To cite this article:
Marco Donnarumma. Against the Norm: Othering and Otherness in AI Aesthetics. *Journal of Digital Culture & Society*, vol. 8, no. 2, 2022, pp. 39-66. https://doi.org/10.14361/dcs-2022-080205.

Please note: some parts of the last section of this personal copy differ from the published version of the article. Here some information has been added or updated for accuracy.

Against the Norm:
Othering and Otherness in AI Aesthetics

Marco Donnarumma
Intelligent Instruments Lab, Iceland University of the Arts, Rejkyavik
m@marcodonnarumma.com | https://marcodonnarumma.com

What are the links between science, aesthetics and episteme sustaining the large deep learning models known as AI image generators, such as Dall-e, StableDiffusion and Midjourney? Here I argue that these software systems are an expression of the normative power wielded by the industrial-cultural complex that funds, produces and disseminates them. This form of expression is what I call *corporate AI aesthetics* or, in short, AI aesthetics. My argumentation begins by defining the particular kind of aesthetics that AI image generators yield and then journeys through the enforcement of neoliberal knowledge they contribute to. Through the lens of cultural criticism and anticolonial scholarship, I try to develop a critique of AI aesthetics as soft propaganda for the Global North. In this sense, AI aesthetics disseminates a techno-deterministic view where anything that is or can be made countable, like human creativity, can be predicted, hence, controlled. To this end, I discuss the cultural and technical production of AI artefacts, emphasising its dependence on an abstraction of labour, the reinforcement of biases in visual culture and

the deceit of art market speculations. The discussion leads to observe what AI aesthetics detracts attention from: an understanding of artistic intention as a form of *collective otherness*; that is, artistic intention not as the will of an inspired individual or a powerful AI system, but as the flow of relations among human and non-humans, existing across generations, cultures and geographies. Such a relational ecosystem becomes manifest in the analysis of two artworks, by Jonathan Chaim Reus and artist group AIseeds Project, that are representative of an affirmative, critical and culturally situated approach to deep learning.

Keywords: AI art, aesthetics, capitalism and art market, deep learning, othering, otherness

Aesthetics of Power

This text is concerned with an analysis of some components of what I call *corporate AI aesthetics*. For brevity and clarity, I will refer to it as AI aesthetics. This is a widespread and dominating visual aesthetics that has been popularised by the major players in the artificial intelligence and deep learning business, US-based corporations Google, Nvidia, Meta, OpenAI, and more recently by smaller actors, such as the US company Midjourney and the British Stability AI.[i] AI aesthetics is highly recognizable and lies at the core of what some identify as 'AI Art', the production of images by means of very large algorithmic systems, called AI image generators. Many generators exists today, but their core theory and implementation is tied to a handful of corporate and independent organizations with access to large financial and computational resources, sometimes in collaboration with university research groups that receive funds from them. Here, my focus will be on what kind of artistic expression AI aesthetics yields and how that expression exists in relation with complexes of what turns out to be normative power. In order to tackle these questions, I will discuss how AI aesthetics is produced technically *and* culturally. This will mean situating the technical operations of AI image generators within the cultural-industrial milieu where they thrive: an assemblage of industry business plans, art market operations and technocratic ideologies that, through deep learning, produces a particular type of reality. I will refer to this milieu as the *AI complex*.

Why the need to link aesthetics and power? Some may argue that deep learning is a technological breakthrough and AI art is proof of the novelty of such a computational paradigm. Others would argue that deep learning image generators may be a lightning trend, fated to fade out of public view (just like other unfortunate inventions) or to transmogrify into commercial video plug-ins. I in part agree with both views. Yet, a careful look reveals a more complex picture: from profiling, microtargeting, media recommendation and predictive policing to finance, climate modelling, autonomous weapons and art, deep and probabilistic learning technologies have come to structure and sustain contemporary capitalist societies. While they also made possible important technological breakthroughs, for example in the study of climate change and biodiversity loss, the widespread adoption of deep learning across disciplines has accelerated a drastic societal change. It has materialised in a technocratic ideology of prediction as a scaffold of political, social, moral and cultural life. By looking at the controversies surrounding deep learning and prediction methods it shall be evident how problematic and dangerous this approach can be; think of the role of Cambridge Analytica in both the Leave.EU campaign for Brexit (UK Parliament 2022) and the Trump US election campaign in 2016 (Hu 2020); the entanglement of US military and Google in Project Maven, where Google's machine learning library, TensorFlow, was used to enhance warfare drones and analyse surveillance data (Hoijtink/Planqué-Van Hardeveld 2022); the capacity of forecasting algorithms and agent-based systems to destabilise already volatile financial markets – which became apparent in the flash crash of 2010 (Sornette/Von der Becke, 2011: 15) and is still being debated since then (Blyth 2018; Olorunnimbe/Viktor 2022); the automated exploitation of labour from Amazon and Netflix to Uber, Spotify and Airbnb (Casilli/Posada 2019); and the negative, even deadly, psychological impact of Meta's Instagram on children (Yearwood 2022; Milmo 2022). Acknowledging AI art as deeply embedded in the ramified, capital-driven regime of deep learning begs the question: what relation exists between the aesthetics of deep learning imagery and the techno-deterministic drive of the AI complex?

The AI complex can be understood as an intricate configuration of different powers; it entangles corporations, rich individuals, universities, governments, culture and people in interacting feedback loops. The AI complex does not only create technologies; it engineers platforms to deploy those technologies and then enthrals users to inhabit them. Platforms have rules, encourage particular ways of being and modes of understanding. Because these

loops involve large networks of people, what emerges is the normalization of one particular world view and the obfuscation of those differing from it. What becomes normalized is a view that posits probabilistic computation and the platforms through which it inhabits the world as key to predict – hence, control – any type of phenomena, from elections to shopping and from self-image to creativity. Deep learning and Big Data are instrumental to the process. In this article, I argue that the very functioning of AI art generators is determined by the power dynamics of the AI complex and, therefore, the kind of aesthetics that AI art expresses is submitted to those powers. Ultimately, it *is* an expression of those powers. Thinking, with Sylvia Wynter (2015), about the link of science, aesthetics and episteme: as long as these types of AI systems remain bound to the AI complex, the aesthetics they afford will continue enacting a sublimation of the current order of knowledge, an order necessary to the "present neoliberal/neo-imperial [...] global order of worlds and things" (ibid: 30). I will come back to this in the penultimate section of the article.

It should be noted that I write this text from the perspective of a media and performance artist, programmer and scholar who has worked with AI, among other technologies, for the last ten years of his career, usually in collaboration with scientists and research laboratories and always with open source systems.[ii] I do not despise AI technologies - quite on the contrary, I am continuously intrigued by their aesthetic and epistemic potential (Caramiaux/Donnarumma 2021). My critique harbours, at its core, an affirmative wish for AI technologies and their cultural instantiations to be increasingly and critically shaped by artists. It is my hope that by contemplating the links between aesthetics and power in AI art, I can formulate that objective with a stimulating sense of urgency. The critique that follows is, therefore, not directly addressed at artists, users or researchers. Many artists currently use deep learning tools, and the aesthetics of their work range widely in terms of quality and expression. An enquiry into those works would require a different framing from the one I adopt here. In closing the article, I will, however, enter into a generative dialogue with two particular artworks that I see representative of an affirmative, critical and culturally situated approach to deep learning. This dialogue will show some ways in which artistic uses of deep learning can encourage artists and public to explore alternative ways of thinking about artificial intelligence, ways that discard capital-driven norms of creativity and thus question the operating modes and platforms of the AI complex. To release such

potential, I suggest, these technologies should be used according to an old tenet of media art practice that seems to have lost currency lately: to 'pervert technological correctness' (Lozano-Hemmer 1996), that is, to enact interventions within the guts of a machine, to creatively feed on criticism about technological novelty, speed and beauty and to turn technological limits and flaws into aesthetic strategies.

AI is not a Thing

Before diving into my enquiry, it is helpful to clarify some issues of terminology and navigate a brief historical overview of the field. Today's most advanced AI systems are still far from achieving a general intelligence. In this sense, AI is less a definite concept fully manifested in a machine and more a general idiom that signifies many things at once according to the context where it is spoken. The most precise way to think of AI is as a discipline of study focusing on how machines, software and hardware, can 'learn' particular tasks. Despite its ubiquity, the term 'learning' is strictly a euphemism; a machine does not learn in the conventional sense of the term, it does not craft a skill *and* gather transferable knowledge by being embedded in the world and in relation with others.[iii] An AI system finds recurrent patterns in data, which allows it to classify items according to categories or to map the image of a dog to a label that reads 'dog'. Rather than learning, more appropriate terms would be pattern matching, feature mapping or data averaging. This leads us to another issue of terminology, the very label artificial intelligence. This, arguably, is an ill-chosen name, for it is haunted by the reductionist and anthropocentric notion of intelligence that was predominant in the 1950s, when the term began being used (McCarthy et al. 1955). Today, thanks to a broader and interdisciplinary scope, scientific understanding of intelligence provides insight into the varied forms in which it is manifested across plants, trees, insects, invertebrates and other living beings. While the Western use of the idiom 'AI' is largely still alien to this knowledge, Indigenous epistemologies offer serious alternate perspectives (Lewis et al. 2018).

The arguably reductionist understanding of intelligence at the core of much AI research can be traced along its two main study traditions: "the logic-inspired and the neural-network-inspired paradigms for cognition" (Lecun et al. 2015: 441), also known as, respectively, symbolic and connectionist approach. Each approach represents a different way of understanding how human intelligence works (Smolensky 1987) and both largely

neglect the role of embodiment in the process of learning. Symbolic AI is the classical approach and posits that humans make intelligent decisions by manipulating symbols through logic, for example: *If COLD OUTSIDE then WEAR COAT*. Symbolic AI uses large, linear sequences of logic operations to manipulate a pre-known set of symbols, and can, therefore, excel at logical inference and description as long as its field of action is precisely defined. The symbolic approach also proves to be too coarse and rigid when applied to computation with fuzzy data or unknown variables. Here, knowledge is enshrined in symbols and there is a limit to how many symbols a system can handle, especially since a symbol lacks an internal informational structure that defines its features. Once considered the best approach to reconstructing human-like intelligence, with time, research on symbolic AI lost its appeal. A particular thorny problem for researchers was that it could not fully explain how the material construction of the brain, its neurons and networks, scales up to a symbol management system.

This question was more successfully addressed by the connectionist approach. This method is directly inspired by the material arrangement of animal neural networks and their multithreaded mode of operation. That is, a neuron performs a simple and small operation, but it does so in unison with other million neurons, as part of a network where they are all connected to one another. According to the connectionist school of thought, within such a system – be it living or machinic – inference happens not through logic, but through statistical operations. For Lecun et al. (2015: 441) this is a form of "fast 'intuitive' inference that underpins effortless common sense reasoning." Such cognitive model can be represented as a non-linear system where a rational choice emerges from the multiple, simultaneous operation of groups of neurons. They do not hold intelligence in themselves, it is their number and their statistical relations that allow data to be manipulated from a granular level up to higher scales that eventually yield an intelligent decision. Artificial neural networks are systems that attempt to approximate the human brain from a strictly connectionist point of view. As such, they require a vast amount of computational power and data. For this reason, research on neural networks was abandoned in the mid 1970s, and for the same reason it was revived around 2012 by corporations such as Google, Microsoft, Amazon and Meta in the US (Sudmann 2018) and Baidu (Kai 2013) and Tencent in China (Zou 2014). Data and computational power are, in fact, their monopoly.

The symbolic and connectionist approaches generated an innumerable amount of algorithmic techniques that can be broadly grouped into supervised, unsupervised and semisupervised learning, reinforcement learning, ensemble learning, instance-based learning and neural networks. Deep learning and probabilistic learning rely on a connectionist approach, yet may borrow components from a diverse range of techniques. The two methods are interrelated and lie at the core of so called AI image generators. In my view, this is another misnomer, for, as we will see in the next section, these AI systems operate more like filters than generators. They use gargantuan neural networks containing billions of artificial neurons and parameters to guide filtering of visual noise[iv] in the shape of a coherent image. Because they follow the principles of the connectionist approach they are void of symbolic logic, manipulate numerical variables instead of symbols and excel at statistical inference from gigantic datasets. This also means that their functioning eschews a scientific ground truth or an established set of concepts. Rather, they *find* meaning within data. Establishing whether that meaning is truthful or useful is an unwelcome responsibility for those who own and use the systems. As a result, deep learning research lives in the shadow of the black-box issue, a serious problem of interpretability and ethics afflicting in particular the implementation of these systems in predictive policing and medicine.

In light of this, it should become clear that to casually refer to 'AI' is to be rather vague. Willingly or not, when spelling the letters 'AI' one refers to a specific approach to simulating human cognition, implemented by means of a particular type of machine learning algorithm. In less common but more powerfully evocative cases, some people refer to 'AI' to conjure up an immaterial, more-than-human, fictional agent that plays out disparate roles in the human imagery, from forthcoming god to existential threat, from prodigal child to innovative creator. Other times, especially in the snappy conversations saturating the online space, these two definitions may mix up to various degrees, propelled by sketchy media headlines and savvy-sounding CEO's tweets. On the other hand, scientists and researchers tend to adhere to precise naming, for there exist a myriad of different machine learning and probabilistic algorithms, each with particular affordances and limitations. Throughout this text, I will use the term *AI* to refer to deep learning algorithms or systems made out of those algorithms. When using inverted commas, as in

'AI', I will be referring to the popular definition of the technology as a myth. When speaking of a particular algorithm I will specify the name and type.

Navigating the Manifold

The emergence of corporate AI aesthetics is a result of recent research and production of particular deep learning models for natural language processing (GPT-3 by OpenAI, BERT[v] by Google), computer vision and image synthesis (Dall-e by OpenAI, Stable Diffusion by Stability AI, and Midjourney). For the sake of clarity, in this article I will focus on image generators based on probabilistic machine learning. Yet, the analysis may be useful in other case studies, for the theoretical scaffold of image-based systems is shared by other generative and predictive AI systems. In the past ten years, the perceived quality of AI-based image synthesis has increased greatly. But, possibly, what transformed a quirky way of manipulating images into a popular form of supposed 'art' is, ironically, something that has little to do with art or creativity. Recent image generators can be prompted with short text to 'generate' a matching image; and they do so by duplicating and overlapping features of existent artworks. This trick makes them appear to 'understand' semantic relations in text and to express them via images. Suddenly, an image generator model acquires a new form of alterity, moving 'AI' closer to the myth of singularity. It is less clear what art has to do with it and, hopefully, what follows can aid this line of enquiry.

In order to set up an aesthetic analysis of AI artefacts, it is important to understand how this kind of models create images. Contrary to popular opinion, a model does not create an image out of thin air; it amalgamates abstract features of existing artworks into pastiches. It is useful to recall here that image generators function according to a connectionist approach. So, rather than creating an image using logic, symbols and proto-intent as a symbolic approach would require, the generators derive images from datasets using probabilistic functions. The mechanics of this process is extremely intricate, for it depends on arbitrarily defined interactions between diverse algorithmic components. In contrast to the myth of 'AI' as a singular agent, these systems are crowded assemblages of algorithms and models that are compounded most often according to empirical experiments. Below, I describe the image generation process of a latent text-to-image diffusion model, which is the most successful and widely adopted model at the time of this writing. My description, I hope, is accessible to a general reader, while staying true to the technical operations and

dispensing with some details for sake of space. For the curious reader, I provide references to the implementation of the relevant algorithms. Analysing how these models work unveils a labyrinthine dramaturgy of modern computation, weaving together probability theory, visual culture and labour exploitation.

First is the data collection phase. Images of artworks and random content are scraped from the web together with the Alt-text that describe them (cf. Schuman et al. 2021). Combining images with captions into data pairs is crucial to these systems, for it is recurring captions that enable the model to track the contents of an image and thus establish links between textual semantics and visual representation. 'Daringly' challenging the established notions of authorship, images and Alt-text are collected without authors' consent. In the process, millions of artworks by living and dead artists are expropriated, in particular those working with two-dimensional imagery. To understand the scale of the scraping consider that LAION5,[vi] the database used by the Stable Diffusion model, includes 5.85 billion image-text pairs (LAION 2022). A subset of this database, LAION-Aesthetics, contains a collection of 600 million images[vii] selected – by another machine learning model – for the reason of being 'aesthetically pleasing images' (ibid). A powerful backlash from living artists and artistic communities has unfolded through increasingly publicly actions, but, so far, it has not been 'persuasive' enough in the eyes of the models' owners.

Second, the system establishes relations between the semantic content of Alt-text and the visual content of images. Both Dall-e and Stable Diffusion perform this phase using the CLIP[viii] neural network designed by OpenAI (Radford et al. 2021). Images and Alt-texts are encoded, meaning that they are described using mathematical representations. Practically, images and captions are encoded into separate numerical vectors where numbers describe some of their features; these are called 'embeddings'. What is key to my analysis here is that the model's engineers have little agency over which features the model recognizes or disregards. The model itself does not 'know' either, for it statistically manipulates probabilistic distributions in ways that are too complex to trace. "[I]t's not actually clear what is making the AI models work well. It's not [...] clear what parts of the data are actually giving [the model] what abilities", stated David Holz, Midjourney's founder (Claburn 2022). The incomprehensibility of the model's operation increases as the process progresses. The neural network works the whole dataset so as to compute the probability

that a given text embedding is the description of a given image embedding.[ix] Once it completed its training, CLIP is capable of returning the most fitting caption for a given image.

The last step before generating an image is to link a user's live prompt to particular visual contents found in the dataset. Using another neural network, called Prior, a user's prompt is mapped to an image embedding and then fed to a decoder. The decoder uses the image embeddings to condition its search of visual components to sample across the dataset. Put simply, the decoder uses words to navigate the labels in the datasets and find the required image components. At the moment, the most used decoder is a diffusion model called Glide, also developed by OpenAI (Nichol et al. 2021). The diffusion process is rather convoluted, but in essence it consists of digitally manipulating visual Gaussian noise until its mathematical representation matches the text and image embeddings created by CLIP after the user's prompt. The diffusion process is perhaps better understood as filtering rather than generation. The model filters Gaussian noise iteratively so as to make it as similar as possible to selected samples of other artworks.

Thinking in a more abstract fashion can help grasp the mechanism of 'learning' and 'generation'. What happens, conceptually, is that the model groups text-image vectors by similarity or co-variance and positions the groups in a high-dimensional space. This is a virtual space called 'manifold' and can be thought of as a geometrical representation of a curvilinear and continuous space. Filling the space of the manifold is akin to constructing a cartography of the dataset, where bits of images and texts are distributed at particular locations according to particular probability densities. The manifold, therefore, contains all the image combinations that are possible with the data at hand. As a crude example, imagine the following (which I document with images in Fig. 1-2-3 below). Multiple images of a painting of a dog by Francis Bacon are grouped at one location in the manifold; multiple images of a flower by Georgia O'Keefe are grouped at another location.

But a point in the manifold exists where Bacon's dogs and O'Keefe's flowers meet. When a user prompts the model to generate 'a dog by Francis Bacon in a flower by Georgia O'Keefe', the model uses the text as directions to find that particular location where dogs and flowers live next to each other. Then, it progressively samples some of the visual

features stored at this location and uses them to incrementally filter Gaussian noise in the form of a matching image. The sampling is stochastic, meaning that the samples are randomly selected from the relevant data available; this is why a model prompted with the same text will always generate a different result.

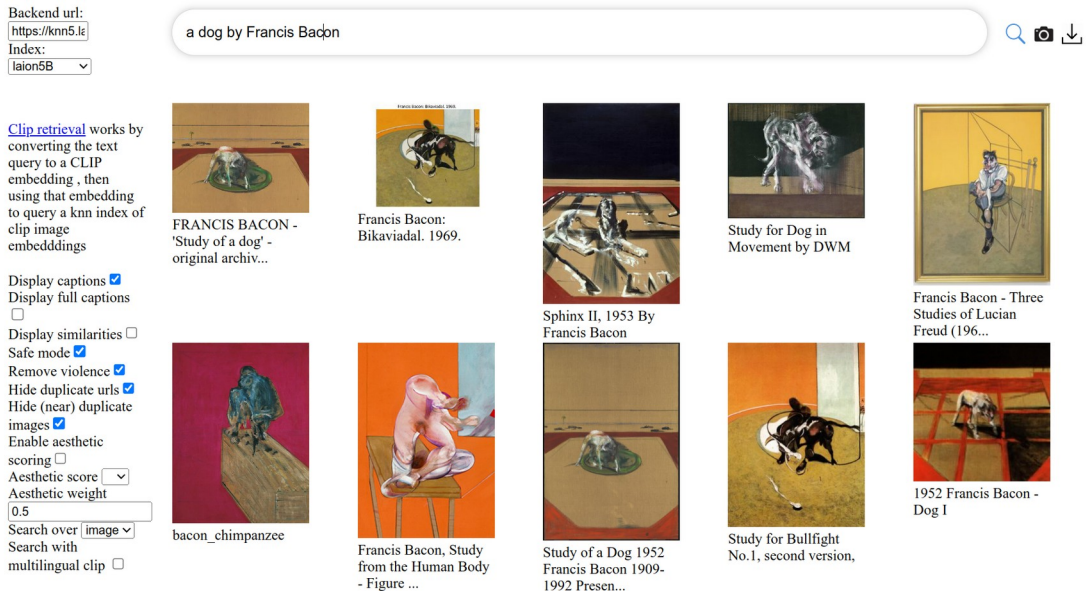

Figure 1: Exploring the LAION5B dataset. Top search results for 'a dog by Francis Bacon'

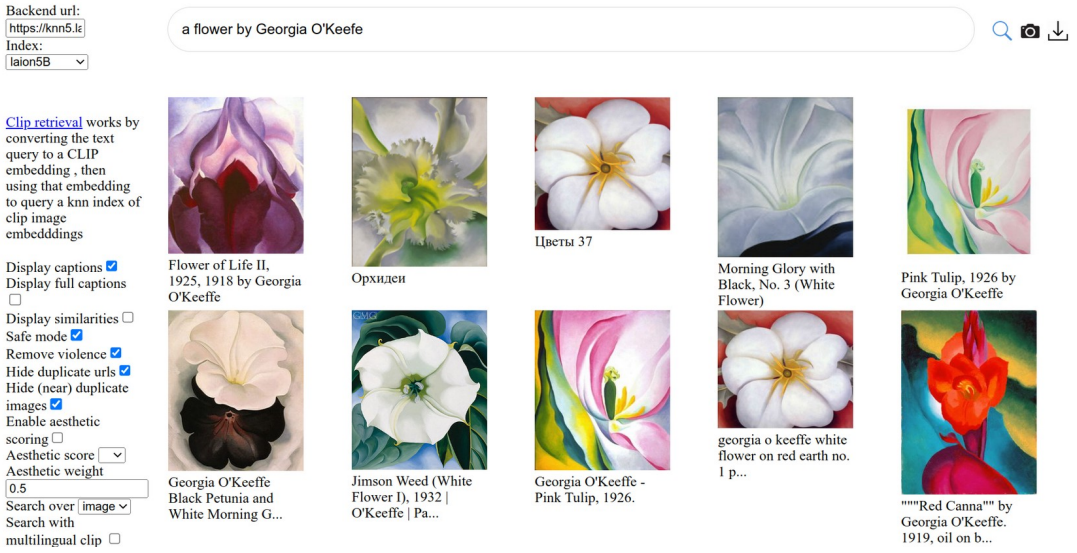

Figure 2: Top search results for 'a flower by Georgia O'Keefe'

Figure 3: First four output images produced by StableDiffusion upon the prompt: 'a dog by Francis Bacon in a flower by Georgia O'Keefe'.

Indifference through Repetition

Unsurprisingly, given their mode of operation as described above, the image artefacts produced by AI generators exhibit sophisticated mimicry of human-made artworks or imagery. The output of these system ranges wildly in quality. Most appealing images are obtained only through numerous trials with different ways of phrasing prompts and iterating variations on initially crude outputs. Curating the output of AI generators in order to select useful images is more human-labour intensive than producing the images. Eventually, a refined AI image may show a control of composition, style and colour palettes, coupled with a polished look and delicate visual intricacies. In these cases, the models' capacity for mimicry is so uncanny that, at first, it may even suggest that the system is actually creating something new. Having demystified the mechanics behind the process, it is evident that there is no artistry at play, at least not in the way it is commonly understood in the cultures of human animals. While it is true that, in view of such a technology, the historical definition of artistry may be expanded, it is important to remain focused on the exploitation that these systems rely on; what these models perform, technically and conceptually, is a brute appropriation and a chancy exploitation of cultural capital and cultural ecosystems. Their capacity for mimesis is closer to a happy accident than to a creative stimulus or an emergent quality of software, for these algorithms are designed to recreate an amalgam of what already exists. In these AI pastiches, what may be perceived as artistic meaning is a by-product of a stochastic sampling of thousands of artists' works. The features extracted from the original artworks are drastically decontextualised, displacing, thus, the artworks' aspirations and modes of expression. Alienated remains of artistic expression are then repeated, superimposed on themselves times and times over, until they loose expressivity and become mere signs, ghostly traces of artistic intention. As signs devoid of intention, they bear an eerie detachment from human creativity. It is a process that creates, freely paraphrasing Gilles Deleuze (1968 [1994]), indifference through repetition.

In fact, such a eerie detachment from artistic intention may be one of the factors that makes AI artefacts popular. The eerie, as Mark Fisher (2016) elaborated, can be understood as a 'failure' of presence or absence. It is the feeling that something that should be there is absent, or, vice versa, something that should not be there is present. When observing an artwork, one expects, consciously or not, the artist's intention to be present in the piece.

The presence of artistic intention is the *causa sine qua non* art making happens. Intention, in my view, is independent from artistry and skill levels and it emerges from one's interaction with others – living beings, materials, cultures. It can be more or less manifest and, still, be sensed in amateur artworks and less artistically successful pieces, for it is the expression of an ecosystem of relations; ultimately, artistic intention may be understood as the expression of a *collective otherness*. I will return to this in the last section when discussing two AI-based artworks whose approaches exemplify the importance of relations and otherness in art making. What is important here is that because of the way in which neural networks wrestle encoded meaning out of actual artworks, AI aesthetics rests on a mutilation of those ecosystemic creative relations that make an artwork what it is. AI aesthetics rests on a forced absence of collective otherness, for the material, cultural and personal relations that contribute to it are discarded. Where there used to be a creative goal, an artistic tradition or a rebellion against it, a chain of entangled influences (and misappropriations) across histories and cultures, nothing is left. Instead, AI image generators rely on the tired myth of the individual artist as a lonely genius. If *the* artist's intention is missing, what kind of agency created the artefact? As speculation and curious doubt are stimulated in the observers, they risk forgetting about the art itself and becoming lost in yet another old myth, that of AI singularity with its clumsy baggage exploding with anticipation over machine domination or salvation. The less romantic reality is that among the different agencies at play – computational, human, material and cultural – the one with less control over the expressivity of AI aesthetics is the computational. For it is the AI complex, with its human, material and cultural agencies, its networks and interests that conceives, designs and fosters how current AI image generation works, how it is disseminated, understood and exploited. The algorithmic system itself functions as a combinatorial agent, enmeshing features of artistic expression with the desires, beliefs and investments of the AI complex. Paradoxically, the algorithmic system does embody an ecosystem of relations, that of the AI complex. This, admittedly, has not much to do with creativity.

The absence of a creative vision in AI artefacts, as well as the presence of normative power in the AI complex, can also be observed in the way generative AI systems cautiously mimic traditional and trending aesthetic criteria.[x] What they offer is *little enough* variation to tickle curiosity, but never enough to unsettle standards. The automated process of mimesis

performed by AI generators conflates aesthetic languages, vocabularies, forms and structures of disparate artworks into a single form, while leaving them unchanged and unchallenged at the same time. What this operation offers, artistically and aesthetically, is the safety of what one already knows, combined with the thrill of a more or less riskless adventure. It recalls trophy hunting, where entitled individuals can experience the thrill of killing magnificent animals in the wild through the safety glass of a four-wheel drive. In this sense, AI aesthetics exists comfortably within the reign of what Ursula Le Guin (2001: XV) once described as "commodified fantasy", a form of cultural production that

> takes no risks: it invents nothing, but imitates and trivializes. It proceeds by depriving the old stories of their intellectual and ethical complexity, turning their truth-telling to sentimental platitude. [...] Profoundly disturbing moral choices are sanitized, made cute, made safe.

The supposed novelty heralded by AI image generators is a kind of change that ends in and with itself; it is sanitized, safe and claustrophobically self-referential. In reproducing ad infinitum the stuff of existent artworks recombined in infinitely slight variations, AI generators make reference only to their own capacity to do so. Following Deleuze (1981: 33) in his analysis of Bacon's paintings, one can see that the self-referentiality of AI generators is alien to *sensations:* those granules of expression or sensory particles traced across an artwork by the embodied act of making, effectively a form of corporeal expression that AI generators expunge from the source artworks they use. It is sensations that guide an observer's perception of the rhythms and intensities expressed by a particular artwork. These, in turn, help reveal how the artwork's concepts, symbols, desires or ways of dealing with them configure themselves into a dynamics of relation – engaging distant living beings, cultures and ideas (Donnarumma 2020). Being void of sensation, AI generators fall short of the aesthetic and sensorial force of art to bring about a questioning or a consequence. A close look at most AI artefacts, in fact, shows them to be neither drastically new forms of digital painting, illustration or generative art nor manifestly innovative ways of conceiving visual representation. These artefacts do not require the creation of a new vocabulary to be described; on the contrary, they can only be characterised by means of well-known, pre-existent tropes and references. Their awe-inducing familiarity and their unchallenging mode of expression is what makes them so conceptually appealing and

financially attractive to the AI complex. I expand on this point below, when discussing the art market.

As far as the vexed question goes: is AI capable of making art? I side with Joanna Zylinska (2020: 49) to state that this is a "misguided question". As I wrote in the introduction to this article, these kinds of questions distract from deeper issues about the AI complex and the way it operates through and onto culture. That said, I think it is helpful to analyse how and where the question of creativity arises, for it may help reveal another principle of corporate AI aesthetics. So, who poses those misguided questions and how do they benefit from it? The deep learning science community is generally weary of mystification. While it could greatly benefit from further developing its working definitions of art, labour and culture, and this impetus is growing within the discipline, a cursory reading of the technical papers referred to earlier quickly shows that researchers are above all interested in models' functionality, operation and optimization. It is rare to see AI scientists making claims of superhuman intelligence or human-like creativity. As expected, these kinds of claims are more often tied to the hype cycles across the industry PR and the mass media landscape. Less expected is to find, as sometimes happens, such claims among artists, curators and producers in media art. For, arguably, this artistic practice was born out of a paradigm that challenged the consumer-ready myth of endless technological progress sold by large part of the industry.

Tellingly, the trope of human-like creativity or singularity-level intelligence is most often found among those working directly or in close contact with the art market, an important agent within the AI complex. The most recent sale operation by auction house Sotheby's involves "the world's first intelligent NFT [...] coded with its own personality", "a modern day Alice [in Wonderland]" that ushers in "the age of living, self-learning artwork" (Sotheby's 2022). Or so we are told. Another sale operation by Sotheby's describes an AI generated video work as using an "AI brain" which is a "self-contained creative agent" enthusing collectors by offering them "the opportunity to watch an AI brain 'think' in real time" (Sothebys, 2018). Leaving aside the courageous copy-editing, this precisely crafted language suggests something intriguing about what actually may be the focus of attention. It signals that perhaps when it comes to art as a commodity and source of financial capital, corporate AI aesthetics has indeed a unique value: the allure of a human-made machine

that can "think". Tragic as it is, in this context, deceit and speculation are the real commodities, while art remain a marginally relevant, haunted shell. But there is another almost antithetical side to the art market exploitation of the AI race. On occasion of the first sale of an AI-generated artwork at the auction house Christies, the organiser of the sale, Richard Lloyd, was asked whether the art market saw a future for AI-generated images. He replied, 'It may not have been painted by a man in a powdered wig, but it is exactly the kind of artwork we have been selling for 250 years' (Christies 2018). It is strikingly ironic, albeit unsurprising, that a supposed ground-breaking artwork owns its fame and price to the fact that it replicates the most conventional kind of painting in the art market. Sameness pays.

As a form of expression of the power and capital dynamics of the AI complex, corporate AI aesthetics has quickly come to dominate the cultural and artistic landscape. This is the most seriously troubling aspect to it. If AI generators were only marketed as plug-ins, as tools that excel at a particular kind of image manipulation, critical appraisal may not be needed. The technology would even perhaps become an empowering tool, for it can certainly provide a gentle and playful entry point to algorithmic art. The problem, more important than the one concerning what this specific technology does, is who owns it or funds it and who produces it. AI generators are, at the time of this writing, a monopoly of a few extremely wealthy corporations or individuals who can afford the resources to pay for research, computational power, legal protection, viral marketing and appealing testimonials. The aesthetics these tools express quickly becomes, therefore, culturally dominant; in other words, it becomes the norm.

Normalizing the Future by refracting the Past

> [A]lgorithms act, but they do so as part of an ill-defined network of actions upon actions, part of a complex of power-knowledge relations, in which unintended consequences, like the side effects of a program's behavior, can become critically important. (Goffey 2008: 19)

Because AI systems are human-made constructs that can create consequences in the world, they are always active agents. Rather than being neutral tools or independent agents decontextualised from social life, they are imbued with a precise and pluripotent agency.

This draws from the particular world views of the people – and their arrangements, such disciplines of study, epistemes or capitalist abstractions - who fund, design, program and implement them. In today's capitalistic societies, deep and probabilistic learning systems are a material substrate (as in surveillance infrastructures, server farms, data networks) and an organizational scaffold (as in the structuring of policies, business plans, trends and trade strategies). Therefore, the agency of those systems helps materialise the very same worldviews that originated them. In other words, predictive systems construct the world by acting in it. Chun (2021) elaborates at length, for example, on how predictive policing systems act in the intimate lives of people by creating polarization and augmenting the vulnerability of those already vulnerable. Similarly, I argue, image generators act within the cultural texture of societies by reinforcing biases in visual culture, abstracting artistic labour and chipping at the idea of creativity as a collective, intergenerational and trans-cultural process until nothing remains of it.

It is not a coincidence that the advent of deep and probabilistic learning in artistic communities is taking place at a time when nuances have become the poorest currency in the cultural arena. Arguably, this may partly be the result of the widespread use of probabilistic and deep learning models across structural components of capitalistic societies; a feedback of sorts. The probabilistic ideology behind the architecture of AI art generators is, in fact, the same that powers predictive policing systems (ibid: 153), surveillance systems by Palantir (Iliadis/Acker, 2022), parts of the financial market, social media timelines from Tik Tok to Facebook and film and music recommendation systems such as Spotify and Netflix. By means of these platforms, social and cultural life becomes largely and silently influenced, and sometimes even driven, by machinic predictions based on past information. The effect is a tight enclosing of social and cultural groups and their isolation from one another according to data from their past activities, clicks and likes. More than just a way of doing computation, it seems that the probabilistic learning view is turning into an invisible dogma. It pervades and regulates societies at a political, social and cultural level.

What is the potential of this societal shift and which role does AI aesthetics play in it? A perspective I want to suggest, especially in view of the recent dissemination of highly disruptive chatbot generators such as ChatGPT, Bing and Bard, is that AI aesthetics

functions as soft propaganda. It would make sense, from the viewpoint of the AI complex, to stir public opinion, to domesticate it by means of a playful version of an AI probabilistic system, before raising the stakes with untested, unstable and 'hallucinating' language systems. It is such an uncanny coincidence that image generators inundated the cultural, media and social discourse just a little in advance of the massively public launch of large language models, which had been in development since several years. What is certain is that corporate AI aesthetics has been mediatically and culturally transfigured into a proxy for human-like creativity. While criticism of these systems is widespread, the work of large AI models is now commonly understood as a friendly, familiar, innocuous and innovative way of opening up creativity. At closer inspection, however, it is clear that the aesthetics of probabilistic systems is a closure for creativity rather than an opening. It can be understood as an operational schemata that is parasitic of cultural output: it uses repetition and imitation to bring about an aesthetic and conceptual normalization of a new that is a cheap duplicate of the present - and past. Echoing Fisher (2009) once more: this copy-paste of the present onto itself repeated infinite times brings about a forgetfulness of the injustices of the past and a foreclosure of the possibilities of the future. It normalizes a future of sameness by refracting the sanitized past ad nauseam.

For critical disability scholar Margrit Shildrick (2002: 71), "we are all implicated on an everyday level in a process of discursive othering that serves to establish and perpetuate standards of normativity". Discursive othering can be understood as a multilayered form of discrimination against those 'other' than the productive white Western able male: BIPOC, elderly, disabled (by society), queer and non-binary, poor, lower class people. This kind of discrimination is discursive because it propagates the enforcement of normativity through words, images, patterns and systems that are *implicitly* discriminatory. The historical connection between eugenics, statistics and AI is an excellent illustration of how ideological racism can turn into discursive othering. This is a widely known history, surprisingly simple to trace yet little talked about. Theoretical and mathematical concepts elaborated by prominent eugenists, including two founders of eugenics, Francis Galton and Karl Pearson, provided the scaffold for modern statistics (MacKenzie 1981; Davis 1995: 30; Chun 2021: 59). These concepts include notions such as linear regression, correlation and principal component analysis[xi] found today at the core of deep learning techniques that, as mentioned earlier, are heavily dependent on statistical methods.

Incidentally, the nemesis of data regression and discrimination – the latter is the technical term for classification – is something named 'outlier', a piece of data that does not fit the norm and therefore corrupts statistical reasoning. The outlier stays in the way of statistical truth, whereas the ideal lies in an average of probabilities. The more a piece of data or a person deviates from a norm the more they enter the domain of deviancy, of the extreme. Those people, thing or data, which can only exist at the edges surrounding the average become, thus, synonymous with the untrue, irrelevant, something that can be safely discarded. Following this logic, deep learning algorithms used for data classification must have layers that "amplify aspects of the input that are important for discrimination and suppress *irrelevant variations*" (Lecun et al. 2015: 436 [emphasis added]). Which are the criteria according to which a variation is deemed irrelevant? Far from being a semantic problem, the key question is, who has the power to establish relevance and irrelevance, to define the norm?

These questions become more poignant when considered in conjunction with the fact that, historically, the development of foundational AI theories and infrastructural AI technologies has been largely at the hand of Western, white males (Katz 2020) – with the exception of robotics, a field that owes much to diverse research communities, in particular those in Japan. Today, researchers in the field come from more varied background and geographical locations and yet, a techno-deterministic and reductionist mentality seems to remain dominant. As a result, parts of that field apply and amplify rooted biases and more or less latent forms of discrimination that, with or without malicious intent, have plagued deep learning since its re-emergence – although have not perturbed its pervasiveness. This reflection is not meant to claim the existence of a suppressed or militant racism at the core of probabilistic theory or deep learning. Many researchers in those fields are well intentioned if not actively striving to counter forms of computational othering (Bender et al. 2021) or the ambiguous creative power of Big Data (Vigliensoni et al. 2022) and their work must be supported and disseminated. More subtly, my observation is meant to focus attention on how deep learning, following statistics with its roots in eugenics, creates and disseminates a world view where people, images, sounds, money, lives - anything that is and is made countable - can be precisely understood and predicted. What is required is 'only' a large amount of data and related computational capacity coupled with a clear assumption of what a usefully functioning, a newly normative pattern will look like. This new kind of

norm still often is, or refers to, what Wynter (2015: 19) aptly defines as a "single genre-specific Western [...] bourgeois model of being", or, in other words, a capitalist, conformed, productive and consumer individual; a blueprint that is not only adopted in the West, Global North and most capitalist societies but also imposed on the others.

It is only from a perspective wherein anything can be algorithmically predicted because it is ascribed to only one particular mode of being – that of the capitalist, conformed and productive worldview – that it may become reasonable to think that an AI system with access to billions of images and artworks, but entirely lacking in intention, ideas, a sense of self and, most importantly, embodied relations to others, can attain creativity. Further, only within that episteme and a form of capitalism that upholds enormous concentration of power in the hands of the few, it may become acceptable to collect thousands of artworks without consent in order to train an AI system. Despite strong backlash from individual artists, whole communities and large companies,[xii] the very fact of exploitation and the disregard for consent slips within the creaks of 'innovation', becoming a passing thought that can only be detrimental to this quest of control. Yet, while an ideology of prediction may help research protein folding or climate modelling, culture and art making require no ideology, but interpersonal and interspecies relations, artistic urge, socio-cultural context, outliers and outsiders, unreachable edges and indescribable thoughts.

What I find important to stress, then, is that AI aesthetics functions as an effective means of enforcement of knowledge, in particular of the system of knowledge that defines neoliberal societies; an order based on the repetition and normalization of the Western understanding of cognitive and aesthetic criteria. It is telling in this sense that the parameters of scientific evaluation of the generative capability of the popular diffusion models are limited to the degree of realism or photorealism, samples diversity and similarity to users' caption (Ramesh et al. 2021, 2022). When considered scientifically, aesthetic quality evaluation is automated (ibid: 13) and when it is not, it is based on individual, unmotivated assumptions of what beauty is or should be; as Holz illustrates "I do think the world needs more beauty. Basically, if I create something that allows people to make beautiful things, and there are more beautiful things in the world, that's what I want *by default*." (Claburn 2022 [emphasis added]). What Holz speaks about is a vision of a stale, normalised beauty that can and *should be* universalised through technology. What is

missing from the picture is that beauty is far from universal and that, depending on who is looking, the world may already be full of it, while in the process of actively destroying it. With each new seeming or actual breakthrough rippling through media outlets, the AI complex normalises a universalist view of human cognition, machinic agency and aesthetic values, involving scientists, artists and workers across generations in the process. It provides scientific proof for a reductionist understanding of cognition and embodiment, all the way to notions of intelligence, beauty and creativity. To validate itself, the AI complex outputs scientific research and artistic artefacts with little space for counterarguments or criticism; as shown by the routine of frantically publishing research lacking peer review and hyperbolic descriptions of artworks.

AI aesthetics, encroached as it is between techno-determinism and capital, has users, viewers and customers implicated in the perpetuation of the illusion of exponential progress, a view oblivious to the significant contribution of the AI complex to societal polarization, undermining of human and workers' rights and the current climate collapse. By equating sophisticated mimicry to a neoliberal reductionist conception of creativity, the AI complex exploits art as yet another distraction from the dramatic and largely irreversible impact of its technologies on humans and non-humans. Here AI aesthetics functions as an epistemic infinity mirror (Fig. 4). It creates receding reflections of the technocratic order of knowledge, amplifying it to infinity within a small, rigid and inescapable frame.

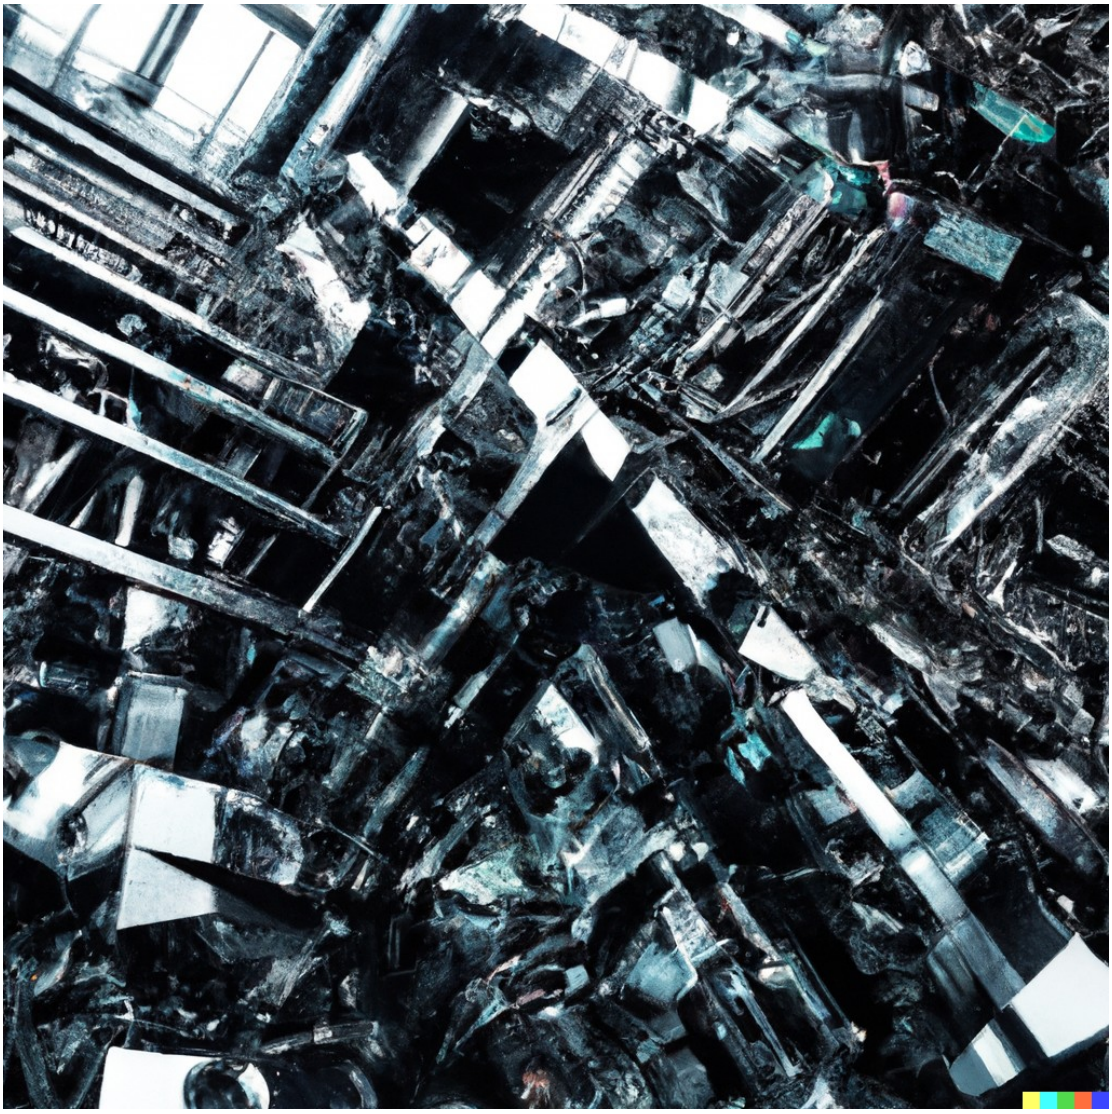

Figure 4: First output image produced by Dall-e with the prompt: 'capitalism reflected in an infinity mirror'

Against the Norm

> Machines, the entire technology of the West, is just that, the technology of the West. Nothing has to look or function the way it does. The West man's freedom, unscientifically got at the expense of the rest of the world's people, has allowed him to expand his mind – spread his sensibility wherever it go, & so shaped the world, & its powerful artifact-engines. (Baraka, 1970)

Deep learning, and artificial intelligence at large, do not have to be conceived and designed in the way they are. As elaborated throughout this text and as Baraka's powerful and prescient statement above reinforces, the present embodiment of deep learning and its expressive capacities are defined by specific powers. These, that I gather under the term AI complex, are the powers owning the financial, cultural and political means of computational production, and therefore those set out to gain unhinged profits from it. What art practice needs to counter this narrative, in addition to awareness and literacy, is the willingness to imagine a radically different present and therefore an alternative future. As illuminated by Fisher (2009), capitalism does not limit itself to the regulation of politics and finance. Rather, it is a pervasive atmosphere, a material and immaterial ambiance that severs imagination from culture, work and education; what Fisher calls "capitalist realism". As a strategy to ensure its own resilience, the doctrine of capitalist realism has erected a "grey curtain" (ibid: 81) over the horizon of possible futures that exists beyond capitalism itself. What is needed in media art and elsewhere, therefore, is rigorous work on imagining alternative ways to perverse the current route of AI technologies. Through this kind of imagination untamed ways of thinking about the meaning of technologically mediated art in times of environmental disruption will emerge. These ways will have little to do with the established canons of Western beauty or unfounded clamours of human-like machine intelligence. These are ways that will call into action the creative work of collectivities, the intelligent non-human entities that live in the world, and the limitations and potentials of deep learning to envision and manipulate polyphonic futures. As brief case studies, in closing this article I want to discuss two artworks, *In Search of Good Ancestors / Ahnen in Arbeit* by Jonathan Chaim Reus (2022) and *Aerodinámicas de las semillas* by the artist group AIseeds Project (2023). Through their imagination, these works are capable of creating new, collective knowledge that can sculpt realities.

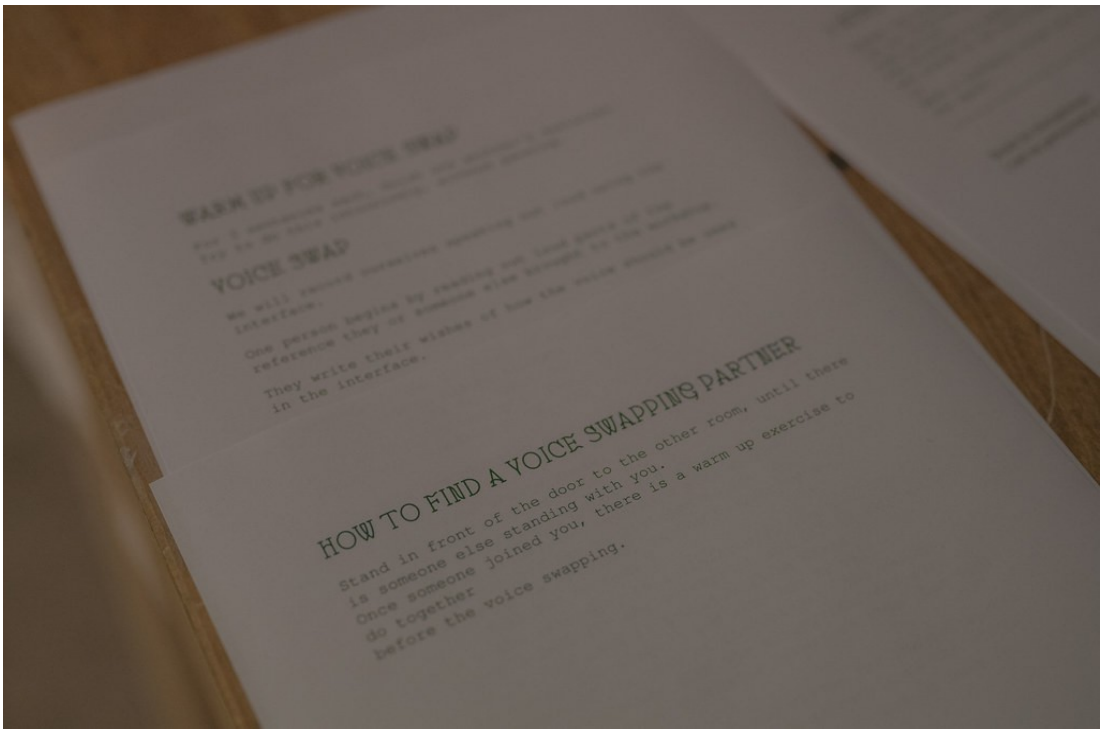

Figure 5: Informal instructions for the participants to one of Reus' workshops on voice and deep learning.

Reus is a transmedia artist and musician, originally from the US and based in the Netherlands. His piece *In Search of Good Ancestors* works technically and socially with the idea of machine learning algorithms as unstable memory. It enquires into the intergenerational implications of voice datasets and it does so by inviting interested strangers to literally seed and participate in the evolution of the work through public workshops (Fig. 5). The work consists of a 24-hour radio stream of generative speech broadcast over a year.[xiii] Speech sounds are generated using some of the most widespread, stock voice models for English speech research at the time of the artwork's creation, including GPT2 and Tacotron2. The system actually began from one voice model and, by the end of the project, it had expanded to multiple, repeatedly fine-tuned models that better suited the workshop process. Acknowledging and creatively repurposing the variety of voice models available in the domain of AI speech technologies is, for Reus (2023), to emphasise the plurality of embodiment that underpins data-driven, voice technologies. In a similar approach, the models were initially trained on two most common research datasets, The Pile and LJSpeech,[xiv] datasets that, as any other, are embedded with foundational,

sociocultural assumptions about what a voice is or means, in terms of gender, race or age, for example. Manipulating these datasets together with the workshop participants means, for Reus, to put into effect a collectively-led subversion of those presuppositions, once again moving away from a singular universal voice towards a shared multiplicity.

New vocal content to be reproduced by the voice models is generated through text predictions based on a lecture by US virologist Jonas Salk - where he calls for Western cultures to emphasise intergenerational responsibility as a high moral imperative. Rather than mulling over endless variations of the original text, throughout the ensuing year the models are gradually fed new, small voice and text datasets that Reus creates from scratch together with the participants of his public workshops. As the model is fine-tuned time after time, it is forced to continuously update its predictions. Machine memory becomes unstable, collective and collectivised. Because the artwork's way of rethinking deep learning technology subverts the strict paradigm of AI image generation discussed so far, *In Search of Good Ancestors* speaks about strategies whereby deep learning becomes a conceptual territory to navigate and shape, rather than a tool to be adopted as is. The workshop participants are neither deep learning nor voice experts, yet together with Reus they familiarise with the technical basis of the means and, importantly, appropriate the technology by developing personal modes of interaction with it. Together, they curate the texts for the model's training, design annotations for expressive text generation, perform with voice clones and record voice data individually and collectively (Fig. 6). As Reus explains, this format enables the diversity of languages and predisposition of the participants to enrich the process by allowing a "plurality of intentions and (literal) voices" to become the work (Reus 2022: 1).

Figure 6: The artist and the participants at one of Reus' project workshop sessions as they prepare to record voice material.

AIseeds Project is a transdisciplinary group of artists and researchers including Gabriela Munguía, Mario Guzmán, Guadalupe Chávez, Ángel Salazar and Erika Torres, whose bases lie across Argentina, Mexico, Ecuador and Colombia. Their work *Aerodinámicas de las semillas* (literally translated as aerodynamics of seeds) is a multifaceted project aiming to generate alternative perspectives on intelligence by looking carefully at plants and their reproduction processes and technologies. The work output takes multiple, interrelated forms: a generative website; a database of real and speculative flying seeds specimen – created using the generative adversarial network StyleGAN2; and an accompanying collection of texts that recombine, through the GPT2 large language model, existing manuscripts on botanics, ecology, philosophy and poetry.[xv] Rather than leveraging AI generators and models to create the kind of beauty that Holz dreamed of, AIseeds Project deploys that technology to create speculative morphologies of some of the flying seeds native to their lands, as well as to catalogue and describe them. Every time a user visits the

website, the model changes the textual descriptions by recombining them in always novel configurations. In doing so, AIseeds Project purposely plays with the core of scientific explanation and visualisation, and, by implication, of scientific knowledge.

Figure 7: AIseeds Project photographing seeds of Jacaranda in order to catalogue them and feed them to a deep learning system to generate further, speculative specimen. Courtesy of the artists.

Similarly to Reus' work, AIseeds Project's interest lies not in the technology per se, but in a collaborative action that generates a form of collective otherness. In this case, it is about an exploration of the territories that the plants inhabit and the artists share with them. Together with biologists and biodiversity researchers, they set out to collect specimen of indigenous flying seeds across their own native lands (Fig. 7). Their travel through those territories soon morphed into a new reading of the land that followed, conceptually and physically, the relations among the non-human lives inhabiting it; a remapping of the land

according to plant intelligence rather than human geopolitics. Through deep learning, the group analysed and compared the various seeds morphologies and found possible kinship relations among them. This enabled speculation on the potential for resilience and dissemination of some morphological features, providing an insight into possible ways in which the plants may evolve their reproductive organs; an urgent topic given the extreme violence to biodiversity provoked by the global environmental changes. The aesthetic and poetic power of the work rests, therefore, on its methodology: as the group puts it (AIseeds Project 2023), they use artificial intelligence to "negotiate otherness". The focus is not on the capability of the machine per se, but on its affordances in aiding the exploration, understanding and negotiation of plant intelligence. Simply put, the machine is a means to facilitate an intimate view into the existence of non-human others, fostering, in the process, strategies of curiosity, care and imagination.

Both works, *In Search of Good Ancestors* and *Aerodinámicas de las semillas*, speak effectively against the norm of corporate AI aesthetics. They do so in different ways and yet share roots in a fundamental embodiedness of the artistic experience. Being in the world as a relational and attentive entity among many others is posited as the spark of audacious relationships: between workshop participants, voices and AI models or between text and memories in the case of Reus' work, and among artists and biologists and land or among seeds and algorithms in the case of AIseeds Project. This is an aesthetic of relationality (Donnarumma 2020: 41), for it embraces the affordances of deep learning from a viewpoint of profound awareness of one's embodiment and interdependence in the world. Thus, rather than masking or mystifying the mechanisms of algorithmic agency, this kind of aesthetics utilizes, excavates and subverts that agency for what it is: codified instructions for patterns matching in a world of living things. This strategy does not trivialise the capacity of deep learning. On the contrary, and perhaps counter-intuitively, by laying naked the bare bones of deep learning systems and making them vulnerable and attuned to other kinds of agencies - human, vegetal and material - algorithmic agency is made an integral part of a new sensory clutch. It becomes a means of grasping the limits of memory, the multiplicities of identity, the intelligence of a plant or the histories of a flying seed (Fig. 8).

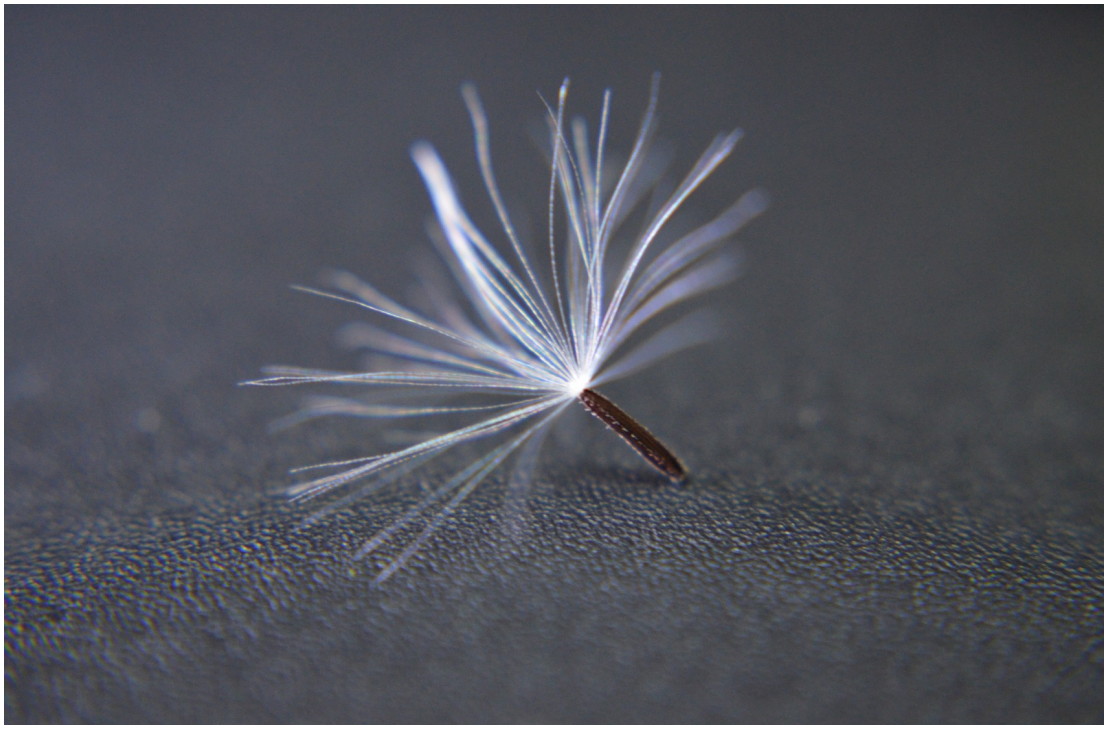

Figure 8: Microscope photography of a seed of a Diente de león. Courtesy of AIseeds Project.

Acknowledgements

This article is the result of hours of conversation with colleagues, who, with their varied expertise, constructive criticism and generous thinking supported my research. I thank the Intelligent Instruments Lab for allowing me the time to work on this manuscript during my residency there, and in particular Victor Shepardson, Jack Armitage and Thor Magnusson for our exchanges. I am grateful to Olga Goriunova for our conversations and her insightful reading of this manuscript. A warm thanks goes also to Minerva H. Trejo, Elizabeth Jochum and Pablo Gobira for their suggestions, Jonathan Chaim Reus and AIseeds Project for their time and thoughts, and the journal's reviewers for their feedback.

List of References

AIseeds Project (2023): Personal conversation with the author.

Baraka, A. (1970): "Technology & Ethos." Vol. 2 Book of Life. In: *Raise Rage Rays Raze: Essays Since 1965*.

Bender, E. M., Gebru, T., McMillan-Major, A., and Shmitchell, S. (2021): "On the dangers of stochastic parrots: Can language models be too big?" In: FAccT 2021 - *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 1, pp. 610-623.

Blyth, S. (2018): "Big Data and Machine Learning Won't Save Us from Another Financial Crisis." *Harvard Business School Cases*. (https://hbr.org/2018/09/big-data-and-machine-learning-wont-save-us-from-another-financial-crisis)

Cambridge Analytica (2022). "Leave.EU: Psychographic Targeting for Britain." (https://www.parliament.uk/globalassets/documents/commons-committees/culture-media-and-sport/BK-Background-paper-CA-proposals-to-LeaveEU.pdf).

Caramiaux, B./Donnarumma, M. (2021): "Artificial Intelligence in Music and Performance: A Subjective Art-Research Inquiry". In: E. R. Miranda (ed.), *Handbook of Artificial Intelligence for Music. Foundations, Advanced Approaches, and Developments for Creativity*. London: Springer, pp. 75-95.

Casilli, A. A./Posada Gutiérrez, J. (2019): "The Platformization Of Labor and Society." In: Graham, Mark & Dutton, William H. (eds.): *Society and the Internet: How Networks of Information and Communication are Changing Our Lives*, (2nd edition). Oxford: Oxford University Press, pp. 293-306.

Chun, W. H. K. (2021): *Discriminating Data: Correlation, neighborhoods, and the new politics of recognition*. Cambridge: MIT Press.

Christie's (2018): "Is artificial intelligence set to become art's next medium?" (https://www.christies.com/features/a-collaboration-between-two-artists-one-human-one-a-machine-9332-1.aspx).

Claburn, T. (2022): "David Holz, founder of AI art generator Midjourney, on the future of imaging." The Register. (https://www.theregister.com/2022/08/01/david_holz_midjourney/).

Davis, L. J. (1995): *Enforcing Normalcy: Disability, Deafness, and the Body*. London: Verso.

Deleuze, G. (1968): *Différence et répétition*. Paris: PUF. Tr. as *Difference and Repetition*, by Paul Patton (1994). New York: Columbia University Press.

Deleuze, G. (1981): *Francis Bacon: The Logic of Sensation*. Paris: Editions de la Difference.

Donnarumma, M. (2020): "Across Bodily and Disciplinary Borders: Hybridity as methodology, expression, dynamic". *Performance Research*, 25(4), pp. 36-44.

Fisher, M. (2009): *Capitalist Realism*. London: Repeater Books.

Fisher, M. (2016): *The Weird and the Eerie*. London: Zero Books.

Goffey, A. (2008): "Algorithm". In: M. Fuller (ed.), *Software Studies - A Lexicon*. Cambridge: MIT press.

Hoijtink, M./Planqué-Van Hardeveld, A. (2022): "Machine Learning and the Platformization of the Military: A Study of Google's Machine Learning Platform TensorFlow." *International Political Sociology*, 16(2), pp. 1-19.

Hu, M. (2020): "Cambridge Analytica's black box." *Big Data and Society*, 7(2), pp. 1-6.

Iliadis, A./Acker, A. (2022): "The seer and the seen: Surveying Palantir's surveillance platform." *The Information Society*, pp. 1-30.

Kai Y. (2013): "Large-scale deep learning at Baidu." In: *Proceedings of the 22nd ACM international conference on Information & Knowledge Management* (CIKM '13). New York: Association for Computing Machinery, pp. 2211-2212.

Katz, Y. (2020): *Artificial Whiteness: Politics and Ideology in Artificial Intelligence*. New York: Columbia University Press.

Lecun, Y., Bengio, Y. and Hinton, G. (2015): "Deep learning". *Nature*, 521(7553), pp. 436-444.

Le Guin, U. K. (2021): *Earthsea: Tales from Earthsea*. Cycle #5. New York:Graphia.

Lewis, J. E., Arista, N., Pechawis, A., and Kite, S. (2018): "Making Kin with the Machines". *Journal of Design and Science*, pp. 1-13.

Lozano-Hemmer, R. (1996): "Perverting Technological Correctness." *Leonardo*, 29(1), pp. 5-15.

Luitse, D. and Denkena, W. (2021): "The great transformer: Examining the role of large language models in the political economy of AI." *Big Data and Society*, 8(2), pp.1-14.

MacKenzie, D. A. (1981): *Statistics in Britain, 1865-1930*. Edinburgh: Edinburgh University Press.

Milmo, D. (2022): "Social media firms 'monetising misery', says Molly Russell's father after inquest". The Guardian. (https://www.theguardian.com/uk-news/2022/sep/30/molly-russell-died-while-suffering-negative-effects-of-online-content-rules-coroner).

McCarthy, J., Minsky, M. L., Rochester, N., and Shannon, C. E. (1955): "A proposal for the Dartmouth summer research project on artificial intelligence." White paper.

Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., and Chen, M. (2021): "GLIDE: Towards Photorealistic Image Generation and Editing with Text-Guided Diffusion Models." White paper.

Olorunnimbe, K./Viktor, H. (2022): "Deep learning in the stock market — a systematic survey of practice, backtesting, and applications." *Artificial Intelligence Review*, 56, pp. 2057-2109.

Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., and Sutskever, I. (2021): "Learning Transferable Visual Models From Natural Language Supervision." White paper.

Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., and Chen, M. (2022): "Hierarchical Text-Conditional Image Generation with CLIP Latents." White paper.

Reus, J. C. (2022): "In Search of Good Ancestors / Ahnen in Arbeit". In: *Nordic Human-Computer Interaction Conference* (NordiCHI '22), October 8–12, 2022, Aarhus, Denmark. New York: ACM.

Reus, J.C. (2023): Personal conversation with the author.

Schuhmann, C., Vencu, R., Beaumont, R., Kaczmarczyk, R., Mullis, C., Katta, A., Coombes, T., Jitsev, J., and Komatsuzaki, A. (2021): "LAION-400M: Open Dataset of CLIP-Filtered 400 Million Image-Text Pairs." White paper.

Shildrick, M. (2002): *Embodying the Monster: Encounters with the Vulnerable Self*. London: Sage Publications.

Smolensky, P. (1987): "Connectionist AI, symbolic AI, and the brain." *Artificial Intelligence Review*, 1(2), pp. 95-109.

Sotheby's (2018): "Contemporary Art Day Auction / Lot 109." (https://www.sothebys.com/en/auctions/ecatalogue/2019/contemporary-art-day-auction-l19021/lot.109.html?locale=en).

Sotheby's (2022): "Meet the world's first intelligent NFT. The age of living, self-learning artwork has arrived." (https://thefirstinft.com/).Sornette, D., and Von der Becke, S. (2011): "Computer trading: crashes and high frequency trading." UK Government Office for Science, 1-26. (https://www.gov.uk/government/publications/computer-trading-crashes-and-high-frequency-trading).

Sudmann, A. (2018): "On the Media-political Dimension of Artificial Intelligence." *Digital Culture & Society*, 4(1), pp. 181-200.

Vigliensoni, G., Perry, P., and Fiebrink, R. (2022): "A Small-Data Mindset for Generative AI Creative Work." *Generative AI and HCI* - CHI 2022 Workshop, May 10, 2022, Online, 1(1), pp. 1-5.

Whittaker, M. (2021): "From ethics to organizing: getting serious about AI." Distinguished Speaker Series, Hariri Institute for Computing, Boston University, Boston, MA, United States. (https://www.youtube.com/watch?v=_BzUobDoIcs).

Wynter, S. (2015): "Yours in the Intellectual Struggle: Sylvia Wynter and the Realization of the Living." In: K. McKittrick (ed.): *On Being Human As Praxis*. Durham and London: Duke University Press.

Yearwood, E. L. (2022): "A whistleblower confronts social media". *Archives of Psychiatric Nursing*, 37(PA1).

Zylinska, J. (2020): *AI Art. Machine Visions and Warped Dreams*. London: Open Humanities Press.

Zou, Y., Jin, X., Li, Y., Guo, Z., Wang, E. and Xiao, B. (2014): "Mariana: tencent deep learning platform and its applications". In: *Proceedings of the VLDB Endowment*, 7/13, pp. 1772-1777.

i Midjourney is owned by David Holz, a Silicon Valley entrepreneur who previously headed the Leap Motion technology startup. Stability AI is owned by Emad Mostaque, a hedge fund manager.

ii My artworks can be viewed at https://marcodonnarumma.com.

iii It should be noted that research on reinforcement learning for robotics often focus on robots' 'learning' in an environment or in partnership with a human. While several successful case studies exist, the results are incomparable to the behavioural plasticity of an animal.

iv More precisely, Gaussian noise is used. This is a form of signal noise where, from a statistical viewpoint, the values that noise can take are 'normally distributed'. This means that the most common values are near the mean, while less common ones are farther away from it.

v Bidirectional Encoder Representations from Transformers. It is used for Google's search engine and is widely adopted as basic tool for natural language processing research.

vi The dataset can be explored using this demo https://rom1504.github.io/clip-retrieval/?back=https%3A%2F%2Fknn.laion.ai&index=laion5B-H-14&useMclip=false&query=a+dog+by+francis+bacon. Last accessed on 08.11.2023.

vii According to a recent survey of a subset of the latter collection, most images are scraped from Pinterest (8.5%) and Wordpress-hosted websites (6.8%), while the rest originates from varied locations including artists-oriented platforms like DeviantArt, Flickr, Tumblr, as well as art shopping sites, including Fine Art America (5.8%), Shopify, Squarespace and Etsy.

viii Contrastive Language-Image Pre-training

ix This is done by maximizing higher probability distributions and minimizing the lower ones.

x The idiom "trending on ArtStation" is one of the most used parts of textual prompts by users of Midjourney.

xi Other concepts include standard deviation, p-values, chi-squared mean, and more exists.

xii Well known illustration artists Simon Stalenhag and RJ Palmer have been publicly vocal about the exploitation issue, while administrators of the online artistic communities Fur Affinity and Newgrounds rewrote their policies to explicitly ban "AI art". Two among the most popular communities, DeviantArt and ArtStation are facing increasing pressure from their users to enact a

similar ban. More recently Getty Images has sued StabilityAI for unlawfully copying and processing millions of their copyrighted images.

xiii The broadcast ran continuously from January 2022 until the end of March 2023 at ahnen.in, while fragments played intermittently on German Public Radio. Further information and documentation are available at the website. Last accessed on 08.11.2023.

xiv As explained by Reus (2023), LJSpeech is the voice of a single person, Linda Johnson, a voice actor that actively participated in the LibriVox project, a group of international volunteers who read and record public domain texts, creating free public domain audiobooks. Her recordings were made into a speech research dataset by researcher Keith Ito. Her voice is now possibly the most widely known within the entire world of English language machine learning voice research.

xv The work currently lives at https://aerodynamics-of-seeds.netlify.app/. Last accessed on 08.11.2023.