

# The Paradox of Anthroponormative Restriction: Artistic Artificial Intelligence and Literary Writing

---

Hannes Bajohr  
(Collegium Helveticum, Zurich)

## **Abstract:**

Artificial intelligence, in the shape of stochastic machine learning models, has seen an increased use in artistic production in recent years. However, it makes an immense difference if such an ‘artistic Artificial Intelligence’ is conceived of as an autonomous agent or only as a tool in the context of a human-machine assemblage. This article introduces the distinction between a *strong* and a *weak* artistic AI, and suggests that each invites a specific aesthetic. The former is inherently anthropocentric, strives for the reduplication of existing artforms, and reproduces concepts of a postromantic tradition such as expression, genius, and creativity. The latter, on the other hand, allows for an experimental approach towards genuine artistic novelty unhampered by human models through, paradoxically, keeping a human in the loop. This point is illustrated by discussing Ahmed Elgammal’s ‘Creative Adversarial Network’ and the digital poetry of Allison Parrish and Zach Whalen.

**Keywords:** Artificial Intelligence, computational creativity, electronic literature, language models, experimental literature, anthropocentrism.

German novelist Daniel Kehlmann has not written a book with an Artificial Intelligence; he has now written a book about that fact. Kehlmann, who might be most aptly described as Germany’s answer to Jonathan Franzen,

*CounterText* 8.2 (2022): 262–282

DOI: 10.3366/count.2022.0270

© Edinburgh University Press

[www.eupublishing.com/count](http://www.eupublishing.com/count)

is best known for his 2005 historic novel *Die Vermessung der Welt* about Alexander von Humboldt and Carl Friedrich Gauss, which was translated into English in 2006 as *Measuring the World* (Kehlmann 2005, 2006). In 2021, he turned to an entirely different subject. In the essay *Mein Algorithmus und ich* (My Algorithm and Me), Kehlmann travels to Palo Alto, where he gets access to an AI system at a start-up, the language model CTRL (Kehlmann 2021a). By entering into a dialogue with the model, he wants to test whether there is literature to be made with AI. Kehlmann is disappointed: the results are too narratively incoherent and too absurd for him, even if, here and there, an interesting sentence appears. Yet the failure of his excursion into the realm of machine learning turns out to be the implicit point of his book: humans need not worry that literature will soon be taken away from them by AIs.

Kehlmann does not approach the matter in a technically naïve way. He prefaces his book by saying that ‘AI’ is actually a misnomer, and that what sails under this moniker has neither consciousness nor intelligence in any real sense, but is a statistical model that merely makes predictions about likely states based on learned data (Kehlmann 2021a: 29). In so-called ‘large language models’, both the data learned and the predictions made have the form of text. These models do not work in a fundamentally different way than a smartphone’s autocomplete function: ‘Good’ is probably followed by ‘morning’, ‘idea’, or ‘heavens’.<sup>1</sup>

The fact that such an ‘intelligence’ has little to do with our own and should better be called ‘artificial rationality’, as Kehlmann remarked in a panel discussion (Kehlmann 2021b, my translation), nevertheless does not tempt him to examine this difference aesthetically. He regretfully admits: ‘I don’t have a story to show to you that I wrote with CTRL and that would seem good enough to me to be published as an artistic work rather than merely as the product of an experiment’ (Kehlmann 2021a: 31).

But what does ‘good enough’ mean? Measured against what aesthetics? When Kehlmann speaks of ‘experiment’, he seems not to have experimental literature in mind, but rather the scientific meaning of the word: a controlled observation whose outcome supports, weakens, or refines a hypothesis. But it does so, according to Thomas S. Kuhn, always only within the framework of an existing paradigm – new paradigms are precisely not what scientific experiments establish.<sup>2</sup> Experimental literature, on the other hand—at least according to its avant-garde

self-image—does not want mere refinement, but ideally questions the paradigm of literature itself.

Seen this way, it is quite possible that it is not Artificial Intelligence which has failed literature, but, as I will argue, Kehlmann who has failed Artificial Intelligence—and perhaps literature, too. For in his dichotomy between fully-fledged ‘artistic work’ and mere ‘experiment’ it becomes apparent how little it occurs to him that one can, or perhaps even must, write literature differently with machines instead of making them jump through the hoops of one’s own poetics. To him, the aberrations and absurdities that CTRL spews out are obviously a bug, not a feature. Moreover, he has a preformed and rigid idea of what literature is and what aesthetics it is supposed to follow. For Kehlmann, a novelist, literature’s perennial core is one thing above all: narrative—coherent and with broad, sweeping plot arcs that ultimately point to a complex authorial intentionality. For the novelist, then, not even poetry—language’s self-reflectivity—comes close to the anthropological need to tell stories; the machine lacks this capacity, and so he considers the experiment a failure.<sup>3</sup>

Never mind that the language model used, CTRL, was already hopelessly outdated when the book came out. Even GPT-3, which made a splash in 2020 as the state of the art in text AI, is a hundred times larger and would have produced much better results in terms of coherence; and the most recent models, like WuDao 2.0 or Google’s PaLM, are even more comprehensive.<sup>4</sup> More interesting than such technical quibbles is a paradox that is behind Kehlmann’s disappointment and that can frequently be discerned in discussions about art-making AI: the more one expects from Artificial Intelligence, the more human it is thought to be, but the less it is appreciated as a phenomenon in its own right. One may call this the *paradox of anthroponormative restriction*. A truly powerful artistic AI would not extend Kehlmann, but actually replace him—and it would not necessitate new aesthetics, but merely repeat the old ones. This paradox is evident in both the theory and practice of artistic AI.

### *Strong and Weak Artistic AI*

Almost all discussions about art and Artificial Intelligence fall under one of two, mostly unarticulated, conceptions of what an artistic AI actually is or should be. They differ immensely in their aspirations and hinge primarily on

the autonomy conceded to the art-producing system. Perhaps the best way to illustrate this difference is to use the parallel of John Searle's canonical notion of 'strong' and 'weak' AI:

According to weak AI, the principal value of the computer in the study of the mind is that it gives us a very powerful tool. For example, it enables us to formulate and test hypotheses in a more rigorous and precise fashion. But according to strong AI, the computer is not merely a tool in the study of the mind; rather, the appropriately programmed computer really *is* a mind, in the sense that computers given the right programs can be literally said to *understand* and have other cognitive states. (Searle 1980: 417; emphasis in the original)

For Searle, then, strong AI refers to the production of an artificial consciousness including all the properties that are constitutive of it (for Searle, this involves, above all, intentionality). Weak AI, on the other hand, is a mere aid for modelling consciousness. Thus, if strong AI means the functional reduplication of the target domain, weak AI is at best a partial simulation of this domain and has at most a heuristic, a 'tool' function, as Searle puts it.

If we move away from consciousness as a target domain, we can analogously speak of *strong* and *weak artistic AI*. The strong conception would see its task as reduplicating the entire production process of art. The weak conception would regard technologies – such as neural networks – as mere assistance systems in this process that take on only partial tasks. This may go quite far, but not to the point of complete independence as imagined in the concept of strong artistic AI. Yet it is precisely in the strong model – which would only be satisfied with a second Kehlmann, that is, an AI that produces an output one might expect from a human author – that a number of difficulties arise.

The possibility of strong artistic AI stands or falls with the question of how to operationalise the concept that is applied in the target domain: the concept of art (or, transitively, literature). Already the title under which the strong model usually operates – 'artificial creativity' or 'computational creativity' – shows that it is easier to circumvent the vagueness of the term 'art' by replacing it with that of 'creativity'. There are various strategies for doing this. The philosopher Margaret Boden defines creativity from the object side as the production of something that is 'new, surprising and

valuable' (Boden 2004: 1). Alternatively, the neuroscience approach aims at the subject side, the creative brain process (Abraham 2018). In both cases, art-as-creativity becomes something that can be schematised and ultimately be simulated by computers (Boden 2009, 2016).

In both cases, however, it is anything but clear whether there is not something about the term 'art' that gets lost once it is equated with creativity. This identification is reductive, partly because it does not attempt to define any criterion that distinguishes the production of a work of art from a technical innovation or a particularly disruptive business strategy. Nor is it clear how creativity, which is often conceptualised as problem solving, actually relates to the aesthetic.<sup>5</sup> At least since the 1960s, if not much earlier, creativity is largely detached from the concept of art and, more recently, literature (Goldsmith 2011) and is instead transferred to the 'creative industry' (Florida 2012) or weighs, as a 'creativity dispositif', on every neoliberal subject as the need to constantly prove one's adaptability in the marketplace (Reckwitz 2018). Art, in any case, can hardly be reduced to the concept of creativity, nor creativity to art.

Nevertheless, the strong model, which replaces the notion of art with that of creativity, still relies on an implicit concept of art. Both definitions of creativity, from the object side and from the subject side, are immanentist in nature. On the one hand, they subtract art from any sociological and historical context and thus posit it as an eternal, never-changing phenomenon.<sup>6</sup> On the other hand, they see art as produced by an isolated actor. Implicitly behind the idea of 'computational creativity' is an aesthetics of autonomy and genius that should raise suspicions in the context of a contemporary aesthetics.<sup>7</sup> Its contradictions become unavoidable when this theory is put into practice.

### *Creativity Machines*

In December 2020, Ahmed Elgammal, a computer scientist at Rutgers University, received a US patent for a 'Creative Adversarial Network' (CAN). The Network is explicitly designed to 'generate art', as suggested in the full patent title: 'Creative GAN Generating Art Deviating from Style Norms' (Elgammal 2020). Elgammal, too, understands art as creativity, defining it, with reference to behaviourist psychology, as an 'arousal' that can be measured in the brain. Triggers of such 'arousal' potentials are

surprise, confusion, complexity and semantic ambiguity, whereby both too little stimulus (boredom) and too much (reluctance) are to be avoided (Elgammal 2020: 6).<sup>8</sup>

Elgammal's CAN is the attempt to implement these novelty factors as a further development of a well-established AI architecture, the so-called 'Generative Adversarial Network' (GAN). A GAN combines two neural networks, where one, the 'generator', initially produces random images that are evaluated by the other, the 'discriminator', which is trained on a specific data set of images. In an iterative optimisation process, the generator adjusts its output according to the discriminator's scores, so that it eventually outputs images that have a statistical similarity to the training set (see Goodfellow et al. 2014). Trained on a set of portraits, the GAN could now produce new, deceptively real faces.<sup>9</sup> If in a GAN the only criterion is that the output image is statistically close to the training set (the operative distinction being similar/dissimilar), Elgammal's CAN introduces two criteria: trained on a canon of digital reproductions of real paintings and thus having obtained a probabilistic model of what counts as art, the discriminator first decides whether an output generated by the generator is art or not; then, using learned metadata about the 'styles' of these paintings from the training set, it evaluates whether the output image the CAN produced matches any of these styles (the distinctions thus being art/non-art and known style/unknown style; see Fig. 1, Nos 116 and 120/130).<sup>10</sup> These operations of distinction form a feedback process in which the discriminator steers the generator further and further towards higher 'art-ness' (by increasing the statistical similarity of its outputs to that of its inputs), and in which, in order to avoid mere pastiche and ensure the output's novelty, the discriminator has the generator avoid known styles and encourages styles that do not fit any of the learned ones.

Here, art-making is a function of deviation from individual structural features in the context of a canon that serves as a framework within which this deviation is permitted. In the CAN, both the idea of genius and the idea of autonomous art are technically implemented: not only is the *Bildungsroman* of an artist recreated in the machine (its 'aesthetic sensibility' is the result of an 'education' that encompasses already existing artworks), but art history is reduced to a sequence of decontextualised, digitised, and dehistoricised training data. As one of the most advanced designs to date within the paradigm of strong artistic AI, the CAN not




Figure 1. Ahmed Elgammal, Schema of the Creative Adversarial Network, in Ahmed Elgammal, 'Creative GAN Generating Art Deviating from Style Norms', US 10,853,986 B2, issued 2020, p. 1.

only takes on auxiliary tasks in the production process, but also creates art 'independently'.<sup>11</sup> This, at any rate, is Elgammal's explicit claim: in an article written together with the art historian Marian Mazzone, he distinguishes the CAN from merely assisting AI systems with the argument that its invention actually proceeds 'intentionally' because it acquires the rules for art production itself instead of already being fed them. Therefore, Elgammal insists, the CAN is not only 'inherently creative', but also an 'autonomous artist' (Mazzone and Elgammal 2019, 4, 3).

Yet this assertion of autonomy is dubious, and not only because here, too, a human has to choose which of the limitless generated images they actually want to release into the world. The claim of a strong artistic AI, which emphasises genius and autonomy, also renounces the concept of art that has been current in aesthetic discourses for more than half a century. At least since the 1960s, art is not seen to lie primarily in the object, but to encompass a social process of negotiation in the context of historical developments and institutional framings.<sup>12</sup> Thus, aesthetic thought in modernity is always measured by the possibility of stepping out of a given paradigm and declaring entirely new domains as art. It implies, as German media theorist Dieter Mersch writes, 'in every act and artifact, a

transformation of the aesthetic itself' (Mersch 2019: 73).<sup>13</sup> This also applies to literature, especially experimental literature. Its most minimal definitions do not make use of any immanent properties – 'literariness' as measurable on the object level – but rely solely on the gesture of declaring a text to be literature.<sup>14</sup>

The CAN, however, has no outside and does not allow for one. As a result, the strong model is *structurally conservative*. Instead of enabling new aesthetics, it reproduces the old ones. It is true that the CAN wants to simulate judgement in the broadest sense, and produce 'styles' that do not match any of the ones it has learned. But because this is conceived as a statistical model, it can only produce average art in the literal sense. For what art is, is already a foregone conclusion, since the CAN's concept is merely derived from past data; the added notion of 'style' does little to liberate this more fundamental preliminary determination. Technically formulated: since the CAN models a vector space ('art') from the distribution of features in the training set, it can *interpolate* arbitrary states in it, but cannot *extrapolate* any that lie beyond this space – it cannot, in other words, expand the concept of art.<sup>15</sup> The gesture of framing – declaring something to be art – has been essential to contemporary art since Marcel Duchamp's readymades, more than a century old, and it appears to be impossible for AI.<sup>16</sup> The output of CAN is correspondingly dull, and could, in its inoffensive abstraction, decorate the lobby of any corporation.

No experiments: instead of allowing for innovation, the strong model results in a *re-traditionalisation of art*. More than that, it is too anthroponormative at its core, despite its assertion of machine autonomy, because it relies on the mere duplication of human art production and appreciation. A *perfect* literary CAN would perhaps be narratively coherent, but unlikely to produce new literary forms. This is the paradox inherent in Kehlmann's desire to build an AI to his literary taste: there is autonomy only at the price of repetition. But it is questionable if anyone needs a second Kehlmann, probably not even Kehlmann himself.

### *Experiments in vector space*

What I have called *weak* artistic AI has more modest aspirations, but possibly produces the more interesting and experimentally daring art. Instead of thinking of the machine as 'creative' and 'autonomous' in one



way or another, the weak model advocates for a much more complex human-machine entanglement. Because of this, it can also accommodate a more nuanced notion of art, in which historical and social contexts are not simply reduced to a free-floating training set. In this paradigm, the degree of interconnectedness between human and AI is almost secondary, and can range from cyborg-like human-machine assemblages to a merely instrumental tool-use, in which AI would be but a better paintbrush or word processor.<sup>17</sup> In any case, the result is that all actors within this entanglement, be they human or machinic, influence and change each other, which almost necessarily produces new aesthetics rather than repeating existing ones.

If Elgammal were to give up his claim to truly independent art production, there would be no reason why a neural network like the CAN could not be used productively in a weak model.<sup>18</sup> For the problem of restricting art lies not at all with the technology used, artificial neural networks, of which GAN and CAN are only two of many subtypes. Neural nets are not in themselves hostile to art or literature, and I do not want to put forward any anti-technological argument here. In fact, these nets are the field in which the most interesting artistic experiments with AI can be observed at the moment – ‘experiments’ in the sense of the historical and neo-avant-gardes, as efforts to explore and create new forms. That such new forms are also necessary in the greater history of *electronic* experimental literature is not least due to the fact that the previous tradition of computer-generated literature and art cannot not simply be absorbed into the paradigm of neural networks; rather, neural nets demand a new poetics.

Until about ten years ago, this lineage was determined by the *sequential* paradigm – the algorithm as a series of formalised but human-readable rule steps. Because these steps can be understood by readers of the code, which is often (but not always) published alongside the output it creates, many of these works are committed to an aesthetic of transparency: the otherwise hidden operations in the artistic process are revealed and documented at the level of its production.<sup>19</sup> Artificial neural networks, however, which follow the *connectionist* paradigm and which are based on the (highly abstracted) model of synapses and neurons in the brain, are no longer programmed in such a stepwise manner.<sup>20</sup> Instead, they learn statistically, as does the CAN, by being fed a large number of inputs and tasked to produce similar

outputs. Since the resultant ‘weight model’ is simply a complex list of numbers determining the activation strengths of its neuron layers, their inner workings are neither easily readable by humans nor translatable into explicit rules.<sup>21</sup>

Language models such as CTRL and GPT-3, too, are neural networks. Trained on gigabytes of text—and it is rarely clear exactly where this training data comes from—they are likewise impenetrable systems to their users. But where the older aesthetics of transparency no longer holds and strong artistic AI seems aesthetically restrictive, two alternative tendencies can be observed in the contemporary landscape of experimental literary practice.

One tendency seizes upon the inscrutability of the language model and takes up the ‘hallucinatory’ strand of modernism, which in Surrealism, for example, focused on the exploration of the unconscious (O’Gieblын 2021; Roloff 2021). Here the model is seen more as a ‘medium’ in a quasi-spiritualist sense than as an autonomous creator. Thus, programmer, artist, and founder of Google’s ‘Artists + Machine Intelligence’ program K Allado-McDowell engaged in a kind of ‘co-creative writing’ for their novel *Pharmako-AI*. Writing directly with GPT-3 in a dialogical and improvisational fashion, they describe the work’s origin in, as Nietzsche would have it, Dionysiac rather than Apollonian terms—as a burst of enthusiastic ego-loss rather than as a work of cerebral rationality (see Nietzsche 1999).<sup>22</sup> In an ‘iterative writing process, between the generation of responses and the “trimming” of output’, a circular, hallucinatory act of language discovery took place: ‘Clusters of concepts emerged from our conversation. Images persisted from session to session. They entered my thoughts and dreams, and I fed them back into GPT-3. In this process, a vocabulary was born: a mapping of space, time and language that points outside of all three.’ (Allado-McDowell 2020: xi)

Nevertheless, Allado-McDowell has no interest in the phantasm of strong artistic AI. On the one hand, they use GPT-3 very much like a Tarot deck, which more often serves as a tool for self-inquiry than for communication with higher powers. On the other hand, they enter into a human-machine assemblage with the AI system, distributing authorship across it and blending their own vocabulary with that of the language model. This is also the difference to Kehlmann: Allado-McDowell engages with GPT-3 as a collaborative partner, and does not reject narrative breaks and

inconsistencies as errors that contradict one's own aesthetic preferences, but understands them as elements of an aesthetic to be developed cooperatively.

The second tendency to react to the unfamiliarity of neural networks lies in the exploration of their media-specific affordances (Bajohr 2022). For *VAUDn oc HORRRR*, programmer and artist Zach Whalen trained a GAN on 4800 comic panels of the horror genre and had it output new ones (Fig. 2). Since the panels contain both characters and dialogue, the neural network processes both according to the same logic—that of the pixel image. Instead of recognising discrete elements of a character system in the text, it treats the text in these images just as image information like the rest of the drawings. The task of outputting statistically similar images to the input not only results in panels with monstrously distorted faces (the small training set prevents the final images from becoming too good); it also produces speech bubbles containing a muddled mixture of quasi-characters that mimics the shape of text without containing any known words. By subjecting text and image to the *same* operation, Whalen illustrates the statistical data processing of neural networks by way of the collapse of the semiotic process, thus adding the uncanniness of a symbolic category confusion to the horror of the origin stories (Whalen 2020b, the discussion of the work, and a link to the code, can be found at 2020a).

In 'Compasses', programmer and poet Allison Parrish investigates how text *as* text—as strings of characters rather than as raster image—is processed by neural networks (Parrish 2019). Language models encode words in 'word embeddings' as high-dimensional vectors. This can be imagined as, in a first step, converting words into unique sets of numbers; in a second step, through a process called 'dimensionality reduction', these sets are then projected into a multidimensional vector space in which their relative probability of occurring in a context alongside each other is retained (Alpaydin 2021: 133–5; in the context of digital literature, see Heflin 2020). In this way, it is possible to model complex relationships between them, so that terms of similar meaning are close to each other in vector space. Even when performing operations on these words, the dependencies between them are preserved (as in the well-known example: 'King – Man + Woman = Queen') (see Allen and Hospedales 2019).

For 'Compasses', Parrish encoded words not by their meaning but by their phonetic value. Her system then was able to output the phonetic 'intermediate states' within this vector space. For example, between the



Figure 2. Zach Whalen, page from *VAUDn oc HORRRR* (2020),  
<http://www.zachwhalen.net/pg/horrrrr/book.pdf>, p. 1.

north			google		
woerth	earthe		augle	agolzen	
west	eaurth	east	apple	aasbol	amazon
waust	seauet		pacebul	aace-bown	
south			facebook		

Figure 3. Allison Parrish, detail from 'Compasses', *Sync*, 2:27 (2019),  
[www.sync.abue.io/issues/190705ap\\_sync2\\_27\\_compasses.pdf](http://www.sync.abue.io/issues/190705ap_sync2_27_compasses.pdf).

phonetic values for the word 'north' and 'west' lies the inferred word 'woerth'. This can be done for more than two-word vectors: the phonetic value of the combination between the four major tech companies – Google, Facebook, Apple, and Amazon – yields 'aasbol' (Fig. 3). Parrish's work (and her teaching, see Parrish 2017) plays with the fact that language can also be thought of as non-discrete, as a vector space that can be traversed continuously. This opens up a different self-understanding for literature and entirely new access to its material.

These avant-garde experiments, which work within the weak paradigm of artistic AI, seem, at least to me, more aesthetically promising and theoretically sophisticated than any attempts at or hopes for a strong model. Instead of, like Elgammal and Kehlmann, making artistic AI (despite all the emphasis on its non-intelligence) into a *simulation of artist subjectivity*, they rely on collaborative practices between humans and machines that generate their own languages. In this, however, these experiments also have an emancipatory character in the context of corporate data extractivism (see Crawford 2021).

For just as the CAN tends toward a statistical standard, large language models such as GPT-3 are also large, privately controlled levellers of difference. On the one hand, they extract publicly available language we all create by being active on the internet – Twitter posts, Wikipedia entries, comments and conversations all feed into the dataset of large language models owned by private corporations (Joler and Pasquinelli 2020). On the other, such models aggregate trends and overall tendencies at the expense of outliers and individual difference: all idiosyncrasies are averaged out in the mass of training data, so that their outputs tend toward a conventional

treatment of language (see Bender, Gebru, McMillan-Major et al. 2021; Manovich 2021: 49). This also applies, Kehlmann notwithstanding, to narration. An AI that narrates coherently and thus performs a standard function of language is precisely not unthinkable, but most likely only a matter of time. Especially serial and genre literature, which already permutes plot elements combinatorically, could plausibly be generated in this way, perhaps in conjunction with older, sequential techniques that include narratological schemas.<sup>23</sup>

Only those who write from the outset in a position from beyond the (vector) space in which language models interpolate their results will escape this averaging out. And that is more likely to be Allison Parrish than Daniel Kehlmann – the experimental avant-garde rather than the more or less conventional narrative literature. If Kehlmann opined that ‘language-experimental literature is what can be most easily algorithmised’ (Kehlmann 2021b, my translation), the very practice and ambition of experimental writing contradicts him: ‘Part of what I want to do as a poet’, Parrish says, ‘is invent forms of language so new that even GPT-[3] can’t predict them’ (Parrish 2020). In an age of large language models, avant-garde is literary self-defence; only by writing with the AI against its levelling tendencies will literature be anything other than the future repetition of its past states.

## Notes

A version of this text first appeared in German as ‘Keine Experimente: Über künstlerische Künstliche Intelligenz’, *Merkur* 75: 864 (2021), 32–44. The present article is a considerably expanded and updated English translation.

1. A useful introduction into the operation and the politics of large language models is Luitse and Denkena 2021.
2. According to Kuhn, paradigms are bound up with the notion of ‘normal science’, that is, the dominant and canonised form of ‘doing science’ at a given historical moment. ‘Paradigm’ relates to the fact ‘that some accepted examples of actual scientific practice – examples which include law, theory, application, and instrumentation together – provide models from which spring particular coherent traditions of scientific research’ (Kuhn 2012: 11). Once a paradigm has been established, experiments have the role of (1) clarifying the paradigm’s basic assumptions, (2) testing these assumptions against empirical evidence, and (3) refining the paradigm as to (i) its mathematical constants, (ii) the articulation of laws, and (iii) possible transfer of its findings onto other realms (Kuhn 2012: 25–9). Kehlmann’s use of the word ‘experiment’ only seems to refer to (1) and (2); he is not interested in (3) or the establishment of new paradigms.

3. Indeed, Kehlmann's collaboration with Bryan McCann, the founder of CTRL, was presented under the heading of 'AI Storytelling', and Kehlmann himself stated that the main result of the collaboration was to 'think deeper about the mechanisms of storytelling' (Kehlmann and McCann 2021). Kehlmann's use of CTRL was, one is led to speculate, in no small part a PR campaign of McCann on behalf of CTRL; given the virtually total insignificance of CTRL compared to other large language models today, the collaboration appears to have been not just an artistic but also a business failure.
4. CTRL has 1.6 billion parameters—or 'neurons' in its neural network—while GPT-3 boasts 175 billion; CTRL was trained on 140 gigabytes of text, GPT-3 on 570 gigabytes; PaLM, introduced by Google in April 2022, has 540 billion parameters and was trained on 780 gigabytes of text. For CTRL, see Keskar, McCann, Varshney, et al. 2019; for GPT-3, see Brown, Mann, Ryder, et al. 2020; for PaLM, see Chowdhery, Narang, Devlin, et al. 2022. The race for ever larger language models is now being criticised ethically and politically: the models reproduce discriminatory language, are no longer transparent and correctable in their size, and are responsible for immense CO<sub>2</sub> emissions. For a prominent example of this discussion, see Bender, Gebru, McMillan-Major, et al. 2021.
5. Unlike Boden, the neuroscientist Anna Abraham defines creativity as novelty plus appropriateness, the latter being understood as a problem-solving or optimisation issue (Abraham 2018: 7–8, 12–13). This certainly helps in the automation of creativity, but the question remains as to which problem a work of art actually solves and in which domain 'appropriateness' would then have to be sought.
6. Margaret Boden distinguishes personal from historical creativity ('P-creativity' and 'H-creativity'), but measures the historicity of the latter solely in terms of whether the creatively produced object objectively represents a novelty in the world (Boden 2004: 43–8). This is ultimately a catalogue model of history in which genealogies, influences, and asynchronous developments are irrelevant. For a critique of this notion of history, see Blumenberg 1958; Kracauer 1995.
7. This is made explicit in two popular discussions of AI art by Marcus du Sautoy (2019) and Arthur I. Miller (2019). For both authors, creativity is essentially a characteristic of genius, forming such series as: Bach, Picasso, Steve Jobs. That the concept of genius has migrated to Silicon Valley is also confirmed by Adrian Daub (2020).
8. In his discussion of arousal as basic measure for novelty and thus creativity (a concept of art thrice removed), Elgammal in his patent refers to the work of behaviourist psychologist Daniel E. Berlyne from the late 1960s and early 1970s (Berlyne 1967, 1971; Elgammal 2020: 5–6). In a co-authored paper, instead of Berlyne psychologist Colin Martindale is cited, whose framework similarly uses a measurable (and thus operationalisable) arousal potential (Elgammal, Liu, Elhoseiny, et al. 2017: 2; Martindale 1990).
9. This is illustrated, for example, by [www.thispersondoesnotexist.com](http://www.thispersondoesnotexist.com), which uses the StyleGAN2 model. For the philosophical background, see Bajohr 2021.
10. The training set consists of 75,753 paintings; the idiosyncratic list of styles that were used as metadata includes: 'Abstract Expressionism, Action Painting, Analytical

Cubism, Art Nouveau Modern, Baroque, Color Field Painting, Contemporary Realism, Cubism, Early Renaissance, Expressionism, Fauvism, High Renaissance, Impressionism, Mannerism/Late Renaissance, Minimalism, Naive Art/Primitivism, New Realism, Northern Renaissance, Pointillism, Pop Art, Post-Impressionism, Realism, Rococo, Romanticism, Synthetic Cubism' (Elgammal 2020: 14). Apart from the reductions that are going on in this model discussed below – its total historical decontextualisation and its identifying art with creativity—it is worth pointing out that the CAN equates artworks with their digitisation as pixel representations. This makes the image the paradigm of art as such, ignores all of its material qualities, and flattens any three-dimensional characteristics into two dimensions. Finally, even as representations, these images cannot be very good: given that the CAN's outputs are only 256x256 pixels large (Elgammal 2020, col. 12), one must speculate that the input images are of the same size.

11. It is precisely this claim to strong artistic AI—to actual *machine authorship*—that make the CAN stand out vis-à-vis newer systems that may yield more impressive results (such as DALL·E 2, see note 18) but content themselves with, even if implicitly, a notion of weak artistic AI.
12. This is also where the definition of art as creativity fails. If conceptual artist Sherrie Levine, whose practice encompasses the appropriation of other artists's works, reproduces Walker Evans's photographs without any alteration and declares them to be her own artworks, the result is, in the framework of Margaret Boden, neither new nor surprising on the object level. Their value is only measured in terms of a rationale that lies *outside* this object, which is precisely not the subject of Boden's definition. Such a concept of art, which locates the designation of an artwork in a context of socio-historical recognition, is rather captured by George Dickie's and Arthur C. Danto's 'institutional theory' (Danto 1981; Dickie 1974) and its contemporary corollaries, such as Sherri Irvin's notion of the 'artist's sanction' (Irvin 2005). For institutional theories, the main problem of postconceptual art—that the same object may be either a mundane object or a work of art (or, in the case of Levine, one artist's work or that of another), depending on the deictic gesture of declaring it to be art—can no longer be answered by reference to skill, beauty, a catalogue of possible and necessary forms, or any metaphysical definition of art as such. The institutional theory reduces art to the act of deixis itself, which is provided by institutions such as critics, museums, and the 'artworld system' in general. It is the most sceptical, or negative, theory of art in that it refrains from any positive notion of what art is, and can thus still capture its social reality best, even at the price of a certain *petitio principii* in that artworld and art being definitionally codependent (Dickie 1984: 82).
13. I agree with Dieter Mersch's critique insofar as it refers to strong rather than weak artistic AI.
14. These positions range from Austrian poet H. C. Artmann's 'poetic act' (Rühm 1969: 10) to the appropriation literature of the present. For an excellent overview of the latter's conception of literature, which is analogous to the 'institutional theory' in art, see Gilbert 2022.



15. This confirms Italo Calvino's intuition about the possibilities of a literary AI system: namely, that 'its true vocation would be for classicism' (1986: 12) – that is, the repetition of existing forms.
16. Would a strong artistic AI be conceivable that does justice to the institutional theory? I think it would – if such an AI had the status of a *social agent*. Interestingly, it wouldn't even have to be strong in Searle's sense and possess consciousness for that to happen, as long as it was *socially* accepted as a communication partner that could make the post-Duchamp framing gesture: 'This is art.' Anna Franková's Twitter bot @this\_is\_art\_bot, which declares all sorts of things to be art incessantly and without any consequences, shows that we are not there yet. The shift from a model of intelligence championed by Alan Turing's 'imitation game' (best known as the 'Turing test') that is primarily based on deceiving a human interlocutor (Turing 1950; Natale 2021) to one of social agency decoupled from intelligence as suggested by Susan Leigh Star's 'Durkheim test' (1989) would open up the possibility of *post-artificial texts*, for which the standard assumption about its authorship being of human origin is suspended in favour of a more agnostic position: in such a situation it may no longer be important whether a text has been written by a human or a machine.
17. See Henrickson 2021: 53 and Schäfer 2010 for accounts of computer-human interaction that are inspired by Actor-Network Theory.
18. One example of such an artistic assistance system is OpenAI's DALL·E and its more powerful successor DALL·E 2, both of which can produce illustrations and designs by description alone ('an armchair in the shape of an avocado') and will certainly soon find its way into professional graphics software practice (see Radford, Kim, Hallacy, et al. 2021; Ramesh, Dhariwall, Nichol, et al. 2022); instead of a GAN, they use 'diffusion models', which are based on the mathematics of thermodynamics (Dhariwal et al. 2021). These text-to-image AIs are already the tool of choice for an active scene of enthusiasts who explore its artistic dimensions, see for instance the Twitter accounts @images\_ai and @advadnoun.
19. This can be exemplified by Nick Montfort's generative Beckett pastiche *Megawatt*, whose code, when executed, not only outputs the text, but also includes this code itself (Montfort 2014: 241–6). As in Lawrence Weiner's strand of conceptual art, the production rule of the work and the work itself are identical. For the connection between conceptual and code literature, see Montfort's own reflections (2018).
20. For a more in-depth discussion of the 'sequential' and the 'connectionist' paradigms, see Bajohr 2022. Nota bene: since the concept of 'AI' itself is agnostic about the technology through which it is pursued – earlier AI models were sequential while today most research is connectionist – the distinction between weak and strong AI does not correlate with that of sequential and connectionist; each technology can be used for each goal, and has been: see Mitchell 2019.
21. Matthew Kirschenbaum has tried to 'read' a neural network, but could do so only at the *output* level (2021), while in his earlier work, he has delved deeply into its code and even, 'forensically', its physical substrate (see Kirschenbaum 2008).

22. The contributions in *Pharmako-AI* are clearly attributed to Allado-McDowell and GPT-3 through roman and bold typeface (Allado Mc-Dowell 2020). This is no longer the case in their most recent work, *Amor Cringe* (2022), also co-written with GPT-3, which forgoes such identifiability.
23. Coherence in language models has so far been limited by their small ‘context window’, that is, they can only ever keep track of and refer to a limited section of a text (for GPT-3 this was initially about 500–1000 words). However, the context window improves with each new published model, and so does its coherence. Of course, there are also objections: the fact that narrativity cannot be simulated as long as AIs can encode correlations but not causalities is an argument put forward by the literary scholar Angus Fletcher (2021). He refers to considerations by the computer scientist Judea Pearl (Pearl and MacKenzie 2018).

## References

- Abraham, Anna (2018), *The Neuroscience of Creativity*, Cambridge: Cambridge University Press.
- Allado-McDowell, K (2020), *Pharmako-AI*, London: Ignota.
- Allado-McDowell, K (2022), *Amor Cringe*, New York: Deluge.
- Allen, Carl, and Timothy Hospedales (2019), ‘Analogies Explained: Towards Understanding Word Embeddings’, *arXiv* <<https://arxiv.org/abs/1901.09813>> [accessed 17 April 2022].
- Alpaydin, Ethem (2021), *Machine Learning: Revised and Updated Edition*, Cambridge, Mass: MIT Press.
- Bajohr, Hannes (2022), ‘Algorithmic Empathy: Toward a Critique of Aesthetic AI’, *Configurations*, 30: 2, 203–31 <doi: 10.1353/con.2022.0011>.
- Bajohr, Hannes (2021), ‘The Gestalt of AI: Beyond the Atomism-Holism Divide’, *Interface Critique*, 3: 3, 13–35 <doi: 10.11588/ic.2021.3.81304>.
- Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, et al. (2021), ‘On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?’, *EAccT ’21* <doi: 10.1145/3442188.3445922> [accessed 17 April 2022].
- Berlyne, Daniel E. (1971), *Aesthetics and Psychobiology*, New York: Appleton-Century-Crofts.
- Berlyne, Daniel E. (1967), ‘Arousal and Reinforcement’, in David Levine (ed.), *Nebraska Symposium on Motivation 1967*, Lincoln: University of Nebraska Press.
- Blumenberg, Hans (1958), ‘Epochenschwelle und Rezeption’, *Philosophische Rundschau*, 6: 1–2, 94–120.
- Boden, Margaret A. (2004), *The Creative Mind: Myths and Mechanisms*, 2<sup>nd</sup> edn, London: Routledge.
- Boden, Margaret A. (2009), ‘Computer Models of Creativity’, *AI Magazine*, 30: 3, 23–34.
- Boden, Margaret A. (2016), *AI: Its Nature and Future*, Oxford: Oxford University Press.
- Brown, Tom B., Benjamin Mann, Nick Ryder, et al. (2020), ‘Language Models are Few-Shot Learners’, *arXiv* <<http://arxiv.org/abs/2005.14165>> [accessed 17 April 2022].

- Calvino, Italo (1986), 'Cybernetics and Ghosts', *The Uses of Literature*, San Diego: Harcourt Brace Jovanovich, pp. 3–27.
- Chowdhery, Aakanksha, Sharan Narang, Jacob Devlin, et al. (2022), 'PaLM: Scaling Language Modeling with Pathways', *arXiv* <<http://arxiv.org/abs/2204.02311>> [accessed 17 April 2022].
- Crawford, Kate (2021), *Atlas of AI. Power, Politics, and the Planetary Costs of Artificial Intelligence*, New Haven: Yale University Press.
- Danto, Arthur C. (1981), *The Transfiguration of the Commonplace: A Philosophy of Art*, Cambridge, MA: Harvard University Press.
- Daub, Adrian (2020), *What Tech Calls Thinking: An Inquiry into the Intellectual Bedrock of Silicon Valley*, New York: Farrar, Straus, and Giroux.
- Dhariwal, Prafulla, and Alex Nichol (2021), 'Diffusion Models Beat GANs on Image Synthesis', *arXiv* <<http://arxiv.org/abs/2105.05233>> [accessed 20 May 2022].
- Dickie, George (1974), *Art and the Aesthetic: An Institutional Analysis*, Ithaca: Cornell University Press.
- Dickie, George (1984), *The Art Circle: A Theory of Art*, New Haven: Yale University Press.
- du Sautoy, Marcus (2019), *The Creativity Code: How AI is Learning to Write, Paint, and Think*, London: Fourth Estate.
- Elgammal, Ahmed (2020), Creative GAN Generating Art Deviating from Style Norms, United States Patent US 10,853,986 B2.
- Elgammal, Ahmed, Binchen Liu, M: Creative Adversarial Networks of the Creative Class, 2<sup>nd</sup> edn, New York: Basic Books.
- Gilbert, Annette (2022), *Literature's Elsewheres: On the Necessity of Radical Literary Practices*, Cambridge, MA: MIT Press.
- Goldsmith, Kenneth (2011), *Uncreative Writing: Managing Language in the Digital Age*, New York: Columbia University Press.
- Goodfellow, Ian J., Jean Pouget-Abadie, Mehdi Mirza, et al. (2014), Generative Adversarial Networks, *Advances in Neural Information Processing Systems*, January, 2672–80.
- Heflin, Judy (2020), *AI-Generated Literature and the Vectorized Word*, (MA Thesis), Cambridge, MA: Massachusetts Institute of Technology.
- Henrickson, Leah (2021), *Reading Computer-Generated Texts*, Cambridge: Cambridge University Press <doi: 10.1017/9781108906463> [accessed 17 April 2022].
- Irvin, Sherri (2005), 'The Artist's Sanction in Contemporary Art', *The Journal of Aesthetics and Art Criticism*, 63: 4, 315–26.
- Joler, Vladan, and Matteo Pasquinelli (2020), 'The Noosphere Manifested: AI as Instrument of Knowledge Extractivism', *Noosphere.ai* <<https://noosphere.ai>> [accessed 17 April 2022].
- Kehlmann, Daniel (2005), *Die Vermessung der Welt*, Reinbek bei Hamburg: Rowohlt.
- Kehlmann, Daniel (2006), *Measuring the World*, New York: Pantheon.
- Kehlmann, Daniel (2021a), *Mein Algorithmus und ich: Stuttgarter Zukunftsrede*, Stuttgart: Klett-Cotta.
- Kehlmann, Daniel (2021b), *Mein Algorithmus und Ich: Stuttgarter Zukunftsrede*, 1 September 2021 <<https://www.literaturhaus-stuttgart.de/event/mein-algorithmus-und-ich-4842.html>> [accessed 10 January 2021, no longer online].

- Kehlmann, Daniel, and Bryan McCann (2021), eVe Award Ceremony 2021, *YouTube*, <<https://www.youtube.com/watch?v=Gtl8zV0Ofz0>> [accessed 17 April 2022].
- Keskar, Nitish Shirish, McCann Bryan, Varshney Lav R., et al. (2019), CTRL: A Conditional Transformer Language Model for Controllable Generation, *arXiv* <<https://arxiv.org/abs/1909.05858>> [accessed 17 April 2022].
- Kirschenbaum, Matthew G. (2008), *Mechanisms: New Media and the Forensic Imagination*, Cambridge: MIT Press.
- Kirschenbaum, Matthew G. (2021), 'Spec Acts: Reading Form in Recurrent Neural Networks', *ELH*, 88: 2, 361–86.
- Kracauer, Siegfried (1995 [1969]), *History: The Last Things Before the Last*, Princeton: Markus Wiener.
- Kuhn, Thomas S. (2012 [1962]), *The Structure of Scientific Revolutions*, Chicago: The University of Chicago Press.
- Luitse, Dieuwertje, and Wiebke Denkena (2021), 'The Great Transformer: Examining the Role of Large Language Models in the Political Economy of AI', *Big Data & Society*, 8: 2 <doi: 10.1177/20539517211047734> [accessed 17 April 2022].
- Manovich, Lev (2021), *Cultural Analytics*, Cambridge: MIT Press.
- Martindale, Colin (1990), *The Clockwork Muse: The Predictability of Artistic Change*, New York: Basic Books.
- Mazzone, Marian, and Ahmed Elgammal (2019), 'Art, Creativity, and the Potential of Artificial Intelligence', *Arts* 8: 1, 1–9 <doi:10.3390/arts8010026>.
- Mersch, Dieter (2019), 'Kreativität und Künstliche Intelligenz: Bemerkungen zu einer Kritik algorithmischer Rationalität', *Zeitschrift für Medienwissenschaft*, 11: 2, 65–74.
- Miller, Arthur I. (2019), *The Artist in the Machine: The World of AI-Powered Creativity*, Cambridge: MIT Press.
- Mitchell, Melanie (2019), *Artificial Intelligence: A Guide for Thinking Humans*, New York: Farrar, Straus, and Giroux.
- Montfort, Nick (2014), *Megawatt: A Novel. Computationally, Deterministically Generated Extending Passages from Samuel Beckett's Watt*, Cambridge: Bad Quarto.
- Montfort, Nick (2018), 'Conceptual Computing and Digital Writing', in Andrea Andersson (ed.), *Postscript: Writing after Conceptual Art*, Toronto: University of Toronto Press, pp. 197–210.
- Natale, Simone (2021), *Deceitful Media: Artificial Intelligence and Social Life after the Turing Test*, Oxford: Oxford University Press.
- Nietzsche, Friedrich (1999), *The Birth of Tragedy and Other Writings*, ed. Raymond Geuss and Ronald Speirs, Cambridge: Cambridge University Press.
- O'Gieblyn, Meghan (2021), Babel: Could a Machine Have an Unconscious?, *n+1*, 40: 1 <<https://www.nplusonemag.com/issue-40/essays/babel-4>> [accessed 17 April 2022].
- Parrish, Allison (2017), Understanding Word Vectors, *GitHub* <<https://gist.github.com/aparrish/2f562e3737544cf29aaf1af30362f469>> [accessed 17 April 2022].
- Parrish, Allison (2019), Compasses, *Sync*, 2: 27 <[www.sync.abue.io/issues/190705ap\\_sync2\\_27\\_compasses.pdf](http://www.sync.abue.io/issues/190705ap_sync2_27_compasses.pdf)> [accessed 17 April 2022].

- Parrish, Allison (2020), 'Q&A with Allison Parrish', *Artists + Machine Intelligence* <medium.com/artists-and-machine-intelligence/q-a-with-allison-parrish-895a72727a4> [accessed 17 April 2022].
- Pearl, Judea, and Dana MacKenzie (2018), *The Book of Why: The New Science of Cause and Effect*, New York: Basic Books.
- Radford, Alec, Jong Wook Kim, Chris Hallacy, et al. (2021), 'Learning Transferable Visual Models from Natural Language Supervision', *arXiv* <<http://arxiv.org/abs/2103.00020>> [accessed 17 April 2022].
- Ramesh, Aditya, Prafulla Dhariwal, Alex Nichol, et al. (2022), 'Hierarchical Text-Conditional Image Generation with CLIP Latents', *arXiv* <<https://arxiv.org/abs/2204.06125>> [accessed 17 April 2022].
- Reckwitz, Andreas (2018), *The Invention of Creativity: Modern Society and the Culture of the New*, Malden, MA: Polity.
- Roloff, Simon (2021), 'Halluzinierende Systeme', *Merkur*, 75: 864, 73–81.
- Rühm, Gerhard (ed.) (1969), 'Vorwort: Die Wiener Gruppe', in *Die Wiener Gruppe: Achleitner, Artmann, Bayer, Rühm, Wiener. Texte, Gemeinschaftsarbeiten, Aktionen*, 2nd ed., Reinbek bei Hamburg: Rowohlt, 7–38.
- Schäfer, Jörgen (2010), 'Reassembling the Literary: Toward a Theoretical Framework for Literary Communication in Computer-Based Media', in Jörgen Schäfer and Peter Gendolla (eds), *Beyond the Screen: Transformations of Literary Structures, Interfaces and Genres*, Bielefeld: Transcript, pp. 25–70.
- Searle, John R. (1980), 'Minds, Brains, and Programs', *The Behavioral and Brain Sciences*, 3: 3, 417–57.
- Star, Susan Leigh (1989), 'The Structure of Ill-Structured Solutions: Boundary Objects and Heterogeneous Distributed Problem Solving', in Les Gasser and Michael N. Huhns (eds), *Distributed Artificial Intelligence*, London: Pitman, 37–54.
- Turing, Alan M. (1950), 'Computing Machinery and Intelligence', *Mind*, 59: 236, 433–60.
- Whalen, Zach (2020a), 'This Comic Does Not Exist', *GitHub* <<https://github.com/nanogenmo/2020/issues/55>> [accessed 17 April 2022].
- Whalen, Zach (2020b), *VAUDn oc HORRRR* <<http://www.zachwhalen.net/pg/horrrrr/book.pdf>> [accessed 17 April 2022].