. Dawson and E. S. Cross. Dance experience sculpts aesthetic perception and related brain circuits. Annals of the New York Academy of Sciences, 1337.1, (2015), 130-139.

[5] L. B. Alberti, (1972) On Painting and Sculpture: the Latin Texts of De Pictura and De Statua, C. Grayson, C. ed. and transl.), New York: Phaidon Press

[6] J. E. T. Taylor, J. K. Witt, and P. J. Grimaldi. Uncovering the connection between artist and audience: Viewing painted. Cognition 125 (2012) 26–36.

[7] H. Leder, S. Bär, and S. Topolinski. Covert painting simulations influence aesthetic appreciation of artworks. Psychological Science 23.12 (2012): 1479-1481.

[8] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, & M. Chen. Hierarchical text-conditional image generation with clip latents. (2022) arXiv preprint arXiv:2204.06125.

[9] M. Eckman. How to Create a Hand-Drawn Pointillism, 2014. URL:https://design.tutsplus.com/articles/how-to-create-a-hand-drawn-pointillism-portrait--vector-23626
[10] S. Patterson. Photoshop Pointillism Painting Effect, 2022. URL: https://www.photoshopessentials.com/photo-effects/photoshop-pointillism-effect/