Art, Creativity, and the Potential of Artificial Intelligence

Marian Mazzone [1,*] and **Ahmed Elgammal** [2,*]

1   Department of Art & Architectural History, College of Charleston, Charleston, SC 29424, USA
2   Department of Computer Science, Rutgers University, New Brunswick, NJ 08901-8554, USA
*   Correspondence: mazzonem@cofc.edu (M.M.); elgammal@cs.rutgers.edu (A.E.)

Abstract: Our essay discusses an AI process developed for making art (AICAN), and the issues AI creativity raises for understanding art and artists in the 21st century. Backed by our training in computer science (Elgammal) and art history (Mazzone), we argue for the consideration of AICAN's works as art, relate AICAN works to the contemporary art context, and urge a reconsideration of how we might define human and machine creativity. Our work in developing AI processes for art making, style analysis, and detecting large-scale style patterns in art history has led us to carefully consider the history and dynamics of human art-making and to examine how those patterns can be modeled and taught to the machine. We advocate for a connection between machine creativity and art broadly defined as parallel to but not in conflict with human artists and their emotional and social intentions of art making. Rather, we urge a partnership between human and machine creativity when called for, seeing in this collaboration a means to maximize both partners' creative strengths.

Keywords: artificial intelligence; art; creativity; computational creativity; deep learning; adversarial learning

---

1. AI-Art: GAN, a New Wave of Generative Art

Over the last 50 years, several artists and scientists have been exploring writing computer programs that can generate art. Some programs are written for other purposes and are adopted for art making, such as generative adversarial networks (GANs). Alternatively, programs can be written that intend to make creative outputs. Algorithmic art is a broad term that points to any art that cannot be created without the use of programming. If we look at the Merriam-Webster definition of art, we find "the conscious use of skill and creative imagination especially in the production of aesthetic objects; the works so produced". Throughout the 20th century, that understanding of art has been expanded to include objects that are not necessarily aesthetic in their purpose (for example, conceptual art), and not created physical objects (performance art). Since the challenges of Marcel Duchamp's practice, the art world has also relied on the determination of the artist's intention, institutional display, and audience acceptance as critical defining steps to decide whether something is "art".

The most prominent early example of algorithmic art work is by Harold Cohen and his program AARON (aaronshome.com). American artist Lillian Schwartz, a pioneer in using computer graphics in art, also experimented with AI (Lillian.com). However, in the last few years, the development of GANs has inspired a wave of algorithmic art that uses Artificial Intelligence (AI) in new ways to make art (Schneider and Rea 2018). In contrast to traditional algorithmic art, in which the artist had to write detailed code that already specified the rules for the desired aesthetics, in this new wave, the algorithms are set up by the artists to "learn" the aesthetics by looking at many images using machine learning technology. The algorithm only then generates new images that follow the aesthetics it has learned.

Figure 1 explains the creative process that is involved in making this kind of AI art. The artist chooses a collection of images to feed the algorithm (pre-curation), for example, traditional art portraits. These images are then fed into a generative AI algorithm that tries to imitate these inputs. The most widely used tool for this is generative adversarial networks (GANs), introduced by Goodfellow in 2014 (Goodfellow et al. 2014), which have been successful in many applications in the AI community. It is the development of GANs that likely sparked this new wave of AI Art. In the final step, the artist sifts through many output images to curate a final collection (post-curation).
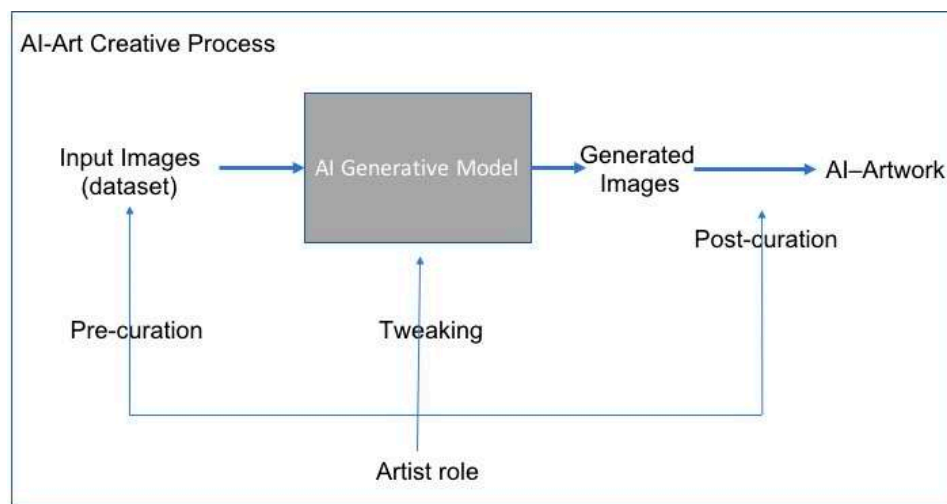
Figure 1. A block diagram showing the artist's role using the AI generative model in making art. Diagram created by author A. Elgammal.

In this kind of procedure, AI is used as a tool in the creation of art. The creative process is primarily done by the artist in the pre- and post-curatorial actions, as well as in tweaking the algorithm. There have been many great art works that have been created using this pipeline. The generative algorithm always produces images that surprise the viewer and even the artist who presides over the process.

Figure 2 is an example of what a typical GAN trained on portrait paintings would produce. Why might we like or hate these images, and should we call them art? We will try to answer these questions from a perception and a psychology point of view. Experimental psychologist Daniel E. Berlyne (1924–1976) studied the basics of the psychology of aesthetics for several decades and pointed out that *novelty*, *surprisingness*, *complexity*, *ambiguity*, and *puzzlingness* are the most significant properties in stimulus relevance to studying aesthetic phenomena (Berlyne 1971). Although there are several alternative newer theories than Berlyne's, we use it in our explanation for its simplicity as the explanation does not contradict other theories. Indeed, the resulting images with all the deformations in the faces are novel, surprising, and puzzling to us. In fact, they might remind us of Francis Bacon's famous deformed portraits such as *Three Studies for a Portrait of Henrietta Moraes* (1963). However, this comparison highlights a major difference, that of intent. It was Bacon's intention to make the faces deformed in his portrait, but the deformation we see in the AI art is not the intention of the artist nor of the machine. Simply put, the machine fails to imitate the human face completely and, as a result, generates surprising deformations. Therefore, what we are looking at are failure cases by the machine that might be appealing to us perceptually because of their novelty as visual stimuli compared to naturalistic faces. However, these "failure cases" have a positive visual impact on us as viewers of art; only in these examples, the artist's intention is absent.

Figure 2. Examples of images generated by training a generative adversarial network (GAN) with portraits from the last 500 years of Western art. The distorted faces are the algorithm's attempts to imitate those inputs. Images generated at Art & Artificial Intelligence Laboratory, Rutgers.

So far, most art critics have been skeptical and usually evaluate only the resulting images while ignoring the creative process that generates them. They might be right that images created using this type of AI pipeline are not that interesting. After all, this process just imitates the pre-curated inputs with a slight twist. However, if we look at the creative process overall and not simply the resulting images, this activity falls clearly in the category of conceptual art because the artist has the option to act in the choice-making roles of curation and tweaking. More sophisticated conceptual work will be coming in the future as more artists explore AI tools and learn how to better manipulate the AI art creative process.

2. Pushing the Creativity of the Machine: Creative, Not Just Generative

At Rutgers' Art & AI Lab, we created AICAN, an almost autonomous artist. Our goal was to study the artistic creative process and how art evolves from a perceptual and cognitive point of view. The model we built is based on a theory from psychology proposed by Colin Martindale (Martindale 1990). The process simulates how artists digest prior art works until, at some point, they break out of established styles and create new styles. The process is realized through a "creative adversarial network (CAN)" (Elgammal et al. 2017), a variant of GAN that we proposed that uses "stylistic ambiguity" to achieve novelty. The machine is trained between two opposing forces—one that urges the machine to follow the aesthetics of the art it is shown (minimizing deviation from art distribution), while the other force penalizes the machine if it emulates an already established style (maximizing style ambiguity). These two opposing forces ensure that the art generated will be novel but at the same time will not depart too much from acceptable aesthetic standards. This is called the "least effort" principle in Martindale's theory, and it is essential in art generation because too much novelty would result in rejection by viewers. Figure 3 illustrates a block diagram of the CAN network where the generator receives two signals, one measuring the deviations from art distribution and the second measuring style ambiguity. The generator tries to minimize the first to follow aesthetics and maximize the second to deviate from established styles.
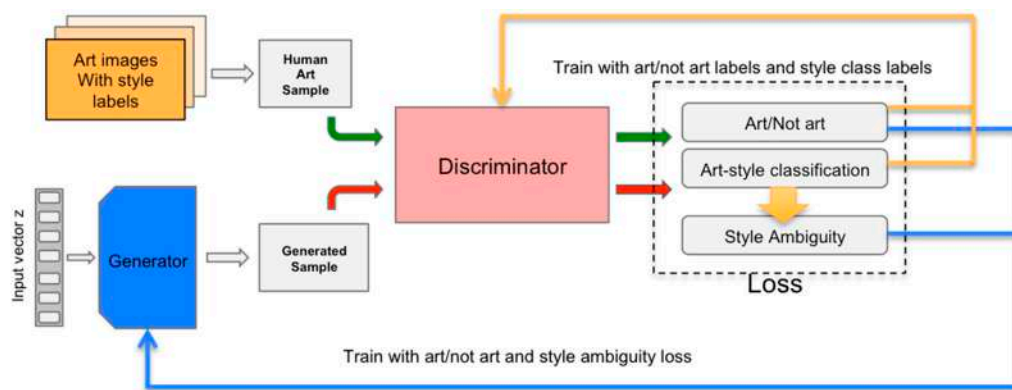
Figure 3. A block diagram of a creative adversarial network. The generator explores the creative space by trying to generate images that maximize style ambiguity while minimizing deviation from art distribution. Diagram by author A. Elgammal.

Unlike the generative AI art discussed earlier, this process is inherently creative. There is no curation on the dataset; instead, we fed the algorithm 80K images representing 5 centuries of Western art history, simulating the process of how an artist digests art history, with no special selection of genres or styles. The generative process using CAN is seeking innovation. The outputs surprise us all the time with the range of art AICAN generates. Figure 4 shows the variety of AICAN-generated art.
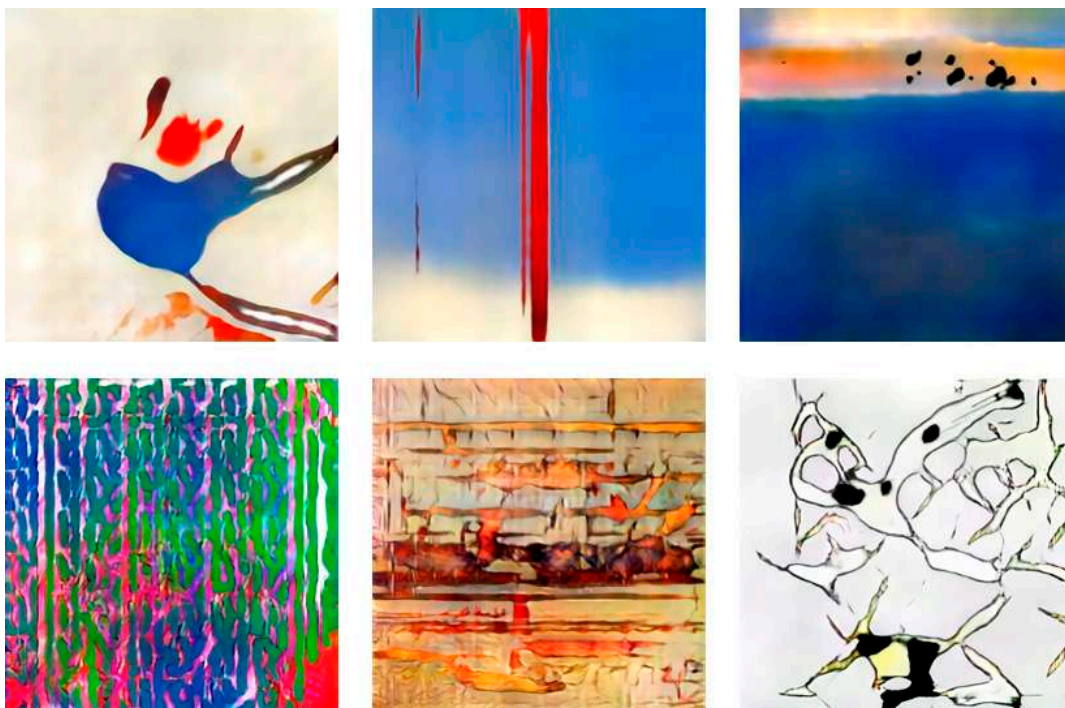
Figure 4. Examples of images generated by AICAN after training with images from all styles and genres from the past 500 years of Western art. Images courtesy of the Art & Artificial Intelligence Laboratory, Rutgers.

We devised a visual Turing test to register how people would react to the generated images and whether they could tell the difference between AICAN- or human-created art. To make the test timely and of high quality, we mixed images from AICAN with works from Art Basel 2016 (the flagship art fair in contemporary art). We also used a set of images from abstract expressionist masters as a baseline. Our study showed that human subjects could not tell whether the art was made by a human artist or by the machine. Seventy-five percent of the time, people in our study thought the AICAN

generated images were created by a human artist. In the case of the baseline abstract expressionist set, 85% of the time subjects thought the art was by human artists. Our subjects even described the AICAN-generated images using words such as "intentional", "having visual structure", "inspiring", and "communicative" at the same levels as the human-created art.

Beginning in October 2017, we started exhibiting AICAN's work at venues in Frankfurt, Los Angles, New York City, and San Francisco, with a different set of images for each show (Figure 5). Recently, in December 2018, AICAN was exhibited in the SCOPE Miami Beach Art Fair. At these exhibitions, the reception of works was overwhelmingly positive on the part of viewers who had no prior knowledge that the art shown was generated using AI. People genuinely liked the artworks and engaged in various conversations about the process. We heard one question time and again: Who is the artist? Here, we posit that the person(s) setting up the process designs a conceptual and algorithmic framework, but the algorithm is fully at the creative helm when it comes to the elements and the principles of the art it creates. For each image it generates, the machine chooses the style, the subject, the forms, and composition, including the textures and colors.

Figure 5. Photographs from AICAN exhibition held in Los Angeles in October 2017. Photographs by author A. Elgammal.

3. AI in Art and Art History

The CAN study provoked a number of concerns about AI as a threat or rival to art made by human beings. Yes, the study is interested in the process of art creation, and the more abstract problem of what creativity is and does. However, AI focuses on developing a *machine* process and *machine* creativity, not merely aping and trying to pass as human-made. Our work is focused on understanding the *process of creativity* such that a means can be found to model that process to generate a creative result. One way to do this, and what this study has chosen, is to model the process by which art is taught and then stimulate AICAN to synthesize that style information and next create something new. To do this, the machine was trained on many thousands of human-created paintings in a process parallel to a human artists' experience of looking at other artists' works, learning by example. The AICAN system was then designed to encourage choices that deviate from copying/repeating what had been seen (the GAN function) to encouraging new combinations and new choices based on a knowledge of art styles (the CAN function). If the creation process is modeled successfully, art may result.

One barometer of whether art has been successfully created through the chosen process is whether human beings appreciate it as art and do not necessarily recognize it as AI-derived. AICAN was tasked with creating works that did not default into the familiar psychedelic patterning of most GAN-generated images as a test of its creativity function. Our inclusion of viewer surveys to gauge peoples' responses did not aim to prove that the AICAN artifacts were better than human creations, but rather to gauge whether the AICAN works were aesthetically recognizable as art, and whether human viewers liked the AI-generated works of art. It seemed most pertinent to have viewers assess the AICAN images in a group with other contemporary images rather than historical ones, hence the

choice to select these from Art Basel. The objective was to learn whether AICAN can produce work that is able to qualify or count as art, and if it exhibits qualities that make it desirable or pleasurable to look at. In other words, could AICAN artifacts be recognized as quality aesthetic objects by human beings? Because we used Berlyne's theory of arousal potential, the response of human beings to the images was a necessary check to evaluate the quality level of AICAN creativity.

There may always be a number of artists and art lovers who resist the idea of AI in art because of technophobia. For them, the machine simply has no place in art. In addition, many lack understanding of what AI actually is, how it works, and what it can and cannot be made to do. There is also an element of fear at work, resulting in an imagined future in which AI will commandeer art making and crank out masses of soulless abstract paintings. However, as we discuss throughout this article, AI is really very limited and specific in what it can do in terms of art creation, and it was never our goal is to supplant the role of the human artist. There is simply and profoundly no need to do that. It is an interesting problem in machine learning to model the process of image creation and to explore what creativity might mean within the confines of computation, but these are issues separate and apart from how a human being makes art, and they are not mutually exclusive in any way. The very best outcome we can imagine is a fruitful partnership between an artist and a creative AI system. However, we are in the very early days of developing algorithms for such AI systems.

A comparison with photography is useful because both forms of technology first encountered resistance in the art world based on the use of a machine in the art-making process. This comparison has been discussed widely, including in this issue of *Arts* (Hertzmann 2018) (Agüera y Arcas 2017), so we will not elaborate on it here. A hopeful sign for AI art is that eventually, some photography was fully accepted as art. A key path towards its acceptance was the dialogue that developed between two mediums: Photographers worked to incorporate some of the formal and aesthetic characteristics of painting, while painters were closely looking at photography and shifting painting in response. Painters were inspired by the compositional flatness, capture of movement, and summary edges of the photographic viewfinder. Photographers shifted their approach to lighting, focus, and subject matter as inspired by the aesthetic criteria of painting. Thus, a feedback loop was established between the practices of painting and photography. In both cases, creators began to see *differently* based on their experiences with the other medium. Perhaps this can happen between AI and painting in turn. Currently, most AI systems are trained on thousands of paintings made by Western European and American artists over the last several hundred years. In turn, the AIs create images that speak the language of painting (color choices, form elements, arrangement of forms on a 2-D surface) and depose their elements before the eyes of viewers in a way similar to how we look at paintings. Already, we have contemporary practitioners such as Jason Salavon or Petra Cortright, whose practice demonstrates a lively exchange between the processes of painting and those of computation. Photography did threaten to supplant some of the functions of painting, particularly in those instances when a high degree of naturalistic representation was desirable, such as in portraiture or in topographical representations. Consequently, photography largely did replace painted portraits and most forms of topographical imagery, for example. We imagine that AI-produced art could usefully replace some mass-produced imagery such as decorative art or tourist scenes where repetition of a few pleasing characteristics is desirable. Consumers would be the drivers of this market, electing for the machine-derived images or preferring those created by a human.

Another sound point of comparison is the replicative process of image production employed by both the camera and the computer. Like the camera, the computer provides its user with a range of repetitive and reproductive means to generate multiple images. As noted by Walter Benjamin in the early 20th century (Benjamin [1936] 1969), the impact of the mass production and reproduction of imagery has changed how we think about the originality and the legitimacy of reproductions of works of art, and our viewing experience of art. Most people's experience of art is now soundly in the realm of reproductions, and we ascribe meaningfulness to the experience of the reproduction. Although the singular, original work of art is a paradigm still operational in painting, it is markedly less so in

print making or photography, and completely absent in computational art. Computers can produce many more and varied versions of an image through parameterization, randomizing tools, and other generative processes than can nondigital photography or prints, but the theoretical principle of the multiple still applies. The contemporary art world is well able to theorize and accept multiples or reproductions as legitimate works of art, we believe even at the rate and level of complexity produced by generative computational systems.

There is, however, one profound difference between AI computer-based creativity versus other machine-based image making technologies. Photography, and the similar media of film and video, are predicated on a reference to something outside of the machine, something in the natural world. They are technologies to capture elements of the world outside themselves as natural light on a plate or film, fixed with a chemical process to freeze light patterns in time and space. Computational imagery has no such referent in nature or to anything outside of itself. This is a profound difference that we believe should be given more attention. The lack of reference in nature has historical implications for how we understand something as art. Almost all human art creation has been inspired by something seen in the natural world. There, of course, may be many steps between the inspiration and the resulting work, such that the visual referent can be changed, abstracted or even erased by the final version. However, the process was always first instigated by the artist looking at something in the world, and photography, film, and video retained that first step of the art-making process through light encoding. The computer does not follow this primal pattern. It requires absolutely nothing from the natural world; instead, its "brain" and "eyes" (its internal apparatus for encoding imagery of any kind) consist only of receptors for numerical data. There are two preliminary points to elaborate here: The first relates to issues in contemporary art, the second to the distinction from human creativity.

First, the lack of referent in the natural world and the resulting freedom and range to create or not create any object as a result of the artistic inspiration aligns AI and all computational methods with conceptual art. Like with photography, the comparison with conceptual art has frequently been made for AI and computational methods in general. In conceptual art, the act of the creation of the art work is located in the mind of the artist, and its instantiation in any material form(s) in the world is, as Sol Lewitt (Lewitt 1967) famously declared, "a perfunctory affair. The idea becomes a machine that makes the art." Thus, the making of an art object becomes simply optional. And although contemporary artists in the main have not stopped making objects, the principle that object making is optional and variable in relation to the art concept still remains. We believe this is at the heart of the usefulness of the comparison with conceptual art: The idea or concept is untethered from nature, being primarily located in the synapses of the brain and secondly disassociated from the dictates of the material world. Most AI systems use some form of a neural network, which is modeled on the neural complexity of the human brain. Therefore, AI and conceptual art coincide in locating the art act in the system network of the brain, rather than in the physical output. The physical act of an artist, either applying paint or carving marble, becomes optional. This removes the *necessity* of a human body (the artist) to make things and allows us to imagine that there could be more than one kind of artist, including *other* than human.

4. AI Art: Blurring the Lines between the Artist and the Tool

Many artists and art historians resist seeing work created with AI as art because their definition of art is based on the modern artist figure as the sole locus of art creation and creativity. Therefore, the figure of the artist is necessary to their definition of art. But understanding art as a vehicle for the personal expression of the individual artist is a relatively recent and culturally-specific conception. For many centuries, across many cultures and belief systems, art has been made for a variety of reasons under a wide range of conditions. More often created by groups of people rather than an individual artist (think medieval cathedrals or guild workshops), art is often made to the specifications of patrons and donors large and small, made to order, funded by a wide variety of groups, civic organizations, or religious institutions, and made to function in an extraordinary range of situations. The notion

of a work of art being the coherent expression of the individual's psyche, emotional condition, or expressive point of view begins in the Romantic era and became the prevailing norm in the 19th and 20th centuries in Western Europe and its colonies. Although this remains a common motivation for many artists working today, it does not mean it is the only and correct definition of art. And certainly, it is not a role that any AI system will ever be able to fulfill. Clearly, machine learning and AI cannot replicate the lived experience of a human being; therefore, AI is not able to create art in the same way that human artists do. Thankfully, we are not proposing that it can in our work. Humans and AI do not share all of the same sources of inspiration or intentions for art making. Why the machine makes art is intrinsically different; its motivation is that of being tasked with the problem of making art, and its intention is to fulfill that task. However, we are asking everyone to consider that a different process of creation does not disqualify the results of the process as a viable work of art. Instead consider that without the *necessity* of the individual expressive artist in our definition of art, how we conceptualize art and art making is greatly expanded.

AI is a set of algorithms designed to function as parallel to human intelligence actions such as decision-making, image recognition, language translation/comprehension, or creativity. Elsewhere in this issue of *Arts*, Hertzmann (Hertzmann 2018) makes a point about art algorithms being tools, not artists. As we have argued, we would agree that the algorithms in AI are not artists like human artists. But AI (art generating algorithm in this case) is more than a tool, like a brush with oil paint on it, which is an inanimate and unchanging object. Certainly, artists learn over time and with experience how to better use their tools, and their tools have a role in the physical actions by which they make work in paint. However, the paintbrush does not have the capacity to change, it does not make decisions based on past painting experiences, and it is not trained to learn from data. Algorithms contain all of those possibilities. Perhaps we can conceptualize AI algorithms as more than tools and closer to a *medium*. The word medium in the art world indicates far more than a tool, a medium includes not only the tools used (brush, oil paint, turpentine, canvas, etc.) but also the range of possibilities and limitations inherent to the conditions of creation in that area of art. Thus, the medium of painting also includes a history of painting styles, the physical and conceptual restraints of the 2-D surface, the limits of what can be recognized as a painting, a critical language that has been developed to describe and critique paintings, and so on. Admittedly, we are in the very early days of the medium of AI in art creation, but this medium might encompass tools such as code, mathematics, hardware and software, printing choices, etc., with medium conditions including algorithmic structuring, data collection and application, and the critical theory needed to detect and judge computational creativity and artistic intention within the much larger field of computer science. At this time, a problem is the relatively small number of people able to work creatively in this field or judge the role of the machine in the exercise of creative processes. This will change over time as artists, computer scientists, and historians/critics all become more knowledgeable. For human artists who are interested in the possibilities (and limitations) of AI in creativity and the arts, using AI as a creative partner is already happening now and will happen in the future. In a partnership, both halves bring skill sets to the process of creation. As Hertzmann notes in his article and Cohen discovered in his work with the AARON program, human artists bring capacity for high-quality work, artistic intent, creativity, and growth/change over time. Art is a social interaction. Actually, we think we can argue that AI does a fair amount of this, and it can certainly all be accomplished in a creative partnership between and artist and his or her AI system.

References

Agüera y Arcas, Blaise. 2017. Art in the Age of Machine Intelligence. *Arts* 6: 18. [CrossRef]

Benjamin, Walter. 1969. The Work of Art in Age of Mechanical Reproduction. In *Illuminations*. Edited by Hannah Arendt. New York: Schocken, pp. 217–51. First published 1936.

Berlyne, Daniel E. 1971. *Aesthetics and Psychobiology*. New York: Appleton-Century-Crofts of Meredith Corporation, p. 336.

Elgammal, Ahmed, Bingchen Liu, Mohamed Elhoseiny, and Marian Mazzone. 2017. CAN: Creative adversarial networks, generating "art" by learning about styles and deviating from style norms. *arXiv*, arXiv:1706.07068.

Goodfellow, Ian, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in Neural Information Processing Systems*. Cambridge: MIT Press, pp. 2672–80.

Hertzmann, Aaron. 2018. Can Computers Create Art? *Arts* 7: 18. [CrossRef]

Lewitt, Sol. 1967. Paragraphs on conceptual Art. *Artforum* 5: 79–84.

Martindale, Colin. 1990. *The Clockwork Muse: The Predictability of Artistic Change*. New York: Basic Books.

Schneider, Tim, and Naomi Rea. 2018. Has artificial intelligence given us the next great art movement? Experts say slow down, the 'field is in its infancy. *Artnetnews*. September 25. Available online: https://news.artnet.com/art-world/ai-art-comes-to-market-is-it-worth-the-hype-1352011 (accessed on 3 February 2019).