Essay

Can Computers Create Art? [†]

Aaron Hertzmann

Adobe Research, San Francisco, CA 94103, USA; hertzman@dgp.toronto.edu

[†] This essay expresses my own opinions and not those of my employer.

Received: 4 February 2018; Accepted: 7 May 2018; Published: 10 May 2018

Abstract: This essay discusses whether computers, using Artificial Intelligence (AI), could create art. First, the history of technologies that automated aspects of art is surveyed, including photography and animation. In each case, there were initial fears and denial of the technology, followed by a blossoming of new creative and professional opportunities for artists. The current hype and reality of Artificial Intelligence (AI) tools for art making is then discussed, together with predictions about how AI tools will be used. It is then speculated about whether it could ever happen that AI systems could be credited with authorship of artwork. It is theorized that art is something created by social agents, and so computers cannot be credited with authorship of art in our current understanding. A few ways that this could change are also hypothesized.

Keywords: art and technology; artificial intelligence; photography; painting; computer animation; image stylization

1. Introduction

Artificial Intelligence (AI) research has made staggering advances recently, including many publicly-visible developments in web search, image recognition, conversational agents, and robotics. These developments have stoked fear about Artificial Intelligence's effect on many aspects of society. In the context of art, news media hype presents new image and video creation algorithms as if they are automating the creation of art. Perhaps they will empower everyday users while putting artists out of work . . . and, while they are at it, rob us of our humanity?

Beyond the hype, confusion about how technology influences art pervades serious discussions. Professional artists are often concerned that computers might put them out of work (Nicholas 2017)—a concern I have heard for decades. Some practitioners present their algorithms as themselves, potentially, artists (Colton 2012; Elgammal et al. 2017; Shein 2017), as do some journalists (Dvorsky 2017; Perez 2018; Pogue 2018). And, I was recently contacted by a prominent social psychologist who, inspired by recent results with neural networks, wished to conduct experiments to assess whether ordinary people might be willing to buy artwork made by a computer, and, if so, why. It was assumed that computers were already happily making their own artwork.

On the other hand, when I have informally asked friends or colleagues the question of whether computers can create art, the answer is sometimes a decisive “No.” Art requires human intent, inspiration, a desire to express something. Thus, by definition, there is no such thing as art created by a computer, so why would anyone worry? The concepts of art and inspiration are often spoken of in mystical terms, something special and primal beyond the realm of science and technology; it is as if only humans create art because only humans have “souls.” Surely, there should be a more scientific explanation.

In this essay, I tackle the question of “Can Computers Create Art?” This might seem like a simple question. Sometimes people ask it as a question about technological capabilities, like asking if a new car can go 100 miles per hour. For this sort of question, someone who understands the technology

ought to be able to give a simple yes-or-no answer. But it is not a given that any computer will ever be widely considered as an artist. To date, there is already a rich body of computer-generated art, and, in all cases, the work is credited to the human artist(s) behind the tools, such as the authors or users of the software—and this might never change.

A more precise statement of this essay's question is: "Could a piece of computer software ever be widely credited as the author of an artwork? What would this require?" This is a question of the psychology and philosophy of art; I do not describe the existing technology in much detail. I also discuss the related question of "Will AI put artists out of jobs?" and, more generally, of whether AI will be beneficial to art and artists.

Before directly addressing these questions, I discuss the history and current state of automation for art. I begin with some historical perspectives: previous moments in history when new technologies automated image and film creation, particularly the invention of photography. In each case, we see that these new technologies caused fears of displacing artists, when, in fact, the new technology both created new opportunities for artists while invigorating traditional media.

I argue that new technologies benefit art and artists, creating new tools and modes of expression, and new styles of expression. Sadly, art and science are often viewed as being separate, or, even, in opposition (Snow 1959). Yet, technological development stimulates so much of the continued vitality of art, and new artistic technologies create new job opportunities. Our new AI technologies follow this trend, and will for the foreseeable future: new AI algorithms will provide new tools for expression and transform our art and culture in positive ways, just as so many other technologies have in the past.

Computers do not create art, people using computers create art. Despite many decades of procedural and computer-generated art, there has never been a computer widely accepted as the author of an artwork. To date, all "computer-generated art" is the result of human invention, software development, tweaking, and other kinds of direct control and authorship. We credit the human artist as author, acknowledging that the human is always the mastermind behind the work, and that the computer is a simple tool.

I then discuss whether this could ever change: would we ever agree to assign authorship to a computer? I argue that artistic creation is primarily a *social* act, an action that people primarily perform as an interaction with other humans in society. This implies that computers cannot create art, just as people do not give gifts to their coffee makers or marry their cars. Conversely, any human can make art, because we humans are social creatures. However, the boundaries of art are fluid, and, someday, better AI could come to be viewed as true social agents. I discuss this and other scenarios where AI algorithms could come to be accepted as artists, along with some of the dangers of too eagerly accepting software as artists.

This essay expresses my point of view, as someone who has developed various kinds of technology for art, has followed the development of new technologies affecting art and culture over the years, and, in a few instances, exhibited my own work as artworks. I have written in a style that is meant to be accessible to a general audience, since these are topics that many different people care about.

2. History and Current Practice

Before making general claims about how computer tools relate to art, I begin by first describing several useful historical examples, including the inventions of photography, film, and computer animation. These examples show previous situations in which a new technology appeared poised to displace artists, but, in reality, provided new opportunities and roles for artists. I also discuss how procedural art and computer art employs automation. In so many of these fields, outsiders view the computer or the technology as doing the work of creativity. Yet, in each of these cases, the human artist or artists behind the work are the true authors of the work.

I argue that, throughout history, technology has expanded creative and professional opportunities for artists dramatically, by providing newer and more powerful tools for artists. The advent of new technologies often causes fears of displacement among traditional artists. In fact, these new tools

ultimately enable new artistic styles and inject vitality into art forms that might otherwise grow stale. These new tools also make art more accessible to wider sections of society, both as creators and as consumers. These trends are particularly visible in the past two centuries since the Industrial Revolution.

2.1. How Photography Became an Artform

“From today, painting is dead!”

—Paul Delaroche, painter, at a demonstration of the daguerreotype in 1839¹

For lessons from the past about AI and art, perhaps no invention is more significant than photography.² Prior to the invention of photography, realistic images of the world could only be produced by artists. In today’s world, we are so swamped with images that it is hard to imagine just how special and unique it must have felt to see a skillfully-executed realistic painting (Figure 1a). The technical skills of realism were inseparable from other creative aspects. This changed when photography automated the task of producing images of the real world.

In 1839, the first two commercially-practical photographic processes were invented: Louis-Jacques-Mandé Daguerre’s dagguereotype (Figure 1b), and William Henry Fox Talbot’s negative-positive process. They were mainly presented as ways to produce practical records of the world. Of the two, the daguerreotype was more popular for several decades, because Talbot’s process was restricted by patents. Improvements to Talbot’s method eventually made the daguerreotype obsolete and evolved into modern film processes.

Portraiture was a main driver for early adoption. Then, as today, people enjoyed possessing pictures of their friends, loved-ones, and ancestors. Portrait painting was only available to aristocrats and the very wealthy. In the 18th century, several inexpensive alternatives were developed, such as the silhouette, a representation of an individual’s outline (Figure 2a), typically hand-cut by an artisan out of black paper. The daguerreotype offered an economical way to create a realistic portrait (Figure 2b). It was very slow and required locking the subject’s head in place with a head brace for several minutes, while the subject tightly gripped their chair, so as not to move their fingers. Nonetheless, numerous daguerreotype studios arose and became commonplace as technologies improved, and many portraitists switched to this new technology. Figure 3 shows one cartoon satirizing the situation for painters. By 1863, a painter-photographer named Henri Le Secq said “One knows that photography has harmed painting considerably, and has killed portraiture especially, once the livelihood of the artist.” Photography largely replaced most older forms of portraiture, such as the silhouette, and no one seems to particularly regret this loss. As much as I appreciate the mystery and beauty of looking at old etchings and portraits, I would rather use my mobile phone camera for my own pictures than to try to paint them by hand.

Another early use for the daguerreotype was to produce souvenirs for tourists: by 1850, daguerreotypes of Roman ruins completely replaced the etchings and lithographs that tourists had previously purchased. As the technology improved, photography became indispensable as a source of records for engineering projects and disappearing architectural ruins, as well as for documentary purposes, such as Matthew Brady’s photographs of the horrors of the American civil war.

¹ The historical information from this section is distilled from two texts: Scharf (1968) and Rosenblum (1984).

² This connection has been made previously (Agüera y Arcas 2017; Hertzmann 2001), but I explore it much more thoroughly here.

Figure 1. The interplay of painting and early photography. (a) By the 19th century, Western painters had achieved dazzling levels of realism. (b) Early cameras took low-quality (though evocative) pictures. This daguerreotype took over 10 minutes to expose. (c) However, camera technology steadily improved, capturing greater and greater realism, in much faster exposures. (d) This challenged painters to create works that were *not* about hyper-realistic depiction, such as Whistler’s Tonalist *Nocturne*. (e) The Pictorialist photographers attempted to establish photography as an art form by mimicking the styles and abstraction of painting. **Works:** (a) *Ophelia*, John Everett Millais, 1851; (b) *Boulevard du Temple*, Daguerre, 1838; (c) *Portrait of Sarah Bernhardt*, Félix Tournachon (Nadar), 1864; (d) *Nocturne in Blue and Gold: Old Battersea Bridge*, James MacNeill Whistler, c. 1872–1875; (e) *Morning*, Clarence H. White, 1908. Public domain images.

“Is photography art?” This question was debated for many decades, coalescing into three main positions. Many people believed that photography could not be art, because it was made by a mechanical device rather than by human creativity. Many artists were dismissive of photography, and saw it as a threat to “real art.” For example, the poet Charles Baudelaire wrote, in a review of the Salon of 1859: “If photography is allowed to supplement art in some of its functions, it will soon supplant or corrupt it altogether, thanks to the stupidity of the multitude which is its natural ally.” A second view was that photography could be useful to real artists, such as for reference, but should not be considered as equal to drawing and painting. Finally, a third group, relating photography to established forms like etching and lithography, felt that photography could eventually be as significant an art form as painting.

Figure 2. (a) A traditional silhouette portrait. (b) Daguerreotype portrait of Abraham Lincoln, 1846. Photographic techniques like this completely displaced previous portraiture techniques. *Public domain images.*

Figure 3. *The Unhappy Painter*, by Theodor Hosemann, 1843, satirizes the painter, a victim of progress, made obsolete by a daguerreotype. *Public domain image.*

Photography ultimately had a profound and unexpected effect on painting. Painters' mimetic abilities had been improving over the centuries. Many painters of the 19th century, such as the Pre-Raphaelites like John Everett Millais and Neoclassicists like Ingres, painted depictions of the world with dazzling realism, more than had ever been seen before. However, cameras became cheaper, lighter, and easier to use, and grew widespread among both amateurs and professionals (Figure 1c). Realistic photographs became commonplace by the end of the 19th century. If photorealism could be reduced to a mechanical process, then what is the artist's role?

This question drove painters away from visual realism toward different forms of abstraction. James McNeill Whistler's Tonalist movement created atmospheric, moody scenes (Figure 1d); he wrote: "The imitator is a poor kind of creature. If the man who paints only the tree, or the flower, or other surface he sees before him were an artist, the king of artists would be the photographer. It is for the artist to do something beyond this." The Impressionists, who sought to capture the perceptions of scenes, were likely influenced by the "imperfections" of early photographs. In contrast, Symbolists and post-Impressionist artists moved away from perceptual realism altogether. Edvard Munch wrote "I have no fear of photography as long as it cannot be used in heaven and in hell ... I am going to paint people who breathe, feel, love, and suffer." Vincent Van Gogh, describing his artistic breakthroughs around 1888, wrote to his brother: "You must boldly exaggerate the effects of either harmony or discord which colors produce. It is the same thing in drawing—accurate drawing, accurate color, is perhaps

not the essential thing to aim at, because the reflection of reality in a mirror, if it could be caught, color and all, would not be a picture at all, no more than a photograph.” Photography continued to influence modern art of the 20th century; one can infer a significant influence of Étienne-Jules Marey’s multiple-exposure photography on Futurism and Cubism, e.g., in Duchamp’s *Nude Descending A Staircase*.

It seems likely, in fact, that photography was one of the major catalysts of the Modern Art movement: its influence led to decades of vitality in the world of painting, as artists were both inspired by photographic images and pushed beyond realism.

Meanwhile, the Pictorialist movement, begun around 1885, was an attempt to firmly establish photography as an art form. Pictorialists introduced much more artistic control over the photographs, often using highly-posed subjects as in classical painting, and manipulating their images in the darkroom. Many of their works had a hazy, atmospheric look, similar to Tonalism, that softened the realism of high-quality photography (Figure 1e). They seemed to be deliberately mimicking the qualities of the fine art painting of the time, and today much of their work seems rather affected. They pursued various strategies toward legitimization of their work as an art form, such as the organization of photographic societies, periodicals, and juried photography exhibitions (Sternberger 2001). Their works and achievements made it harder and harder to deny the artistic contributions of photography; culminating in the “Buffalo Show”, organized by Alfred Stieglitz at the Albright Gallery in Buffalo, NY, the first photography exhibition at an American art museum, in 1910. Photography was firmly established as an art, and free to move beyond the pretensions of Pictorialism.

This story provides several lessons that are directly relevant for AI as an artistic tool. At first, photography, like AI, was seen by many as non-artistic, because it was a mechanical process. Some saw photography as a threat and argued against its legitimacy. Photography did displace old technologies that had fulfilled non-artistic functions, such as portraiture’s social function. Some artists enthusiastically embraced the new technology, and began to explore its potential. As the technology improved, and became more widespread over nearly a century, artists learned to better control and express themselves with the new technology, until there was no more real controversy over the status of photography. The new technology made image-making much more accessible to non-experts and hobbyists; today, everyone can experiment with photography. Furthermore, the new technology breathed new life into the old art form, provoking it toward greater abstraction. Wherever there is controversy in AI as an artistic tool, I predict the same trajectory. Eventually, new AI tools will be fully recognized as artists’ tools; AI tools may stimulate traditional media as well, e.g., the New Aesthetic (Sterling 2012).

2.2. The Technology of Live-Action Cinema

The story of filmmaking and technology has important lessons about how artists and technologists can work together, each pushing the other further. Most of the early photographers were, by necessity, both artists and technologists, experimenting with new techniques driven by their art or to inspire their art. But, in film and animation, this interaction has been much more central to the art form.

The history of film is filled with artist-tinkerers, as well as teams of artists and technologists. The Lumière Brothers created one of the first films, a simple recording of workers leaving their factory (Figure 4a), but also experimented with a wide range of camera technologies, color processing, and artistic ways to use them. The stage magician George Méliès filmed fantastical stories like *A Trip to the Moon*, employing a wide range of clever in-camera tricks to create delightfully inventive and beguiling films (Figure 4b). Walt Disney employed and pushed new technologies of sound and color recording, and drove other innovations along the way, such as the multiplane camera. Many of Orson Welles’ innovative film techniques were made possible by new camera lenses employed by his cinematographer Gregg Toland (Figure 4c). The introduction of more portable camera and audio equipment enabled the experiments of the French New Wave, who, in turn, influenced young American directors like Francis Ford Coppola and George Lucas. George Lucas’ team for Star Wars

was an early developer of many new visual effects on a shoestring budget (think of Ben Burt hitting telephone guy-wires to create the “blaster” sound effect), as well as an early innovation in digital film editing and compositing (Rubin 2005). Digital and computer graphics technology, have, obviously, revolutionized film storytelling since then, with directors like Michel Gondry and James Cameron pushing the technology further into unforeseen directions. In each case, we see technologies rapidly adopted by directors to create new storytelling techniques and styles, transforming the medium over and over.

Figure 4. Technological developments in the art of filmmaking: (a) the first captured film, of workers leaving the Lumiere Brothers’ factory (1895); (b) George Méliès’ 1902 *A Trip to The Moon* filmed like a stage play but with fantastical special effects; (c) *Citizen Kane*, which used numerous experimental camera and lens effects to tell the story. *Public domain images.*

2.3. 3D Computer Animation: A Collaboration

3D computer animation as an artform was pioneered by Pixar Animation Studios, and that success is due to the close collaboration of artists and engineers (Price 2008). It all began with Ed Catmull, an animation enthusiast who received a PhD in computer science in 1974. In his thesis, he invented several core techniques that every major 3D computer graphics system uses today. During his time in graduate school, he quietly set a goal for himself: to make the world’s first computer-animated film (Catmull and Wallace 2014). Consequently, he founded the Graphics Group at Lucasfilm Computer Division, and hired a team of brilliant engineers to invent computer systems to be used for film-making. However, none of this group could animate, that is, bring a character to life through movement. Hence, they recruited John Lasseter, an animator trained deeply in the Disney tradition. Through tight collaboration between Lasseter and the technical staff, they were able to invent new technologies and discover together how computer animation could start to become its own art form (Lasseter 1987). This group, led by Catmull and Alvy Ray Smith, spun out as Pixar, and, over the following years, invented numerous technical innovations aimed at answering the needs set out by Pixar’s artists; in turn, the artists were inspired by these new tools, and pushed them to new extremes. One of their mantras was “Art challenges technology, technology inspires art” (Catmull and Wallace 2014).

Pixar, by design, treats artists and engineers both as crucial to the company’s success and minimizes any barriers between the groups. When I worked there during a sabbatical, despite my technical role, I had many energizing conversations with different kinds of artists, attended many lectures on art and storytelling, sketched at an open life drawing session, watched a performance of an employee improv troupe, and participated in many other social and educational events that deliberately mixed people from different parts of the company. This is the culture that, though it still has some flaws to address, achieved so many years of technical and creative innovation, and, ultimately, commercial and artistic success.

Computer animation is another technology that scared traditional artists. In the early days before they found Lasseter, the Lucasfilm Graphics Group made many attempts to interest Disney animators in their work (Price 2008). Smith later said: “Animators were frightened of the computer. They felt that it was going to take their jobs away. We spent a lot of time telling people, ‘No, it’s just a tool—it doesn’t do the creativity!’ That misconception was everywhere” (Paik 2007). It is a common misconception

that computer animation just amounts to the computer solving everything; a programmer presses a button and the characters just move on their own. In reality, computer animation is extraordinarily labor-intensive, requiring the skills of talented artists (especially animators) for almost every little detail. Character animation is an art form of extreme skill and talent, requiring laborious effort using the same fundamental skills of performance—of bringing a character to life through pure movement—as in conventional animation (Lasseter 1987).

Traditional cel animation jobs did not last at Disney, for various reasons. Disney Feature Animation underwent a renaissance in the early 1990s, starting with *The Little Mermaid*. Then, following some changes in management, the Disney animation began a slow, sad decline. After releasing duds like *Brother Bear* and *Home on the Range*, management shut down all traditional 2D animation at Disney, and converted the studios entirely to 3D computer animation. Many conventional animators were retrained in 3D animation, but Disney's first 3D animation, *Chicken Little* was still a dud. Following Disney's acquisition of Pixar several years later, they revived Disney's beloved 2D animation productions. The result, a charming and enjoyable film called *The Princess and The Frog*, performed so-so at the box office, and, moreover, the animators' creative energy was focused on the newer 3D art form (Catmull and Wallace 2014). Today, traditional 2D animation at Disney is dead.³ Today, computer animation is a thriving industry, and it thrives in many more places than cel animation ever did: at many different film studios, in visual effects for live-action films, in video games, television studios, web startups, independent web studios, and many more. There are now more types of opportunities for animators than ever before. The story here is not the destruction of jobs, but the evolution and growth of an art form through technology. This is another story that contradicts the popular notion of art and technology operating in conflict, when, in fact, the opposite is usually true.

2.4. Procedural Artwork

In the art world, there is a long tradition of procedural artwork. Jean Arp created artworks governed by laws of chance in the 1910s (or so he claimed), and, beginning in the 1950s, John Cage used random rules to compose music. The term "Generative Art" appears to have originated in the 1960s. Sol LeWitt's wall drawings are provided as lists of precise instructions; people are still drawing new versions of his paintings after his death (Cotter 2008). Starting in the 1970s, classically-trained painter Harold Cohen began exhibiting paintings generated by a program he wrote called AARON (Cohen 1995). Since the 1980s, many current artists, such as Karl Sims, Scott Snibbe, Golan Levin, Scott Draves, and Jason Salavon, create abstract artworks by writing computer programs that generate either static images, or create interactive artistic experiences and installation works (Figure 5). In Sims' and Draves' work, the artwork "evolves" according to audience input. The popularity of the Processing computer language for artists speaks to the growth of procedural art.

In each of these cases, despite the presence of procedural, emergent, and/or crowdsourced elements, the human behind it is credited as the author of the artwork, and it would seem perverse to suggest otherwise. The human has done all of the creative decision-making around the visual style, of designing a framework and process, of testing and evaluating alternative algorithms, and so on.

³ Traditional animation styles are still vital in countries like Japan and France that, unlike America, do not believe that animation is "just for kids." Even so, their visual styles have evolved considerably due to computer technology.

Figure 5. Procedural artworks in the fine art world of galleries and art museums. (a) *Electric Sheep*, by Scott Draves, evolves dazzling procedural abstract animations based on thousands of votes. *Creative Commons (CC-BY-3.0 US) image, “sheep” by BrothaLewis*. (b) *The Top Grossing Film of All Time*, 1×1 (2000) by Jason Salavon, shows the average color of each frame of the movie *Titanic*. *Used by permission of the artist*.

2.5. State of the Art in Computer Science Research

Recent developments in computational artistic image synthesis are quite spectacular. But they should not be mistaken for AI artists.

Non-Photorealistic Rendering (NPR) is a subfield of computer graphics research (Rosin and Collomosse 2013) that I have worked in for many years. NPR research develops new algorithms and artistic tools for creating images inspired by the look of a conventional media, such as painting or drawing. Paul Haeberli’s groundbreaking 1990 paper (Haeberli 1990) introduced a paint program that began with a user-selected photograph. Whenever the user clicked on the canvas (initially blank), the system placed a brush stroke with color and orientation based on the photograph. In this way, a user could quickly create a simple painting without any particular technical skill (Figure 6a). In a follow-up paper, Pete Litwinowicz automated the process entirely, by placing brush strokes on a grid (Litwinowicz 1997). My own first research paper arose from experimenting with modifications to his algorithm: the method that I came up with creates long, curved strokes, beginning with large strokes that were then refined by small details (Hertzmann 1998) (Figure 6b). The algorithm was inspired by my experience with real painting, and the way artists often start from a rough sketch and then refine it.

This type of artistic algorithm design reflects the majority of computer graphics research in this area (Hertzmann 2003; Rosin and Collomosse 2013). The algorithms are automated, but we can explain in complete detail why the algorithm works and the intuitions about artistic process it embodies. This mathematical modeling of artistic representation continues the investigations begun in the Renaissance with Filippo Brunelleschi’s invention of linear perspective, a viewpoint I have written about elsewhere (Hertzmann 2010).

Figure 6. Painterly rendering algorithms that process an input photograph, using hand-coded rules and algorithms. (a) Paul Haeberli’s interactive painting system (Haeberli 1990). Images used by permission of the ACM. (b) My automatic painting system (Hertzmann 1998) processes a photograph without user input aside from selecting some parameter settings.

At some point, I found it very difficult to embody richer intuitions about artistic process into source code. Instead, inspired by recent results in computer vision (Efros and Leung 1999), I began to develop a method for working from examples. My collaborators and I published this method in 2001, calling it “Image Analogies” (Hertzmann et al. 2001). We presented the work as learning artistic style from example. But the “learning” here was quite shallow. It amounted to rearranging the pixels of the source artwork in a clever way, but not generalizing to radically new scenes or style (Figure 7). (A related style transfer method was published concurrently by Efros and Freeman (2001).) Since then, other researchers have improved the method substantially, making it much more robust (Fišer et al. 2016).

Figure 7. Our Image Analogies algorithm (Hertzmann et al. 2001), which stylizes a photograph in the style of a given artwork; in this case, Van Gogh’s *Starry Night over the Rhône*. Source photograph by John Shaw, used by permission; Van Gogh image in public domain.

In 2016, Leon Gatys and his colleagues published a new breakthrough in this space, Neural Style Transfer (Gatys et al. 2016). Based on recent advances in neural networks, their method transfers certain neural network correlation statistics from a painting to a photograph, thus producing a new painting of the input photograph (Figure 8). The method is still “shallow” in a sense—there is no “understanding” of the photograph or the artwork—but the method seems to be more robust than the

original Image Analogies algorithm. This paper led to a flurry of excitement and new applications, including the popular Prisma app and Facebook's Live Video stylization, as well as many new research papers improving upon these ideas. This work is ongoing today.

Figure 8. The Neural Style Transfer Algorithm (Gatys et al. 2016), which stylizes a photograph in the style of a given artwork; in this case, Van Gogh's *The Starry Night*. This algorithm has led to numerous new apps and research in stylization. *Van Gogh image in public domain; images by Leon Gatys, used by permission.*

Another development which received considerable attention in 2015 was the invention of DeepDreams by Mordvintsev et al. (2015), who, developing a visualization tool for neural networks, discovered that a simple activation excitation procedure produced striking, hallucinatory imagery of a type we had never seen before. There are many other current projects, particularly those around Generative Adversarial Networks (Goodfellow et al. 2014) and Project Magenta at Google (Metz 2017), that also show promise as new artistic tools. For example, Figure 9 shows images that we generated by visualizing trends learned by a neural network from a large collection of artistic images in different styles. A variety of related images are produced by Creative Adversarial Networks (Elgammal et al. 2017); it is a visual style that seems familiar but not the same as what we are familiar with (and is probably driven, in part, by the biases of the convolutional neural network representation).

In each of these cases, the artworks are produced by a human-defined procedure, and the human is the author of the imagery.

Figure 9. Images that we generated using a neural network trained on different subsets of a database of artistic imagery (Wilber et al. 2017). Each one is meant to typify a single stylistic category or a medium, though the biases of the convolutional neural network architecture are also evident.

2.6. Artificial Intelligence is Not Intelligent

Unfortunately, there has been a considerable amount of media hype around AI techniques. In the news media, algorithms are often anthropomorphized, as if they have the same consciousness as

humans (e.g., (Wilson 2017)), and sometimes they are described as artists (Dvorsky 2017; Perez 2018; Pogue 2018). In fact, we do not really know what consciousness is (despite many theories), or what it would mean to embody it in an algorithm.

Today, the most-successful AI and machine learning algorithms are best thought of as glorified data-fitting procedures (Brooks 2017). That is, these algorithms are basically like fitting a curve to a set of datapoints, except with very sophisticated ways to fit high-dimensional curves to millions of datapoints. When we as researchers speak of “training” an algorithm, or an algorithm that “learns”, it is easy to misinterpret this as being the same thing as human learning. But these words mean quite different things in the two contexts. In general, “training” a model to learn a task involves careful human effort to formulate the problem, acquire appropriate data, and test different formulations. It is laborious and requires considerable expertise and experimentation. When a new task needs to be solved, the human starts over.

Compared to human intelligence, these algorithms are brittle and bespoke. For example, image recognition algorithms have undergone breathtaking breakthroughs in the past decade, and are now widely used in consumer products. Yet they often fail on inputs that suggest bizarre misunderstandings; the existence and robustness of adversarial examples (Szegedy et al. 2014) and procedurally-evolved images (Nguyen et al. 2016) demonstrates that these algorithms have not really learned anything like human-level understanding. They are like tourists in a foreign country that can repeat and combine phrases from the phrasebook, but not truly understand the foreign language or culture. These systems have no autonomy except within the narrow scope for which they were trained, and typically fail-safes must be put in place as well, e.g., Google Photos no longer classifies anything as a “gorilla” because of one high-profile failure (Simonite 2018).

There are some fascinating parallels between human learning and machine learning, and it does seem likely that humans are, in some way, optimized by evolutionary principles (Alexander 1996; Gopnik and Wellman 2012; Koerding and Wolpert 2006; Tenenbaum et al. 2011). But going from these high-level analogies to actual machine intelligence is a problem for which the solution is not even on the horizon.

3. How Technology Changes Art

Based on this history, I now make several specific claims about how technology changes art. Far from replacing artists, new technologies become new tools for artists, invigorating and changing art and culture. These claims apply equally to current developments in AI as they did to previous developments like photography and animation.

3.1. Algorithms Are Artists' Tools

“my watercolor teacher used to say: let the medium do it. true that—so my sketch provides the foundation and then the network does it thing; i don’t fight, just constantly tweak the #brushGAN toolkit.”

—Helena Sarin (@glagolista)

In every technology that we currently employ—whether photography, film, or software algorithm—the technologies and algorithms we use are basic tools, just like brushes and paint.⁴ The same is true for the new AI-based algorithms that are appearing. They are not always predictable, and the results are often surprising and delightful—but the same could be said for the way watercolor

⁴ This view reflects, I think, the conventional wisdom in the computer graphics research field, which is my research background. This field has always had close ties with certain artistic communities, especially computer animation and visual effects, and, based on this experience, the field is often resistant to attempts to automate creative tasks. In contrast, artificial intelligence researchers use terminology much more aspirationally, historically using words like “intelligence,” “learning,” and “expert systems” in ways that are far simpler than the human versions of these things.

flows on the page. There is no plausible sense in which current systems reflect “true” artificial intelligence: there is always a human behind the artwork.

Applying the same standard to the current research in neural networks and neural style transfer, it would seem equally perverse to assign authorship of their outputs to the software itself. The DeepDream software was authored by a human; another human then selected an input image, and experimented with many parameter settings, running the software over and over until good results were obtained. Indeed, in a recent art exhibition meant to promote these methods and their exploration (Tyka 2016), human artists were credited for each of the individual works.

The same process of selection of tools and inputs, adjusting settings and even modifying code, and iterating until a desirable output is produced, occurs in all current and foreseeable computer artworks. Computer-generated artworks result from considerable time and effort from human artists, from conception of the idea, to painstakingly guiding the execution, to selecting from among the outputs.

There have been a few cases where AI algorithms have been presented as artists or potential artists (Colton 2012; Dvorsky 2017; Perez 2018; Pogue 2018; Shein 2017). For the above reasons, I think this claim misunderstands the nature of procedural art.

In short, in our present understanding, *all art algorithms, including methods based on machine learning, are tools for artists; they are not themselves artists.*

3.2. New Technology Helps Art Stay Vital

Rather than being afraid of the new technologies, we should be enthusiastic about the new artworks that they will enable artists to produce. When we think of art as having external influences, we normally think of social or political influences, but ignore the effect of new tools. In contrast, I argue that, especially from the 19th century onwards, technological developments have played a pivotal role in advancing art, in keeping it vital and injecting fresh ideas. The stories I gave of photography and cinema include many examples of this. However, the effect is far more widespread.

One of the most important breakthroughs in the history of Western art was the invention of oil paint by Flemish painters such as Jan van Eyck in the 15th century (de la Croix et al. 1991). Previously, painting had been done primarily with tempera, which lacks subtle coloration, and fresco, which was very cumbersome to work with. Oil paint had existed in some form for centuries, but van Eyck and others found new techniques that gave them a very practical new medium. It was fast-drying and allowed rich colors and tones, sharp edges, and hard surfaces. The rich light and color that we associate with the Northern Renaissance and the Italian Renaissance are due to this technology (Figure 10).

In each decade since the 1950s, many of our culturally-important works used technology that had only been invented within the previous 10 years. For example, most technology used in today's feature films did not exist 10 years ago (e.g., widespread use of HD digital cameras; facial performance capture); the same goes for artworks using smartphones and crowdsourcing; artworks involving white LEDs and Arduino controllers; DJs performing on stage behind their laptops; and so on. Even the most vérité-seeming romantic comedies frequently involve recent digital video editing and digital backdrops.

Conversely, artistic styles that fail to change become stale and lose their cultural relevance; the adoption and exploration of new technology is one of the ways that art stays vibrant. For example, the introduction of synthesizer music into 1980s pop music created a new sound that was exciting and modern. The sound diversified as the tools improved, until grunge became popular and made the 1980s synthpop sound seem superficial and old-fashioned. Nowadays, a recent revival of 1970s and 1980s instruments by bands like Daft Punk and LCD Soundsystem seems most exciting at times when they are creating new types of music using old instruments. In contrast, the swing music revival of the 1990s never went anywhere (from bands like Big Bad Voodoo Daddy and Squirrel Nut Zippers) in my opinion, because the bands mimicked classic styles with classic instruments, without inventing anything particularly original themselves.

Figure 10. The development of the oil painting technology changed painting as an art form. (a) Fresco painting by Michaelangelo on the Sistine Chapel (ca. 1508). The fresco process was difficult and achieved limited tonal range. Today, fresco is defunct as a medium. (b) Oil painting by Jan van Eyck for the Ghent Altarpiece (ca. 1430). Much richer colors and lighting are possible with oil paint. *Public domain images.*

In each era, radical technological innovations are met by artists with both enthusiasm and rejection. For example, when the Moog synthesizer became popular, it was adopted by big-name bands like Emerson, Lake, and Palmer. Other bands felt that twisting knobs to make music was “cheating”: Queen’s album covers proudly state that the band did not use synthesizers. Robert Moog described one New York musician who said of the instrument “This is the end of the world” (Snowden 2012). It now seems silly to imagine that people might have ever categorically objected to synthesized music, or to the scratching and sampling of hip-hop DJs, just as it now seems silly that people once rejected waltzing, the Impressionists, and the Rite of Spring as invalid or immoral.

In addition to stimulating professional artists, new tools make art more accessible to larger portions of society. Photography was once accessible only to the most determined early adopters, but has continually become easier, faster, and more compact, to the point where nearly everyone carries a mobile phone camera in their pocket or purse. The same goes for the tools of cinematography (from hand-cranked to heavy cameras to Steadicams to handycams to iPhones), and so on. Modern computers give nearly everyone access to digital equivalents of darkrooms, mixing studios, painting studios, and so on; these were formerly highly-specialized technologies requiring laborious effort.

3.3. New Technology Does Not Cause Net Unemployment

Concerns of how technology displaces jobs have been around since at least the 19th century, when Luddite protesters destroyed mechanical weaving machines, and, in folk songs, John Henry competed against a steam drilling machine. These fears are real and understandable. Yet, despite centuries of technological disruption, we do not live in a world of massive unemployment. This is because, as old roles are erased, many more arise in their stead. But these fears keep recurring, because, at any given time, it is easy to imagine losing specific jobs but it requires superhuman imagination to forecast what new opportunities will be created by transformative new technologies. Nowadays, most of us do jobs that would be hard to explain to a 19th century worker.

The real workforce concerns should not be about the technology itself, but about whether the economic system shares the benefits of new productivity fairly across society versus concentrating wealth only among the very richest (Stiglitz 2012) and whether machine learning systems are misused to magnify existing forms of inequality (O’Neil 2016). When displacements due to new technology occur, their effects can be eased by social safety nets and better educational foundations (for employment

flexibility and retraining). Conversely, a society which fails to distribute wealth and economic gains fairly has much bigger problems than just the impact of AI.

Fears of new technology seem to be human nature. I suspect many people view the “normal” state of things as being how they were when they came of age, and they view any significant change as scary. Yet, nearly all of our familiar modern technologies were viewed as threatening by some previous generation.

The fear of human-created life has been with us for a long time. Notably, 18th-century scientists discovered electricity. As they searched to understand it, they discovered the life-like effect of galvanism, that the muscles of dead frogs could be stimulated by electrical currents. Had the secret of life been discovered? This inspired Mary Shelley’s novel *Frankenstein; or The Modern Prometheus*, in which an ambitious university student uses modern science to create new life (Shelley 1818). Today, the story is vivid and evocative, but, intellectually, we recognize it as preposterous. The fear of AI is essentially the same irrational fear; SkyNet is Frankenstein’s monster, but with neural networks as the Promethean spark instead of galvanism.⁵ At present, the Terminator’s autonomous AI is only slightly more plausible than its ability to travel backwards in time.

3.4. New AI Will Be New Tools for Artists

Some general trends around the evolution of technology and art seem quite robust. As discussed above, current AI algorithms are not autonomous creators, and will not be in the foreseeable future. They are still just tools, ready for artists to explore and exploit. New developments in these tools will be enthusiastically adopted by some artists, leading to exciting new forms and styles that we cannot currently foresee. Novices will have access to new simplified tools for expression. It is possible that some tasks performed by human artists will fade out, but these will generally be mechanical tasks that do not require much creativity because they fill societal functions other than artistic expression. Some traditional arts may fade simply due to seeming old-fashioned. This is the nature of art: nothing is fresh forever, which is not to be blamed on technology. Artistic technology is an “imagination amplifier” (Cohen 2000) and better technology will allow artists to see even further than before.

Aside from general trends, it is hard to make specific predictions about the art of the future. Les Paul, who invented the solid-body guitar in the 1940s (Figure 11, himself primarily performed light pop, country, and showtunes, and could hardly have predicted how the electric guitar would be used by, say, Led Zeppelin, just as it is hard to imagine Daguerre predicting Instagram. More generally, making predictions about how AI technologies might transform society is very hard because we have so little understanding of what these technologies might actually be (Brooks 2017). Even the science fiction writers of the 1950s and 1960s completely failed to imagine the transformative power of the Internet and mobile computing (Krauss 2017); for them, the computers of the future would still be room-scale monstrosities that one had to sit in front of to operate. But they did predict moon colonies and replicants by 2018.

⁵ In fact, *Frankenstein* is presented as a cautionary tale about the quest for knowledge in general. Victor Frankenstein tells his story as a warning when he learns that Captain Walton is himself driven by an obsessive quest for knowledge that is entirely unrelated to Frankenstein’s.

Figure 11. Les Paul, inventor of the solid-body electric guitar in 1943. His technology transformed popular music in ways they could not have foreseen. *Public domain image.*

In short, we cannot predict what new inventions and ideas artists will come up with in the future, but we can predict that they will be amazing, and they will be amazing because they make use of technology in new, unpredictable ways.

4. What Is an Artist?

So far, I have described how computer technologies are currently accepted as tools for artists, not as artists themselves. Why is this? After all, computers can do other human tasks like speak, search, print, navigate, and, to some extent, drive cars. There are several obvious reasons why computers do not make art, including tradition, the incentives involved, and the relatively predictable nature of existing automation. But, still one could imagine an alternate history in which some machines or computer programs had already been called artists. I believe there is a more fundamental reason, which explains not just why this has not happened, but why it is unlikely to happen anytime soon.

In this section, I theorize about the prerequisites for an entity being an artist, focusing on my hypothesis that art is a social behavior. I will then apply this idea to AI in the next section. I also explore several alternative hypotheses for what could make an AI as an artist. Most of these alternative hypotheses identify some attribute of human artists, and then hypothesize that that attribute is needed for an AI to be an artist.

4.1. Art Is Social

“What an artist is trying to do for people is bring them closer to something, because of course art is about sharing: you wouldn’t be an artist if you didn’t want to share an experience, a thought.”

—David Hockney ([Hockney 1996](#))

Why do we create and consume art? I argue that art is primarily a social behavior: art is about communication and displays between people. For example, people often speak of art as being about personal expression, which is an act of communication.

Any human can make art, because humans are social creatures. A painter cranking out the same conventional landscapes for tourists year after year is still considered an artist. A child’s drawing might only be interesting to their family, but it is still art.

I am directly inspired by the theory, going back to Charles Darwin, that art-making is an adaptive product of our biological evolution. [Dutton \(2009\)](#) sets out a persuasive argument for this theory, which I briefly summarize, though I cannot do it justice. Creating art served several functions for our Pleistocene ancestors. Art-making served as a fitness signal for mating and sexual selection. Art can also be used as a display of wealth and status. Storytelling, music, and dance strengthen social bonds within a group. Storytelling additionally plays a very important role of communicating information that would otherwise be hard to share.

I observe that each of these functions of art is *social*: art arose as forms of communication, displays, and sharing between people. Although art takes many different forms in different cultures today, each of these forms serves one or more of the same basic social functions that it did in the Pleistocene.

I generalize this theory beyond humans to hypothesize: *art is an interaction between social agents*. A “social agent” is anything that has a status akin to personhood; someone worthy of empathy and ethical consideration. Many of our other behaviors are interactions between social agents, such as gifts, conversation, and social relationships like friendship, competition, and romance.

In contrast, while we can get emotionally attached to our computers and other possessions, we feel no real empathy for their needs, and no ethical duty toward them. Possessions can participate only in shallow versions of these interactions. For example, we frequently talk about our possessions with statements like “The brakes on my bicycle were complaining so I gave it a new pair as a gift and now it’s much happier. I love my bike.” This statement indicates emotional attachment but not true empathy with the bike’s feelings, despite the anthropomorphizing language. We do not live in a social hierarchy with our possessions: we do not compete with them for status, or try to impress them. We care about what other people have to say because we care about other people; we care about what computers have to say only insofar as it is useful to us. We generally treat conversational agents (like Siri and Alexa) as user interfaces to software, not like people.

Art-making does also have non-social benefits to the artist. For example, art-making can help practice skills like dexterity and problem solving. Creating art is often pleasurable or meditative in itself. But these benefits are secondary to their social benefits: they are not the reasons that evolution has produced art as a human activity. Similarly, one may also talk or sing to oneself while alone, but talking and singing are still fundamentally social activities.

Note that the evolutionary argument here is optional; one can discuss whether art is fundamentally social without it. But I believe that the evolutionary view of art gives some additional understanding.

4.2. Non-Human Authors

As we have seen, despite many technological advances, current algorithms are not accepted as artists. There are a few other existing examples where objects are created by authors or processes that are not human-driven. These give some support to this theory.

Natural processes. Natural processes, including landscapes like the Grand Canyon or the HuangShan Mountains, are not considered art, even though they may be extraordinarily beautiful and change one’s perspective immensely. Beautiful structures made instinctively by animals, such as honeycombs and coral, are not considered art. This indicates that simply creating complex and beautiful outputs is not itself sufficient for art, since there is no creative social communication in these cases.

Animals. Some higher mammals, including chimpanzees, elephants, and dolphins, have been trained to paint (Desmond 2016; Dutton 2009). Many writers are skeptical of animal-made art. Typically, the animal’s owner or handler steers the process, letting the animal throw paint on the canvas, then stopping the painting when they believe it is done, then selecting which works to show; the animals seem to show no interest in the artwork afterward. Animal artwork has not had any significant cultural impact or popularity; it seems to have been largely the product of media stunts. (People for the Ethical Treatment of Animals have recently tried to claim copyright in favor of a monkey, but failed (Millward 2017), as US copyright law only allows humans to claim copyright.)

The most interesting aspect of this discussion is not whether animals can create art, but *how we decide*. Discussions of whether animal-made artifacts are art are not based on *a priori* rules whether animals can be artists. Instead, they are attempts to study the evidence of the animal’s behavior around the artwork, and, from that, to infer whether the artwork is some form of inner expression (Desmond 2016), or an artifact that the animal has a special appreciation for (Dutton 2009). In other words, we are open to the idea of animals creating art, because they can have social relationships

with us. It is just that we have not found any other creature that satisfies our criteria for creating art, whatever they are.

4.3. Judging the Work Instead

It is tempting to judge whether a computer is an artist based solely on the merits of the work that it produces. In this hypothesis, whether or not a computer can be an artist is a judgement of the quality of the work that it produces, independent of the properties of the computer itself. If an algorithm outputs a continual stream of diverse, stimulating, beautiful, and/or skillful outputs, without many duds, we might be quite tempted to call this algorithm an artist. The better computer-generated art becomes, the more we will hear questions about whether computers are artists.

Skill is clearly not the real requirement for someone or something to be an artist. Any human can make art, including unskilled amateurs and children. Conversely, computers can already be programmed to create infinite sequences of dazzling realistic or abstract imagery, exhibiting technical proficiency way beyond the typical human capacity.

I suspect that, when we look at a computer's output and ask "is this work good enough to call the computer an artist?", we are not actually judging the quality of the work *per se*. Instead, we are really looking for evidence that the system itself is intelligent, conscious, and feeling: traits that we associate with social agents. No matter how skillful and surprising a computer's output is, we will not accept it as an artist until we infer some sort of social being inside.

4.4. An Intent Machine

Another hypothesis is as follows: in the modern art world, the role of the artist is to supply the "intent" and the "idea" for the work; it is not necessary for the artist to execute on the work, other than coordinating in its production. For humans, this is clearly true in numerous examples, such as those in Section 2.4, as well as appropriation works like Duchamp's Readymades and Richard Prince's questionable Instagram reproductions; artists can also employ helpers or crowdworkers, such as in Scott Draves' *Electric Sheep* and Aaron Koblin's *Sheep Market* (Koblin 2009). Consequently, for a computer to be an artist it simply needs to supply an intent.

It is easy to imagine designing a system that creates intent and even coordinates the labor of producing a work. For example, one could write a simple procedural algorithm to generate very basic intents (e.g., "portray ominous landscape"), or sample intents from a Recursive Neural Network trained on artists' statements scraped from the web. Or the method could randomly select some news item, photograph, or historical event, and randomly sample some attitude toward that thing. Starting from this intent, crowdworkers could be used to refine the idea and convert it into a new image, similar to systems like SoyLent (Bernstein et al. 2015). One could also automate steps of the process, e.g., using GANs (Goodfellow et al. 2014) to generate entirely new starting images from scratch. Crowdworkers could also be used to rate and evaluate the outputs of the system, selecting just the best results and discarding the rest. Final steps of the process could also automatically hire professional designers, e.g., from sites like Upwork or 99Designs. This system could then run continuously, generating new images over time (with payment being automatically made to the crowdworkers involved). Workers could group images with common themes and intents, and create separate collections around these themes. Artists' statements could be generated around these themes. The system's preferences could grow and adapt over time as more data are gathered or external data streams (e.g., photography blogs) change.

Suppose someone were to build this system, calling it, say, The Intent Machine, and exhibits its work in an art show or gallery. Suppose, moreover, they convinced the curator or gallery owner to credit The Intent Machine with authorship of the works it had created, but fully disclosed the procedure by which it worked. Would people credit it as the artist who had authored its works, or would they say that the system-builder is the real artist here?

I believe that, in general, the consensus would be that the system-builder is the real artist here, and that this is really an artwork about probing the nature of computer-generated art, or the nature of the

commercial artworld. Note that the procedure used to define the system is not fundamentally different than any other procedural computer-generated art algorithm (as in Sections 2.4 and 2.5). Even if the work itself ended up being quite good, viewers would ask why it is good, and it is doubtful that the computer's own contribution would be judged as significant beyond those of the humans involved.

Are artists just “intent machines”? If art is a social act, then the answer is no.

4.5. Creativity, Growth, Responsiveness, and so on

Finally, three related criteria that have intuitive appeal are the notions of creativity, growth, and responsiveness. They are important for human artists, so perhaps they should be for AI artists as well.

Several authors have proposed energy terms or criteria for *creativity* (Elgammal et al. 2017; Nguyen et al. 2016). These are often attempts to express the idea that the system's output should somehow “surpass” the human programmer and/or the training data. For example, Colton (2008) proposes that we judge the creativity of a system, in part, by whether the system's output *surprises* the system's author. I believe that this is too weak a criterion, since many mechanical or algorithmic phenomena may be surprising to their own discoverer or author at first. For example, the basic algorithm that produces Mandelbrot set images can be specified in a single sentence,⁶ yet produces dazzling animations of infinite complexity.⁷ The Mandelbrot set is very surprising and produces beautiful, unprecedented images, but we do not call its iteration equation creative, or an artist.

All current procedural art systems, such as the Mandelbrot set, have a recognizable style, and, after awhile, lose their novelty. I believe that the same will be true for systems specifically designed with “creativity” objectives (Elgammal et al. 2017; Nguyen et al. 2016). Unlike human artists, these systems do not grow or evolve over time.

Perhaps an AI artist would need to exhibit some form of *growth*. For example, Harold Cohen, the fine-art painter who began, in 1968, to write software to generate art, described the evolution of his views: “Ten years after [1968], I would have said ‘look, the program is doing this on its own’ Another ten years on and I would have said ‘the fact that the program is doing this on its own is the central issue here,’ denoting my belief in the program's potential and growing autonomy over the whole business of art-making It was producing complex images of a high quality and I could have had it go on forever without rewriting a single line of code. How much more autonomous than that can one get? . . . [But] it's virtually impossible to imagine a human being in a similar position. The human artist is modified in the act of making art. For the program to have been similarly self-modifying would have required not merely that it be capable of assessing its own output but that it had its own modifiable worldview” (Cohen 2014).

In any existing system, it is easy to think of trivial ways for the system to evolve and change over time, e.g., subtly change the color palette or the training data over the years. Superficial growth is easy; meaningful growth is hard to even define for a computer AI. If someone could design a system that produces a sequence of art that is meaningful to people and also significantly evolves over time, that would be truly remarkable. Another missing piece from current systems is an artists' ability to *respond* meaningfully to their culture, experiences, world events, responses to their work, and other aspects of their environment. It seems hard to imagine achieving these goals without enormous technological advances—they may not be possible without true AI or social AI in some form.

4.6. Definitions of Art

For guidance about the nature of art, we could have also looked to existing definitions of art. However, there is no Royal Society that prescribes what is and what is not valid art. Instead, art is a

⁶ Color the image location (x, y) proportional to the number of steps the iteration $z_{n+1} = z_n^2 + x + yi$ requires to reach $|z| \geq \tau$, starting from $z_1 = 0$, where τ is a large constant.

⁷ For example: <https://www.youtube.com/watch?v=PD2XgQOyCCK>.

phenomenon that results from the interplay of cultural institutions and the general population that we can analyze, and it changes over time. Philosophers have attempted to devise concise definitions of art that include all existing types (music, dance, painting, etc.) and styles of art. The Institutional Definition, originated in the 1960s in response to conceptual art, states, roughly, that art is anything in a style that is broadly accepted as art (Danto 1964; Dickie 1969; Levinson 1979). A more fine-grained approach is to identify attributes common to many different types of art (Dutton 2000 2009; Gaut 2000 2005). Each of these definitions of art is an attempt to fit the data, and draw a line between those things that we call art (like theatre) and those that we do not (like spectator sports). Understandably, these definitions all assume that the artist is always human, without exploring much whether non-humans can create art, and thus do not provide much guidance for this discussion.

4.7. Attribute Theories in General

Many of the theories in this section have the following recipe: identify some attributes of human artists, and then hypothesize that AIs with these attributes will be considered artists. Artists make high-quality work; artists supply intent; artists are creative; artists grow.

The development of Cohen's thinking illustrates this. Cohen initially thought that a machine that makes high-quality art could be considered an artist. The more he observed his own software, the more something seemed to be missing. He noted that human artists grow over time, and his system did not. If he had added some form of "growth" to his system, would it then be an artist?

Maybe someday someone will develop a system with enough of these attributes to cross some intangible threshold and thus be perceived as an artist. However, it seems very hard to reason concretely about this possibility. How do we define any of these attributes precisely? How much is enough?

Furthermore, many of these attributes are not really required of human artists. Any human can make art, even if it is not very original or surprising; the artist need not grow noticeably or respond to culture or feedback. We do judge the work by these attributes, but there is no minimum requirement for humans to make art.

In contrast, the social theory makes a much more concrete statement. Art is fundamentally a social interaction, and thus can only be made by social agents. The social theory has the additional appeal of being based on a plausible evolutionary hypothesis for the reasons we create art.

5. Will an AI Ever Be an Artist?

With this background, I now turn to speculating about the future. As we have seen, authorship of all current algorithmic art is assigned to the human author behind the algorithms. Will we ever say that an AI itself created art? Will we ever recognize a piece of software as the author of a work of art?

Human-Level AI. If we ever develop AI with human-level intelligence and consciousness, by definition, it would be able to create art, since it would have the same capacity for consciousness, emotions, and social relationships. But, as discussed in Section 2.6, this scenario is science fiction and we have no idea if this is possible or how it would be achieved. Making meaningful predictions about a world with "true AI" is impossible (Brooks 2017), because we have so little idea of how specifically this AI would actually operate. Moreover, this AI would transform society so much as to make it unrecognizable to us. We may as well speculate about what kind of artwork is made by aliens from outer-space—if do we ever meet them, we will have more pressing questions than what kind of music they like.

Hence, the interesting question is whether there could be computer-authored artwork *without* human-level intelligence or consciousness.

Social AI. If, as I have argued, creating art is a fundamentally social act of expression and communication, then it follows that *AI can be granted authorship when we view the AI as a social agent, and it is performing some communication or sharing through art.*

What does it mean for us to view an AI as a social agent? We have to view the AI as deserving of empathy and ethical consideration in some way. However, the AI does not need human-level intelligence, just as we have social relationships with our pets. But we do expect that the AI has something to say socially, something that suggests an inner consciousness and feeling.

Short of true intelligence (the science fiction scenario), I think that the only way this can happen is through “shallow AI” agents. People are sometimes “fooled” by shallow AI. The classic example is Eliza, a simple text-based “psychiatrist” program developed in 1964, based on simple pattern-matching and repetition of what the user types (Weizenbaum 1966). It was meant as a demonstration of the superficiality of the AIs of the time, but, unexpectedly, many people attributed human-like emotions to the machine. Since then, there are many anecdotes of people being fooled by “chatbots” in online settings (Dvorsky 2016; Epstein 2007), including the recent plague of Twitter bots (Kessler 2014). But, once the veil is lifted, it is clear that these chatbots do not exhibit real intelligence.

Some software and robots have been designed to have relationships with their owners, including talking dolls, Tamagotchi, and Paro therapeutic baby seals. A related effect is that people behave toward their computers as if they were social agents in certain ways (Reeves and Nass 1996), even when they do not believe that they are intelligent. For example, dialogue systems like Siri and Alexa all use female voices by default, based on many findings that male and female users both respond better to female voices (Griggs 2011; Reeves and Nass 1996).

Perhaps, for many users, the system does not need to be truly intelligent, it just has to be *perceived* as a social agent, like a Siri or Alexa that you can ask to make you an artwork (Xu et al. 2017). The day may come in which these agents are so integrated into our daily lives that we forget that they are carefully-designed software. One can easily imagine the development of AI that simulates emotion and affection toward the user; it is easy to imagine, for example, a toy doll that paints pictures for its owner, as one of many behaviors designed to display companionship and affection.

Non-social AI. It seems possible that non-social algorithms could be successfully promoted as artists; there have been a few tentative forays in this direction, e.g., (Colton 2012; Elgammal et al. 2017). For the reasons given above, I am skeptical that such methods will be accepted as true artists without some plausible belief about their underlying social and/or conscious attributes.

Perhaps a curator at a well-known museum would download or otherwise acquire various artifacts from software “artists”, and list the software systems as the authors. There would be controversy, and discussion in newspapers and journals. Perhaps other curators and galleries would follow suit. Perhaps people would find enough value in these computer-generated artworks, while also being convinced that no human could be rightly given credit for their works. This sort of process has happened for things like abstract expressionism, and not for chimpanzee art. Could it happen for computer art?

Suitcase words. The term “artist” could come to be used as having multiple meanings, just as words like “intelligence” and “learning” have come to mean something different for humans than they do for algorithms (Brooks 2017). A software program that, say, automatically stylizes your photos, could be called an “artist” in the same way that software applications like “calendar” and “mail” programs have replaced their physical-world namesakes. Unfortunately, the use of the same word to mean different things in different contexts causes endless confusion, as discussed in Section 2.6.

Dangers. A continual danger of new AI technology is that human users misunderstand the nature of the AI (O’Neil 2016). When we call a shallow AI an “artist”, we risk seriously misleading or lying to people. I believe that, if you convince people that an AI is an artist, then they will also falsely attribute emotions, feelings, and ethical weight to that AI. If this is true, I would argue that calling such AIs “artists” is unethical. It leads to all sorts of dangers, including overselling the competence and abilities of the AI, to misleading people about the nature of art.

It seems likely that some companies will not have any scruples about this. For example, Hanson Robotics has promoted (in many contexts) a social robot as a truly intelligent being, even though it is clearly nothing more than a “chatbot with a face” (Gershgorn 2017) or, in Yann LeCun’s words, a “Potemkin AI” (Vincent 2018).

A related concern is that we deprive the AI’s designer(s) of authorship credit. At present, we credit the author of a piece of automatic software with the output of that software. This usually acknowledges the skill and effort required to engineer and iterate with software so that it produces good outputs. Artistic credit is important for understanding the real sources of how something was made.

Outside of science fiction, I can see no positive benefit to calling a computer an artist, but I do see dangers.

6. Conclusions

I do not believe that any software system in our current understanding could be called an “artist”. Art is a social activity. I mean this as a warning against misleading oneself and others about the nature of art. Of course, the ambitious reader could take this as a challenge: I have laid out some of the serious objections that you must overcome if you wish to create a software “artist.” I do not think it can be done anytime soon, but I also know that proving critics wrong is one of the ways that art and science advance.

One of my main goals in this essay has been to highlight the degree to which technology contributes to art, rather than being antagonistic. We are lucky to be alive at a time when artists can explore ever-more-powerful tools. Every time I see an artist create something wonderful with new technology, I get a little thrill: it feels like a new art form evolving. Danny Rosin’s *Wooden Mirror*, Jason Salavon’s *The Top Grossing Film of All Time*, 1×1 , Bob Sabiston’s *Snack and Drink*, Michel Gondry’s *Like A Rolling Stone*, Kutiman’s *ThruYOU*, Amon Tobin’s *Permutation*, Ian Bogost’s *Cow Clicker*, Christian Marclay’s video installations, Íñigo Quilez’s procedural renderings, and Wesley Allsbrook’s and Goro Fujita’s Virtual Reality paintings are a few examples of artworks that have affected me this way over the years. Today, through GitHub and Twitter, there is an extremely fast interplay between machine learning researchers and artists; it seems like, every day, we see new tinkerers and artists Tweeting their latest creative experiments with RNNs and GANs (e.g., @JanelleCShane, @helena, @christophrhesse, @quasimondo, @DrBeef_).

Art maintains its vitality through continual innovation, and technology is one of the main engines of that innovation. Occasionally, the avant garde has tremendous cultural impact: electronic music and sampling was once the domain of experimental electronic and musique concrète pioneers, like Wendy Carlos and Delia Darbyshire. Likewise, at one time, computer-animated films could only be seen at obscure short-film festivals. Today, we are seeing many intriguing and beguiling experiments with AI techniques, and, as artists’ tools, they will surely transform the way we think about art in thrilling and unpredictable ways.

Acknowledgments: Thanks to Shira Katz, Alvy Ray Smith, Craig Kaplan, Shiry Ginosar, and Dani Oore for valuable comments on the manuscript. Thanks to everyone who shared discussion and/or encouragement online, including Aseem Agarwala, Mark Chen, Lyndie Chiou, Michael Cohen, James Landay, Nevena Lazic, Jason Salavon, Adrien Treuille, and many others.

Conflicts of Interest: The author declares no conflict of interest.

References

- Agüera y Arcas, Blaise. 2017. Art in the Age of Machine Intelligence. *Arts* 6: 18. [\[CrossRef\]](#)
- Alexander, R. McNeill. 1996. *Optima For Animals*. Princeton: Princeton University Press.
- Bernstein, Michael S., Greg Little, Robert C. Miller, Björn Hartmann, Mark S. Ackermann, David R. Karger, David Crowell, and Karina Panovich. 2015. Soylent: A Word Processor with a Crowd Inside. *Communications of the ACM* 58: 85–94. [\[CrossRef\]](#)
- Brooks, Rodney. 2017. The Seven Deadly Sins of AI Predictions. *MIT Technology Review* 120: 79–86.

- Catmull, Ed, and Amy Wallace. 2014. *Creativity, Inc: Overcoming the Unseen Forces that Stand in the Way of True Inspiration*. New York: Random House.
- Cohen, Harold. 1995. The further exploits of AARON, painter. *Stanford Humanities Review* 4: 141–58.
- Cohen, Harold. 2014. ACM SIGGRAPH Awards—Harold Cohen, Distinguished Artist Award for Lifetime Achievement. Available online: https://youtu.be/_Xbt8lzWxlIQ?t=13m20s (accessed on 9 May 2018).
- Cohen, Michael F. 2000. Imagination amplification. *IEEE Computer Graphics and Applications* 20: 54–55 [CrossRef]
- Colton, Simon. 2008. Creativity Versus the Perception of Creativity in Computational Systems. Paper presented at the AAAI Spring Symposium: Creative Intelligent Systems, Stanford, CA, USA, March 26–28.
- Colton, Simon. 2012. *The Painting Fool: Stories from Building an Automated Painter*. In *Computers and Creativity*. Berlin and Heidelberg: Springer, pp. 3–38.
- Cotter, Holland. 2008. Now in Residence: Walls of Luscious Austerity. *The New York Times*, March 24.
- Danto, Arthur. 1964. The Artworld. *The Journal of Philosophy* 61: 571–84. [CrossRef]
- De la Croix, Horst, Richard G. Tansey, and Diane Kirkpatrick. 1991. *Gardner's Art through the Ages*, 9th ed. New York: Harcourt Brace.
- Desmond, Jane. 2016. Zoos Make Money Selling Paintings Made by Animals. Are They Art? *The Washington Post*, September 8.
- Dickie, George. 1969. Defining Art. *American Philosophical Quarterly* 6: 253–56.
- Dutton, Denis. 2000. But They Don't Have Our Concept of Art. In *Theories of Art Today*. Edited by Noel Carroll. Madison: University of Wisconsin Press.
- Dutton, Denis. 2009. *The Art Instinct: Beauty, Pleasure, and Human Evolution*. New York: Bloomsbury Press.
- Dvorsky, George. 2016. Computer Science Students Fooled By Artificially Intelligent TA. *Gizmodo*, May 9.
- Dvorsky, George. 2017. This Artificially Intelligent Robot Composes and Performs Its Own Music. *Gizmodo*, June 14.
- Efros, Alexei A., and Thomas Leung. 1999. Texture synthesis by non-parametric sampling. Paper presented at the IEEE International Conference on Computer Vision, Corfu, Greece, September 21–22.
- Efros, Alexei A., and William T. Freeman. 2001. Image Quilting for Texture Synthesis and Transfer. Paper presented at the 28th Annual Conference on Computer Graphics and Interactive Techniques, Los Angeles, CA, USA, August 12–17, pp. 341–46.
- Elgammal, Ahmed, Bingchen Liu, Mohammad Elhoseiny, and Marian Mazzone. 2017. CAN: Creative Adversarial Networks Generating “Art” by Learning About Styles and Deviating from Style Norms. Paper presented at the 8th International Conference on Computational Creativity (ICCC), Atlanta, GA, USA, June 19–23.
- Epstein, Robert. 2007. From Russia, with Love. *Scientific American Mind* 18: 16–17. [CrossRef]
- Fišer, Jakub, Ondřej Jamriška, Michal Lukáč, Eli Shechtman, Paul Asente, Jingwan Lu, and Daniel Šýkora. 2016. StyLit: Illumination-guided Example-based Stylization of 3D Renderings. *ACM Transactions on Graphics (SIGGRAPH)* 35: 92:1–92:11. [CrossRef]
- Gatys, Leon A., Alexander S. Ecker, and Matthias Bethge. 2016. Image Style Transfer Using Convolutional Neural Networks. Paper presented at the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, June 27–30.
- Gaut, Berys. 2000. Art as a Cluster Concept. In *Theories of Art Today*. Edited by Noel Carroll. Madison: University of Wisconsin Press.
- Gaut, Berys. 2005. The Cluster Account of Art Defended. *The British Journal of Aesthetics* 45: 273–88. [CrossRef]
- Gershgorin, Dave. 2017. Inside the Mechanical Brain of the World's First Robot Citizen. *Quartz*, November 16.
- Goodfellow, Ian J., Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. Paper presented at the Advances in Neural Information Processing Systems, Montreal, QC, Canada, December 8–13.
- Gopnik, Alison, and Henry M. Wellman. 2012. Reconstructing constructivism: Causal models, Bayesian learning mechanisms, and the theory theory. *Psychological Bulletin* 138: 1085–108. [CrossRef] [PubMed]
- Griggs, Brandon. 2011. Why Computer Voices Are Mostly Female. *CNN*, October 21.
- Haeberli, Paul. 1990. Paint By Numbers: Abstract Image Representations. *ACM SIGGRAPH Computer Graphics* 24: 207–14. [CrossRef]
- Hertzmann, Aaron. 1998. Painterly Rendering with Curved Brush Strokes of Multiple Sizes. Paper presented at the Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '98), Orlando, FL, USA, July 19–24.

- Hertzmann, Aaron. 2001. *Algorithms for Rendering in Artistic Styles*. Ph.D. thesis, New York University, New York, NY, USA.
- Hertzmann, Aaron. 2003. A Survey of Stroke-Based Rendering. *IEEE Computer Graphics & Applications* 23: 70–81.
- Hertzmann, Aaron. 2010. Non-Photorealistic Rendering and the Science of Art. Paper presented at the Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering (NPAR '10), Annecy, France, June 7–10.
- Hertzmann, Aaron, Charles E. Jacobs, Nuria Oliver, Brian Curless, and David H. Salesin. 2001. Image Analogies. Paper presented at the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '01), Los Angeles, CA, USA, August 12–17.
- Hockney, David. 1996. *That's the Way I See It*. San Francisco: Chronicle.
- Kessler, Sarah. 2014. How Twitter Bots Fool You Into Thinking They Are Real People. *Fast Company*, June 10.
- Koblin, Aaron Michael. 2009. The Sheep Market. Paper presented at the Proceedings of the Seventh ACM Conference on Creativity and Cognition (C&C '09), Berkeley, CA, USA, October 26–30, pp. 451–52.
- Koerding, Konrad P., and Daniel M. Wolpert. 2006. Bayesian decision theory in sensorimotor control. *Trends in Cognitive Sciences* 10: 319–26. [\[CrossRef\]](#) [\[PubMed\]](#)
- Krauss, Lawrence. 2017. Why Science-Fiction Writers Couldn't Imagine the Internet. *Slate*, September 11.
- Lasseter, John. 1987. Principles of traditional animation applied to 3d computer animation. *ACM SIGGRAPH Computer Graphics* 21: 35–44. [\[CrossRef\]](#)
- Levinson, Jerrold. 1979. Defining art historically. *The British Journal of Aesthetics* 19: 232–50. [\[CrossRef\]](#)
- Litwinowicz, Peter. 1997. Processing Images and Video for an Impressionist Effect. Paper presented at the 24th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '97), Los Angeles, CA, USA, August 3–8.
- Metz, Cade. 2017. How A.I. Is Creating Building Blocks to Reshape Music and Art. *The New York Times*, August 14.
- Millward, David. 2017. Monkey Selfie Case: British Photographer Settles with Animal Charity Over Royalties Dispute. *The Telegraph*, September 12.
- Mordvintsev, Alexander, Christopher Olah, and Mike Tyka. 2015. Inceptionism: Going Deeper into Neural Networks. *Google Research Blog*, December 16.
- Nguyen, Anh Mai, Jason Yosinski, and Jeff Clune. 2016. Understanding Innovation Engines: Automated Creativity and Improved Stochastic Optimization via Deep Learning. *Evolutionary Computation* 24: 545–72. [\[CrossRef\]](#) [\[PubMed\]](#)
- Nicholas, Gabriel. 2017. These stunning A.I. Tools Are About to Change the Art World. *Slate*, December 15.
- O'Neil, Cathy. 2016. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown.
- Paik, Karen. 2007. *To Infinity and Beyond! The Story of Pixar Animation Studios*. San Francisco: Chronicle Books.
- Perez, Sarah. 2018. Microsoft's New Drawing Bot Is an AI Artist. *TechCrunch*, January 19.
- Pogue, David. 2018. Is Art Created by AI Really Art? *Scientific American*, February 1.
- Price, David A. 2008. *The Pixar Touch*. New York: Vintage.
- Reeves, Byron, and Clifford Nass. 1996. *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. Stanford: CSLI.
- Rosenblum, Naomi. 1984. *A World History of Photography*. New York: Abbeville Press.
- Rosin, Paul, and John Collomosse. 2013. *Image and Video-Based Artistic Stylisation*. Berlin and Heidelberg: Springer.
- Rubin, Michael. 2005. *Droidmaker: George Lucas and the Digital Revolution*. Gainesville: Triad.
- Scharf, Aaron. 1968. *Art and Photography*. Harmondsworth: Penguin.
- Shein, Esther. 2017. Computing the Arts. *Communications of the ACM* 60: 17–19. [\[CrossRef\]](#)
- Shelley, Mary Wollstonecraft. 1818. *Frankenstein; or, the Modern Prometheus*. London: Lackington, Hughes, Harding, Mavor & Jones.
- Simonite, Tom. 2018. When It Comes to Gorillas, Google Photos Remains Blind. *Wired*, January 13.
- Snow, Charles Percy. 1959. *The Two Cultures*. London: Cambridge University Press.
- Snowden, Don. 2012. Robert Moog: 'I Wouldn't Call This Music'—A Classic Interview to Mark a Google Doodle. *The Guardian*, May 23.
- Sterling, Bruce. 2012. An Essay on the New Aesthetic. *Wired*, April 2.
- Sternberger, Paul Spencer. 2001. *Between Amateur & Aesthete: The Legitimization of Photography as Art in America, 1880–1900*. Albuquerque: University of New Mexico Press.

- Stiglitz, Joseph E. 2012. *The Price of Inequality*. New York: Norton.
- Szegedy, Christian, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2014. Intriguing properties of neural networks. Paper presented at the International Conference on Learning Representations, Banff, AB, Canada, April 14–16.
- Tenenbaum, Joshua B., Charles Kemp, Thomas L. Griffiths, and Noah D. Goodman. 2011. How to Grow a Mind: Statistics, Structure, and Abstraction. *Science* 331: 1279–85. [[CrossRef](#)] [[PubMed](#)]
- Tyka, Mike. 2016. Exploring the Intersection of Art and Machine Intelligence. *Google Research Blog*, June 18.
- Vincent, James. 2018. Facebook's Head of AI Really Hates Sophia the Robot (and with Good Reason). *The Verge*, January 18.
- Weizenbaum, Joseph. 1966. ELIZA—A computer program for the study of natural language communication between man and machine. *Communications of the ACM* 9: 36–45. [[CrossRef](#)]
- Wilber, Michael J., Chen Fang, Hailin Jin, Aaron Hertzmann, John Collomosse, and Serge Belongie. 2017. BAM! The Behance Artistic Media Dataset for Recognition Beyond Photography. *International Conference on Computer Vision* 1: 4.
- Wilson, Mark. 2017. AI Is Inventing Languages Humans Can't Understand. Should We Stop It? *Co.Design*, July 14.
- Xu, Tao, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. 2017. AttnGAN: Fine-Grained Text to Image Generation with Attentional Generative Adversarial Networks. *arXiv*.

© 2018 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).