

AI and Art

Arguments for Practice

Arno Schubbach

Over the past decade, the advances in artificial intelligence (AI) research have been attracting a lot of attention and provoked a broad variety of debates. Especially in the last two or three years, the progress in image generation by ‘generative adversarial networks’ (GAN) or ‘diffusion models’ (like DALLÉ-2 or Stable Diffusion) has been breath-taking—and has perhaps only been overshadowed in the public’s attention by OpenAI’s ChatGPT, which moreover will soon already be part of the everyday life of almost all computer users, if this is not already the case.

Compared with these swift technological advances, the debates they entail seem rather stable and often dominated by the same recurring, quite speculative questions: Can machines have consciousness? Should we fear that the machines we create will master us? Narratives that are as old as the fascination with machines and automatons are thus revived, and age-old fears of the creature surpassing its creator or of autonomous machines replacing human beings are stirred up. Finally, fear that machines will replace humans, their labour or creativity, is often expressed based on the assumption that machines can now do what previously seemed to be an exclusively human activity: machines are therefore supposed to be able to think or create art.

These debates seem insufficient to me since they compare machines and humans and yet, at the same time, tend to blind us to both sides. On the machine side, it does seem essential to take into account the specific approaches, algorithms, and applications that have been developed in recent years and are being used more and more. This is necessary, on the one hand, in order to gain insights into the amazing capabilities, but also the limitations, of this technology. On the other, it is also essential because the successful research of recent years is based on the rather sober approach of machine learning, which does not have much to do with AI research insofar as it has done away with the ideological undertones that were sometimes inherent in traditional attempts to artificially re-create and surpass human intelligence and capabilities.

On the human side, it would, however, be equally important to have a more detailed account of the activities that we may have previously regarded as a human

privilege, but which can potentially be taken over by computers. After all, what activities such as thinking or artistic creation actually are is by no means obvious or undisputed. One traditional and until today influential approach is that such activities should be discussed in terms of the psychological capabilities they may presuppose. Moreover, this approach has been particularly influential whenever AI research has seen its mission as imitating human capabilities and intelligence. As a consequence, the question of whether computers can think or make art is often reckoned against a psychological model that refers to the inner experience of the human being and does not provide a precise criterion. For example, the question of whether machines can make art leads back to the question of whether or not these machines can be creative—at the same time, while we may have the experience of being creative, we also do not have a straightforward understanding or a precise criterion for what being creative means for humans.

Alan Turing already presented a similar argument against a psychological approach to the question of whether computers can think. In his view, the problem is not so much that it would be speculative to attribute psychological capabilities or intelligence to computers. Given the traditional conundrums about the soul, consciousness, et cetera, he instead suggests that whether computers can think is a pointless question, because it is not at all clear what exactly might be meant by this: ‘The original question, “Can machines think?”, I believe to be too meaningless to deserve discussion’ (Turing 2004, 449). Turing thus argues that a verifiable empirical criterion is needed and proposes for this purpose his ‘imitation game’, later called the ‘Turing test’. The key criterion here is whether a human can distinguish the computer from the human in an interaction in which human and machine interact through the same medium, such as a typewriter, thus excluding any immediate perception of the human being based on his or her face, voice, et cetera. From the critique of a psychological understanding of thinking or art-making follows a feasible criterion, which is, however, not particularly revealing: it does not relate to the activities of thinking or creating art and says nothing about the corresponding activities and processes on the part of humans or computers, because it refers exclusively to their results and to their experience through the participants in the ‘imitation game’ or ‘Turing test’.

In the following, I would like to discuss the relation between AI and art with reference to the recent successes in image generation based on machine learning methods. In the first section, I will, however, step back a bit and go back to the beginnings of computer art and A. Michael Noll’s studies in the 1960s. In doing so, I want to show not only that the question of whether a computer can make art was already asked at that time, although the technical means were extremely simple from today’s point of view, which may speak for the question having life of its own, independent of technical developments. I would also like to show on the one hand above all how Noll’s question regarding ‘human or machine’ is inspired by the Turing test,

but, on the other hand, that his own approach simultaneously proves how questionable this either/or is. It shows how much human work had to be invested and made invisible before generated images could be viewed under the assumption of the apparent alternative between human or machine. By making this human work visible again and bringing it into focus, we, however, also gain insight into a practice of picture-making that involves humans, machines, and tools. As I will show in the second section, even recent examples of AI art based on the latest machine learning techniques cannot dispense with human work and should therefore be regarded as part of practices involving humans. In doing so, I will provisionally distinguish different stages or layers of human work that are still, and probably will be for the foreseeable future, indispensable in order for computer-generated images to become part of a creative practice and manifest a claim to art.

Michael Noll's Early Computer Art: 'Human or Machine?'

The question of whether computers can make art is apparently as old as the first experiments to generate images with computers. This observation is surprising from the perspective of today, because these beginnings were technically relatively primitive and would hardly force such a question on us anymore. This suggests that questions like these should apparently not be understood as direct reflections on the concrete state of technology, but instead have a life of their own and seem to be readily revived every time the technology achieves astonishing advances.

It was the engineer A. Michael Noll at Bell Labs in Murray Hill, New Jersey, who was one of the first pioneers to use computers to generate pictures in the early 1960s.¹ He presented the first results of his programs running on an IBM 7090 computer and prints produced with a microfilm printer in a technical report from 1962. Noll cautiously spoke of 'patterns', in order to avoid any 'unintentional debate at this time on whether the computer-produced designs are truly art or not' (Noll 1962, 1). The entire report is, however, written in view of the speculation, which is only formulated in the final sentence, namely that, with a bit more research, 'the programmer-artist might be better prepared to produce not just haphazard patterns, but 'true art'' (Noll 1962, 4).

Noll, in fact, soon began to tackle this challenge head on. His *Computer Composition with Lines* (fig. 1) dates from 1964/65 and obviously imitates Piet Mondrian's 1917 *Composition with Lines* (fig. 2). The significance of this experimental imitation is revealed in Noll's paper 'Human or Machine', which was published in the journal *The Psychological Record* in 1966 (Noll 1966a). Here, Noll first emphasizes Mondrian's art

1 For a retrospective account of the situation and his work at Bell Labs, see Noll 1994, 39. For more detailed accounts of Noll's experiments, see Taylor 2012 and Crowther 2019, 36–41.


historical importance and then briefly explains the technical approach to recreating a painting like Mondrian's *Composition* using a programmable computer.² The focus of the article, however, is on an inquiry into how the two pictures are perceived. More specifically, with the involvement of about a hundred employees of the Bell Labs, Noll conducted a small survey that asked all the participants which picture was 'generated by a machine', which image was 'painted by a human' (Noll 1966a, 4), and which of the two images they preferred. Noll concludes that only 28 per cent of the participants were able to correctly attribute the pictures to the computer or the human, while 59 per cent simultaneously preferred his computer-generated picture to Mondrian's work (Noll 1966a, 4–9).

Figure 1: A. Michael Noll, *Computer Composition with Lines*, 1964/65. Source: A. M. Noll, <http://noll.uscannenberg.org/>



2 See Noll 1966a, 1–4, and, in greater depth, 1966b.

Figure 2: Piet Mondrian, *Composition with Lines*, 1917. Source: Mediathek, Heinrich-Heine-Universität Düsseldorf, Institut für Kunstgeschichte, prometheus.



At first glance, Noll's question of 'human or machine' seems to be aligned with the Turing test. The question would thus be whether computers can make art, and the criterion for determining the answer would be whether people are able to distinguish the computer-generated picture from the human artwork.³ The results of Noll's survey would therefore suggest that computers can make art. Such an approach, however, would require treating the picture 'generated by a machine' as if it were a work by a computer and produced by it autonomously. By contrast, Noll, however, emphasizes his own role, that is, his own work on the conception and production of the computer-generated picture. Noll was very transparent in this regard from the very beginning, when he first explained that he chose Mondrian's *Composition* from 1917 because such a picture, due to its formal simplicity, is within the range of what he can achieve with his computer equipment despite its technical limitations (Noll 1966a, 3). He then wrote, in a trial-and-error approach, a FORTRAN program whose sole purpose was to produce something like Mondrian's picture (Noll 1966a, 3f., and 1966b, 68 and 70). Furthermore, he ran the program several times and obtained different pictures (fig. 3), because the program executed the concrete design depending on random numbers and was thus able to generate new pictures again and again (Noll 1966b, 70f.).

3 For such a transfer of the Turing test to the field of the fine arts and the question of creativity, see Boden 2010.

Figure 3: A. Michael Noll, *Four computer-generated random patterns based on the composition criteria of Mondrian's Composition with Lines*. Source: Noll, *Computers and the Visual Arts*, 71.





Figure 4: A. Michael Noll, *Picture pair as presented on separate sheets to the participants of his survey*. Source: Noll, *Human or Machine*, 5.



In the next step, he selected the picture that, in his eyes, best matched Mondrian's *Composition* and thus promised to be difficult to identify as the computer-generated picture. Finally, he cropped the circular computer-generated picture a little at the top and bottom to match the shape of Mondrian's picture and also copied

both pictures xerographically so that the microfilm print would not be easily distinguishable from the painting and its reproduction due to having different material appearance and properties (Noll 1966a, 4f., 10). Noll thus makes it very clear that before he could ask 'man or machine', he had to invest a great deal of human work not only in the computer-generated picture, but also in staging the visual comparison of his computer-generated picture and Mondrian's painting within his survey (fig. 4).

The similarity of Noll's survey to the Turing test, which Noll does not mention in his article, may thus suggest itself, but it is at best, as Noll himself formulated retrospectively, a 'crude approximation to Turing's experiment': 'In a sense, the computer with its program could be considered creative, although it can be argued that human creativity was involved in the original program with the computer performing only as an obedient tool.' (Noll 1967, 92f.) Consequently, the question of 'human or machine' in Noll's survey can hardly be understood in the sense of whether a human or a machine created a work independently, and there is nothing to suggest that the aim is for the computer to create art, hitherto thought to be the prerogative of humans. Since this approach would presuppose leaving out the programmer on the one hand, in order to present the result as the product of solely the computer, and leaving out traditional tools and aids on the other, as if a human being could produce a painting without making use of brushes and paint. Noll's survey instead inquires whether the participants are able to distinguish between the pictures made by a traditional artist and painter or a new kind of artist who operates a computer. It thus compares the traditional artistic practice without a machine, but with a brush and canvas, and a creative practice with machine, that is, a computer. It is therefore not surprising that Noll occasionally compares the computer to the brush instead of allowing it to replace the artist (Noll 1966b, 71, and 1967, 90).

Noll is therefore not concerned primarily with the question of whether computers can make art, but instead with showing that people can also make art with computers, just as they did before with brushes and oils. In other words, he aims at the computer as an 'artistic' and 'creative medium' (Noll 1967, 89). For this reason, he emphasizes the indispensable role of the new 'programmer-artist' (Noll 1966a, 9) or 'computer-artist' (Noll 1966b, 71) to an equal extent as the role of the computer, its own mode of operation in general, and its possible randomization of the designs executed in particular. Because of the inherent complexity of the computer, he understands it not merely as a new tool or material, but much more as 'assistance' (Noll 1966a, 9) or as an 'intellectual and active partner' (Noll 1967, 89) of a new type that artists should engage with.

Thus, at the beginnings of computer art, two ways of thinking about computer and art together emerge. We can either try, as Noll suggests, to envision and explore creative practices of humans *and* machines, in which 'a tight interaction between artist and computer constitutes a totally new, active, and exciting artistic medium' and allows for 'new art forms and possibly new aesthetic experiences' (Noll 1967,

89f.).⁴ Or we can try to replace the artist with the computer, but we will then probably embrace a questionable model of artistic creation, disguise the actual human work involved, and possibly perpetuate traditional aesthetics.⁵ As I will show in the following section, this conceptual alternative is still relevant despite the impressive technological progress and recent advances in picture generation based on machine learning methods. These new methods are, however, mostly discussed under the latter perspective, which is not only inappropriate, because they still cannot do without human work, but also occludes the artistic practices and potentials for which they might otherwise pave the way.

AI, Art, and Practice

If we follow the headlines of recent years, there seems to be little doubt: computer programs based on AI technologies are making art. Last year, an AI art picture titled *Théâtre d'opéra Spatial* (fig. 5) won an art prize for the first time, at the Colorado State Fair Fine Arts Competition, earning Jason M. Allen, who submitted the picture, 300 dollars in prize money (Roose 2022). In 2018, the *Portrait of Edmond de Belamy* (fig. 6) by the Paris-based artist collective Obvious already made headlines as the 'first work of art produced by artificial intelligence to be sold at auction' (Alleyne 2018), achieving the astonishing price of 432,500 dollars.

Let us leave aside for a moment the question of whether we are dealing here with works of art created by computers based on machine learning methods or, ultimately, with well-staged stunts for the art market and the press. For it is important first and foremost to take note of the amazing and impressive progress that computer-assisted picture generation has made. Until the advances in machine learning for picture generation in the last decade, it was basically necessary to write, as Noll did, a program which specifies how a picture is assembled from elementary graphical operations explicitly and step by step. If the outcome was not always the same and unpredictable, it was only because random values or the like affected the graphical output.

The computer generation of pictures today works very differently because it is based on machine learning methods and, in particular, on artificial neural networks (ANN). To enable such networks to perform certain tasks, they are not programmed, but instead trained on data that specify the desired task.

4 There is no question that Noll's experiments and particularly his re-creation of Mondrian's *Composition* fall short here, as he himself also concluded only a few years later (Noll 1970, 10f.).

5 The perpetuation of traditional aesthetics along this path in the field of literature is scrutinized in Hannes Bajohr 2021.

Figure 5: Jason M. Allen, Théâtre d'opéra Spatial, 2022. Source: Wikimedia Commons.



Figure 6: Obvious, Portrait of Edmond de Belamy, 2018. Source: Wikimedia Commons.



In the case of the generation of pictures, they must therefore be trained on many pictures so that the trained network should eventually be able to produce analogue pictures in turn. In recent years, various methods have been developed in this framework and two types of ANNs in particular have become established. The two pictures discussed above are examples from these two most successful approaches.

The *Portrait of Edmond de Belamy* was generated by a generative adversarial network (GAN). This approach was introduced in 2014 and denotes the first successful artificial neural network able to generate high-quality pictures.⁶ A GAN consists of the combination of two artificial neural networks: The first network, called a 'discriminator', is trained to distinguish elements from a given training dataset, for example, photographic portraits of human beings, from other elements or pictures. The second network, called a 'generator', is then trained to produce pictures that the discriminator is nevertheless unable to distinguish from the original training data and therefore mistakes them, for example, for photographic portraits even if they were computed from scratch.

GANs defined the state of the art in machine learning for picture generation for several years, but were outperformed last year by a new approach, which was also used to generate Jason M. Allen's award-winning *Théâtre d'opéra Spatial*.⁷ So-called diffusion models are based on an idea that may seem curious at first glance. It is simple to diffuse any picture by gradually adding some Gaussian noise. But could we then not attempt to de-noise the blurry picture that results step by step in order to finally reconstruct a picture that resembles the original one as much as possible? This is indeed possible and precisely the task that so-called diffusion models are trained for. At the same time, they need other input to guide the process of de-noising: this guidance is provided by a text prompt describing the target picture. This approach also created such a buzz in 2022 because the technology became available very quickly and is relatively easy to use, from Midjourney by the company of the same name, which was used by Jason M. Allen, to OpenAI's DALL-E 2 and Stability AI's Stable Diffusion, which has become the most popular model: it requires less computing power and was made openly available, so that it was used so frequently that it nearly gave rise to its own flood of images.⁸

Given these unquestionable advances in computer-assisted picture generation over all previous approaches and especially over the simple programs from the 1960s, it may be tempting to take up Noll's question of 'man or machine' and answer it now

6 For GANs used in the field of AI-generated art, see the seminal paper Goodfellow/Pouget-Abadie/Mirza et al. 2014 and the overview by Maerten/Soydaner 2023, 14–17.

7 The seminal paper here is Ho/Jain/Abbeel 2022. See also Marten/Soydaner 2023, 19–22, for an overview of different diffusion models.

8 Decisive advances in efficiency were achieved in Rombach/Blattmann/Lorenz et al. 2022.

without hesitation in favour of the machine. In view of the impressive technological progress, it nonetheless seems perfectly reasonable to hesitate for a moment and take a closer look before believing that the 'programmer-artist' has been replaced by the recent models of machine learning or the 'program as artist'. Is the advance in methods and techniques really so fundamental that we should claim that the machine now makes art? Or should we not ask, following our reflections on Noll's approach in the first section, how much human work had to be put into the machine before we feel pushed to the alternative of 'man or machine' or even think that we have to decide it in favour of the computer?

In fact, a closer look quickly reveals how much human work goes into such computer-generated art. I would like to provisionally distinguish four different steps or layers of human work based on my two examples. A first, quite simple aspect refers to the final stage in the making of the picture. In the two most recent examples of putative AI art, the computer-generated picture was apparently edited by hand. In the case of *Théâtre d'opéra Spatial*, Jason M. Allen 'cleaned up' the image 'by giving one of the female figures in his winning image a head with wavy, dark hair after Mid-journey had rendered her headless' (Metz 2022). Such an intentional and manual intervention can most likely also be found in the *Portrait of Edmond de Belamy*: It is evident that the mathematical formula in the lower right corner of the picture was added manually by the art collective Obvious in order to mimic a signature. A first stage or layer of the involvement of humans in computer-generated art is thus the intentional and manual editing of computer-generated pictures.

Moreover, Obvious's pseudo-signature hints at a second layer or stage of human involvement. The pseudo-signature in the form of a mathematical formula is apparently a means to emphasize and orchestrate the computer's authorship, but it goes back to a manual addition by the human artists. This indicates that it is not the computer that claims to have produced a work of art, for, in these pictures, as always, the computer has simply calculated an output from an input. Read against the grain, then, the manually added pseudo-signature reveals that human involvement is required to transform a computer output into a work of art, that is, to first select suitable outputs, present them as works of art, and thus finally feed them into the art system or the art market. As long as art is primarily addressed to human beings, it hardly seems conceivable that art can do without this second layer or stage of human work, which I would characterize as curatorial practice.

These first two aspects of human involvement take computer-generated pictures as starting point. Human involvement, however, already took place before a picture could be generated by a computer in the first place. Not only do the artificial neural networks presuppose training data that are naturally taken from human culture, which means that the machine learning-based picture generation ultimately incorporates the history, genres, styles, and media of pictorial representation, a process that entails many ethical and legal problems, from machine bias to copyright issues.

Similarly, the procedures employed, from network architectures to training algorithms, while based on a great deal of mathematics, are cultural products through and through: They are ultimately engineering accomplishments developed for a very specific task, such as the generation of pictures, and thus embody cultural values and economic aims in many ways, from the ideals of realistic and coherent depiction, which are thoroughly conventional and normative, to goals of increased efficiency, automation, or new products and services.⁹ At the level of technical procedures and computational processes, we must therefore already presuppose a fundamental layer of human involvement in the form of cultural and economic aims and contexts, without which not even one single picture could be computed.

The fourth layer or stage of human involvement in computer-generated art has to do with the concrete use of programs or machine learning models: Because the computer does not generate output without input, humans are also necessarily involved in the input stage. Depending on the methods, models, and programs used, the input can take on different forms. In the case of diffusion models, which are currently attracting so much attention, it was the standard until recently that they required a short text, a so-called prompt, as input. This text should characterize the target picture and is necessary to guide the technical de-noising process. The power of the new technology is often demonstrated by presenting this text along with the selected and, naturally, convincing result. This form of presentation, however, obscures the fact that, in the use of these tools in practice, the relationship between text input and picture output is anything but trivial.

This denotes, first of all, that it is by no means the case that changes in the picture output following a modification of the text prompt would be readily understandable. As can easily be seen by tinkering around with models like Stable Diffusion or in experiments accessible on the Internet, many surprises lurk here, and the addition of even one or more phrase points at the end of a text prompt, which may be rather meaningless for us, can significantly affect the output.¹⁰ In order to actually use such a model or tool productively, some practical experience on the part of the human user is thus definitely required. In addition, even with a more intuitive input/output behaviour, any use of such tools requires some kind of intermodal translation from text to image: one paints and draws with words, as it were, which would

9 I am unable to address this argument in detail here for reasons of space. But the approach of relating culture to its tools and technologies and simultaneously considering the latter within the context of cultural conditions and purposes has a long and varied tradition, which can be linked, for example, to the history of technology and the journal *Technology and Culture* or to the popular account by Arnold Pacey (1983), which includes institutional and organizational points of view. See also Damisch 1963 for a similar approach with reference to the technical picture of photography.

10 For example, the 'Stable Diffusion Prompt Guide' on the youtube channel of Nerdy Rodent, accessible via the link <https://www.youtube.com/watch?v=c5dHlZoRyMU>.

be a completely new situation for designers and artists and simultaneously open up undreamt-of possibilities for creative laymen to produce pictures. At the same time, text input allows only limited control over image output. It therefore comes as no surprise that since the beginning of this year, newer tools and models have been emerging (such as M-VADER or ControlNet), which permit people to characterize the target picture not only by means of a text prompt, but simultaneously through a picture input and its various algorithmically analysable properties.¹¹ What might considerably facilitate the integration of such tools into design practice, however, simultaneously indicates that human involvement is required here, from a practice of interaction and interplay between the user and the computer in general to the person of a designer who is experienced in handling these tools and knows how to use their input/output behaviour productively for his or her creative practice.¹²

Conclusion

The recent advances in machine learning-based picture generation are certainly impressive. Nevertheless, my argument is that we are not in a fundamentally new situation compared with Noll's early experiments with computer-assisted picture generation, even though the latest models and tools are technically far more complex than Noll's simple and short programs: Before we can finally ask the question of 'human or machine', we are nonetheless investing a lot of human work in computer-generated pictures or artworks, from the technical development of the methods and models and identifying suitable input, to the selection and editing of the output and feeding the result into the art system and art market. In other words, this alternative is only conceivable because we have made a lot of human work invisible, so that it can appear as if such a picture was generated solely by a computer or machine.¹³ The fear that the human artist will be soon replaced by the art-creating computer therefore seems to me to be a misguided and short-sighted conclusion influenced by the stunning new pictures.

This argument is by no means about preserving for humans the privilege of creativity, from which machines would remain excluded and would therefore be barred from making art themselves. For just as little as machines, human beings also do not

11 See Weinbach/Bellagente/Eichenberg 2022 and Zhang/Agrawala 2023.

12 To put it another way, it takes human work to make the computer and specific tools a part of a human and social practice. This type of argument was, to my knowledge, first developed with reference to the pocket calculator and its integration into arithmetic practices in Collins/Kusch 1999, 121–24, and in greater detail in Collins 1990, 62–71.

13 By invisible human work, I do not primarily mean the clickworkers who often label the datasets needed to train the machine learning models, Irani 2016; Gray/Suri 2019; Crawford 2021, 63–69.

make art on their own. We should instead take into account the creative practices in which, alongside human beings, instruments and tools, techniques and materials have always played a central role, and in which intensive engagement with the latter has always also been a source of creativity.¹⁴ Since Noll's time, the computer, various applications, and software suites have thus transformed creative practices in many ways and the new machine learning-based tools will bring further sweeping changes.

The danger here is therefore not so much that the computer will replace the human being, but that artists and designers who use these new tools will probably replace the artists and designers who do not work with them. The real challenge, then, is to use the new tools not merely to automate standard workflows and reproduce traditional aesthetics, but also to engage productively and creatively with the new technology, which in no way precludes critical scrutiny. For the purpose of leveraging the creative potential of the new tools, we can and must develop creative practices in art and design in the context of and in confrontation as well as collaboration with these new technologies. Given the rather primitive beginnings of the computer-assisted generation of pictures in the 1960s, Noll's hope for a 'new creative partnership and collaboration between artist and the computer' (Noll 1967, 93) seems to have a simply counter-factual and, in this sense, utopian character. With respect to the new machine learning tools, this hope is perhaps a bit more realistic, since these tools do not simply execute programmed instructions. They are much more complex and operate on the basis of what they are able to extract from the training data, in this case usually a huge set of images, and thus certainly open up new possibilities for interaction and collaboration.¹⁵ But, even today, Noll's hope for a 'new creative partnership' can barely conceal its utopian touch, if only because the new technology and its widespread application will not only bring about many changes in practice, but will also entail economic and legal conflicts and necessitate political debates and regulatory measures.

References

Alleyne, Allyssia (2018). A Sign of Things to Come? AI-Produced Artwork Sells for \$433K, Smashing Expectations. CNN.com, 25 October 2018. Available online at <https://edition.cnn.com/style/article/obvious-ai-art-christies-auction-smart-creativity/index.html> (all URLs here accessed in August 2023).

14 Pointing in this direction is also the treatment of Obvious's *Portrait of Edmond de Belamy* in the context of a critical discussion of the concept of creativity in Stephensen 2019, esp. 27f.

15 In an epistemological context, I have already argued in a similar but more detailed way in Schubbach 2021.

- Bajohr, Hannes (2021). Keine Experimente. Über künstlerische Künstliche Intelligenz. *Merkur* 75 (864), 32–44. Available online at <https://hannesbajohr.de/wp-content/uploads/2021/04/Bajohr-2021-Keine-Experimente-U%CC%88ber-ku%CC%88nstlerische-Ku%CC%88nstliche-Intelligenz.pdf>.
- Boden, Margaret A. (2010). The Turing Test and Artistic Creativity. *Kybernetes* 39 (3), 409–13. <https://doi.org/10.1108/03684921011036132>.
- Collins, Harry M. (1990). *Artificial Experts: Social Knowledge and Intelligent Machines*. Cambridge, MA, The MIT Press. <https://doi.org/10.7551/mitpress/1416.001.0001>.
- Collins, Harry M./Kusch, Martin (1998). *The Shape of Actions: What Humans and Machines Can Do*. Cambridge, MA, The MIT Press. <https://doi.org/10.7551/mitpress/6200.001.0001>.
- Crowther, Paul (2019). *Digital Art, Aesthetic Creation: The Birth of a Medium*. New York and London, Routledge. <https://doi.org/10.4324/9780429467943>.
- Damisch, Hubert (1978). Five Notes for a Phenomenology of the Photographic Image. October 5, 70–72. <https://doi.org/10.2307/778645>.
- Goodfellow, Ian J./Pouget-Abadie, Jean/Mirza, Mehdi et al. (2014). *Generative Adversarial Nets: Proceedings of the 27th International Conference on Neural Information Processing Systems (NIPS'14)*, 2672–80. Cambridge, MA, The MIT Press. <https://doi.org/10.48550/arXiv.1406.2661>.
- Gray, Mary L./Suri, Siddharth (2019). *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Boston, Houghton Mifflin Harcourt.
- Ho, Jonathan/Jain, Ajay/Abbeel, Pieter (2020). Denoising Diffusion Probabilistic Models. *arXiv:2006.11239*. <https://doi.org/10.48550/arXiv.2006.11239>.
- Irani, Lily (2016). The Hidden Faces of Automation. *XRDS* 23 (2), 34–37. <https://doi.org/10.1145/3014390>.
- Maerten, Anne-Sofie/Soydaner, Derya (2023). From Paintbrush to Pixels: A Review of Deep Neural Networks in AI-Generated Art. *arXiv:2302.10913*. <https://doi.org/10.48550/arXiv.2302.10913>.
- Metz, Rachel (2022). AI Won an Art Contest, and Artists Are Furious. *CNN.com*, 3 March 2022. Available online at <https://edition.cnn.com/2022/09/03/tech/ai-art-fair-winner-controversy/index.html>.
- Noll, A. Michael (1962). *Patterns by 7090: Technical Memorandum of the Bell Phone Laboratories*. Available online at <http://noll.uscannenberg.org/Art%20Papers/BTL%201962%20Memo.pdf>.
- Noll, A. Michael (1966a). Human or Machine: A Subjective Comparison of Piet Mondrian's 'Composition with Lines' (1917) and a Computer-Generated Picture. *The Psychological Record* 16, 1–10. <https://doi.org/10.1007/BF03393635>.
- Noll, A. Michael (1966b). Computers and the Visual Arts. *Design Quarterly* 66/67, 64–71. Available online at <http://noll.uscannenberg.org/ComputerArt.htm>.

- Noll, A. Michael (1967). The Digital Computer as a Creative Medium. *IEEE Spectrum* 4 (10), 89–95. Available online at <http://noll.uscannenberg.org/ComputerArt.htm>.
- Noll, A. Michael (1968). A Subjective Comparison of Piet Mondrian's 'Composition with lines' 1917. In: Jasia Reichardt (Ed.). *Cybernetic Serendipity: The Computer and the Arts*. London and New York, Studio International Foundation, 74. Available online at https://monoskop.org/images/2/25/Reichardt_Jasia_ed_Cybernetic_Serendipidity_The_Computer_and_the_Arts.pdf.
- Noll, A. Michael (1970). Art Ex Machina. *IEEE Student Journal* 8 (4), 10–14. Available online at <http://noll.uscannenberg.org/Art%20Papers/Art%20ex%20Machina.pdf>.
- Noll, A. Michael (1994). The Beginnings of Computer Art in the United States: A Memoir. *Leonardo* 27 (1), 39–44. Available online at <http://noll.uscannenberg.org/Art%20Papers/Memoir.pdf>.
- Pacey, Arnold (1983). *The Culture of Technology*. Cambridge, MA, The MIT Press. Available online at <https://hdl.handle.net/2027/hebo1153.0001.001>.
- Rombach, Robin/Blattmann, Andreas/Lorenz, Dominik et al. (2022). High-Resolution Image Synthesis with Latent Diffusion Models. *arXiv:2112.10752*. <https://doi.org/10.48550/arXiv.2112.10752>.
- Roose, Kevin (2022). An A.I.-Generated Picture Won an Art Prize: Artists Aren't Happy. *The New York Times*, 2 September 2022. Available online at <https://www.nytimes.com/2022/09/02/technology/ai-artificial-intelligence-artists.html>.
- Schubbach, Arno (2021). Judging Machines: Philosophical Aspects of Deep Learning. *Synthese* 198, 1807–27. <https://doi.org/10.1007/s11229-019-02167-z>.
- Stephensen, Jan Løhmann (2019). Towards a Philosophy of Post-Creative Practices?—Reading Obvious' Portrait of Edmond de Belamy'. In: *Proceedings of POM Beirut 2019*, 21–30. <http://dx.doi.org/10.14236/ewic/POM19.4>.
- Taylor, Grant (2012). Routing Mondrian: The A. Michael Noll Experiment. *Media-N* 8 (2). Available online at <https://median.newmediacaucus.org/routing-mondrian-the-a-michael-noll-experiment/>.
- Turing, Alan (2004). Computing Machinery and Intelligence. In: B. Jack Copeland (Ed.). *The Essential Turing: Seminal Writings in Computing, Logic, Philosophy, Artificial Intelligence, and Artificial Life plus The Secrets of Enigma*. Oxford, Oxford University Press, 441–64. <https://doi.org/10.1093/oso/9780198250791.003.0017>.
- Weinbach, Samuel/Bellagente, Marco/Eichenberg, Constantin et al. (2022). M-VADER: A Model for Diffusion with Multimodal Context. *arXiv:2212.02936*. <https://doi.org/10.48550/arXiv.2212.02936>.
- Zhang, Lvmin/Agrawala, Maneesh (2023). Adding Conditional Control to Text-to-Image Diffusion Models. *arXiv:2302.05543*. <https://doi.org/10.48550/arXiv.2302.05543>.