Computers Do Not Make Art, People Do

By Aaron Hertzmann
Published in *Communications of the ACM*, May 2020

The continually evolving relationship between artistic technologies and artists.

We live in an age of amazing new visual art created with artificial intelligence (AI) technology. The recent wave began with neural stylization apps and the trippy, evocative DeepDream. Many fine artists now work with neural network algorithms, creating high-profile works appearing in major venues.[1]

Together with these new developments comes the hype: technologists who claim that their algorithms are artists and journalists who suggest that computers are creating art on their very own. These discussions usually betray a lack of understanding about art, about AI, or both.

This column explains why today's technologies do not create art; they are tools for artists. This is not a fringe viewpoint; it reflects mainstream understanding of both art and computer science. There is a long tradition of computer-driven procedural art, and all of it is ultimately made by people, even when they use software branded as AI. It is possible this could change someday, that our software gets so good that we assign it authorship of its own works. As I will explain, I believe this is unlikely.

Generative artwork by Tom White, made with aesthetic considerations and to fool image recognition algorithms. Most current algorithms identify this image as "trombone."

The First Artistic Machine

Art has a long history of evolving in response to new technologies. In the past century, many of these technologies have led to debates and misunderstandings about the role of the artist. Tools that seemed at first to make artists irrelevant actually gave them new expressive opportunities.

The invention of photography is, arguably, the most important invention in the modern history of art, but it did not seem so positive at first. The les-sons from this story repeat periodically when new artistic technologies arise.

In 1839, Louis-Jacques-Mandé Da-guerre described the first practical photographic technique.[9] Public interest was immediate and widespread; practical applications of photography were immediately evident, and new developments and uses quickly followed in the subsequent decades. The first major impact of photography was on portraiture, where it soon became dominant, largely replacing portrait painting, silhouette cutouts, and printmaking.

The status of photography as art was more controversial. In one of the early presentations of the Daguerreotype in 1839, classical painter Paul Delaroche was quoted as saying "From today, painting is dead."

To understand why, it is helpful to Imagine the time before photography, when realistic images of the world could only be produced by skilled artists. Today we are so swamped with images that it is hard to imagine just how special and unique it must have felt to see a well-executed, realistic painting. The technical skills of realism were in-separable from other creative skills in making images. This changed when photography mechanized the task of producing images of the real world.

Some critics, like the poet Charles Baudelaire, saw photography as a corrupt and dangerous attack on true artistic genius. On the other side, photographers developed and advocated for their own art form. They argued that the artist's ability to control image creation to express their vision made it an art form in itself. Some early photographers, like Henry Peach Robinson, tried to elevate their work by mimicking the themes and composition techniques of painting. Subsequent generations explored the unique qualities of photography. By 1910, after years of work by photographers and critics, mainstream muse-ums and other institutions began to recognize photography as an art form in its own right. Today, major photographers are included in the art history canon, and photography continues to be a significant art form.

When the camera was first invented, it looked like a machine that auto-mated the creation of art. It required no skill and would destroy high-quality art. What actually happened? A new art form was invented, with its own unique qualities. Portraiture technologies became largely obsolete and portrait artists did need to learn the new technology. Image-making became more available to hobbyists; nowadays, anyone with a mobile phone can take a picture.

Unexpectedly, photography had a profoundly positive effect on painting. As cameras became cheaper, lighter, and easier to use, 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? Many artists of the era, including Whistler, Munch, and Van Gogh, wrote that true art was not about reproducing reality, because that was "just photogra-phy"[9]—true art was about something beyond realism. It seems the Modern Art movement came about because of photography. Rather than killing painting, photography spurred decades of innovation in painting.
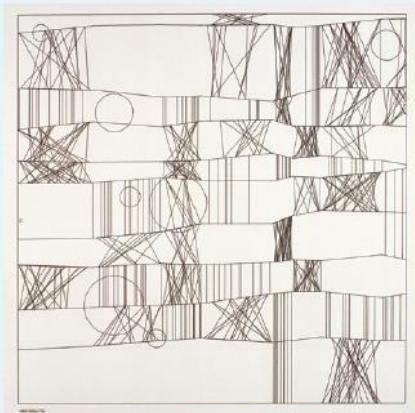
Important steps in fine art technology. (a) Oil paint technology, (b) Early artistic photography, mimicking conventional painting, (c) Rule-based, procedural computer-generated art, (d) Evolutionary computer-generated art, (e) Current neural-network GAN art.

Artworks:

(a) "The Emperor Napoleon in His Study at the Tuileries," by Jacques-Louis David, 1812.

(b) "Fading Away" by Henry Peach Robinson, 1858.

(c) "Homage á Paul Klee" by Frieder Nake, computer-generated graphic, 1965, China ink on paper.

(d) "Electric Sheep" by Scott Draves, individual "sheep" by BrothaLewis. CC BY 3.0 US. https://creativecommons.org/licenses/by/3.0/us/

(e) "AmI Dali Yet?" by Helena Sarin, 2018. (collection of Jeremy Howard).

The Artist is the Mastermind

In the Modern Art era that followed, the definition of art broadened significantly. Marcel Duchamp's "Fountain" was a landmark: he (or possibly a friend of his) found a urinal, flipped it over, signed it, and submitted it as a sculpture in 1917. Later artists, like Robert Smithson and Yves Klein, removed the requirement that the artwork be an object at all. This, and other work like it, ultimately set a precedent that an artist creates an art-work simply by naming it as such.

Throughout the arts, the artist is the mastermind behind the work, no matter how much or how little they contributed to its actual execution (see the figure images in this column). Since auteur theory, the director is credited as author of a film, despite the recognized contributions of many other artists and artisans. An architect is credited for their buildings, even when large teams of artisans, technicians, and builders all contribute. A DJ who samples and remixes sounds is the artist behind a new track. (Copyright law treats ownership differently, but that is an entirely separate topic.)

The same applies to software art, of which there is now a long tradition. The first software art was created in the 1960s, by artists including A. Michael Noll, Georg Nees, and Frieder Nake.[7] Harold Cohen's AARON soft-ware generates paintings based on a set of randomized procedural rules.[1] Karl Sims'[10] and Scott Draves'[4] evolutionary artworks involve automatic creation of images, virtual creatures, and procedural animations from user feed-back. Many artists have created lovely abstract interactive artworks that respond to the viewer's movements, including the work of Daniel Rozin, Camille Utterback, and Golan Levin. New Media arts programs typically have entire courses of study around software and interactive art.

In each of these cases, the practice of creating artistic software—or making art with software—involves iteration, experimentation, and refinement. The artist does not simply write a program and let it go. The artist writes a piece of software and then tinkers and refines the algorithms over a long period of time, continually judging and evaluating the imagery produced by the system. The final results we see exhibited come from many hundreds of hours of hard work from the artist.

In popular art, computer animation is widely accepted as an art form. But progress in computer animation frequently renews the old fears. For example, in the 1980s, Ed Catmull and Alvy Ray Smith made many trips to Disney to convince them to fund computer graphics research.[8] The animators always resisted, afraid that computers would take their jobs, and management was too conservative to take the risk. So Catmull and Smith ultimately founded Pixar instead, and created an environment where artists and technologists could develop the art form together. Now, computer-animated films employ hundreds of animators and other kinds of artists; high-quality computer animation is both dependent on human creativity and is extremely labor-intensive.

Even though software, crowdworkers, and/or artisans may have executed on the artist's instructions, the artist is the person who instigated and coordinated the work. None of these software systems is called "an artist."

Current AI-Based Artwork

All of this leads to the inevitable conclusion that AI-based artwork is still artwork made by a human. Our current "AI" software is just software,[2] despite the fancy branding, and there is a long precedent of art made with software.

Earlier computer-generated artwork, such as AARON and Sims' evolved virtual creatures, employed classic AI methods, that is, handcoded rules and numerical optimization. The same goes for procedural image stylization algorithms.

The recent neural stylization algorithms use data-fitting as one step, but they are each ultimately the result of a human writing software, and then experimenting with and improving the soft-ware algorithms, parameters, and training data until they get results they like.

We can get a window into some of today's Generative Adversarial Network (GAN) artists' experimentation via their Twitter feeds. Artists like Mario Klingemann (@quasimondo) and Helena Sarin (@glagolista) regularly tweet about their experiments using the latest image transformation software to create art. They experiment with code, parameters, and datasets, tinkering with the tools until they get great results. Assigning authorship of their art to software is perverse, dismissing the value of the artists' own hard work and creativity.

In a few cases, computer scientists have claimed their software is (possibly) the artist. In each case, they are writing the code, running an optimization, tuning the algorithm and the optimization, and selecting the favored results—just as in all previous artworks, like those listed here. Perhaps these authors claim this out of ignorance of the history of computer-generated art, or perhaps based on a desire to be provocative. If these software packages are artists, then so is the font rendering and page layout package that renders your Word or PowerPoint documents, and so is a game engine that beautifully renders procedural 3D environments.

People sometimes talk about the possibility of "collaborating" with an AI. We do collaborate with our tools in that they can automate tasks, produce unexpected results, and push us in ways that we would not have otherwise considered. But "collaboration" also often implies co-ownership and joint high-level decision-making. In this sense of the word, one does not collaborate with software any more than one collaborates with watercolor paints or Photoshop.

But the lessons from history are ultimately positive: new technology gives new tools to artists, who in turn invent new visual styles and new artistic media. The infusion of new technology into art is one of the main ways that art remains vital. New AI technology will continue to invigorate art and empower artists in the future.

Will An AI Ever Be An Artist?

The definition of art changes over time.[11] In the distant future, will we start to accept software systems as themselves artists, independent of their creators?

To think about this question, one may first ask why we create art. Evolutionary theory provides a compelling answer, saying that artmaking is the product of our evolution.[5] According to this theory, art emerged as a way for our Pleistocene ancestors to strengthen their social ties and social status. For example, art can serve as gifts, as fitness signals for mating, and as displays of status and tribal affiliation. In each role, the fundamental purpose of art is to affect peoples' relationships with each other, where the relationships are themselves important. We have many behaviors for establishing and maintaining personal and group relationships, like gifts, competition, conversation, games, and romance, and making art is one of these behaviors.

Hence, I hypothesize that art can only be created by people (or other independent actors) capable of these kinds of social relationships. In contrast, while we can get emotionally attached to our computers and other possessions, we feel no real empathy for their emotions, no ethical duty toward them, and no need to demonstrate our feelings toward them. This means computers cannot be credited as artists until they have some kind of personhood, just as people do not give

gifts to their coffeemakers or marry their cars. If there is ever such a thing as human-level AI, with thoughts, feelings, and moral status comparable to ours, then it would be able to create art. But "human-level AI" is pure science fiction right now, and we are no-where near achieving it. We do seem open to the idea that animals (such as our pets) could create art—it is just that, while they have social relationships, they lack interest in creative objects or performance.

Perhaps someday we will believe that social, shallow AIs are artists. There are many anecdotes of people being fooled into thinking that chatbots are real people, including the recent plague of Twitter bots. But, once the veil is lifted, it is clear that these chatbots do not exhibit real intelligence, and we feel cheated if we had thought they were "real." Maybe agents like a Siri or Alexa will someday be treated like junior members of the family, who answers questions, raises children, and makes artwork. This seems unlikely.

Some people say that only humans can create art because art requires intent, or, it must express something, such as an emotion. However, it would be easy to build artificial systems that do this. For example, in many artworks, the intent can be summarized by a short sentence: "depict a specific beautiful landscape," or "convey an experience of the horrors of war." It would be straightforward to build systems that generate "intents" like these, and then create artwork from them. But most people would probably agree that this system is not "an artist;" it is still a human-engineered system, and the authorship really belongs to the author and/or user of the system.

Even if we could someday develop an algorithm that autonomously pro-duces an endless stream of artworks that are original, beautiful, surprising, provocative, expressive, and culturally relevant, as long as we under-stand the software as just executing the instructions it has been given, it will continue to be a dumb machine, and not an artist.

Conclusion

I do not believe any software system in our current understanding could be called an "artist." Art is a social activity, and our "AI" software is still just software, mechanically following the instructions we give it.

Moreover, calling a software system an artist is irresponsible, because it is misleading: it could make people think that the software has human-like intelligence, autonomy, and emotions.

Art maintains its vitality through continual innovation, and technology is one of the main engines of that innovation. We are lucky to be alive at a time when artists can explore ever-more powerful tools. Today, through GitHub and Twitter, there is an extremely rapid interplay between machine learning researchers and artists; it seems like, every day, we see tinkerers and artists tweeting new creative experiments with the latest neural networks. Seeing an artist create something wonderful with new technology is thrilling, because each piece is a step the evolution of new forms of art. As artists' tools, AI soft-ware will surely transform the way we think about art in thrilling and unpredictable ways.

References

1. Bailey, J. The tools of generative art, from Flash to neural networks. Art in America (Jan 8, 2020).
2. Brooks, R. The seven deadly sins of AI predictions. Technology Review (Oct. 6, 2017).
3. Cohen. H. The further exploits of AARON, painter. Stanford Humanities Review 4, 2 (1995).
4. Draves, S. The electric sheep screen-saver: A case study in aesthetic evolution. In Proceedings of EvoWorkshops. (2005).
5. Dutton, D. The Art Instinct: Beauty, Pleasure, and Human Evolution. Bloomsbury Press, 2009.

6.   Hertzmann, A. Can computers create art? Arts 7, 2 (2018).
7.   Nake, F. Computer art: A personal recollection. In Proceedings of C&C. ACM, New York, NY, USA, (2005), 54–62; https://doi.org/10.1145/1056224.1056234
8.   Paik, K. To Infinity and Beyond! The Story of Pixar Animation Studios. Chronicle Books, 2007.
9.   Scharf, A. Art and Photography. Penguin, 1968.
10.   Sims, K. Artificial evolution for computer graphics. In Proceedings of SIGGRAPH. ACM, New York, NY, USA, (1991), 319–328. https://doi. org/10.1145/122718.122752
11.   Warburton, N. The Art Question. Routledge, 2002.

Aaron Hertzmann (hertzman@dgp.toronto.edu) is a Principal Scientist at Adobe Research in San Francisco, CA, USA, and an ACM Fellow.