

The Network That Painted This: Latour, Verbeek, and Generative AI

Blagoj Naumov

(16182162)

Abstract:

This thesis examines how generative artificial intelligence mediates human relations to art, asking how GenAI reshapes the conditions under which art is produced, encountered, and valued. Bringing together postphenomenology and Actor-Network Theory, it argues that generative systems should be understood as mediating technologies embedded in networks of human and non-human actants that co-constitute artistic outcomes. The thesis first establishes a conceptual framework through the work of Peter-Paul Verbeek and Bruno Latour, showing that technologies shape perception, action, interpretation, and moral orientation, while also emerging through distributed relations among datasets, code, hardware, institutions, and users. It then clarifies what counts as GenAI and GenAI art for the purposes of the study, distinguishing contemporary probabilistic, data-trained systems from earlier rule-based generative art and excluding purely operational AI outputs that do not produce meaningful expressive artefacts. A further theoretical foundation is developed through a postphenomenological philosophy of images, arguing that images already mediate selfhood, embodiment, comparison, and social relation, and that GenAI intensifies these dynamics by increasing the scale, flexibility, and opacity of visual mediation. These arguments are applied to the cases of Obvious's Portrait of Edmond de Belamy and the work of Mario Klingemann, through which the thesis explores authorship, black-boxing, deskilling, aesthetic value, and the changing status of the artist. The central conclusion is that GenAI art exposes a broader transformation in contemporary culture, one in which artistic creation and interpretation are increasingly distributed across relational networks rather than grounded in a single sovereign human subject.

Introduction	4
Chapter 1: Framework	7
<i>1.1 Postphenomenology</i>	7
<i>1.2 Actor-Network Theory</i>	11
Chapter 2: Situating the thesis and Theoretical Foundations	14
<i>2.1 What counts as GenAI and GenAI Art</i>	14
<i>2.2 A postphenomenological philosophy for images</i>	17
Chapter 3: Visual Cases	22
<i>3.1 Obvious's Portrait of Edmond de Belamy</i>	22
<i>3.2 Mario Klingemann</i>	26
<i>3.3 A further discussion on the problems of GenAI</i>	30
Conclusion	33
<i>Limitations and Future Research</i>	33
<i>Conclusion</i>	34
Bibliography	36

Introduction

In October 2018, a portrait rendered completely using generative AI, was sold at an auction for \$432,500. The work itself was titled 'Portrait of Edmon de Belamy' and was signed off with '[Ex~pdata(x)log D(x) + Ez~pz(z)log(1 - D(G(z)))]'. This signature represents a fragment of the loss function, which is a mathematical expression used during model training to measure how far the system's output is from the desired result and to guide the network in adjusting its parameters. In this case, the fragment symbolically encoded the architecture of the generative adversarial network (GAN) that produced the image. The whole project was undertaken by a French collective called Obvious, who had trained the model on a dataset of historical portraits and allowed the system to generate outputs that no individual human had previously drawn. Fast-forward seven years and Regan Morris of the BBC reports that 'AI was enemy No. 1 during Hollywood strikes. Now it's in Oscar-winning films' (Morris, 2025). All of this is to say that generative AI, whether its GANs for images or LLMs (Large Language Models) for scripts, has slowly penetrated the art world and seems to show no signs of slowing down.

Generative artificial intelligence, a family of probabilistic models that includes generative adversarial networks, large language models, and similar systems, has in a remarkably short time become a constitutive element of cultural production. Some estimates have up to 71% of social media images as AI generated (Elad, 2025). Text-to-image AI tools such as General Adversarial Networks, or GANs, or DALL-E, Midjourney and others, allow users to generate imagery and illustrations from language based prompts. Large Language Models are able to co-write novels, generate poetry, produce screenplays, and engage in creative collaboration with their prompters. In my thesis, this is exemplified by Portrait of Edmond de Belamy, and by Mario Klingemann's AI-generated portraits, an artist who uses neural networks to generate artwork using AI. Potential exists for these practices to get more and more accessible and dominant in the cultural sphere and existing debates tend to circle around the question of authorship (whose work is it?) and/or value (is it any good? Is it ethical and desirable?), however a question that seems to be underexplored is what is the actual relationship between the human and the generative AI system and more specifically how are these relationships influenced by the technological tool itself and in what ways does this technology shape our way that we interact with it in the future and with art in general? I will refer to this research question as how does generative AI (GenAI) technology mediate human relations to art and to the technologies itself.

Thus, the guiding question of this thesis becomes: How does generative AI mediate human to art relations and what are its repercussions? In other words, how do GenAI tools affect the way we see GenAI art, as viewers, and how do they affect the way we do art, as artists. This is what is meant by how human to art relations are mediated by GenAI. In order to answer this, I will take the work of Bruno Latour and Peter-Paul Verbeek to establish a framework to analyse this mediation, or in other words to analyse how the technologies influence and affect the way humans understand the world, or in this specific case, understand our relationship to the world of art and the use of said technologies. I will then look at specific examples from the art world and analyse how they specifically mediate human to art relations within the sphere of postphenomenology and Actor-Network Theory (ANT) and what can be learned from such an analysis. The thesis serves as an exploration of what a postphenomenology and ANT analysis can say for GenAI and art as a whole, and where these theories lead us if we use them analyse the contemporary debates surrounding GenAI. This thesis limits itself to images because visual GenAI provide a case for how generative systems mediate perception within human-art relations and human-technology relations. An account of what counts as GenAI and thus what counts as GenAI art will be made later, however while things like poetry or books or screenplays can be generated using GenAI or large language models (LLMs), this thesis will specifically focus on examples from image generation in particular. This is partly due to the fact that an analysis of large language models cases of GenAI might require a slightly different framework as before delving into the specific case studies I also provide an overview of how images as a whole can mediate human-technology relations, and it is this analysis combined with the philosophical framework of Verbeek and Latour that will be used to explore the case studies. While the framework can be applied to

LLMs and literary cases, that would require a slightly different angle that for the purposes of depth is not included here.

Debates about generative AI and art have expanded rapidly in recent years, and much of the existing discussion has concentrated on questions of authorship, originality, copyright, labour displacement, and artistic value, often treating generative systems either as technical tools or as threats to human creativity, and often from a legal perspective (McCarthy and Yoon, 2023; Caldwell, 2023; He, 2026). While these debates are important and some of them will be addressed later, they do not focus on the more foundational question of how these systems mediate the relation between humans, artworks, and the creative practice itself. A smaller but growing body of work has begun to address these questions such as studies of computational authorship and creativity in AI art (Manovich and Arielli, 2024; Di Dio et al., 2023). While media-theoretical work on AI art demonstrates how prompting and curation reorganise creative practice and authorship (Manovich and Arielli, 2024), there is still relatively little sustained analysis of how such systems mediate artistic experience. This thesis is therefore relevant because it moves beyond the binary of human versus machine authorship and instead asks how generative AI reorganises the conditions under which art is made, encountered, and valued. In doing so, it treats GenAI as a mediating technology. I will first start by outlining the framework, and then move on the definitions and criteria I will use for remainder of the thesis. This will be followed by two specific case studies analysed from the perspectives of postphenomenology and ANT and an exploration and discussion of some of the consequences that have emerged as a result of these technologies and said analysis, and an exploration of where this theoretical framework applied onto these questions lead. Before moving onto the framework however it is important to establish why I have specifically chosen postphenomenology and ANT as the framework itself, or in other words: Why Latour and Verbeek?

Both the work of Verbeek and Latour is necessary for exploring the thesis question at hand because both of them approach a way to explore mediated relations in the human experience between humans and non-human things. The two theories go together hand in hand quite well, and in fact, in his work, Verbeek makes constant references and almost builds upon Latour, making his work intertwined with his. Verbeek's postphenomenology is necessary because it provides an account for describing how GenAI is encountered in day to day life. It also helps explain why these relations mediated by technologies are never ethically neutral, since technologies shape what users notice, value, and do. Latour's actor-network theory is necessary because GenAI art is a marriage between interface plus a user, a network of actants, models, datasets, code, hardware, and institutions all contribute to what the artwork is and how it comes to appear. In other words, for Latour a production of a GenAI image is the work of the prompter, the specific code that is built into the AI system, the dataset which the system is trained on (because these models are often trained on a variety of data), the actual computer hardware used to type and generate the prompt and any physical servers used to support the system, and the institutions which actually own and allow access to these tools. What ties the two together is a shared rejection of the idea that there is a distinct divide between human and technology. The work of both treats relations as primary. Verbeek is strongest on phenomenological description of experience and moral mediation, while Latour is strongest on tracing the distributed production of outcomes. That combination is especially suitable for this thesis because the research question asks both how GenAI art is experienced and how that experience is made possible by networks of mediation.

Verbeek asks how technologies shape human intentionality and moral orientation, while Latour asks how associations among human and non-human actants make outcomes possible in the first place. Used together, they let the thesis answer both halves of the research question: postphenomenology explains how GenAI art is experienced in the human-art relation, and ANT explains how that relation is produced by distributed networks of technical, social, and institutional actors, which is why both are necessary rather than interchangeable. Finally, of important note is to mention that Don Ihde writes 'Verbeek...continued his interest in comparing ANT and Postphenomenology' (Rosenberger and Verbeek, 2015) and in the preface titled *Positioning Postphenomenology in Postphenomenological Investigations: Essays on Human-Technology Relations*, he writes extensively about the relations of ANT to Postphenomenology and their emergence from the same school of thought in sociology of science and science and technology studies (Rosenberger and Verbeek, 2015). In fact, in the same essay, a brief summary

of the major comparative features of the two theories is provided which notes that “Both ANT and Postphenomenology employ inter-relational ontologies...both ANT and Postphenomenology are materially sensitive...Both methodologies have abandoned early modern subject-object dichotomies and substituted multiple transition and structure notions” (Rosenberger and Verbeek, 2015). As demonstrated, both the work of Verbeek and Latour is complimentary to each other and thus I shall move on to outline the work of both philosophers and establish a framework which will then be applied to analyse three different visual, image based case studies and this will be used to explore the question of how exactly GenAI tools shape and affect the relationship humans have with art and with artistry.

Finally, to further answer the research question, the thesis proceeds in three stages. First, it establishes the theoretical framework by outlining postphenomenology and Actor-Network Theory as the two complementary approaches through which generative AI mediation will be analysed. Second, it clarifies the definitional scope of the thesis and specifies what counts and does not count as GenAI and GenAI art, and goes over the demarcation criteria for the case studies. This is followed by a theoretical foundation on the framework being applied to images. Third, it applies this framework to the visual case studies of Portrait of Edmond de Belamy and Mario Klingemann, analysing each as a mediated image that is produced through relations among human and non-human actants, and through which questions such as those of authorship, deskilling, and artistic value become visible. In this way, the case studies serve as concrete sites through which the thesis traces how GenAI reorganises the conditions under which art is made, encountered, and understood.

Chapter 1: Framework

First, we can start by establishing the conceptual and methodological framework through which the remainder of the thesis will analyse generative AI art. The first question, addressed in this chapter, is how postphenomenology and Actor-Network Theory can be brought together to explain the mediating role of generative AI in human-art relations. Rather than treating GenAI as a neutral tool or as an autonomous creative agent, the chapter approaches it as a relational technology whose meaning emerges through its use, its design, and the networks in which it is embedded. The aim here is to set out the theoretical vocabulary needed to describe how GenAI shapes perception, interpretation, and artistic practice. The chapter begins with postphenomenology, in order to explain how technologies structure human-world relations through mediation. This section outlines the basic categories through which technologies can be understood, and it introduces the idea that technologies should not be seen as neutral tools, but should be seen as active participants in shaping how phenomena appear to human users. It then moves to Peter-Paul Verbeek's development of this framework, especially his account of technological mediation and the ethical dimensions of human-technology relations. This allows the chapter to show why generative AI should be analysed as something that shapes how art is seen, made, and valued. The chapter then turns to Actor-Network Theory and the work of Bruno Latour in order to examine how generative images are produced through networks of human and non-human actants. This section explains the importance of tracing associations rather than presuming a single source of agency, and it introduces concepts such as black-boxing and inscription to show how GenAI systems become stabilised and operationalised in practice. Together, these frameworks provide the basis for the case studies that follow, since they allow the thesis to analyse not only what GenAI art is, but also how it comes to appear as art in the first place.

1.1 Postphenomenology

The first question postphenomenology asks is, "How do technologies shape our choices, our actions, and our experience of the world?" (Rosenberger and Verbeek, 2015). It can therefore be understood as the study of the "relations that develop between users and technologies", and of the ways in which those relations help to constitute human experience (Rosenberger and Verbeek, 2015). Hence, the choice of postphenomenology as a framework for this thesis is because it allows us to ask what generative AI is and how it mediates our relation to images, art, creativity, authorship, and the world more broadly. In other words, how it shapes, forms, and changes our relation, or our outlook towards the aforementioned concepts. It also gives us a framework for describing the fact that technologies are more than tools and can be seen as co-constitutive partners in our experience of the world.

Before moving to the work of Peter-Paul Verbeek, however, it is useful to give a little background as to the emergence of postphenomenology and to the work of Don Ihde, since postphenomenology begins with Ihde's attempt to provide an empirical and phenomenological analysis of specific technologies and of the ways in which they shape our experience and perception of the world through our interaction with them (Ihde, 1990). Ihde started from looking at technologies and asked what kinds of relations they produce between human beings and the world. The core of this philosophy is made up of four different human-technology-world relations, each describing a distinct structure of what we can call mediated intentionality. Mediated intentionality is when our way of perceiving or acting in the world is shaped through a technology, so that we experience things through it rather than simply alongside it. These are not fixed categories into which all technologies must fit, but rather starting points for describing how technologies participate in experience. In no particular order, they include embodiment relations, hermeneutic relations, alterity relations, and background relations (Rosenberger and Verbeek, 2015). It is important to stress from the beginning that these relations are not mutually exclusive. A single technology can occupy different relations depending on context, use, and situation. This is referred to as multistability (Rosenberger and Verbeek, 2015), meaning that technologies are not

stable in meaning or function in any absolute sense, but can become configured differently in different practices. A GenAI interface, for example, may function as an alterity relation when the model is experienced as an interlocutor, as a hermeneutic relation when its output has to be interpreted, and as a background relation when it is quietly embedded in an otherwise invisible design workflow. The point is not that every technology must be squeezed into a rigid typology, but rather that the typology helps us see how technologies shape and restructure human-world relations. To expand on this, we can delve into the four different relations.

An embodiment relation is when a technology is incorporated into the body so that perception can extend through it. The example often given is glasses, whereby the wearer perceives the world anew through the lens and, in this way, the glasses become an extension of embodied experience. The relation can be represented as: (Human-Technology) → World. In embodiment relations, the technology becomes quasi-transparent (Rosenberger and Verbeek, 2015). What this means is that the technology does not disappear completely, but becomes so integrated into action that the user's attention is directed through it rather than toward it. The technology is felt less as an object than as a medium of access. A broader argument as to how exactly GenAI can function as an embodied relation will be made and applied later in the thesis, but to set the stage for such an argument I would first also like to expand on this using Maurice Merleau-Ponty's notion of the flesh. Merleau-Ponty's notion of the flesh describes this shared relationship that connects body and world, and tries to show that perception or experience happens through an underlying intertwining of the two rather than between the two separate things. It means that the perceiver and the perceived are not fully distinct, but belong to the same relational field of experience (O'Brian, 2017). The flesh is not meant literally as meat or tissue; rather, it is Merleau-Ponty's way of naming the shared, lived fabric in which perceiver and perceived are already bound together. This helps explain why technologies can become so closely integrated into experience. Ihde's embodiment relations can then be understood as a more concrete account of this same intimacy, a camera, or a cane, or glasses can become part of how the body reaches toward and engages the world. In other words "His postphenomenological take on embodiment relationships fundamentally proposes that a human's experiences are mediated *through* a technological artefact, which can be as simple as a piece of chalk." (O'Brian, 2017). Merleau-Ponty gives the broader philosophical idea of bodily intertwining, while Ihde shows how that intertwining is mediated through specific technologies (O'Brian, 2017). If we transpose this onto GenAI, a GenAI tool (such as DALL-E or Midjourney) can become a sort of quasi-embodied extension of the human body. An argument has been made before through postphenomenological readings of film and cinema that "film has, and is, a body (O'Brian, 2017)" most notably by Jennifer Barker. This is the same argument Daniel O'Brian draws upon in his postphenomenological analysis of cinema, which I won't delve into here. In the same way, we can use this sort of narrative to argue that a GenAI tool such as GANs can act as a body, which produces a separate body of work, and through the fact that in order to produce this separate body of work it needs a human prompter or collaborator, in this way, it can act as a metaphorical embodied extension of the human artists' vision and act as the human's extension of their own body in the production of the work. This argument put forth by Jennifer Barker is an extension of Merleau-Ponty's work on the flesh, as described by O'Brian:

"Jennifer Barker's work by route of Merleau-Ponty's flesh explores the liminal space in which film and viewer can emerge as co-constituted, individualised but related, embodied entities" (O'Brian, 2017).

In this same way, a human artist in collaboration with a GAN can be seen as a co-constituted, individualised but related, embodied entities, however as mentioned, this will be further explored later. This is also not intended as a claim that AI possesses a body in the same sense as a human subject, nor as a rejection of the critique that AI systems are often disembodied. Rather, it reflects a broader argument in the literature that generative systems can function as extensions of human artistic practice, in which embodied agency is distributed across the artist, the interface, and the model itself (Zeilinger, 2021; Zylinska, 2020). In this sense, the GAN does not become a body in itself, but may be understood as a metaphysical extension of human intentionality and creative action, shaping the production of the artwork without replacing the embodied artist (O'Brian, 2017; Zeilinger, 2021; Zylinska, 2020). This distinction is important for the present thesis because it allows GenAI to be analysed not as an autonomous body, but as a mediating technology through which artistic embodiment is reconfigured and extended.

A hermeneutic relation is when a technology produces a representation that requires interpretation. Here the technology becomes something like a text, which is to be read and interpreted. Other examples include a thermometer, an ultrasound image, or a weather map. All of these mediate the world by translating it into a representational artefact. The relation can be represented as follows: Human → (Technology-World). The world becomes encountered through the representational output, not directly (Rosenberger and Verbeek, 2015). This is helpful for understanding many image-based technologies. A health dashboard, for instance, offers a representational account of the body in the form of data visualisation. Likewise, a GenAI output can be an artefact that must be interpreted in relation to the prompt, the training set, the interface, and the assumptions built into the system.

An alterity relation is when a technology itself becomes the focus of quasi-intentional experience. The technology appears as quasi-other. An ATM, a chess-playing machine, and Large-Language Models could be examples. These are neither fully transparent as in embodiment, nor fully representational as in hermeneutic relations, but are engaged as a kind of interlocutor (Rosenberger and Verbeek, 2015). They call forth responses, expectations, and adjustments. In the context of GenAI, this is relevant because the model is often experienced as something that “answers,” “suggests,” “generates,” or “responds,” even though it is not a subject in the human sense. This quasi-otherness appears to meet the user halfway, but it also exceeds the user’s intentions, introducing unforeseen patterns, styles, or associations. Alterity, in this sense, names the phenomenological experience of encountering technology as if it had a limited agency of its own.

A background relation is a technology that conditions experience without being the object of any explicit intentional act. A thermostat that maintains room temperature is a common example. Such technologies are present in the environment in which they function, but they remain in the background of attention (Rosenberger and Verbeek, 2015). This relation matters because not all technological mediation is foregrounded. Some technologies shape what is possible while remaining largely unnoticed. In the case of GenAI, background relations could be things like a model that may be embedded in systems of image production. For instance, if someone is using a platform to make a flyer and the system automatically offers ‘smart’ image generation or background removal, the model is functioning as a background relation. The user may not consciously think about the AI model itself, yet it is shaping what kinds of images are available, how quickly they can be produced, and what visual choices feel natural or easy. It quietly conditions the production process without becoming the direct focus of interpretation.

It is important to note that these relations are not meant as exhaustive or fixed properties of artefacts. A particular technology can sustain multiple relational configurations depending on context (Rosenberger and Verbeek, 2015). This is a useful feature of postphenomenology, it encourages us to ask what relation a technology enacts here, in this situation, for this user, and for this purpose. The same GenAI tool can therefore be an interlocutor in one moment, a representational device in another, and a hidden infrastructure in a third. This helps explain why postphenomenology is so appropriate for analysing generative AI art: such works are often multi-stable in their mode of engagement, shifting between tool, partner, and output. Not only are they multi-stable but Ihde introduces this idea of amplification and reduction. Ihde’s ideas of amplification and reduction are central to postphenomenology because they show that technologies never just add something to experience; they always reshape it. By amplification, Ihde means that a technology makes certain features of the world more visible, more precise, or more accessible than they would be without it. By reduction, he means that the same technology also pushes other features into the background or makes them less available to us (Ihde, 1990; O’Brian, 2017). A dental probe, for example, amplifies small differences in the surface of a tooth or board that a hand would not reveal as clearly. At the same time, it reduces other forms of contact, because the object is no longer experienced in the same broad, tactile way. So amplification and reduction always happen together: a technology gives us one kind of access while limiting another. This shows that technologies actively shape what can be sensed, noticed, and interpreted (O’Brian, 2017).

Peter-Paul Verbeek extends Ihde’s framework. He does this partly by expanding the relation types and partly by developing the ethical and normative implications of technological mediation.

Verbeek first extends Ihde's postphenomenology by emphasising that human-technology relations can also take forms such as fusion and immersion. Fusion relations describe cases in which the boundary between body and technology becomes especially close or partially dissolved, as in implants or other deeply integrated devices, while immersion relations describe situations in which the user is surrounded by a technological environment that actively shapes experience from within. In Verbeek's account, technologies also help configure how we live, choose, and become. In particular, Verbeek is interested in how technologies shape not only perception and action, but also morality and subjectivity. One of his key contributions is the idea that technologies are not neutral with respect to values. Instead, they actively mediate how we perceive, judge, and inhabit situations. This is what he means by the term technological mediation. This is the ways in which technologies shape and reshape both human beings and the world through specific mediating effects (Rosenberger and Verbeek, 2015).

Verbeek's article *To-Do Is to Be: Foucault, Levinas, and Technologically Mediated Subjectivation* is especially useful here because it shows how technological mediation can be understood ethically (Bergen and Verbeek, 2021). In the article, Verbeek examines a gamified to-do list app in order to show how a seemingly mundane technology can confront the user with themselves, and how this confrontation can involve motivation, shame, discipline, and self-care. The article presents technological mediation as something that goes beyond just assistance or tool. The to-do list app helps the user organise tasks as a primary tool but it also goes a step further, and by helping the user organise tasks it also shapes how the user experiences responsibility. In this way, Verbeek demonstrates that technologies are active mediators of subjectivation, meaning that they help shape who we become through the practices they enable and constrain (Bergen and Verbeek, 2021). He exemplifies and explains this by looking at Foucault and Levinas.

Verbeek draws on Foucault's notion of the "care of the self" in order to explain how technologies may participate in practices of self-formation. Foucault's point is that the self is made through techniques, disciplines, and forms of reflection (Bergen and Verbeek, 2021). Verbeek uses this to show that technologies can become part of those techniques of self-formation (Bergen and Verbeek, 2021). In the case of the to-do app in his case study, what we have is a device that prompts the user to confront their habits, delays, and failures. It structures self-evaluation and produces a form of motivational pressure. This is important because it reveals that technologies can mediate subjectivity through subtle forms of normativity, for example in the case used above, by turning unfinished tasks into visible reminders, translating procrastination into shame, and making self-improvement appear as the proper response to technological feedback.

At the same time, Verbeek brings in Levinas to explain the ethical dimension of this confrontation. Levinas is important because he thinks subjectivity begins in responsibility toward the other. Verbeek adapts this idea to show that technologies can confront us with a kind of quasi-otherness that calls us into relation and responsibility (Bergen and Verbeek, 2021). In the to-do app, the user is confronted with unfinished tasks, reminders, and accumulating obligations. This may produce shame. This can also function as a form of ethical provocation, an interruption that invites self-care. Verbeek's use of Levinas allows him to argue that technologies can mediate ethical subjectivation by structuring the conditions under which responsibility is felt and enacted. This is relevant for generative AI art because it allows us to think about these technologies as more than simple tool use. A GenAI system can shape the field of possible actions, possibilities of judgment, and conditions of artistic self-understanding. The model, interface, prompt structure, and dataset help configure what the artist notices, what kinds of choices appear available, what forms of feedback are amplified, and what kinds of responsibility are obscured.

To extend this idea we can also look at the concept of affordance. An affordance is what a technology makes possible, invites, or constrains within a given situation. This idea comes from James J. Gibson's concept of affordance, which he used to describe the action possibilities offered by an environment relative to an organism's capabilities (Gibson, 1979). In the context of postphenomenology, affordances are important because they show that technologies mediate action without fully controlling it: a tool may invite certain forms of use, while discouraging others, and this shaping effect becomes part of how the technology is experienced and understood (Gibson, 1979; Ihde, 1990; Verbeek, 2005). In other words, affordances describe the practical possibilities that emerge when a person engages a technology in a particular way. In the context of technological mediation, affordances are relevant because they show that technologies do not

determine action absolutely, but they do shape the horizon of action. Any technology affords certain uses, and for this reason technological mediations are constrained by the affordances of a technology (Rosenberger and Verbeek, 2015). Taking this idea further, the affordances of certain technologies can also affect and shape our actions and behaviours in relation to that technology and the world (Rosenberger and Verbeek, 2015). This is to say that there is always a constitutional element to technological mediation (Rosenberger and Verbeek, 2015). What this means in practice is that technologies participate in forming the conditions under which action, perception, and interpretation become possible. A paintbrush affords certain gestures; a camera affords certain angles; a GenAI interface could afford prompting, refining, and selecting. Each of these shapes not only what can be done, but also how doing is experienced. In GenAI art, affordances matter because they show that the model and interface are not neutral channels through which pre-existing intentions pass. They actively structure the creative process. They make some aesthetic paths easier, some decisions more immediate, and some possibilities more visible than others. A prompt interface where one can write whatever prompt they want, for instance, affords a certain kind of compositional thinking: the user is encouraged to think in terms of description and instruction as opposed to a prompt interface which gives the user a choice of preselected prompts. The model then returns outputs that may open up further possibilities or close down others. This is not a free and unconstrained creative field, but neither is it a completely determined one. It is a mediated field of action, and that mediation has consequences for how the technology is experienced.

Verbeek's also introduces the idea of moral mediation. Moral mediation means that technologies help shape what is experienced as desirable, responsible, appropriate, or valuable (Verbeek, 2021). Verbeek gives the example of the speed bump. The speed bump shapes driver behaviour in normative ways by encouraging drivers to slow down. It says, "slow down here," without speaking. In the same way, generative AI systems can encode aesthetic norms, cultural biases, and assumptions about what counts as a good image or a successful output. An AI system trained on western data will produce western centric images consistently, with the implication that these are what is wanted. They can also normalise particular forms of visual style, again if trained on western data, they will normalise the western style as that is the style they will constantly produce. This will become more clear in the case studies that follow later in the thesis. In *Portrait of Edmond de Belamy*, for example, the system mediates how authorship is perceived, attributed, and valued. The work invites the viewer to confront the ambiguity of a machine-generated portrait whose aesthetic force is inseparable from its technological and institutional framing. The important point is that the technology participates in configuring the practical and ethical environment within which art is made and interpreted.

For this reason, postphenomenology is central to this thesis. It provides the conceptual resources to show that GenAI is a collaborator of perception, action, interpretation, and value. Ihde's early work helps us describe the different ways in which technologies enter into human-world relations, while Verbeek's work deepens this by showing how those relations are also ethically and normatively formative. Together, they allow us to see generative AI as a site where human agency, technological affordance, and ethical subjectivation are in constant conversation with each other. The framework gives us a way to analyse what GenAI art looks like, what it does to the conditions of artistic practice and to the ways in which we relate to images, to ourselves, and to the world.

1.2 Actor-Network Theory

We can synthesise the work of Verbeek, which builds upon the work of Ihde, by looking at Bruno Latour's Actor-Network Theory. Actor-Network Theory (ANT) emerges under the rejection of what Bruno Latour refers to as the 'social explanation' of phenomena, which is essentially the idea that events can be explained by pointing to a pre-given social force such as culture, or society, as a cause (Latour, 2005). Latour argues that instead of this, we ought to follow the 'actors' themselves (in the event) and analyse how the outcomes have emerged through the combination of associated actors (Latour, 2005). To put it simply, an actor or 'actant' is anything that modifies (or in a way affects), another actant, regardless of whether it is human or nonhuman. For example, anything like a machine, a text even, a building, an organism, legal codes perhaps, and humans themselves can act as actants if they make a difference to what happens in a network (Latour, 1996). This leads rise to the methodological principle of 'general symmetry' (Latour, 1996) that human and nonhuman entities should be described using the same analytical vocabulary, without

assuming that humans are the only genuine sources of action. In *Reassembling the Social*, Latour argues that sociology should be understood as the study of associations, attachments, and mediations through which collective life becomes assembled in the first place (Latour, 2005). ANT therefore asks the analyst to resist the temptation to explain too quickly, and instead to describe how heterogeneous elements hold together, how they translate one another, and how their relations stabilise into something that can later appear self-evident.

To build on this, when a network (which is composed of actants acting together on each other, as previously implied) becomes complex and stable enough that its internal complexity is no longer obvious, it can be seen as a 'black box' (Law, 1992). A black box is essentially an entity that we treat as simple and unified when in fact it is the product of many entangled relations. The concept of black-boxing is important for the analysis of GenAI art. Since a black box is a network element whose internal workings are rendered invisible, in GANs or diffusion models such as DALL-E, the model itself can be seen as a black box since users see prompts and outputs but not the billions of parameters and the training corpus that mediate between them. This introduces asymmetries of power: those who can open the black box (engineers, institutions) hold different relational positions than those who cannot. Part of the task of an ANT analysis is to open these black boxes and trace the actants concealed within them.

For this thesis, that emphasis is relevant because generative AI art is rarely the product of a single identifiable source. When a text-to-image model produces an image, or when a generative system produces a poem, a video, or a portrait, it is not enough to say that 'the AI made it,' or alternatively that 'the artist used AI as a tool.' Such formulations are too reductive, because they conceal the more complex network of associations through which the work emerges. An ANT analysis instead suggests that a generative artwork should be understood as the effect of interactions among multiple actants, including but not limited to the artist, the dataset, the model architecture, the prompt structure, the interface, the software, the hardware, the platform, the institutional setting, and the audience. Each of these participates in shaping what the work becomes. The final output is can therefore be seen as both an expression of a human intention, a product of the machine, and as the result of a relational process in which different elements act upon one another. It does not have to be a black and white outlook that is either 'the human made this' or 'the machine made this'. It can instead be read as 'this was made by both a human and machine actants in a complex network of actants'.

We can also turn to the work of Madeleine Akrich and her account of inscription to further expand on this. In *The De-scription of Technical Objects*, Akrich argues that technical objects are not neutral devices because they can embody or carry within them particular forms of human intention or expectations in terms of how they are used, which influences behaviour (Akrich, 1992). Innovators, in her account, "inscribe" visions of the world into technical artefacts, and these inscriptions help define the roles, behaviours, and possibilities available to users when using those same artefacts. She refers to this as the script or scenario of the object (Akrich, 1992). A script is the set of anticipated actions, assumptions, and norms built into the design of the object, which then orient how the object will be used and what kinds of behaviour it will solicit or encourage. Akrich's point echoes that of Verbeek in that that technical objects can affect user behaviour and therefore shape user behaviour, and the scripts implemented are the hidden things in the network that influences said behaviour. To exemplify what is meant, Akrich gives the example of a door-closer. A door-closer performs the instruction 'close the door' on behalf of the designer, whether or not a human enforcer is present. It carries an intention, acting as a representative of the designer's anticipated world (Akrich, 1992).

In text-to-image systems such as DALL-E or Midjourney, it can be said that aesthetic scripts are embedded at multiple levels of the system. We can start by looking at the training data, which if for example only encompasses a very specific style of art, then the system will only generate an art form in that specific style. Say you train a system only on Van Gogh paintings, then it will try and emulate similar paintings. We can also look at the filtering mechanism or the specific code which instructs it what sort of images to produce as an output, which may carry scripts by the people who have designed the code. The platform's moderation policies and prompt interface might allow certain prompts and disallow others, limiting what the user is able to do and the generation possible. In this way, such systems can encode and reproduce particular aesthetic norms, often in ways that remain invisible to the user. If, for instance, the model has been trained

predominantly on Western, contemporary, and commercially circulated imagery, then that training material becomes part of the script through which new outputs are generated. The system may therefore naturalise certain visual preferences, compositional styles, body types, or representational norms, all while appearing to present itself as flexible or open-ended. This account allows us to see that the technical artefact is not innocent. It is a scripted object whose internal design already shapes the field of possible creativity. This point is necessary to emphasise because it shows that generative AI does structure the range of what can appear as a plausible image in the first place. By scripting certain visual preferences and representational norms into the system, the technical artefact participates in shaping creativity before the user even begins to prompt.

Before moving onto the criteria and definitions, it is useful to touch a little more on Latour's concept of black-boxing, because it allows us to see how technologies come to appear as stable and unified even though they are in fact composed of many hidden relations and decisions. In the case of GenAI, this is especially important, since users typically encounter only the prompt and the output, while the training data, code, and institutional conditions that make the output possible remain concealed within the system's apparent simplicity. As previously mentioned, in ANT and related science and technology studies, the term refers to the way the internal workings of a system disappear from view once it functions reliably, so that users engage with its outputs without needing to see the many elements that make it work (Law, 1992). The point is not that the system is simple, but that it becomes taken as simple because its complexity is successfully concealed by routine use and stability.

The concept comes from studies of science and technology, especially work in actor-network theory, where scholars such as John Law and Bruno Latour used it to show how scientific facts, technical devices, and social arrangements become stabilised over time. Latour's broader account of ANT also treats black-boxing as a crucial feature of networks: when associations hold together well enough, their internal negotiations and controversies recede from visibility, and the resulting object or process appears unified and finished (Latour, 2005). In this sense, black-boxing is both descriptive and analytical: it names how systems become opaque in practice, and it gives the researcher a method for reopening them and tracing the relations inside. In relation to GenAI, this matters because a user typically sees only a prompt and an output, while the model's training data, architecture, filtering, and interfaces remain hidden within the black box (Law, 1992; Latour, 2005). ANT also helps clarify why black-boxing is so important. As a network becomes stable and complex enough, its internal relations cease to be visible in ordinary use, and the network can come to appear as a simple, unified entity. This is what John Law describes as black-boxing (Law, 1992). A black box is thus a network whose internal complexity has been rendered invisible by its own successful stability.

From a Latourian perspective, a generative artwork is best understood as the effect of an assemblage of numerous actors together in a network. The artist may initiate the process, but the dataset constrains what can be learned; the model architecture shapes how those patterns are transformed; the interface governs how the user can intervene; the hardware conditions the speed and scale of generation; the platform imposes rules and limits; and the audience contributes interpretation and reception. The final image, or artefact is therefore the product of the outcome of many interacting actants. This has direct consequences for how the case studies in this thesis should be read. A work such as *Portrait of Edmond de Belamy* is the work of all of a multitude of actants together. Moreover, it must be understood through the network that includes the historical portrait corpus, the mathematical architecture of the model, the code that generated the output, the auction house that framed it as a high-value artwork, and the public discourse that received it as evidence of machine creativity. To analyse these works is to trace how those relations are assembled, and how some of them become invisible through black-boxing. A finished artwork is, from this perspective, a networked event. ANT provides a way of thinking beyond both technological determinism and human exceptionalism. It neither reduces the artwork to the intentions of the artist nor attributes creative agency wholly to the machine. Instead, it treats creativity as distributed across a network of actants, each of which makes a difference to what the artwork becomes. In this sense, ANT helps describe the actual conditions under which GenAI art is produced, circulated, and interpreted, and it allows us to see that what appears as a single image is in fact the visible surface of a far more complex assemblage.

Chapter 2: Situating the thesis and Theoretical Foundations

2.1 What counts as GenAI and GenAI Art

Before we start unpacking what is meant by GenAI art, let us first start by specifying what counts as a generative AI system in the first place, for the purposes of this thesis. Generative AI systems can refer to large-scale probabilistic models trained on substantial corpora of human-produced data (such as text, images, audio, or multimodal combinations) that generate novel outputs by learning and reproducing the statistical structure of their training distribution (Buckner, 2024). This includes, but is not limited to, diffusion models (such as DALL-E), generative adversarial networks (GANs), or large language models (such as Claude) which model the probabilistic structure of language to generate texts (Buckner, 2024), and multimodal models that can operate across text, image, audio, and other modalities simultaneously.

Generative AI can be very broadly defined as anything “created with an autonomous system” (Ghosh and Fossas, 2022). However, we can specify what is meant by this even further. In *Artificial Aesthetics: Generative AI, Art and Visual Media*, Lev Manovich also first starts off by trying to define what is meant by generative AI. The first distinction he makes is that “The terms generative media, synthetic media, AI media, generative AI, GenAI are all interchangeable” (Manovich and Arielli, 2024). He then offers the definition of generative AI being anything that involves “the process of creating new media artifacts with software tools that use certain types of artificial neural networks (e.g., AI models) trained on vast collections of media objects already” (Manovich and Arielli, 2024). For the purposes of the thesis we will start with this definition and build upon it slightly to define what is meant by generative AI.

There are however, several important exclusions from this specification. Early computer-generated art such as the algorithmic and rule-based work of Vera Molnár, or Harold Cohen’s AARON system. Rule-based examples such as these are excluded because the structure is fundamentally different to that of a neural network based system. Rule-based systems execute explicit human-authored instructions, while the generative models that are highlighted above instantiate a statistical approximation of a training distribution whose internal logic is not transparent. The case can be made that the problems posed by both are slightly different and should thus be treated as distinct as conflating them might obscure the specific novelty of the current moment. Similarly, AI systems which are used for purely classificatory or optimisation tasks such as fraud detection algorithms, recommendation engines, logistics systems, are excluded even in cases where they involve some sort of machine learning, neural network, or generative aspect, because their outputs are not expressive or communicative artefacts in the relevant sense (more on this below), or in other words, no human encounters them as meaningful objects to be interpreted.

With that established, we can move on to trying to define what counts as GenAI art. To do so we can look at the work of previous others who have tried to do so. Lev Manovich defines generative AI art as a new creative medium with its own logic, arguing that it “thinks” through pattern recognition, recombination, and the generation of variations rather than singular objects (Manovich and Arielli, 2024). In his work on AI image media and generative media more broadly, he emphasises that these systems are defined by prediction, dataset logic, and the ability to produce many possibilities from a prompt, which means that GenAI art acts as a form of computational aesthetic production rather than function as a pure machine imitation (Manovich and Arielli, 2024).

In *What is Generative Art?*, Boden and Edmonds argue that “generative art” can be a range or a family of related practices that they then separate into distinct categories. According to them, their “taxonomy distinguishes eleven types of art. We call them Ele-art, C-art, CA-art, D-art, G-art, CG-art, Evo-art, R-art, I-art, CI-art, and VR-art.” (Boden and Edmonds, 2009). For example, while computer-assisted art or CA-art is when the computer is mainly a tool under direct human control, generative art or G-art is when some part of the process is no longer fully under the artist’s direct control. They then narrow this further to computer-generated art or CG-art for cases where a program is left to run with minimal or no human interference. They also separate out evolutionary art, robotic art, interactive art, computer-interactive art, and virtual-reality art (Boden and

Edmonds, 2009). The categories differ mainly by how much control is handed over to the system, the audience, or the environment. In computer-assisted art, the artist remains in charge; in generative and computer-generated art, the system takes on more of the decision-making; in interactive and computer-interactive art, the audience becomes part of the artwork's production; and in virtual reality, the participant is immersed in a computer-generated world. Even though these categories differ, however, the authors argue that the aesthetic criteria are often surprisingly similar. What changes most is the locus of creativity. Sometimes creativity lies mainly with the artist who sets the rules; sometimes it is distributed across artist, program, and audience; and sometimes the system itself seems to contribute something genuinely creative through autonomy, emergence, or interaction.

Boden and Edmonds emphasise that the key difference between the categories lies in how much decision-making is handed over to the system, the audience, or the environment (Boden and Edmonds, 2009). This is useful for the present thesis because it helps explain why different forms of AI art cannot simply be treated as equivalent. It also suggests that questions of authorship, agency, and aesthetic value may need to be framed differently depending on the role the system plays in the production process. A legal or philosophical account of authorship, for example, may need to distinguish between AI used for stylistic assistance, AI used for co-creation, and AI used as a largely autonomous generator.

For this reason, AI systems used for purely classificatory or optimisation tasks are also excluded from the present definition. Systems such as fraud detection algorithms, recommendation engines, or logistics systems may involve machine learning, neural networks, or even generative components, but their outputs are not expressive or communicative artefacts in the relevant sense. This distinction is made because the focus of this thesis is not on AI in general, but on those generative systems that shape encounters with expressive artefacts. A recommendation engine may influence human behaviour, but it does not ordinarily produce an artwork, image, or text that is taken up as an aesthetic or communicative object. By contrast, a text-to-image model may generate an artefact that is then interpreted, judged, shared, or critically discussed. It is this latter kind of encounter that is central here. Lev Manovich's broader definition also helps refine this point further. He notes that the terms generative media, synthetic media, AI media, generative AI, and GenAI are often used interchangeably, and that they refer to the process of creating new media artefacts with software tools using artificial neural networks trained on large collections of pre-existing media objects (Manovich and Arielli, 2024). The artefacts these models can produce include images, video, animation, songs, music, text, 3D models, scenes, code, and synthetic data (Manovich and Arielli, 2024).

The definition of GenAI art adopted here refers to those artefacts produced through a direct human-model collaboration in which the generative system is constitutive rather than incidental to the work. This excludes cases where AI is used only as a stylistic filter or as a minor post-production aid, since in such cases the model is not structurally active in shaping the artefact in the same way. At the same time, the output must be encountered as meaningful. This includes things like images produced through DALL-E. Accordingly, the case studies chosen for this thesis are not intended to exhaust the field of GenAI art. The definition for what counts as GenAI for this thesis will be the same one proposed by Lev Manovich (with certain exceptions which I will cover later), which is anything that involves "the process of creating new media artifacts with software tools that use certain types of artificial neural networks (e.g., AI models) trained on vast collections of media objects already" (Manovich and Arielli, 2024).

The other question that needs to be answered before analysing how GenAI art mediates human to art relation then is to first discern what is meant by GenAI *art*, and by extension what sort of case studies and examples will be used to supplement this thesis. The category 'Generative AI art' can be too narrow or too broad depending on the way we define it. If we restrict ourselves to mean just art produced within art institutions like galleries with the term then we exclude a large majority of expressive and communicative artefacts produced using generative models such as graphic design, AI-assisted journalism, and everyday uses of large language models that are probably the most numerically significant instances of humans interacting with these systems. On the other hand, it becomes too broad if we include every single instance of an output of generative AI systems because this would include things like background operations which might not have a direct interaction between human and technology. Therefore I propose GenAI Art to be understood to mean anything or any form of an object or artefact that is produced using a

generative artificial network in direct collaboration with a human, this is based on the same definition proposed by Lev Manovich (Manovich and Arielli, 2024). It makes a specific distinction to have a direct collaboration with a human user involved in its creation, rather than through a fully autonomous creative process (Manovich and Arielli, 2024; Hertzmann, 2018). This definition is intentionally limited to works in which the human remains part of the creative loop, because current debates on AI art increasingly focus on collaboration, distributed agency, authorship, and the role of human intention in generating the final work (Cetinic and She, 2023; Di Dio et al., 2023). Thus, it would include images generated by systems such as DALL-E, even if for personal use, and exclude any background artefacts produced with no direct human interaction and with no purpose to form meaning out of said artefact.

A case belongs to this study's field when it instantiates a distinct and somewhat impactful configuration of human-GenAI-world relations, and when it produces an expressive or communicative artefact that a human subsequently encounters as meaningful. There are two main components to this aforementioned idea (for what substantiates something to be used as a case study of GenAI art):

The first component is that the generative model must be (co)-constitutive rather than incidental to the production of the work. This separates out cases where a generative model is a structurally active participant in producing the output (where things like its architecture, training data, and probabilistic next-token prediction actively shape what the output is) from cases where it functions as a filter or just a stylistic overlay applied to a fully-realised human-produced object. This follows from debates on AI co-creation and distributed authorship, which argue that generative systems are active participants in artistic production through their architecture, training data, and probabilistic generation (Manovich and Arielli, 2024; Di Dio et al., 2023). For example, a photographer who passes an image through an AI style-transfer application to add a painterly texture is using a generative model incidentally where as an artist who uses a text-to-image interface to generate a composition from a prompt is using the model constitutively. The distinction is one of degree rather than kind, and borderline cases will exist but the criterion provides a starting point to go from. Even though borderline cases will exist, the purpose of this is not to demarcate some sort of framework and pass definitive normative judgement on what counts as a valid case study and what does not, but justify the use of the case studies I have particularly chosen and provide an insight into what I consider to be GenAI Art that mediates human - art relations.

The second component is that the output must be an artefact that a human encounters as meaningful, as carrying significance that exceeds just functional information. The second component, that the output must be an artefact a human encounters as meaningful, is grounded in postphenomenological work on mediation and in broader discussions of AI-generated images as expressive and interpretive artefacts rather than purely functional outputs (Ihde, 1990; Verbeek, 2005; Manovich and Arielli, 2024). This criterion is designed to be deliberately broader than the term 'art' in any institutional sense (such as art in museums or galleries). A commercial illustration generated by a designer using Adobe Firefly for a product campaign is an expressive artefact because it is produced with communicative intent, encountered by viewers as an image making visual impressions on the viewer, and embedded in a network of human meaning-making practices. A legal document drafted with LLM assistance is a communicative artefact. A therapeutic journal produced collaboratively with ChatGPT is an expressive artefact for the person who produces it. The 'meaningful encounter' that this criterion requires is not restricted to trained aesthetic appreciation. It instead encompasses any engagement in which a human brings interpretive attention to an object as an object, as something to be read, perceived, felt, or reflected upon. These first two criteria can also be applied to further explain what is meant by the term GenAI art, while the final one merely serves to justify which case studies can be looked at to analyse how this art shapes our relation to art as a whole and how this relation is mediated.

The question might arise of 'Why not use an aesthetic lense to discern what is meant by GenAI art as opposed a postphenomenological one?' An aesthetic based lense would imports the assumption that the relevant philosophical category is 'art' rather than 'mediated expressive art encounter'. This thesis asks how Generative AI mediates the human-to-art relation, but the term 'art' in that question is understood broadly, as shorthand for any encounter with an expressive or communicative artefact generated by AI in which the human brings interpretive attention.

Having defined the field positively through the three criteria above, it is useful to specify a set of explicit exclusions. The first is that early and rule-based generative art such as the work of Molnár or Cohen is excluded, as noted above, because the mediating structure differs fundamentally from that of probabilistic large-scale models (Galanter, 2003; Hertzmann, 2018). In this example, Molnár's and Cohen's work relied on algorithmic, rule-based procedures and early computational systems, rather than on the data-driven training and probabilistic image generation characteristic of contemporary GenAI (Galanter, 2003; Hertzmann, 2018; Manovich and Arielli, 2024). In other word, Molnár's and Cohen's work was based on rules written in advance. The artist or programmer told the computer what steps to follow, and the computer then carried out those instructions. The system did not 'learn' from large amounts of existing images in the way contemporary GenAI does. Instead, it followed a fixed set of commands, producing results within a defined structure. By contrast, contemporary GenAI models are trained on very large datasets and generate images by calculating probabilities, which means they predict what visual form is most likely to come next based on patterns in the data (Manovich and Arielli, 2024; Hertzmann, 2018). So the difference is that Molnár and Cohen used rule-based systems, while GenAI uses training-based, probabilistic systems. The second is that purely operational AI outputs such as fraud detection results, recommendation scores, and logistical route optimisations, are also excluded because no expressive or communicative artefact is produced and no meaningful human-artefact encounter of the relevant kind occurs (Rosenberger and Verbeek, 2015). Even where these systems involve sophisticated machine learning, they do not generate objects that humans encounter as demanding interpretation. The third is that scientific visualisation produced by AI systems such as protein structure images, climate model outputs, medical imaging processed by neural networks. These are clearly meaningful to the humans who work with them, and they are encountered as artefacts requiring interpretation. However, the interpretive framework within which they are encountered is technical and operational rather than expressive (Manovich and Arielli, 2024). In those contexts, the human typically asks whether the output is accurate rather than what it means or how it is felt, which places it outside the scope of this thesis's focus on mediated expressive encounters (Ihde, 1990; Verbeek, 2005). While scientific visualisation raises important postphenomenological questions, it would significantly expand the thesis's scope beyond what I am aiming to achieve or argue. It is therefore explicitly excluded, with the acknowledgement that this exclusion is practical rather than principled.

The case studies analysed in the next chapters were selected to vary the dimensions along which Generative AI art differs. This range is not representative of Generative AI art practice in any statistical sense as a case study does not aspire to representativeness but to analytical coverage. The goal is to test the postphenomenological framework across a range of configurations. It bears repeating, finally, that the demarcation offered here is not permanent or exhaustive. Generative AI systems are developing at a pace that regularly produces new relational configurations, so the criteria offered are not the end all and be all, and an acknowledgement is made that there can indeed be borderline cases, but these cases shall not be dwelled on. Lastly, for the sake of depth of analysis the cases focused on here will be two different visual based cases as opposed to writing based cases such as those generated by an LLM. Text based cases such as those generated by an LLM such as Chat-GPT do fit the demarcation criteria, however I leave those cases to be explored at another time, either by myself at a future date, or someone else entirely as for the rest of the thesis I have focused on image based cases.

2.2 A postphenomenological philosophy for images

Before we can analyse specific GenAI produced artefact, it is important to understand how images are mediated and shape human-technology-world relations in the first place (and by extension how most forms of art, be it music, visual images, or writing, do so too). Once we have understood and explained how images themselves can mediate and shape human-technology-world relations, we can then move onto analysing specific examples created by GenAI and see how those specifically mediate and shape human-technology-world relations and what problems emerge as a result.

Images structure how the world becomes apparent to us, how we see ourselves in it and within it, how we orient ourselves within it, and how we come to understand ourselves through it. From a postphenomenological perspective, images are mediating artefacts that shape perception,

attention, self-relation, and social interaction. As an example this is especially true in digital environments such as dating apps, health apps, and social media profiles, where images operate simultaneously as representations of the self and as technologies through which the self is extended, curated, measured, and judged (Ihde, 1990; O'Brian, 2017; Verbeek, 2005; Rosenberger and Verbeek, 2015; Smith, 2014). With this context, images can be seen as actively participating in how we see ourselves and how we portray ourselves.

Before we explore a postphenomenological and ANT analysis of GenAI art and how that mediates human-art relation, an analysis of the image must be made. From a postphenomenological lense, we can start this by rejecting the assumption that vision is a neutral channel through which the world is viewed. Images can mediate. Images can mediate human-world relations by guiding perception, organising meaning, and reframing what is taken to be immediately visible (Ihde, 1990; Verbeek, 2005; Bolter and Grusin, 1999). We can reject the assumption that vision is a neutral channel through which the world is simply given to us. Images mediate reality by shaping what becomes visible, how it becomes visible, and how it is subsequently interpreted. To say that images can mediate is to say that they are not passive surfaces but active relational artefacts that structure our encounter with the world, guiding attention, filtering perception, and organising meaning (Ihde, 1990; Rosenberger and Verbeek, 2015). For example, a photograph on a dating app depicts a person, but it also shows them off in a way they themselves want to be seen, it tries to establish a connection, and communicate certain things to the viewer of the image, most predominantly in this case, some sort of desirability. What we see is already organised by the image, and what we do afterward is influenced by that visual organisation. If I swipe right on a tinder profile, it is because the image in front of me has been successful in proving the other persons' desirability and influenced my decision to swipe right. This claim also resonates with a broader media and film theoretical tradition that treats images as active mediators rather than passive reflections. Flusser, for example, describes images as mediations between humans and the world, while later media theory and visual culture studies emphasise the agency of images in organising perception, meaning, and visibility (Flusser, 2000; Hansen, 2015). Read through this tradition, the image represent desirability on a dating app and participates in producing the conditions under which desirability becomes legible in the first place. Postphenomenology shares this broad family resemblance, but gives it a more specific vocabulary for human-technology-world relations.

This feels closely tied to a sort of quasi-embodiment relation. What I mean by this is not the same sort of embodiment that you would get with glasses but a metaphysical embodiment of ones self to the online sphere. Online, the self often feels extendable and editable. People can curate different versions of themselves through photos, captions, and profiles (Belk, 2013; Turkle, 1995). Health apps turn the body into numbers, and goals, while social media can be a vessel for users to show what they look like, and also how they want to look like (Lupton, 2016). The body becomes projected online, it is externalised as an image. That can be freeing because it allows identity to be explored, but it also makes the self vulnerable to comparison, evaluation, and commodification (Marwick, 2013; Illouz, 1997). This can affect things like self-esteem. When images are liked, rejected, shared, or ignored, self-worth can start to feel unstable and dependent on feedback (Vogel et al., 2014; Ward, 2017). On dating apps, matches and messages can feel like measures of value. On health apps, visible progress can feel rewarding, while failure can feel deeply personal. Images can encourage self-care and reflection, but they can also increase anxiety, inadequacy, and obsessive monitoring (Lupton, 2016; Vogel et al., 2014; Ward, 2017; Lanzing, 2018; Lanzing, 2016). In these contexts, the body is externalised as an image and judged through feedback such as likes, matches, comments, or silence, which can make self-worth feel unstable and dependent on visible validation (Vogel et al., 2014; Ward, 2017). At the same time, these images mediate how we relate to others, since they often precede direct interaction and filter social encounters through rapid visual judgment (David and Cambre, 2016). Our relationship with ourselves and with the world can become affected through these images and is actively mediated through them. They also change how we relate to others. On dating apps, the image usually comes before the person, so visual representation becomes the filter through which relationships begin. This can make interaction easier, but it can also reduce people to quick visual judgments based on attractiveness, youth, confidence, or style. The point is that images are part of how we inhabit the world. They guide attention, shape values, structure comparison, and influence identity. By extension, this same argument for mediation can be made not just for the images that we see, but the text we can see online. Thus both images and written text online act

as active actants in how we see ourselves and the world and therefore are active participants in our relation to the world (Ihde, 1990; Rosenberger and Verbeek, 2015). What this shows is that images can be mediating forms that shape how people perceive, compare, and relate to one another. This matters because it shifts the focus from what an image depicts to what it does in practice, and that provides the bridge for our next step which is to examine how Generative AI intensifies and complicates these mediations by changing how images are produced, encountered, and understood. Thus, we now turn to link this to the way GenAI can affect this, and to studies from psychology and neuroscience as evidence that these technologies do have a mediating impact. Once we have done this, we will have established how both images and GenAI images do impact how human relations change to art and to the technology itself.

The rise of Generative AI makes the philosophical question of the image more urgent because it alters not only the production of images, but also the conditions under which images mediate human to image relations. As discussed, digital images can already function as things that shape how subjects appear to others, how subjects interpret themselves, and how subjects orient their desires, insecurities, and aspirations in relation to the social world. Dating apps, health apps, and social media all demonstrate that the contemporary image is something that influences our relation to the world. Generative AI intensifies this condition because it expands the scale, speed, and accessibility of these images while simultaneously obscuring the relation between image and reality (Hansen, 2015). For example, a text-to-image system can generate countless variations of a person, a landscape, or a product without any direct photographic relation to an original scene, which means that images no longer need to function as traces of something that was physically encountered. This becomes prominent in the case of profile images, advertising, and social media content, where the image may appear realistic enough to carry social meaning while nevertheless depicting a person who does not exist, a product that was never photographed, or a situation that never occurred. In this sense, GenAI also destabilises the epistemic status of the image by weakening the assumption that visual representation is anchored in a determinate physical reality. The result is that images participate in shaping what is taken to be real, desirable, or credible, even as their relation to the world becomes more opaque (Flusser, 2000). This shows that generative AI actively extenuates the conditions under which images mediate perception, self-understanding, and social relation.

Postphenomenology teaches us that technologies can mediate our view of ourselves and the world. In digital image environments, this is especially clear: dating apps, photo-editing tools, and health platforms become part of everyday self-presentation and self-assessment. On dating apps, the profile image functions as a kind of public body, projecting the self into social space before any direct interaction takes place. Approval or rejection then feels personal, because the image feeds back into self-worth. Generative AI intensifies this process by making self-images easier to edit, idealise, and reshape. Users can now create enhanced portraits, avatars, and alternative versions of themselves that are less tied to the lived body. The result is a more synthetic and flexible form of embodiment, where identity can become rehearsed through images before it is fully lived or recognised. Generative AI also raises a darker set of concerns, especially where bodily images are manipulated without consent, as in nudification apps, deepfake abuse material, and the non-consensual circulation of sexualised synthetic images of women (Henry et al., 2018; Citron and Franks, 2014; Chesney and Citron, 2019). In these cases, the same image-based mediation that can support self-presentation and self-fashioning becomes a mechanism of harm, since control over one's image is displaced into systems that can strip, alter, or redistribute the body in ways the subject does not authorise. This makes the question of mediation inseparable from questions of vulnerability, power, and loss of control over visual identity (Henry et al., 2018). This shows that GenAI mediates human-to-image relations in two contrasting ways: it can extend self-presentation and self-fashioning, but it can also intensify vulnerability by enabling forms of image-based harm and loss of control. It is therefore relevant to the research question because it demonstrates that GenAI does actively reshape the conditions under which images structure subjectivity, agency, and social relation.

The psychological implications of this image regime are already well documented, and they become even more significant in the context of Generative AI. Research on social media and body image shows that visual self-presentation and image comparison are closely tied to self-esteem, self-objectification, and perceived attractiveness. Ozimek and al. found that photo editing on social networking sites was negatively associated with self-perceived attractiveness and self-

esteem, with this relationship mediated by self-objectification and appearance comparisons (Ozimek et al., 2023). This is important because it shows that visual manipulation does reorganise how users relate to themselves. The more subjects modify their images, the more they may come to view their bodies from an externalised, evaluative perspective. Other studies reinforce this pattern. Research on Instagram use and social comparison indicates that upward comparisons decrease body-esteem, while downward comparisons can temporarily raise it. More broadly, it has been found that there is a significant association between social comparison on social media and body image concerns, lower positive body image, and eating-disorder symptoms (Fardouly and Vartanian, 2016; Vogel et al., 2014; Linardon et al., 2025). Generative AI may intensify these problems in at least three ways. First, it increases the quantity of idealised images in circulation. Second, it allows those images to be personalised and iteratively optimised according to platform feedback. Third, it makes it possible for users to compare themselves to both edited photographs of other humans, but also to fully AI generated visual ideals that may be unconstrained by ordinary bodily limits. If the online visual environment becomes increasingly saturated with AI-enhanced or AI-generated ideals, then the benchmark for attractiveness, health, or lifestyle may shift. This creates a paradox. Generative AI may appear democratising because it gives more users access to tools of visual enhancement and self-fashioning. Yet the same tools may also deepen inequality at the level of affect and self-worth by intensifying pressures to appear polished, attractive, or optimised.

Neuroscientific research helps explain why these effects are not superficial. Studies in neuroaesthetics suggest that meaningful visual experience can engage systems involved in emotion, reward, and self-referential thought. Vessel and al. found that intense aesthetic experience activates the default mode network, a set of brain regions associated with self-related cognition and introspection (Vessel et al, 2012). Related work has further linked aesthetic experience to self-referential processing, suggesting that powerful images become integrated into the subject's own reflective architecture. Reviews of EEG and fMRI work also indicate that aesthetic processing involves multiple systems associated with attention, emotion, and reward rather than a single isolated 'beauty centre' (Rui and Gu, 2021; Chatterjee and Vartanian, 2014; Cela-Conde et al., 2011). These findings show that images can affect how subjects feel and think about themselves. In the context of GenAI, this suggests that synthetic or AI-enhanced images may have real affective and cognitive effects regardless of their ontological status. A generated ideal body, a hyper-polished self-portrait, or an AI-curated feed can still become psychologically and neurally salient. The fact that an image is artificial does not prevent it from entering into the subject's reflective life. If images already help structure self-esteem, desire, and social recognition, then more adaptive and synthetic images will likely intensify that structuring role.

What GenAI adds to this picture is not just more image editing, but a more radical loosening of the relation between self-image, embodiment, and visual normativity. We are now in a position where GenAI can generate bodies, faces, styles, and lifestyles from prompts rather than from lived photographic events, which means that while users already compare themselves to existing ideals, they can now compare themselves to ideals which are completely arbitrary as GenAI can also produce ideals that are infinitely adjustable, personalised, and detached from ordinary bodily limits (Hansen, 2015; Manovich and Arielli, 2024). This highlights the broader argument of this chapter that images are mediating relations that help configure how subjects perceive themselves, evaluate others, and imagine what counts as a desirable body or life and that GenAI therefore intensifies this world or sphere by making the production of ideals faster, and cheaper, while also making the standards themselves less stable (Flusser, 2000; Verbeek, 2005). In that sense, the relevance of GenAI is that it can deepen the mediating power of images over subjectivity, comparison, and self-relation.

Across these sections, we have established a framework that defines GenAI as a specific class of probabilistic, data-trained systems and GenAI art as a relational practice in which such systems are constitutive of expressive or communicative artefacts. We then narrowed the field by excluding earlier rule-based computer art and purely operational AI, and by stressing that the thesis is concerned less with whether AI outputs count as art in the abstract than with how they mediate human experience, authorship, and interpretation. On that basis, I have argued that GenAI art should be studied through postphenomenology and ANT because it reorganises the conditions under which images, texts, and other artefacts are encountered, judged, and made meaningful. The philosophical account on images extends this claim by showing that images are

mediating artefacts that shape selfhood, social comparison, embodiment, and desirability, especially in digital environments. We will now move on to looking at two distinct case studies that incorporate GenAI to produce artwork, and analyse how these affect our outlook towards art and these tools through a postphenomenological and ANT lens.

Chapter 3: Visual Cases

This chapter turns from framework and definition to application by examining two visual case studies that make the thesis question concrete. Having established that generative AI art is best understood as a mediated thing, and that images themselves already shape selfhood, perception, and social relation, the task now is to show how these dynamics appear in specific works. The first case, *Portrait of Edmond de Belamy*, is especially useful because it raises questions about authorship, and exemplifies the concept of black-boxing and the distribution of agency across a network of human and non-human actors. The second case, Mario Klingemann's AI-generated portraits, develops this discussion further by showing how generative systems mediate aesthetic experience, embodiment, and the relation between artist, model, and viewer. Taken together, these cases allow to show how GenAI reorganises what it means for art to be made, encountered, and valued (Latour, 2005; Verbeek, 2005; Rosenberger and Verbeek, 2015). This chapter finishes with a short discussion and note surrounding other contemporary debates relating to GenAI.

3.1 *Obvious's Portrait of Edmond de Belamy*

Edmond de Belamy is an AI-generated portrait created in 2018 by the Paris-based collective Obvious. It became famous after selling at Christie's for \$432,500, drawing attention as one of the first major auctioned works made with generative AI. The image is a blurred, old-fashioned-looking portrait meant to resemble a fictional aristocratic family of AI-made paintings. Obvious's *Portrait of Edmond de Belamy* foregrounds authorship, and black-boxing. The image was signed by a fragment of the model's loss function. As a reminder, this is a mathematical expression used during model training to measure how far the system's output is from the desired result and to guide the network in adjusting its parameters. The work's sale price and the controversy surrounding it helped establish it as an early case in the public debate over AI art. An ANT reading shows us that the actants in the network includes the Obvious collective, the GAN architecture, the WikiArt dataset of historical portraits, Robbie Barrat's open-source code, Christie's institutional framing, and the art market that converted the work into a highly visible event. This case is important to highlight because it destabilises the singular author model. There was controversy that emerged when it was revealed that Obvious used code developed by another artist-programmer without initial credit. An actant that had been rendered invisible suddenly reappeared, forcing a redistribution of credit and responsibility. This highlights Latour's idea of black-boxing and shows that it can have real consequences. The more a work is framed as the product of "AI," the easier it becomes to obscure the layers of human labour, prior artistic work, and infrastructural dependence that make the output possible. As such, this case will be used to illustrate two things, the first being that GenAI art can raise questions surrounding authorship and copyright, and that if we use an ANT reading, it can reveal itself as a black-box which when opened causes further questions regarding its authorship and copyright questions. However, let us first start with a postphenomenological reading to analyse how this case can serve as an example for how GenAI mediates relations and then return to the key questions after.

Postphenomenology distinguishes between different types of relations: embodiment relations, hermeneutic relations, alterity relations, and background relations. *Edmond de Belamy* does not primarily function as an embodiment relation in the strict sense, because the viewer is not wearing or using it as an extension of bodily action. It is closer to a hermeneutic relation, in which the image must be interpreted. The viewer is led to ask what it means, how it was made, and why it matters. On the other hand, the work has an alterity dimension for the artist, because the AI system appears as a quasi-other whose output is not fully predictable and whose logic remains partially opaque. The work therefore stages a relation in which the machine is encountered by the artist as something partly alien and partly collaborative. Here the AI system appears as a quasi-other, which is why alterity is a useful term. In Ihde's sense, alterity relations occur when a technology is encountered as something with a kind of partial otherness, something that can respond, resist, surprise, or exceed intention (Ihde, 1990; Rosenberger and Verbeek, 2015). That is what happens in generative art, the artist prompts, sets parameters, or curates outcomes, but the model produces an output that is not fully predictable and not entirely under direct control. The machine is therefore experienced as a collaborator that is not fully human and also not fully independent. This is important because it means that GenAI mediates the human-to-technology relation by changing what technology feels like in practice. The system is no longer just a tool in

the instrumental sense, but a partner in a relation of co-creation (Manovich and Arielli, 2024; Di Dio et al., 2023).

The work also illustrates inscription, especially in the Latourian sense of a technical process that records, transforms, and stabilises phenomena into a visible output. The GAN trained on historical portrait data inscribes the statistical regularities of that dataset into a new image, but it does so indirectly and opaquely. The output is therefore not a transparent mirror of human intention; it is the material trace of a chain of transformations involving dataset selection, code, architecture, training, curation, and institutional framing. In this sense, the portrait is not “made” by Obvious, but inscribed by a network of human and nonhuman agencies. This inscription changes how we should come to understand GenAI tools. Instead of viewing them as neutral instruments, we should begin to see them as systems that carry prior cultural material forward in transformed form. If we look at the concept of affordances that is highlighted in the first chapter, we can also ask at what this tool that has been used to generate this portrait allows us to do, and in this case, it is to generate similar images to the images that have been used as training data. All of these things allows to have a case that shows us how humans can come to view GenAI tools more generally and how this is affected by the tool itself. In this way, Edmond de Belamy mediates the image-viewer relation and also the human-tool relation: people are invited to see GenAI as something that participates in making meaning and the artist is invited to see GenAI as something that can participate in co-creating said meaning. It mediates the relation to art by making the viewer interpret the artwork as a technologically produced object and it mediates the relation to technology by making the artist encounter the model as a quasi-other that shapes the work in ways the artist cannot fully predict or command (Latour, 2005; Verbeek, 2005; Rosenberger and Verbeek, 2015).

In turn, one of the questions raised by this and by the work in general is what kind of relation humans now have to artworks whose production is technologically co-produced and therefore whose production is distributed across datasets, code, institutions, and markets. As previously mentioned, this case embroiled itself in controversy surrounding the authorship and copyright of its code and this provides us the basis for the copyright discussion. The features that it has such as unclear authorship, hidden contributors, and a western dataset make it legally difficult to fit into existing copyright models. Even if we remove law from the equation, the work already asks a philosophical question of authorship when a work is co-constitutive with an GenAI model. While this case is one example, in any case regarding GenAI art, the question surrounding copyright and authorship can emerge and this question is a part of how GenAI mediates human-world relations as it can asks us to reconsider the way we treat copyright and authorship. Generative artificial intelligence has produced one of the most significant challenges to existing legal and philosophical frameworks for art. Systems such as Midjourney, DALL-E, and Stable Diffusion allow users to produce images through text-based prompts. As such, the question of who, if anyone, owns the result has become the subject of intense legal and critical debate, as was the case for our chosen case study. Drawing on Bruno Latour’s actor-network-theory (ANT), I will aim to argue that generative AI images or texts are not the product of a single bounded human author but of a distributed network of human and nonhuman actors. On that basis, the classical framework of individual copyright is philosophically inadequate for AI-generated art, and an ANT reading opens the possibility of rethinking AI-generated images as commons-like, shared cultural products.

It is very easy to look at recent image generated software such as DALL-E and simplify the process as a person types in a prompt and image comes out. This oversimplification would be what can be referred to as a black box. We can turn to ANT to analyse how different actants actually come together to generate and produce the image. In doing so, we begin to realise that the image generated emerges from an assemblage of different actants as opposed to a single source. Let us consider an image generated by DALL-E or by any Generative Adversarial Network (GANs). The process starts by the user providing a prompt for a desired image, however the prompt alone does nothing without the series of actants that are involved in the entire process. The prompt must be interpreted by the model which itself is built on a training set of billions of image-text pairs that have essentially been scraped from the internet (see Kalpokiene And Kalpokas, 2022 for a more detailed explanation). The very dataset itself is the product of the labour of thousands of human photographers, painters, illustrators, and designers, whose work has been collected, usually without consent or compensation, and used to shape the model’s

output (Kalpokiene and Kalpokas, 2022). This in itself is the product of multiple actants, but building on this, the model architecture itself is designed by computer scientists and engineers, informed by prior scientific work and code. The model itself runs on hardware such as GPUs, servers, and data centres (which require so much energy and material infrastructure that the ethics of which deserve a separate analysis). Continuing from this, the platform through which the user accesses the system has its own interface design, moderation policies, and terms of service, all of which shape what outputs are possible and permissible. The image is stored, rendered, and delivered through further systems of software and infrastructure (Kalpokiene and Kalpokas, 2022). In ANT terms, every element in this description is an actant. The model shapes the image; the training data inscribes aesthetic tendencies, biases, and visual vocabularies into the output; the hardware determines latency, resolution, and accessibility; the terms of service determine ownership and licensing; the prior artists whose work was scraped contribute to the generative capacity of the model even though they are invisible in the final product. The user's prompt is one actant among many and while not necessarily unimportant (my aim is not to diminish or remove the impact of the prompter itself), but certainly not the sole origin of the work. By seeing this process as a black box rather than what is actually is, what we essentially say is that 'the user made this', as opposed to the effort of a series of actants coming together. We treat the image as if it has emerged from a single source of human creativity, when this is not necessarily the case.

And yes, on the other hand, we can provide the same analysis for almost anything. Even a painter who paints using a brush on a canvas can only do so with the producer of the brush, the paint, the canvas, and acts according to their own inspirations and creative process, which is often not the product of solely their own imagination. However, let us think of it this way: let us imagine that there is person A and person B. Person A is a skilled painter and person B hires person A and tells them what to paint. While person B decides on the 'prompt', it is person A that is actually responsible for having the talent and expertise to make this prompt into a reality onto a canvas. This example is not a stretch, throughout history, many artists have been commissioned to produce specific paintings, and their work compensated and rewarded, it is often Person A, the actual artist who is considered the artist behind the work, and not person B, the 'prompter'. We can liken this example to the case of GenAI, if person B is the prompter, then GenAI is person A, the artist itself. The problem with GenAI however is that there is no such thing as person A, because person A is really an assemblage of various actors (coders, engineers, other artists whose work has been used in datasets. Etc.). Thus the issue arises: who do we attribute copyright to? Who owns the actual intellectual property of the image generated and who should be compensated for it?

The traditional framework of copyright is built on humanist premises (Kalpokiene and Kalpokas, 2022). At its core, copyright assumes that there is a human author who originates a work through an act of individual creative effort, and that this origination gives them an exclusive right to control and profit from that work. This model has been extended and qualified in many legal jurisdictions, but its foundational logic remains: authorship, originality, and individual ownership are the pillars on which copyright rests. However, by turning to ANT and analysing how authorship is actually distributed across a network, then the obvious conclusion is that copyright cannot be attributed to a single person. If creativity emerges from a relational assemblage, then it cannot be treated as the expression of an individual mind. If the image is an outcome of translation, inscription, and enrolment across many actants, then the concept of a determinate human author does not apply to a single human contributor. Margaret Davies' work on property provides a theoretical supplement here. In "Can Property Be Justified in an Entangled World?" (Davies, 2023), Davies argues that the very logic of property, the division of the world into owners and owned, subjects and objects, fails to capture the relational, entangled character of the world. Davies argues that ownership depends on a prior and philosophically dubious separation between a self-sufficient human subject and an external world of ownable things. Once we take seriously the idea that subjects and objects, humans and nonhumans, are constitutively entangled, that separation becomes impossible to sustain. Davies develops the concept of "material subjects and vital objects," arguing that legal personhood and property relations need to be rethought in light of the relational ontology that posthumanist theory proposes (Davies, 2016). Applied to generative AI art, Davies' argument suggests that copyright's assignment of the image to a single owner is both legally dubious and philosophically exists in a grey area. It seems difficult to impose a clean ownership relation on a product whose genesis was irreducibly entangled. This has significant implications for how we ought to see AI Generated images (or any artistic generative AI work, be it

AI used for poetry, or filmmaking purposes). If GenAI art is not the product of a single author (which a ANT view suggests), then they do not have ownership, or cannot be owned in the traditional sense (that foregrounds our legal framework). They are produced by and from a collective process, from the work of previous artists encoded in training data, from engineering knowledge embedded in the model, from the material infrastructure on the internet, and the prompter themselves. This collective process suggests that AI-generated art should be more appropriately thought of as a kind of quasi shared cultural production, which does not belong to a single individual, but to a wide network of contributors.

An objection to this is that it could be argued that copyright does not need to be grounded in any particular theory of authorship or human nature. An argument can be made that copyright exists as a pragmatic legal mechanism to incentivise creative production and ensure that creators are rewarded for their work (Merges 1989). With this view in mind, the logic follows that it does not matter whether or not the image is really the product of an individual author or multiple actants but what matters is that assigning copyright to the user gives them an incentive to generate and distribute images, and in turn this supports a creative economy. While this argument is fair, it raises a lot of potential problems and ethical questions if we use it as the basis or as the reason to give copyright privileges to the prompter. The incentive rationale does not work the context of generative AI art. The human artists whose work was scraped to train the model receive no incentive under copyright as currently applied. The engineers whose code made the output possible receive no per-image reward. The labourers who annotated and filtered training data are invisible to copyright law. The only party incentivised is the user, who often invested very little effort in the process (Cohen, 2024). If copyright is applied to the prompter as an incentive tool, it is a poorly aimed one as it rewards the least-invested participant and ignores the most-invested contributors to the network (Latour, 2005; Davies, 2023). This is not an argument for abolishing copyright, but for recognising that individual authorship is the wrong frame for distributing the rewards of AI art production.

The bottom line is that this analysis aligns with a broader posthumanist critique of anthropocentrism (for a more in-depth analysis on this, again, see Kalpokiene and Kalpokas, 2022). Rosi Braidotti argues that the human must be understood as a relational becoming always produced through material and discursive encounters with others, human and nonhuman alike (Braidotti, 2013). In the context of art, this means that creativity is not the exclusive property of an individual human but an event produced through networks of actors. Lastly, an argument can be made that by distributing agency across networks, it makes it impossible to assign moral and legal responsibility. If everyone and everything made the image, then no one is responsible for it. This worry has practical consequences as for example if an AI-generated image reproduces an artist's style without credit, or produces harmful or biased content, we need to be able to hold someone accountable. However, in the same way that copyright needs to be rethought in the context of networks, so does responsibility. If a generative AI platform builds a system on scraped data, the platform is an actant with specific relations and specific capacities for harm or benefit. If an engineer builds bias into the model, that design decision is an inscription with consequences. If a corporation writes terms of service that appropriate all outputs, that legal document is an actant with effects on power distribution. Making the network visible does not dissolve responsibility but it reveals that responsibility is borne by multiple actors, including nonhuman ones (institutional structures, technical architectures, legal instruments), and that a legal framework focused only on the individual prompter misses most of what is ethically at stake (Davies, 2016; Kurki, 2019).

The argument of this here can be stated as follows: generative AI art is not made by a single human author, and any legal framework that assumes it is misrepresents both the nature of the object and the distribution of contributions that produced it. In any case, "Currently, the EU lacks specific rules on the copyrightability of AI-generated works" (Karttunen, 2025) and this seems to be the case for the US as well as for other countries (for an detailed analysis of country to country copyright laws as relating to generative AI, see Karttunen, 2025). It seems to be the case that the traditional copyright model is inadequate to the reality that GenAI exists as the product of a multiple actors in a network. What should follow is not necessarily a world without ownership, but a rethinking is needed of what ownership means for objects produced through distributed, more-than-human networks. This has been exemplified by the case of Edmond de Belamy primarily. The case has showed us how GenAI art in general can mediate human-technology relations by

encountering GANs and other networks used in image generation as quasi-other alterity relations, and how human-art relations change as humans are now encountered with art generated by a technological tool in direct collaboration with a human which raises questions about how exactly this art should be understood, namely from the authorship and copyright perspective.

Nevertheless, if we take the copyright question aside for the time being, the fact that the painting sold for such a large amount is also significant because it shows that GenAI art is slowly becoming accepted in the world. It is becoming accepted by the collective human zeitgeist. Why else would someone pay such a large amount for it? There is evidence to suggest that the more people use AI or get exposed to AI then the more they are accepting of AI and have stronger positive attitudes towards it. This can be extended to GenAI art. There are two studies that point to this. The first one is A study on students' attitudes toward LLMs suggests that students with more exposure/use tend to report more favourable attitudes and positive user experiences with LLMs (Yeh and Siah, 2025). The second is a study on the frequency of generative AI tool use found that reliability and satisfaction were important for attitude formation, and that usage frequency interacts with these perceptions in shaping how people feel about the tools (Fakhri et al., 2024). The more we are exposed to AI (and GenAI), the more we come to accept it and value it positively. In this way, GenAI actively shapes our attitude towards. The MBS Global research summary on public attitudes toward AI reports that people who intentionally use AI regularly are also more likely to trust, accept, and feel positive about it, and it explicitly states that AI literacy is associated with greater use, trust, and acceptance (Gillespie et al., 2025). This is perfectly exemplified by Edmond de Belamy, it is the perfect case study to illustrate that GenAI art has become apart of the art world and is here to stay. The more we interact and get exposed and use such tools, the more we accept them and positively value them. In this way, in a postphenomenological sense, the usage actively shapes our attitudes towards them. This point, will be expanded upon with the next example, that of Mario Klingemann.

3.2 Mario Klingemann

Mario Klingemann is an artist that works with neural networks and generative models to produce his art. His work is somewhat similar to the Portrait of Edmond de Belamy as it also deals in portraiture, for example in works such as Memories of Passersby I, Klingemann trained a neural network on thousands of historical portraits and then allowed the system to generate an endless sequence of new faces. The result is a digital portrait that is similar to the first case study that we looked at. While his work is less controversial and can be seen as having a very specific prompt, and therefore a less complex network of actors regarding authorship and copyright, it is still useful as an exploration as to how these GenAI tools shape art, directly with the artist. Mario Klingemann's work is a useful second case because it draws from the same argument made in the case of Portrait of Edmond de Belamy, however the questions that I want to discuss in relation to it are different and deal with deskilling and the shifting public perception. Using the example of Klingemann allows the thesis to move from the question of authorship to the question of deskilling in relation to human-technology mediated relation through GenAI and shifting cultural norms in relation to human-world relations mediated through GenAI.

From a postphenomenological perspective, Klingemann's work can be understood as a case in which technology enters into the constitution of experience itself. Postphenomenology emphasises that technologies mediate the relation between humans and the world, amplifying some aspects of perception while reducing others, and shaping how reality becomes available to us (Rosenberger and Verbeek, 2015). Klingemann's AI-generated portraits exemplify this kind of mediation because they display an image in which the face, traditionally one of the most immediate and socially legible forms of human appearance, becomes in his work ambiguous and partly opaque. A portrait produced by a generative model looks at once recognisable and foreign, human and artificial, specific and generic. I will borrow a quote from Diego Bernaschina in *'Visual Art and Artificial Intelligence: Tensions between Ethics, Aesthetics, and Authorship in the Algorithmic Age'*, in which he posits:

"The aesthetics generated by AI challenge traditional categories of the beautiful and the disturbing. This phenomenon—such as the creation of images that blend the familiar with the unknown—produces a sense of uncanny familiarity or the uncanny valley, which influences the viewer's perceptual and emotional response. In this aesthetic, as seen in the visual study of the

personalised art installation, the coexistence of recognizable human faces with their deformations and fragmentations materializes, evoking both empathy and unease. However, this study expands the reflection from the perspective of algorithmic aesthetics, where the machine not only replicates existing visual patterns but also generates entirely new forms that elude human logic. This ability to produce unpredictable and dynamic images—such as those that change with the viewer's position within the installation—redefines the traditional aesthetic experience and places the audience in an active and participatory role.” (Bernaschina, 2025).

This is perfectly exemplified by the work of Klingemann, in the sense that the portraits generated look different to what we would expect normal portraits done by a human to look like. This does not mean that a human artist cannot recreate or create such portraits if they choose to, they surely can, it just exhibits a very particular form of portraiture which is why GenAI art such as that of Klingemann can be said to “redefine[s] the traditional aesthetic experience and places the audience in an active and participatory role” (Bernaschina, 2025). The active and participatory role comes from the fact that ultimately, it is the audience that will come to accept or reject this art and ultimately give it value. This is what I meant earlier with the idea of shifting cultural norms or public perception. The more these artworks enter the zeitgeist, the more public perception will shift in one direction or the other and the more cultural norms will slowly change. In the previous example, the portrait of Edmond de Belamy, having sold at an auction, places value upon the generated artwork as someone has shown willingness to pay for it. This would for example be evidence that such art is beginning to be accepted in the public zeitgeist. Thus, GenAI can be seen as mediating human-world relations as it allows the uniquely GenAI aesthetics to become slowly accepted and valued by society. I used the term valued here very generously and very loosely, almost metaphorically as it does have to be noted that AI generated images or AI mediated art, is received less favourably. “When an image is identified as the product of an algorithmic system, it tends to receive a less favorable evaluation and generate a weaker emotional connection, regardless of its formal or expressive attributes. This phenomenon reveals a persistent cognitive bias, reflecting the tensions between the humanist tradition of art and contemporary technological innovation. Although AI-generated works exhibit high levels of visual complexity or expressiveness, public perception hinges on the idea that subjective intention overrides aesthetic value” (Bernaschina, 2025). Despite this however, the more GenAI art becomes part of the zeitgeist, the more we are likely to see these cognitive biases change.

The Eiffel Tower at one point too was aesthetically rejected but has since become a symbol of Paris and is “now an icon of modernism and avant-garde design” (Zarevich, 2023). The point I am trying to make is that normative aesthetic judgements change. Certain things which are initially seen as uncanny, or ugly, or avant-garde, or against the status quo can be dismissed at first, but through time become accepted and sometimes even valued highly. The case of the Eiffel Tower being a prime example for this, or even if we take someone like Van Gogh, whose post-impressionist paintings were often dismissed until later in his life, and only really valued, after his death. I am not trying to make an aesthetic argument as to what should be valued and what shouldn't. This aim is not to engage in an aesthetic debate nor to try and pass normative judgement on which aesthetic criteria should be valued, the aim is to point out the fact that over time, societal taste can change, and aesthetic taste, and by extension what is accepted and valued, can change. Whether it is for better or for worse there are examples that this does happen, and the reason for that are irrelevant to the discussion at hand. All of this is to serve to say that even though at this very moment in time surveys and studies that have been aforementioned might show a slight distaste towards AI created, or mediated art, this might not be the case for the future. The constant exposure to works of GenAI, the exposure and use of these tools, will only result in more and more acceptance of GenAI art (whether in cases like this one or Obvious) and the valuation of said art. This is not to say that art created using AI tools will be valued on the same playing field as photography, or film, or classic portraiture, or any art not utilising AI tools for that matter. It will probably be valued in relation to its own field, in relation to other artworks that have also used said tools. However, the point is to say that the more a user, or person sees the artworks, and the more they become accustomed to the uncanniness or the novelty of the artworks produced by GenAI, the more they will be accepted. Further, the more users try and engage in GenAI art themselves by using these tools, the more they will be accepted. This is the main point of what postphenomenology is trying to say and analyse. The more we integrate GenAI as embodied, hermeneutic, alterity, fusion, or background relations, the

more they become part of our world and by extension they become a part of how we experience and interact with the world.

To quote Iris Murdoch, "Art attacks art and thrives by doing so. The artist derives energy from loving, but also from hating, the art of the past" (Murdoch, 1992). Photography and oil paintings are two different means of artistic outlet or expression, but when photography first emerged there was a fear that it would replace classical portraiture (Gershon, 2022;). This clearly did not materialise as we see photographs in their own sphere of art, and portraits have continued to be painted to this day. Photography has only given us a new art form to discover and a new set of rules for aesthetic judgement to apply when it comes to judging photography specifically. Equally, photography has meant the emergence of new and specific skills, different to that of painting a portrait with oil on a canvas. Photography requires knowledge and use of camera, how to set the correct aperture, what the exposure time should be, how to properly light the subject, and in older cameras, how to develop the film which are vastly different skills than that of a portrait painter who would need to know how to mix different colours together, how to blend them, how long to let the first layer or coat of paint dry before moving on to the next, and so forth. In this same way, would it not be safe to assume that GenAI art is just another medium of art that is slowly emerging? Much like photography, it will not replace art (although no one seems to be claiming this in the first place, but the more important point to be made here is that (and this refers to the art world specifically), it will lead to the emergence of new skills and a new distinct criteria for aesthetic judgement. The point about the new skills that I bring up here, is also relevant with regard to the deskilling issue that is also covered in this thesis, later on in this chapter. Nevertheless, for now the point is that if we judge a photographer by how they light their subject, how they play around with exposure or aperture, then we might judge the GenAI artist as how they interact with the specific generative network or how they use prompts and what ideas and themes they explore. In other words, all of this can be summed up by saying "AI art and its generation affords opportunities to create images that no other medium truly could, meaning it is well poised to become its own medium or style" (Ghosh and Fossas, 2022). In the same article, Ghosh and Fossas conclude that:

"Art as a medium has always molded to fit the society that creates it, from cave drawings created by early humans to fine art painted with quality acrylics, to contemporary technology powered society. That being said, we strongly suggest that the inclusion of AI into art should be performed responsibly in order to avoid the legal, financial, and ethical impacts on human artists. We also believe that AI art is fully capable of becoming a large, unique medium of its own, like photography and digital art did" (Ghosh and Fossas, 2022).

While it is important to make this note, and it is not controversial to push for an ethical, and legal, and responsible approach to the inclusion and implementation of GenAI into the art world, the article does not make any direct recommendations for what an ethical and responsible implementation looks like. This is not meant as criticism, as this very thesis also does not make direct recommendations for what an ethical and responsible implementation looks like. What is needed however, before such recommendations can be made, is an analysis of how GenAI art actively interacts and shapes our outlook towards it and art in general before specific recommendations can be made. In order to make specific recommendations and in order to know what legal actions, or ethical judgments are needed, we must first understand how GenAI affects us, how it shapes us, and what consequences emerge as a result. By identifying and analysing the problem first, we lay the foundation for a more fruitful conversation around solutions in the future. Ghosh and Fossas, among others, have identified some of the earlier problems. This has been the aim of this thesis. To expand on these problems and analyse them from a postphenomenological and ANT lense in order to expand the discourse around the subject and get a step closer to a potentially viable normative outlooks on the ethical dimension surrounding GenAI and its relation and implementation into the art world.

Stepping away from this for a minute, we can also read this case study through the concept of embodiment. In postphenomenology, technologies can become quasi-transparent extensions of perception and action, meaning that they are not merely external objects but are folded into the user's way of being in the world. Klingemann's work extends this insight into the domain of AI-generated visuality. In his images the viewer is looking at an image but also engaging with a technically mediated environment in which machine perception, training data, and aesthetic selection co-produce the image that is seen. Essentially, the artwork is not the vision of a human

artist by themselves, and their experience painted (or in this case generated) onto a canvas but the experience of the neural network used, and this experience then transposed onto a canvas, or a digital canvas if you will, with the guidance of a human artist. The experience I am talking about here is the training data set that the neural network has been trained on. The network would have used the particular data it was given as a basis to generate the portrait. Thus, this very data becomes something we can see as a form of 'inspiration' for the neural network. In the same way that a Van Gogh, or a Mark Rothko, or a Rembrandt painting might inspire a human painter to paint something in its likeness, the particular paintings used to train the neural network, and so the particular training data set, would function as a similar sort of inspiration for the neural network. This is admittedly anthropomorphising the neural network in a way but it still serves as a useful comparison to explain exactly what the user is seeing. At the same time, this also means that the neural network, or tool used, is in a way a co-author as previously discussed, which means that the human 'artist', is not doing 100% of the work. Apart from the question of authorship that we explored earlier, the other point that needs to be mentioned is that examples such as these bring up the question of deskilling.

This isn't meant as a critique of the artists or to take away the artistry from the artist, but who truly makes the art is one question, and do people retain the necessary skills to make said art valuable is another question. The rise of generative artificial intelligence has revived a longstanding philosophical anxiety: that machines slowly erode the capacities that make certain forms of skill meaningful. In the case of GenAI, this anxiety is usually described as deskilling, by which what is meant is the weakening, displacement, or disappearance of abilities that were once acquired through sustained practice (Brophy, 2008; Zuboff, 1988; Noble, 2018; Appiah, 2025). If we accept the idea that GenAI leads to deskilling then we can say that one of the ways that it mediates human-world and human-technology relations is that by constant reliance on these technologies, our skills decrease and therefore we have less skills, assumably as a collective, in this world, compares to previously, and I suppose in a way become even more reliant on these technologies. From an ANT perspective, deskilling can be read as an emergent property of an arrangement of actors. The relevant actants include the model itself, the interface through which it is used, the platforms that distribute it, the markets that reward speed over craft, the educational institutions that may no longer insist on apprenticeship, and the users who adapt their practices to what the system makes easiest. This means that GenAI does not remove skill in an abstract sense; rather, it redistributes labour within a socio-technical network. A writer who relies heavily on AI to draft, revise, and style their work may appear to retain authorship, but the network has changed: the machine now performs tasks that were once central to learning how to write. A designer who uses image generators may still make aesthetic decisions, yet the long process of sketching, failing, revising, and developing visual intuition can become less central to the practice. This is how an ANT analysis can explain this loss of skill. Furthermore, I am purposefully not mentioning any artist by name here to avoid being critical of their work and will instead remain abstract in my examples for the remainder of this chapter.

To support this, Peter-Paul Verbeek's postphenomenological work shows that technologies can mediate action as they actively shape moral and perceptual relations. By his account, technologies script behaviour by inviting some actions and discouraging others, while also shaping how the world appears to users (Rosenberger and Verbeek, 2015). In the case of GenAI, many AI systems are designed to make the fastest path also the easiest path. The interface invites a user to generate rather than struggle, to prompt rather than draft. I mean this from an artistic perspective. Using DALL-E invites me to generate an image using a verbal prompt, not to design the image myself. The point is, and the point that has been made before, is that technologies are not neutral, they influence what people do and how they relate to their own capacities. In the case of deskilling, this means that the design of a technological system matters ethically because it can either support human learning or quietly replace it. In the case of GenAI it seems self-evident that the technology at hand seems to replace or displace skills and therefore lead to deskilling, or the loss of certain skills. If painters stop longer practice painting and rely on DALL-E, it is safe to assume that painting as a skill will not be as prominent as it is now.

This issue could become especially concrete in the arts. Artistic skill, like most skills, are cultivated through constant practice. GenAI can intervene in these practices by making it possible to bypass intermediate stages that are normally formative. The result is a faster production, sure, but also a redefinition of what counts as artistic competence. The fact that many artists have had

their work used to train systems that then imitate or approximate their styles without compensation could be seen as a cultural problem as it can create a situation in which the outputs of years of practice are absorbed into a machine system that can reproduce surface effects without participating in the embodied history from which those effects emerged.

On the other hand, it would be unfair to say that GenAI only destroys skill. Some artists and writers use it as a collaborative or exploratory tool, and in some cases it can open new aesthetic possibilities. Kwame Anthony Appiah has argued that the loss of some forms of skill should not automatically be treated as a catastrophe. Human beings have always externalised abilities into tools, from writing to arithmetic to digital search. One might say that the history of culture is partly a history of strategic offloading. What matters, on this view, is not whether some old skill disappears, but whether new forms of intelligence, judgment, or creativity emerge in its place (Appiah, 2025). This is the primary objection to the deskilling panic, and it should not be dismissed. Many people do not need to memorise maps because they have navigation software; many do not need to hand-write first drafts in order to produce meaningful prose. In some contexts, automation does remove drudgery and create room for higher-level reflection. Going back to the example I gave regarding photography, to take a great photograph requires a different set of skills than to paint a great portrait. Thus, to generate a great portrait will also take a different set of skills than to paint a great portrait, especially if eventually generated images develop an aesthetic criteria for normative judgement regarding their distinct style of their own.

However, in the case of GenAI the potential problem with this counterargument is that while tasks are being offloaded and the skills change, the result is that this rewards those who control the infrastructures rather than those who practice the craft. ANT helps reveal this asymmetry because it forces us to follow the network all the way down to its material and institutional bases. Who owns the model? Who trains on whose work? What are the consequences of a platform that can generate images at near-zero marginal cost?

The work of Verbeek can push this further by arguing that design itself is a site of moral responsibility. If GenAI systems are designed in ways that produce outputs, they may become instruments of deskilling. If they are designed as mediating tools that support transparent reasoning, iterative learning, and user reflection, then they can participate in a more constructive form of human–technology relation. In this sense, the solution maybe is not more regulation or more protective laws but a movement towards the deliberate embedding of values into the models and infrastructures themselves. This is especially important for educational and artistic contexts.

To conclude, Klingemann’s work can be used as an example that shows that GenAI mediates both the cultural status of the artwork, or GenAI artworks in general, and also the conditions under which artistic practice itself is understood, valued, and learned, which are all encompassed under the issue of deskilling. The case shifts the argument from authorship toward deskilling and changing cultural norms, while also showing how repeated exposure to generative images can gradually normalise new aesthetic expectations and make AI-mediated art more legible within the wider art world. This therefore leads us into the next section, which broadens the discussion from individual case studies to the wider problems raised by GenAI. The following section will consider a range of different issues and will ask what the two case studies together reveal about the broader consequences of GenAI for art, labour, and cultural production.

3.3 A further discussion on the problems of GenAI

Before moving onto the conclusion I want to quickly address other problems or question surrounding GenAI and that could equally arise from these case studies that also show a different side to how GenAI mediates relations and explore what the theoretical framework would have to say about them. The rise of GenAI has introduced a new family of tools into cultural life. We earlier went over issues surrounding authorship. The legal disputes around training-data use show how contentious the issue has become. It is not just Obvious’ case that ran into legal disputes, in *Andersen v. Stability AI*, artists including Sarah Andersen, Kelly McKernan, and Karla Ortiz argued that AI image companies used copyrighted works to train image-generation systems without permission (Williams, 2024). The case demonstrates that the question of copyright has entered the legal sphere and carries with it serious consequences that can affect the livelihood of certain

people. When an artist's work is absorbed into a training set and used to generate new outputs, the relation between original and derivative is not exactly straightforward. ANT could say that this is a network-level transformation, meaning that images are detached from the institutional and personal settings that once gave them meaning and are re-enrolled in another system of circulation. The original artist's control is weakened because the network reassigns its function as a training set as opposed to a stand alone original work.

A second problem of GenAI is that generative systems can be used for things like fraud, deepfakes, sexual exploitation, and the production of content that implies endorsement where none exists. For example, there is damage that can be caused when an artist's style becomes indistinguishable from a machine-generated imitation, or when synthetic content is falsely associated with them. From an ANT frame, reputation is not a property of a person alone; it is a network effect, sustained by institutions, platforms, audiences, and symbolic recognition. GenAI destabilises that network by allowing styles, faces, and voices to be copied at scale. The result is that trust in images and texts weakens, and the moral economy of authorship becomes harder to maintain. The issue here is that moral harm can come from the design of a system that makes misrepresentation cheap, fast, and scalable. This is one reason why artists have responded so strongly (Ghosh and Fossas, 2022). When a recognisable style can be one of the main ways an artist earns commissions, develops a following, and secures cultural legitimacy. The more easily style can be simulated, the more fragile that form of identity can become.

A third problem is the tendency toward homogenisation. UNESCO has issued a warning that GenAI is already reshaping cultural markets, with projected revenue losses for creators and concerns about cultural sovereignty, pluralism, and equitable economies (Tsimitakis, 2026). In UNESCO's recent reports, the emphasis is that algorithmic systems (which include things like LLMs and GANs) may intensify platform dominance and narrow the range of visible cultural production. An ANT perspective sharpens this concern by showing how models are trained on datasets that disproportionately reflect dominant cultural forms, especially those that are already highly visible online. In *Facing Gaia*, Latour insists that the political order must be reassembled to include nonhuman actors and long-term material consequences (Latour, 2017). GenAI systems, by standardising creative production around the most statistically probable outputs, risk diminishing the irregularity, friction, and local specificity that make cultural ecosystems diverse. An ethical response could therefore be something like to reconstruct the collectives through which cultural diversity is maintained such as funding structures, training environments, public institutions, and platform regulations.

The problem is intensified by market concentration and platform gatekeeping. GenAI is being built, owned, and distributed by a small number of powerful firms that control the infrastructure, the models, the interfaces, and the terms of access. This means that creative labor increasingly depends on proprietary platforms whose rules are set by corporate actors such as Anthropic and OpenAI. The platform in question here is an assemblage of different codes, capital, servers, legal agreements, moderation policies, and user interfaces. These actants jointly shape what can be created, who can afford to create it, and whose output gets visibility. Market concentration therefore dictates where the platform and thus its options go.

If we accept that the rise of GenAI will lead to shifting cultural norms, such as accepting that GenAI can be valued as art, valued at upwards of \$400,000 as is the case of Edmond de Belamy, then the more creative power is concentrated in a few places, the less capacity there is for public deliberation over the shape of cultural futures. In his later work, Latour repeatedly argues that political representation must be rethought in light of the agencies that modern institutions have ignored (Latour, 2004; Latour, 2005; Latour, 2017). GenAI's gatekeeping makes the issue relevant in his argument because platform owners can determine which models are trained, which users are privileged, and what counts as acceptable content. This gives them a quasi-sovereign role over cultural production.

A plausible counterargument is that GenAI encourages and increases creativity by lowering barriers to entry. Many people who cannot draw or write at a professional level can now experiment with expressive forms. Artists can also use GenAI productively as a sketching or drawing tool. However, the fact that more people can generate content does not mean more people can meaningfully participate in shaping the system. Another counterargument is that all

technologies produce loss, and artists simply need to adapt. This too contains a truth. Photography transformed painting, digital editing transformed photography, and desktop publishing transformed print culture. However, simply because something new comes along that needs adapting to, it does not necessarily mean that the new tool should be accepted or is ethical in the first place. Furthermore, this does not escape the previously made point that it still risks power to be concentrated in very specific places, Ghost and Fossas make the very same point when they conclude:

“Certainly, an argument can be made that AI text to image models are just another tool in the artists’ arsenal, however, reality is much more nuanced. Most of the models...are owned and maintained by private companies, and these entities charge for model usage. It costs \$15 to buy 115 credits on DALL.E... This essentially means that not only are the artists, whose original works were scraped without permission to train these models, not getting profits from this business model, but under the garb of democratization of art, the middlemen, who in this case are companies like OpenAI or Midjourney, are siphoning off money that would have otherwise gone to original art creators. In plainer terms, these companies are profiting off the artistic endeavors of others, and possibly supplanting those human artisans.” (Ghosh and Fossas, 2022).

GenAI, then, should be understood as two sided. On one side it changes how meaning is controlled, how reputation is sustained, how diversity is preserved, and how value is distributed. It can certainly be useful, and it will likely remain part of artistic and cultural life. Nevertheless, the questions remains of what kinds of associations are being stabilised, which actors are being made more powerful, and which forms of life are being quietly rendered less possible.

Conclusion

The preceding chapters have argued that Generative AI art is best understood as a field of mediated relations in which human and non-human actants co-produce visual and literary works. Once this relational structure is made visible, several debates surrounding GenAI art can be considered in view of the established framework. The questions of copyright and deskilling are the first that come to mind, among others. This chapter addresses some of these debates through the postphenomenological and ANT framework established earlier.

A postphenomenological analysis shows that the real issue is not whether AI stands outside human creativity and threatens it from the exterior, but how specific technologies mediate, redistribute, and sometimes obscure the relations through which creativity, ownership, and responsibility are constituted. ANT further adds to this point by showing that what appears as a single act such as prompting, is in fact the outcome of networks of translation involving datasets, interfaces, infrastructures, institutions, and markets. The task, then, is to explore what conclusions can be drawn from all this.

Limitations and Future Research

A limitation of this thesis is that its analytical scope has been fairly selective, focusing on a small number of exemplary cases rather than on the full breadth of generative AI practice. This means that some important domains like commercial design, platform-based image generation, and rapidly changing multimodal systems are not addressed. Similarly, text-based generative cases such as those from Large Language Models or cases of generative music are also not covered. Future research could extend this framework through comparative studies across different genres, platforms, and cultural contexts, as well as through closer attention to the institutional, economic, and technical conditions that shape generative output. It would also be fruitful to examine how these mediations develop over time as generative systems become more embedded in artistic training, curation, and reception, thereby testing whether the relations identified here remain stable or require revision.

In addition, the thesis has focused primarily on systems understood through postphenomenology and actor-network theory, which has made it possible to foreground mediation and distributed agency, but it has not fully engaged with other frameworks that could have sharpened the account, such as different economic theories, legal studies, critical race theory, feminist STS, or political economy. Future research could apply a different framework or lense onto GenAI art and produce different conclusions or at least approach the question from a different way and potentially lead to new conversations.

There are also limitations stemming from the rapid pace of technological change. Generative systems are evolving so quickly that any philosophical account risks becoming temporally bounded, describing a set of relations that may already be shifting by the time the thesis is completed. This is especially true in relation to multimodal models, agentic systems, and platform-integrated tools, which this thesis has not mentioned but seem to be becoming more and more common, and which may alter the human-artwork relation in ways that go beyond the cases discussed here. Likewise, the thesis does not fully address the legal and regulatory dimensions of generative AI, including the legal frameworks of copyright, consent, dataset provenance, and institutional policy, except insofar as they intersect with the philosophical problem of mediation. Nor does it provide a comprehensive empirical study of how artists, audiences, or institutions actually experience these systems in practice, since its main concern is conceptual rather than ethnographic or statistical.

A further area left outside the scope of the present thesis is the global and geopolitical dimension of generative AI. Questions concerning the extraction of labour and data, unequal access to computational resources, and the cultural dominance of certain linguistic and aesthetic norms are highly relevant to the subject, but they would require a broader analysis than the one undertaken here. Future research could therefore build on this thesis by combining philosophical analysis with empirical methods, cross-cultural comparison, and political critique. Such work would test the

applicability of the framework developed here and also clarify how generative AI is being shaped by broader structures of power, inequality, and cultural mediation.

A further limitation is that the thesis has treated its case studies primarily as philosophical examples rather than as fully historicised art objects with distinct exhibition histories, reception contexts, and curatorial framings. While this was necessary in order to maintain conceptual clarity and keep the analysis focused on mediation, it also means that the historical specificity of each artwork has not been explored to the degree it could be. Future research could therefore supplement the framework developed here with more detailed art-historical and curatorial analysis, paying closer attention to the reception of specific works across different publics, institutions, and time periods. This would help clarify not only how GenAI mediates human-art relations in principle, but how those mediations are stabilised, contested, or transformed within concrete art worlds. Equally, future work could benefit from interviews with artists, curators, and audiences in order to test whether the conceptual relations identified here correspond with lived experience in practice. Such an expansion would not replace the theoretical account developed here, but rather strengthen it by situating mediation more explicitly within concrete institutional and social settings.

Conclusion

The thesis has argued that generative AI should be seen as a mediating technology that influences the relation between human and art. By bringing postphenomenology and actor-network theory into the dialogue, it has served as an exploratory thesis to show how generative systems actively shape relations while simultaneously participating in the production of the works to which they give rise, and what are the potential consequences of this, analysed through the established framework. In this sense, the central contribution of the thesis has been to shift the discussion from the question of whether AI-generated works are ‘really art’ to the more fundamental question of how such works mediate human engagement with art, creativity, and value.

The theoretical framework developed in the opening chapters clarified why this shift is necessary. Postphenomenology, especially in the work of Ihde and Verbeek, enable an account of how technologies structure experience. Actor-network theory, by contrast, made visible the distributed character of cultural production by tracing the role of datasets, interfaces, infrastructures, institutions, and users in the stabilisation of generative outputs. Taken together, these frameworks reject the assumption that agency belongs exclusively to a sovereign human subject. They instead show that artistic meaning emerges through relations that are relational, contingent, and often opaque.

The main claim of this thesis is that GenAI changes art changes the terms on which art is encountered, interpreted, and produced. When generative systems enter artistic practice they affect what can be seen, what kinds of decisions become available to artists, how viewers interpret authorship, and how value becomes attached to a work. This is why the concept of mediation has been so central throughout the thesis. It offers a way of understanding GenAI not as an external threat to art, nor as a neutral enhancement of it, but as a force that reorganises the relations through which art is made legible and significant.

This has been especially clear in relation to the philosophical account of images developed in Chapter 2. There, the thesis argued that images shape things like attention, comparison, and selfhood. In digital environments, images already mediate how subjects see themselves and others, often in ways that become tied to self-worth, and vulnerability. Generative AI intensifies this condition by making images easier to produce, easier to personalise, and less tethered to any determinate physical reality. The result is images become more synthetic, and more adaptable, while still retaining psychological and social force. This matters for art because visual artworks do not exist apart from broader image cultures. They are encountered by subjects already shaped by digital mediation, and GenAI inserts itself directly into that field.

One of the implications of the thesis is that debates about authorship need to be reframed. Much of the public debate still asks whether the human or the machine should count as the real author, but the analyses undertaken here suggest that this binary is inadequate. Once a work is

understood as the work of multiple actants in a network, the search for a single author appears to be futile. This means that intention and responsibility must be located within distributed arrangements rather than isolated individuals alone. The value of ANT in this regard has been to show that a work can be irreducibly networked without becoming meaningless or authorless in a trivial sense. Instead, authorship becomes something more relational, more negotiated, and more dependent on infrastructures than the traditional humanist model allows. A similar point can be made about the issue of deskilling. The analysis has shown that GenAI redistributes skill, often in ways that make some capacities less central while elevating others. At the same time, when generative systems are owned and controlled by a small number of powerful firms, the transition from one set of skills to another can benefit platforms and infrastructure owners more than practitioners. Thus, the questions become what new possibilities GenAI opens, and also whose practices are displaced, whose labour remains invisible, and whose agency is amplified by the design of the system.

Finally, the broader significance of this study lies in its attempt to redirect the philosophical discussion of GenAI away from binaries and toward a more relational vocabulary. My point is to ask a new question. While questions and issues such as whether machines can replace artists, or whether AI-generated objects deserve the label of art are equally relevant in philosophy, I also want to ask how relations between humans, images, tools, institutions, and values are being reorganised through generative systems. In turn, the follow up question is what does that reorganisation have in terms of consequences for things like authorship, embodiment, aesthetics, labour, cultural diversity, and responsibility amongst other things.

If there is one final conclusion to be drawn, it is that generative AI art should be approached as a field of mediation whose consequences are still unfolding. My aim, and the aim of this thesis has been to try to demonstrate that philosophy can act as an exploration as to as to what these consequences can be. My aim has not been to deliver a final verdict on whether GenAI is good or bad for art, but to clarify the structures through which it changes what art can be, how it is experienced, and who or what gets to participate in its production.

Bibliography

- Akrich, M. (1992) 'The de-scription of technical objects', in Bijker, W.E. and Law, J. (eds.) *Shaping Technology/Building Society: Studies in Sociotechnical Change*. Cambridge, MA: MIT Press, pp. 205–224.
- Al-Omari, O. and Al-Omari, T. (2025) "Artificial Intelligence and Posthumanism: A Philosophical Inquiry into Consciousness, Ethics, and Human Identity", *Journal of Posthumanism*, 5(2), pp. 458–469.
- Appiah, K.A. (2025) 'The Atlantic'. The Atlantic. Available at: <https://www.theatlantic.com/ideas/archive/2025/10/ai-deskilling-automation-technology/684669/>
- Bergen, J.P. and Verbeek, P.-P. (2021) 'To-Do Is to Be: Foucault, Levinas, and Technologically Mediated Subjectivation'.
- Bernaschina, D. (2025) 'Visual Art and Artificial Intelligence: Tensions between Ethics, Aesthetics, and Authorship in the Algorithmic Age', *Contemporary Visual Culture and Art*, 1(1), pp. 68–78. doi:10.63385/cvca.v1i1.189.
- Belk, R.W. (2013) 'Extended self in a digital world', *Journal of Consumer Research*, 40(3), pp. 477–500.
- Boden, M.A. and Edmonds, E.A. (2009) 'What is generative art?', *Digital Creativity*, 20(1–2), pp. 21–46. doi:10.1080/14626260902867915.
- Bolter, J.D. and Grusin, R. (1999) *Remediation: Understanding New Media*. Cambridge, MA: MIT Press.
- Braidotti, R. (2013) *The Posthuman*. Cambridge: Polity Press.
- Brophy, E. (2008) *Language Put to Work: Materializing Ideology in the Service Economy*. Basingstoke: Palgrave Macmillan.
- Buckner, C.J. (2024) *From Deep Learning to Rational Machines: What the History of Philosophy Can Teach Us about the Future of Artificial Intelligence*. Oxford: Oxford University Press.
- Caldwell, M. (2023) 'What is an "Author"? Copyright Authorship of AI Art Through a Philosophical Lens', *Houston Law Review*, 61, pp. 411–457. Available at: <https://houstonlawreview.org/article/92132-what-is-an-author-copyright-authorship-of-ai-art-through-a-philosophical-lens>
- Cela-Conde, C.J., Agnati, L., Huston, J.P., Mora, F. and Nadal, M. (2011) 'The neural foundations of aesthetic appreciation', *Progress in Neurobiology*, 94(1), pp. 39–48.
- Cetinic, E. and She, J. (2023) 'Understanding and creating art with AI: Review and outlook', *ACM Transactions on Multimedia Computing, Communications, and Applications*, 19(3), pp. 1–22.
- Chatterjee, A. and Vartanian, O. (2014) 'Neuroaesthetics', *Trends in Cognitive Sciences*, 18(7), pp. 370–375.
- Chesney, R. and Citron, D.K. (2019) 'Deep fakes: A looming challenge for privacy, democracy, and national security', *California Law Review*, 107(6), pp. 1753–1820.
- Citron, D.K. and Franks, M.A. (2014) 'Criminalizing revenge porn', *Wake Forest Law Review*, 49(2), pp. 345–391.
- Cohen, J.E. (2024) *Between Truth and Power: The Legal Construction of AI*. Oxford: Oxford University Press.

- Crawford, K. and Paglen, T. (2019) 'Excavating AI: The Politics of Images in Machine Learning Training Sets', *AI & Society*, 34, pp. 77–91.
- Crowston, K. and Pentland, B.T. (2024) 'Deskilling and upskilling with generative AI systems'. Available at: https://crowston.syr.edu/sites/default/files/GAI_and_skills.pdf
- Davies, M. (2016) 'Material subjects and vital objects: Prefiguring property and rights for an entangled world', *Australian Journal of Human Rights*, 22(1), pp. 1–19.
- Davies, M. (2022) 'Posthuman legalities: New materialism and law beyond the human' special issue editorial. Available at: <https://orca.cardiff.ac.uk/id/eprint/147446/7/SI000EditorialFINAL.pdf> (Accessed: 29 March 2026).
- Davies, M. (2023) 'Can property be justified in an entangled world?', *Law and Critique* online. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4508581 (Accessed: 29 March 2026).
- David, G. and Cambre, C. (2016) 'Screened intimacies: Tinder and the swipe logic', *Social Media + Society*, 2(2).
- Di Dio, C., Ardizzi, M., Schieppati, S.V., Massaro, D., Gilli, G., Gallese, V. and Marchetti, A. (2023) 'Art made by artificial intelligence: The effect of authorship on aesthetic judgments', *Psychology of Aesthetics, Creativity, and the Arts*. doi:10.1037/aca0000602.
- Elad, B. (2025) 'AI in Social Media Tools Statistics 2025: Powerful Trends', *SQ Magazine*. Available at: <https://sqmagazine.co.uk/ai-in-social-media-tools-statistics/> (Accessed: 6 June 2026).
- Fakhri, M.M., Ahmar, A.S., Isma, A., Rosidah, R. and Fadhilatunisa, D. (2024) 'Exploring Generative AI Tools Frequency: Impacts on Attitude, Satisfaction, and Competency in Achieving Higher Education Learning Goals', *EduLine: Journal of Education and Learning Innovation*, 4(1), pp. 196–208. doi:10.35877/454RI.eduline2592.
- Fardouly, J. and Vartanian, L.R. (2016) 'Social media and body image concerns: Current research and future directions', *Current Opinion in Psychology*, 9, pp. 1–5.
- Flusser, V. (2000) *Towards a Philosophy of Photography*. London: Reaktion Books.
- Galanter, P. (2003) 'What is generative art? Complexity theory as a context for art theory'. Available at: http://www.philipgalanter.com/downloads/ga2003_paper.pdf
- Gershon, L. (2022) 'Did Photography Really Kill Portrait Painting?', *JSTOR Daily*. Available at: <https://daily.jstor.org/did-photography-really-kill-portrait-painting/> (Accessed: 6 June 2026).
- Ghosh, A. and Fossas, G. (2022). *Can There be Art Without an Artist?*
- Gibson, J.J. (1979) *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.
- Gillespie, N., Lockey, S., Ward, T., Macdade, A. and Hassed, G. (2025) *Trust, attitudes and use of artificial intelligence: A global study 2025*. The University of Melbourne and KPMG. doi:10.26188/28822919.
- Ginsburg, J.C. and Budiardjo, L.A. (2023) 'Authors and machines', *Columbia Journal of Law & the Arts*, 46(2), pp. 1–45.
- Hansen, M.B.N. (2015) *Feed-Forward: On the Future of Twenty-First-Century Media*. Chicago: University of Chicago Press.
- Haraway, D. (1991) 'A Cyborg Manifesto: Science, Technology, and Socialist-Feminism in the Late Twentieth Century', in Haraway, D. *Simians, Cyborgs, and Women: The Reinvention of Nature*. New York: Routledge, pp. 149–181.

He, Y. (2026) 'Ethical dilemmas in artificial intelligence-generated art: authorship, ownership, and the blurring of creative boundaries', *Digital Scholarship in the Humanities*. Available at: <https://doi.org/10.1093/llc/fqag035>

Hearn, A. (2008) 'Meat, mask, burden: Probing the contours of the branded "self"', *Journal of Consumer Culture*, 8(2), pp. 197–217.

Henry, N., Flynn, A. and Powell, A. (2018) 'Image-based sexual abuse: A study on the causes and consequences of non-consensual nude image sharing', *Violence Against Women*, 24(13), pp. 1576–1596.

Hertzmann, A. (2018) 'Can computers create art?', *Arts*, 7(2), 18.

Ihde, D. (1990) *Technology and the Lifeworld: From Garden to Earth*. Bloomington: Indiana University Press.

Illouz, E. (1997) *Consuming the Romantic Utopia: Love and the Cultural Contradictions of Capitalism*. Berkeley: University of California Press.

Kalpokiene, J. and Kalpokas, I. (2022) 'Creative encounters of a posthuman kind – anthropocentric law, artificial intelligence, and art'. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4264704 (Accessed: 29 March 2026).

Karttunen, S. (2025). Briefing Copyright of AI-generated works: Approaches in the EU and beyond. [online] Available at: [https://www.europarl.europa.eu/RegData/etudes/BRIE/2025/782585/EPRS_BRI\(2025\)782585_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/BRIE/2025/782585/EPRS_BRI(2025)782585_EN.pdf).

Kurki, V.A.J. (2019) *A Theory of Legal Personhood*. Oxford: Oxford University Press.

Landes, W.M. and Posner, R.A. (1989) 'An economic analysis of copyright law', *Journal of Legal Studies*, 18(2), pp. 325–363.

Lanzing, M. (2018) "Strongly Recommended" Revisiting Decisional Privacy to Judge Hypernudging in Self-Tracking Technologies. *Philos. Technol.* **32**, 549–568 (2019). <https://doi.org/10.1007/s13347-018-0316-4>

Latour, B. (1987) *Science in Action: How to Follow Scientists and Engineers through Society*. Cambridge, MA: Harvard University Press.

Latour, B. (1992) *We Have Never Been Modern*. Cambridge, MA: Harvard University Press.

Latour, B. (1996) 'On actor-network theory: A few clarifications'. Available at: <http://www.bruno-latour.fr/sites/default/files/P-67%20ACTOR-NETWORK.pdf> (Accessed: 29 March 2026).

Latour, B. (1999) *Pandora's Hope: Essays on the Reality of Science Studies*. Cambridge, MA: Harvard University Press.

Latour, B. (2004) *Politics of Nature: How to Bring the Sciences into Democracy*. Cambridge, MA: Harvard University Press.

Latour, B. (2005) *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford: Oxford University Press.

Latour, B. (2017) *Facing Gaia: Eight Lectures on the New Climatic Regime*. Cambridge: Polity.

Law, J. (1992) 'Notes on the theory of the actor-network: Ordering, strategy, and heterogeneity', *Systemic Practice and Action Research*, 5(4), pp. 379–393.

- Linardon, J. et al. (2025) 'The association between social comparison in social media, body image concerns and eating disorder symptoms: A systematic review and meta-analysis', *Body Image*, 52, 101841. doi:10.1016/j.bodyim.2024.101841.
- Lupton, D. (2016) *The Quantified Self*. Cambridge: Polity.
- McCarthy, J. and Yoon, K. (2023) 'Generative AI and copyright: principles, priorities and practicalities', *Journal of Intellectual Property Law & Practice*, 18(12).
- Manovich, L. and Arielli, E. (2024) *Artificial Aesthetics: Generative AI, Art and Visual Media*.
- Marwick, A.E. (2013) *Status Update: Celebrity, Publicity, and Branding in the Social Media Age*. New Haven: Yale University Press.
- Merges, R.P. (1989) 'Are you making fun of me?: Notes on the copyright incentive theory', *Columbia Law Review*, 89(2), pp. 313–320.
- Morris, R. (2025) 'AI helped cause Hollywood strikes. Now it's in Oscar-winning films', *BBC News*. Available at: <https://www.bbc.com/news/articles/ce303x19dwgo> (Accessed: 6 June 2026).
- Murdoch, I. (1992) 'Conceptions of Unity: Art', in *Metaphysics as a Guide to Morals*. London: Vintage, pp. 1–24.
- Noble, S.U. (2018) *Algorithms of Oppression*. New York: New York University Press.
- O'Brien, D.P. (2017) *Technological embodiment and haptic narrative: postphenomenology in cinema, interactive art and computer gaming*.
- Ozimek, P., Lainas, S., Bierhoff, H.-W. and Rohmann, E. (2023) 'How photo editing in social media shapes self-perceived attractiveness and self-esteem via self-objectification and physical appearance comparisons', *BMC Psychology*, 11(1). doi:10.1186/s40359-023-01143-0.
- Rosenberger, R. and Verbeek, P.-P. (eds.) (2015) *Postphenomenological Investigations: Essays on Human-Technology Relations*. Lanham, MD: Lexington Books.
- Rui, Z. and Gu, Z. (2021) 'A Review of EEG and fMRI Measuring Aesthetic Processing in Visual User Experience Research', *Computational Intelligence and Neuroscience*, 2021, Article ID 2070209. doi:10.1155/2021/2070209.
- Tsimitakis, M. (2026) 'AI set to slash creators' revenues by up to 24% by 2028, UNESCO warns', *Creatives Unite*. Available at: <https://creativesunite.eu/article/ai-set-to-slash-creators-revenues-by-up-to-by-unesco-warns> (Accessed: 6 June 2026).
- Turkle, S. (1995) *Life on the Screen: Identity in the Age of the Internet*. New York: Simon & Schuster.
- UNESCO (n.d.) *ReShaping Policies for Creativity and related reports on AI, cultural diversity, and creative economies*.
- Valkenburg, P.M., Peter, J. and Schouten, A.P. (2006) 'Friend networking sites and their relationship to adolescents' well-being and social self-esteem', *CyberPsychology & Behavior*, 9(5), pp. 584–590.
- Verbeek, P.-P. (2005) *What Things Do: Philosophical Reflections on Technology, Agency, and Design*. University Park, PA: Pennsylvania State University Press.
- Verbeek, P.-P. (2011) *Moralizing Technology: Understanding and Designing the Morality of Things*. Chicago, IL: University of Chicago Press.

Vessel, E.A., Starr, G.G. and Rubin, N. (2012) 'The brain on art: intense aesthetic experience activates the default mode network', *Frontiers in Human Neuroscience*, 6, 66. doi:10.3389/fnhum.2012.00066.

Vogel, E.A., Rose, J.P., Roberts, L.R. and Eckles, K. (2014) 'Social comparison, social media, and self-esteem', *Psychology of Popular Media Culture*, 3(4), pp. 206–222.

Ward, J. (2017) 'What are you doing on Tinder? Impression management on a matchmaking mobile app', *Information, Communication & Society*, 20(11), pp. 1644–1659.

Williams, S. (2024) 'AI and Artists' IP: Exploring Copyright Infringement Allegations in *Andersen v. Stability AI Ltd.*', Center for Art Law. Available at: <https://itsartlaw.org/art-law/artificial-intelligence-and-artists-intellectual-property-unpacking-copyright-infringement-allegations-in-andersen-v-stability-ai-ltd/> (Accessed: 6 June 2026).

Yeh, S.N. and Siah, C.-J.R. (2025) 'Students' attitudes toward LLMs and its association with metacognitive abilities: A cross-sectional study', *Nurse Education in Practice*, 88, 104567. doi:10.1016/j.nepr.2025.104567.

Zarevich, E. (2023) 'The Artists Who Hated the Eiffel Tower', *JSTOR Daily*. Available at: <https://daily.jstor.org/the-artists-who-hated-the-eiffel-tower/> (Accessed: 6 June 2026).

Zeilinger, M. (2021) *Generative Adversarial Copy Machines*. Available at: <https://culturemachine.net/wp-content/uploads/2021/09/Martin-Zeilinger.pdf>

Zuboff, S. (1988) *In the Age of the Smart Machine: The Future of Work and Power*. New York: Basic Books.

Zylinska, J. (2020) 'CG-Art: demystifying the anthropocentric bias of artistic creativity', *AI & Society*. Available at: <https://www.tandfonline.com/doi/full/10.1080/09540091.2020.1741514>